

Machine Learning Methods for Cross Section Measurements

By

Krish Desai

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Benjamin Nachman, Co-chair

Professor Uros Seljak, Co-Chair

Professor Joshua Bloom

Professor Saul Perlmutter

Summer 2025

Machine Learning Methods for Cross Section Measurements

Copyright 2025
By
Krish Desai

Abstract

Machine Learning Methods for Cross Section Measurements

By

Krish Desai

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Benjamin Nachman, Co-chair

Professor Uros Seljak, Co-chair

Precise differential cross section measurements are indispensable for tests of Standard Model predictions at the energy frontier and for searches for new physics, yet their extraction from collider data is an ill posed inverse problem. *Unfolding*, also known as deconvolution, is the process of removing detector distortions to reconstruct particle level truth from detector level data. Conventional, histogram based, binned unfolding techniques introduce artifacts, impose arbitrary bin edges, and become

computationally prohibitive in high dimensional phase spaces, potentially obscuring underlying physics.

This dissertation develops a unified framework that leverages modern machine learning techniques to surmount these limitations. First, the NEURAL POSTERIOR UNFOLDING (NPU) method demonstrates how conditional normalising flows can serve as differentiable surrogates of detector response, enabling likelihood based unfolding through implicit regularisation. Building on this foundation, the MOMENT UNFOLDING algorithm directly extracts distribution moments without binning, providing precise experimental predictions for effective field theories and phenomenological models. The framework is further advanced by development of REWEIGHTING ADVERSARIAL NETWORKS (RANs), which perform full spectral unfolding using adversarial training to implement particle level reweighting guided by a detector level classifier, offering theoretical and computational advantages over extant methods.

A critical statistical analysis of event correlations in unfolded data reveals systematic misestimation of uncertainties when these correlations are ignored, leading to methodological recommendations that ensure correct coverage for all derived observables. The methods presented in this dissertation are validated using both idealised Gaussian distributions and proton–proton collision simulations of the CMS experiment as realistic particle physics examples, specifically simulations of Z+jets events, demonstrating significant improvements in precision, accuracy, and computational efficiency, reducing computational time by orders of magnitude while maintaining

or exceeding the precision of existing methods for the unbiased recovery of spectral features.

By marrying statistical rigour with powerful machine learning methods, this work establishes a scalable blueprint for precision measurements at current and future high energy physics experiments. The resulting open source software enables more reliable extraction of fundamental physics parameters from complex detector data, advancing the ability to test theoretical models and potentially discover new phenomena.

Publications Resulting from This Dissertation

- [1] K. Desai, O. Long, and B. Nachman. “Unbinned Inference with Correlated Events”. In: *E-print* (Apr. 2025). arXiv: 2504.14072 [physics.data-an].
- [2] F. T. Acosta, K. Desai, V. Mikuni, B. Nachman, and J. Pan. “Multidimensional Deconvolution with Profiling”. In: *NeurIPS. ML4PS*. 177. 2024.
- [3] K. Desai, B. Nachman, and J. Thaler. “Moment extraction using an unfolding protocol without binning”. In: *Physical Review D* 110.11 (2024), p. 116013.
- [4] H. Zhu, K. Desai, M. Kuusela, V. Mikuni, B. Nachman, and L. Wasserman. “Multidimensional Deconvolution with Profiling”. In: *NeurIPS. ML4PS*. 150. 2024. DOI: <https://doi.org/10.48550/arXiv.2409.10421>. arXiv: 0902.0885 [hep-ph].
- [5] K. Desai, B. Nachman, and J. Thaler. “Deconvolving Detector Effects for Distribution Moments”. In: *NeurIPS. ML4PS*. 43. 2022.
- [6] K. Desai, B. Nachman, and J. Thaler. “SymmetryGAN: Symmetry discovery with deep learning”. In: *Phys. Rev. D* 105 (9 May 2022), p. 096031. DOI:

10.1103/PhysRevD.105.096031. URL: <https://link.aps.org/doi/10.1103/PhysRevD.105.096031>.

- [7] K. Desai, B. Nachman, and J. Thaler. “Symmetry Discovery with Deep Learning”. In: *NeurIPS. ML4PS*. 117. 2021.

To Ben Nachman.

Contents

Contents	iv
List of Figures	viii
List of Tables	xii
I Introduction and physics background.	1
I.A Units and conventions.	3
I.B The Standard Model: Theoretical framework.	18
I.C Fundamental role of cross section measurements in particle physics.	25
I.D Cross section measurements: From theory to experiment.	27
I.E Detector response in precision measurements.	37
I.F Challenges at modern experiments.	41
I.G Dissertation scope and physics impact.	43
II Theoretical foundations.	46
II.A Statistical formulation of the unfolding problem.	48
II.B Forward and inverse problems in HEP.	57

II.C	Historical development: From matrix inversion to modern approaches.	61
II.D	Traditional unfolding methods in experimental analyses.	63
II.E	Regularisation: Need, approaches, and limitations.	71
II.F	Unbinned methods: statistical considerations.	75
II.G	Evaluation metrics for unfolding.	83
III	Machine learning for unfolding.	100
III.A	The emergence of machine learning in particle physics.	103
III.B	Introduction to neural networks.	111
III.C	Supervised learning approaches for unfolding.	128
III.D	Deep Learning Architectures	137
III.E	Modern machine learning frameworks.	147
III.F	Comparative analysis of ML based unfolding methods.	183
III.G	Case studies: ML based unfolding in HEP analyses.	190
III.H	ML upstream and downstream of unfolding.	201
IV	NEURAL POSTERIOR UNFOLDING	207
IV.A	Motivation for improved binned unfolding	209
IV.B	Normalising flows for inverse problems.	213
IV.C	The NEURAL POSTERIOR UNFOLDING Algorithm	217
IV.D	Numerical results.	224
IV.E	Beyond binning: the path forward.	239
V	MOMENT UNFOLDING.	241
V.A	Why moments: physics context and QCD calculations.	243

V.B	A GAN like method to unfold moments	253
V.C	Machine learning implementation.	269
V.D	Case studies.	278
V.E	Comparison with alternative methods.	294
V.F	Conclusion.	298
VI	Reweighting Adversarial Networks.	301
VI.A	The need for full spectral measurements.	303
VI.B	From moments to complete differential cross section spectra.	327
VI.C	Methodology and regularisation.	345
VI.D	Machine learning implementation.	362
VI.E	Results	379
VII	Unbinned inference on correlated data.	393
VII.A	Introduction.	395
VII.B	Statistical independence in HEP.	396
VII.C	Violation of statistical independence.	401
VII.D	Consequences for inference.	412
VII.E	Formalism.	416
VII.F	Uncertainty quantification.	426
VII.G	Case studies.	449
VII.H	Conclusions and outlook.	482
VIII	Symmetries in data: connections to unfolding challenges.	487
VIII.A	Symmetries and unfolding.	490

VIII.B	Formalism and importance.	500
VIII.C	Statistical definition of dataset symmetries.	529
VIII.D	SYMMETRYGAN: Discovering symmetries with adversarial learning.	538
VIII.E	Empirical experiments.	550
VIII.F	Towards symmetry inference.	569
VIII.G	Symmetry informed unfolding	571
VIII.H	Improved precision using symmetries	580
VIII.I	Conclusion.	585
IX	Conclusion.	590
IX.A	A synthesis of unfolding methods.	592
IX.B	A unified framework for modern unfolding.	595
IX.C	Advances in cross section measurement.	603
IX.D	Applications to current and future experiments.	610
IX.E	Open challenges and future directions.	616
IX.F	Closing: The broader vision	624
	References	628
A	Glossary	778
B	Theorems and proofs.	791

List of Figures

I.1	η vs. y for various m/p_T	7
I.2	Taxonomy of datasets used	13
I.3	A schematic illustration of the Standard Model [42].	20
II.1	Coverage properties of Tikhonov regularisation and IBU	92
II.2	Weight correlation between event pairs as a function of distance between events.	96
III.1	Neural network depth and width.	113
III.2	Shallow vs. deep NN	114
III.3	Forward pass.	116
IV.1	NPU: 2 bin example	227
IV.2	NPU: Gaussian example setup	228
IV.3	NPU: Gaussian unfolding results and posterior analysis	230
IV.4	Statistical calibration analysis for NPU and FBU	231
IV.4	Jet substructure unfolding results using NPU, FBU, and IBU	233
IV.5	Corner plots for jet substructure observables	237

V.1	Schematic diagram of MOMENT UNFOLDING	257
V.2	Effect of detector resolution on Truth and Generation datasets	280
V.3	MOMENT UNFOLDING: Particle level histograms of Gaussian examples of Truth, Generation, and the prediction from MOMENT UNFOLDING, and the MLC loss landscape.	281
V.4	Jet substructure distributions comparing truth, generation, and MOMENT UNFOLDING results	286
V.5	Loss landscapes showing MOMENT UNFOLDING parameter estimation accuracy for unfolding jet substructure variables	287
V.6	Momentum dependent unfolding of inclusive jet substructure distributions	290
V.7	Momentum dependent differential moments of jet substructure observables	291
V.8	Comparison of unfolding methods for momentum dependent jet moments	294
VI.1	Schematic description of the training setup of a RAN.	351
VI.2	Comparison of activation functions to parametrise a RAN's generator. .	359
VI.3	RAN architecture with generator and critic networks	365
VI.4	Unfolding performance comparison between RAN and OmniFold under increasing detector smearing	383
VI.5	Unfolded particle level distributions for six jet substructure observables, comparing RAN, OMNIFOLD, and IBU.	385
VII.1	Detector and particle level distributions for the study on correlations during inference.	452
VII.2	Detector response effects for smearings 0.00, 0.25, 0.50 and 0.75	453

VII.3	Comparison of uncertainty estimates from asymptotics and bootstraps for binned unfolded distribution parameters as a function of detector resolution.	456
VII.4	Validation of unbiased parameter extraction from unfolded distributions across varying detector resolutions.	457
VII.5	Event wise weight correlations as a function of separation distance for different unfolding methods.	461
VII.6	Covariance structure of binned unfolded distributions for KDE and NN-based methods.	465
VII.7	Comparison between binned and unbinned parameter inference with unbinned and binned unfolding methods as a function of detector resolution.	469
VII.8	Validation of asymptotic uncertainty estimates across dimensionalities in unbinned unfolding.	479
VII.9	Impact of detector resolution on asymptotic uncertainty reliability in 6 dimensional unfolding.	481
VIII.1	Architectural comparison of standard GAN and SymmetryGAN for automated symmetry discovery.	542
VIII.2	SYMMETRYGAN: Analytic loss landscape for $\mathcal{N}((0.5, 1))$	552
VIII.3	Empirical validation of SYMMETRYGAN on 1 dimensional Gaussian data showing convergence to discrete symmetries based on initial conditions.	553
VIII.4	Loss landscapes demonstrating SYMMETRYGAN's ability to discover continuous $SO(2)$ and discrete V_4 symmetries in isotropic and anisotropic 2D Gaussians.	555

VIII.5	Controlled discretization of continuous $SO(2)$ symmetry into cyclic subgroups $\mathbb{Z}_q$ via constraint enforcement.	556
VIII.6	Orthogonal slices of $\text{Aff}(2)$ symmetries for isotropic Gaussian.	557
VIII.7	Discrete symmetry structure of anisotropic Gaussian in $\text{Aff}(2)$ discovered by a SYMMETRYGAN.	558
VIII.8	SYMMETRYGAN's predictions for a bimodal Gaussian mixture with $\mathbb{Z}_2$ reflection symmetry and corresponding loss landscape validation.	559
VIII.9	Discovery of dihedral symmetries D_8 and D_4 in octagonal and square lattice Gaussian mixtures.	560
VIII.10	Symmetries discovered by a SYMMETRIGAN in transverse momentum data of dijet events in a $SO(2) \times SO(2)$ search space, and the symmetry discovery map for the training dynamics.	562
VIII.11	Visual validation of discovered dijet symmetry through transverse momentum preservation.	563
VIII.12	Non-uniform azimuthal distributions from random $SO(4)$ rotations demonstrating symmetry violation, and uniform azimuthal distributions from discovered symmetry confirming invariance of dijet system.	564
VIII.13	Kullback-Leibler divergence analysis showing the transformations identified by the SYMMETRYGAN approach statistical noise floor.	566
VIII.14	Loss values computed by training <i>post hoc</i> discriminators to validate the transformations discovered by the SYMMETRYGAN.	567

List of Tables

I.1	Unit conversions in natural units	4
I.2	Standard notation for common sets and groups	7
II.1	Comparing the MC dependence and uncertainty propagation of tradi- tional unfolding methods.	70
II.2	Distributional distance metrics for unfolding evaluation	88
IV.1	Computational cost comparison between FBU and NPU	231
V.1	MOMENT UNFOLDING: Summary of hyperparameters.	274
V.2	Moments of jet observables at particle level for Truth, Generation, and MOMENT UNFOLDING.	284
VI.1	Comparison of moment-only, binned spectrum, and unbinned spectrum unfolding approaches.	308
VI.2	Comparison of unfolding method characteristics	320
VI.3	Wasserstein distances between unfolded and true distributions	387
VI.4	VLC divergence between unfolded and true distributions	388

VII.1	Sources of correlated uncertainties in cross section measurements. . . .	414
VII.2	Gaussian model parameters for 2D unfolding validation studies. . . .	476
VII.3	Gaussian model parameters for 4D unfolding validation studies. . . .	477
VII.4	Gaussian model parameters for 6D unfolding validation studies. . . .	478
VIII.1	Relationship between unfolding and symmetry discovery in particle physics.	495
VIII.2	Fundamental symmetries and their manifestations in HEP observables.	507
VIII.3	Symmetry breaking induced by detector effects in HEP experiments. . .	519
IX.1	Applications of presented methods to experimental scenarios.	611

Acknowledgments

Like any journey worth taking, a PhD is never traveled alone. This dissertation represents not only years of research, but also the culmination of support, mentorship, and kindness from so many people who have shaped my path as a physicist. My adventure in physics began with an incidental chat with Vipul Mehra over lunch in my first year of high school. Thank you, Vipul, for pushing me harder than I thought I could go, and believing in me when I didn't believe in myself. At Yale, I was fortunate to learn from some of the greatest minds in the field. I want to thank Witold Skiba, for teaching me both how to do physics and how to do academia. I still find myself wishing I was in your classroom furiously keeping up with you before you pause and proclaim "This is not a fit." I also want to thank Laura Newburgh, David Poland, Tom Appelquist and Vincent Moncrief for guiding me patiently through my first research experiences.

Since arriving at Berkeley and LBNL, I have been surrounded by brilliant and generous colleagues. This PhD would not have been possible without the love and support of the Nachman ML Group. To the postdocs, Aishik Ghosh, Dennis Noll, Vinicius Mikuni, Sascha Diefenbachar, Mariel Pettee, Fernando Torales Acosta and Tanvi Wamorkar, you made the hardest days easy. To my co-grad students, Radha Mastandrea and Jordan Dudley, you have always been a source of inspiration. Over the course of this PhD, I have had the privilege of working with some of the most brilliant, generous and thoughtful scientists in the field, and I apologise to all of you who have been part of this journey who I have not thanked here explicitly. To all of my

collaborators—Jesse Thaler, Fernando Torales Acosta, Vinicius Mikuni, Jingjing Pan, Richard Zhu, Mikael Kuusela, Larry Wasserman, and Owen Long—thank you for all the long conversations, your incredibly patience with my strangely timed messages, and your willingness to walk me through your thought processes. I am grateful to my dissertation and qualifying exam committee members, Joshua Bloom, Saul Perlmutter and Hitoshi Murayama, for their guidance in charting this journey and sharing their wisdom and unique perspectives that were instrumental in making it a success. In particular, I would like to thank Uros Seljak for being an invaluable co-advisor.

Words feel inadequate to express my gratitude to Benjamin Nachman. As his first graduate student, I’ve had the privilege of not just an advisor, but a true mentor and friend. Ben has believed in me when I struggled to believe in myself, celebrated my successes as if they were his own, and somehow always known exactly when I needed encouragement or a gentle push forward. His unwavering support, boundless enthusiasm for physics and genuine care have shaped not just this dissertation but who I am as a scientist and person. I have often told new PhD students that should find themselves and advisor who makes them feel like they are the best version of themselves. Ben has been all that and more; I am profoundly lucky to have had him as my advisor.

My sincere appreciation also goes to the administrative and technical staff of the Physics Department at the University of California, Berkeley, and Lawrence Berkeley National Laboratory for their invaluable assistance. Additionally, I acknowledge the support from funding agencies and computational resources provided by the university and national lab. Large language models (Claude 3.5 [92] and 4.0 [19]) were

occasionally consulted during the preparation of this dissertation for brainstorming structural elements, refining transitions between sections and surveying relevant literature. However, the presented ideas, analyses and conclusions are entirely my own. I attest that I stand by all the facts and references presented in this dissertation. I am expecially grateful to the National Energy Research Scientific Computing Center (NERSC) for providing the computational resources that made this work possible. This adventure would not have been possible without the Berkeley physics community, that made a home for me far away from home.

Finally, to my parents, Bijal and Nilesh, and my brother, Dweej, there aren't enough words in this entire disseration to express how grateful I am. And to the friends who have loved me like family, Tyler, Charlie, Kevin, Benjamin, Lenny, Simon, Sam, Narain, Fran, Levi, David, Jeremy, Anna, Gabe and so many more, here's to many more smiles and laughs to come.

This disseration is not the end of a road, it merely marks the beginning of the next adventure. I am grateful to everyone who was part of this journey for helping me reach this milestone, and to everyone who will be part of the next.

Chapter I

Introduction and physics background.

I.A	Units and conventions.	3
I.A.1	Coordinate systems and relativistic geometry.	5
I.A.2	Some common sets and groups.	6
I.A.3	Statistical frameworks and uncertainty quantification.	9
I.A.4	Machine learning architectures and notation.	10
I.A.5	Information theory and optimal observables.	11
I.A.6	Datasets.	12
I.A.7	A note on conventions and clarity.	17
I.B	The Standard Model: Theoretical framework.	18
I.B.1	Fundamental particles and forces.	19
I.B.2	Theoretical framework and symmetries.	22
I.B.3	The Higgs mechanism and mass generation.	23
I.B.4	Limitations and beyond the Standard Model (BSM) physics.	23

<i>Chapter I. Introduction and physics background.</i>	2

I.C	Fundamental role of cross section measurements in particle physics.	25
I.D	Cross section measurements: From theory to experiment.	27
I.D.1	Theory.	28
I.D.2	Experimental measurement.	30
I.D.3	Applications in particle physics.	31
I.E	Detector response in precision measurements.	37
I.F	Challenges at modern experiments.	41
I.G	Dissertation scope and physics impact.	43

I.A Units and conventions.

I.A.1	Coordinate systems and relativistic geometry.	5
I.A.2	Some common sets and groups.	6
I.A.3	Statistical frameworks and uncertainty quantification.	9
I.A.4	Machine learning architectures and notation.	10
I.A.5	Information theory and optimal observables.	11
I.A.6	Datasets.	12
I.A.6.i	Nature versus Monte Carlo.	12
I.A.6.ii	Particle level versus detector level.	14
I.A.6.iii	Four dataset categories.	14
I.A.6.iv	The detector response.	15
I.A.6.v	Operational framework.	16
I.A.7	A note on conventions and clarity.	17

Table I.1: Combined unit conversions in natural units, showing both abbreviated and full natural unit expressions, their SI equivalents and their physical significance.

Quantity	Abbrev.	Full unit	SI (approx.)	Comment
Speed	1	c	$3 \times 10^8 \text{ m s}^{-1}$	Speed of light
Action	1	$\hbar$	10^{-34} J s	Quantum of action
Energy	GeV	GeV	$1.6 \times 10^{-10} \text{ J}$	Binding energy
Momentum	GeV	GeV/c	$5.3 \times 10^{-19} \text{ N s}$	Typical HEP particle momentum
Mass	GeV	GeV/c^2	$1.8 \times 10^{-27} \text{ kg}$	Proton mass
Length	GeV^{-1}	$\hbar c/\text{GeV}$	$2 \times 10^{-16} \text{ m}$	Compton wavelength
Time	GeV^{-1}	$\hbar/\text{GeV}$	$6.6 \times 10^{-25} \text{ s}$	
Charge	1	$e/\sqrt{4\pi\alpha}$	$5.3 \times 10^{-19} \text{ C}$	$e = 1.6 \times 10^{-19} \text{ C} \approx 0.3$
Cross section	GeV^{-2}	$\hbar^2 c^2/\text{GeV}^2$	$4 \times 10^{-32} \text{ m}^{-2}$	barn= $2.6 \times 10^{-9} \text{ GeV}^{-2}$
B field	GeV^2	$\text{GeV}^2/\hbar c^2$	$5 \times 10^{16} \text{ T}$	

Even though nature does not establish preferred units, the physics that describes it requires establishing a common language and framework. In high energy physics, the choice of units reflects not simply an arbitrary choice, but a philosophical stance about the nature of reality. While everyday experience suggests that meters, kilograms, and seconds might be the most suitable units to express the physical world, particle physics reveals that speed, action, and energy form more natural rulers for measuring the universe at its smallest scales.

Definition 1 (Natural units). *Natural Units* are a system of units where $\hbar = c = 1$, effectively setting the speed of light and quantum of action as fundamental measuring sticks. This choice transforms all physical quantities into powers of energy.

Some common physical quantities, their abbreviated natural units, their full natural units, their approximate SI equivalents and their physical significance are described in Table I.1.

I.A.1 Coordinate systems and relativistic geometry.

At a high energy physics (HEP) experiment, the choice of coordinates must respect the underlying symmetries of the physics and the experimental setup. The laboratory frame typically provides a natural Cartesian system. For example, for a circular collider experiment, one can set the z -axis to be along the beam direction (longitudinal), the x -axis to be horizontal, pointing toward the centre of the accelerator ring and the y -axis to be vertical, completing the right handed system.

However, the partons that participate in an interaction can carry unknown fractions of the beam momentum, making the centre of mass frame of each interaction unknowable. This uncertainty motivates a coordinate system that transforms simply under longitudinal Lorentz boosts.

Definition 2 (Pseudorapidity). The *pseudorapidity* η , is defined as

$$\eta = -\ln \tan \frac{\theta}{2} \quad (\text{I.1})$$

where θ is the polar angle from the beam axis [252]. Under a longitudinal boost with rapidity $\beta = \text{artanh } v/c$, a massless particle's η transforms as $\eta \mapsto \eta + \beta$

Expressing θ in terms of momentum,

$$\eta = -\ln \tan \frac{\theta}{2} = \frac{1}{2} \ln \frac{|p| + p_z}{|p| - p_z} = \text{artanh} \frac{p_z}{|p|} \quad (\text{I.2})$$

Figure I.1 shows the relationship between pseudorapidity η and true rapidity y for particles with different mass to p_T ratios. As expected, the massless case ($m/p_T = 0$) lies exactly on the diagonal $y = \eta$, while non-zero m/p_T introduces a systematic deviation that grows at large $|\eta|$. Even a modest ratio ($m/p_T = 0.1$) produces a measurable shift, and by $m/p_T = 2.0$ the rapidity is significantly reduced relative to η . This behaviour must be accounted for when inferring kinematic distributions of heavy particles from detector measurements expressed in pseudorapidity.

For massless particles, differences in η remain invariant under longitudinal boosts, motivating the definition of the angular distance as follows.

Definition 3 (Angular distance). The *angular distance metric* is defined as

$$\Delta R^2 = \Delta\eta^2 + \Delta\varphi^2, \quad (\text{I.3})$$

where φ is the azimuthal angle around the beam axis. This metric approximates the geometric angle between particles in the detector while remaining approximately longitudinally boost invariant.

I.A.2 Some common sets and groups.

The standard notation adopted for frequently occurring sets and groups is listed in Table I.2.




Figure I.1: Comparison of pseudorapidity η and true rapidity $y = \frac{1}{2} \ln \frac{E+p_z}{E-p_z}$ for several mass to transverse momentum ratios. Curves correspond to $m/p_T = 0.0, 0.1, 0.5, 1.0, 2.0$, illustrating how finite mass distorts away from the massless limit $y = \eta$ at large $|\eta|$.

Table I.2: Standard notation for common sets and groups

Notation	Description
<i>Number sets and spaces</i>	

Table I.2 (continued)

Notation	Description
$\mathbb{N}$	Natural numbers (1, 2, 3...)
$\mathbb{N}_0$	Non-negative integers (0, 1, 2, 3...)
$\mathbb{Z}$	Integers
$\mathbb{Q}$	Rational numbers
$\mathbb{R}$	Real numbers
$\mathbb{R}^\times$	Non-zero real numbers
$\mathbb{R}^+$	Positive real numbers
$\mathbb{C}$	Complex numbers
$\mathbb{F}_q$	Finite field with q elements
$\mathbb{S}^n$	n -sphere
<i>Discrete groups</i>	
$\mathbb{Z}_n$	The cyclic group of order $n^{[\alpha]}$
D_n	The dihedral group of order $2n^{[\beta]}$
S_n	Symmetric group on n elements
A_n	Alternating group on n elements
<i>Lie groups</i>	
$\text{Aff}(n)$	Group of affine maps on $\mathbb{R}^{n[\gamma]}$
$\text{GL}(n)$	General linear group on $\mathbb{R}^n$
$\text{SL}(n)$	Special linear group on $\mathbb{R}^{n[\delta]}$

Table I.2 (continued)

Notation	Description
$O(n)$	Orthogonal group on $\mathbb{R}^{n^{[\epsilon]}}$
$SO(n)$	Special orthogonal group on $\mathbb{R}^n$
$U(n)$	Unitary group on $\mathbb{C}^n$
$SU(n)$	Special unitary group on $\mathbb{C}^n$
$Sp(2n)$	Symplectic group on $\mathbb{R}^{2n}$

^[α] May be referred to in other sources as C_n and $\mathbb{Z}/n\mathbb{Z}$.

^[β] Referred to as D_{2n} in other sources. The group of symmetries of a regular n -gon.

^[γ] This and other Lie groups may specify a different base field as $\text{Aff}(n, F)$ if needed.

^[δ] For Lie groups “Special” implies $\det = 1$.

^[ε] $O(p, q)$ denotes the orthogonal group on $\mathbb{R}^{p,q}$ if signature is relevant.

While these conventions are widely used in the mathematical literature, they are listed here to avoid ambiguity and provide a convenient reference.

I.A.3 Statistical frameworks and uncertainty quantification.

Note. Some basic notation:

- $\mathbb{P}(\cdot)$ denotes the probability of an event.

- $\mathbb{E}_p[\cdot]$ denotes the expected value of a random variable under probability density function p . The probability density function may be omitted if it is clear from context.

In modern particle physics analyses, every measurement emerges from millions of interaction events, each carrying both statistical and systematic uncertainties.

Definition 4 (Poisson distribution.). The *Poisson distribution* is the fundamental distribution governing counting experiments. For n observed events with expected value λ ,

$$\mathbb{P}(n \mid \lambda) = \frac{\lambda^n e^{-\lambda}}{n!} . \quad (\text{I.4})$$

In the high statistics limit, ($n \rightarrow \infty$), the Poisson distribution approaches a Gaussian.

Definition 5 (Profile likelihood ratio). The *profile likelihood ratio*

$$\mathcal{L}(\mu) = \frac{L(\mu, \hat{\theta}_\mu)}{L(\hat{\mu}, \hat{\theta})} \quad (\text{I.5})$$

where $\hat{\theta}_\mu$ maximises L for fixed signal strength μ is the optimal test statistics for hypothesis testing under the Neyman–Pearson lemma.

I.A.4 Machine learning architectures and notation.

Neural networks can be specified by their architecture vectors, $[d_0, d_1, \dots, d_\ell]$, which denotes a network with input dimension d_0 , hidden layers of dimensions d_1 through $d_{\ell-1}$, and output dimension d_ℓ .

For a network $f : \mathbb{R}^{d_0} \rightarrow \mathbb{R}^{d_\ell}$ with parameters $\theta = \{W_i, b_i\}$, the forward pass computes

$$h_0 = x \tag{I.6}$$

$$h_i = \sigma(W_i h_{i-1} + b_i) \quad i = 1, \dots, \ell - 1 \tag{I.7}$$

$$f(x; \theta) = W_\ell h_{\ell-1} + b_\ell \tag{I.8}$$

where σ denotes the activation function. The universality theorem guarantees that sufficiently wide networks can approximate any continuous function.

Training proceeds via gradient descent on a loss function $L(\theta)$, with the gradient computed through automatic differentiation.

I.A.5 Information theory and optimal observables.

Information theory provides a rigorous framework through the Neyman-Pearson lemma, which implies that the likelihood ratio $\frac{L(x|S)}{L(x|B)}$ provides the most powerful test for distinguishing signal S from background B at any given significance level.

In practice, this optimal observable can be approximated using machine learning classifiers. A well trained classifier computes

$$f(x) \approx \mathbb{P}(S|x) = \frac{L(x | S) \mathbb{P}(S)}{L(x | S) \mathbb{P}(S) + L(x | B) \mathbb{P}(B)} \tag{I.9}$$

The mutual information $I(Y; f(X))$ between the true labels Y and classifier output $f(X)$ quantifies the information captured

$$I(Y; f(X)) = \iint p(y, f) \ln \frac{p(y, f)}{p(y)p(f)} dy df. \tag{I.10}$$

This connects directly to the area under the ROC curve and provides a model-independent measure of classification performance.

I.A.6 Datasets.

Meaningful distinctions in particle physics data arise from the transformation between the underlying physical processes one seeks to understand and the experimental observations one can actually record, as well as from the distinction between Monte Carlo simulation and measurement in nature. Datasets can therefore be classified according to two principal axes that capture the essential distinctions in particle physics data.

The first axis distinguishes between events that occur in nature versus those generated through Monte Carlo simulation. The second axis separates particle level quantities interest from detector level ones.

This classification yields four distinct dataset categories, each serving a unique role in the measurement process. Understanding the relationships between these categories plays an important role in the development of the unfolding methods discussed in this dissertation.

Figure I.2 illustrates this taxonomy graphically.

I.A.6.i Nature versus Monte Carlo.

The vertical dimension of Figure I.2 distinguishes between natural phenomena and their simulated counterparts. Natural events represent the actual physical processes occurring in particle interactions, while Monte Carlo events are computational repro-




Figure I.2: The four categories of data used in this dissertation organised along two orthogonal dimensions. The vertical axis distinguishes between natural phenomena (Nature) and computational models (MC), while the horizontal axis separates particle level quantities from detector level observations. Arrows indicate the detector response that maps particle level events to their detector level manifestations.

ductions based on theoretical understanding. This distinction is fundamental because although natural events contain true physics, Monte Carlo simulations provide the controlled environment necessary to understand the detector response and develop analysis techniques.

Monte Carlo simulations incorporate theoretical knowledge, including matrix element calculations, parton shower models, and hadronisation prescriptions. While these simulations can achieve remarkable fidelity to natural processes, they remain approximations subject to theoretical uncertainties and modelling assumptions.

I.A.6.ii Particle level versus detector level.

The horizontal dimension captures the transformation from particle level information to detector level signals. Particle level quantities represent the properties of particles immediately after the hard scattering process and subsequent decay chains, before interaction with detector. These quantities form the natural language of theoretical predictions and are directly calculable from field theory.

Detector level quantities, in contrast, represent the actual experimental observables after particles have traversed the detector apparatus. This transformation incorporates numerous physical effects. The detector acts as a complex transfer function that irreversibly transforms the particle level distributions into the measured quantities.

I.A.6.iii Four dataset categories.

The intersection of these two dimensions yields four fundamental dataset categories.

Truth $\sim p(z)$ represents nature at the particle level—the actual physical distributions HEP experiments ultimately seek to measure. These distributions are never directly observable but represent the ground truth that all analyses attempt to approach. In practice, truth distributions are estimated through unfolding procedures from the information contained in the other three distributions.

Data $\sim p(x)$ comprises natural events after detector effects. It is the raw experimental output of all measurements. Data embodies the convolution of the physics of interest with instrumental effects, making it the starting point for any

experimental analysis. The challenge lies in disentangling the physics content from the detector effects.

Generation $\sim q(z)$ comprises Monte Carlo events at the particle level, providing a controlled means to understanding theoretical predictions. Generation allows the study of particle level distributions with perfect knowledge and unlimited statistics, subject to theoretical modelling uncertainties. Generated events serve as proxies for truth in developing and validating analysis methods.

Simulation $\sim q(x)$ represents Monte Carlo events processed through detailed detector modelling, creating artificial detector level datasets with known particle level origins. Simulation enables the study the detector response function by providing matched pairs of particle level and detector level quantities. The fidelity of detector simulation directly impacts the ability to correct for instrumental effects.

I.A.6.iv The detector response.

The arrows in Figure I.2 represent the action of the detector response function, mapping particle level distributions to their detector level counterparts. This mapping is inherently stochastic, incorporating both resolution effects that blur distributions and acceptance effects that create blind regions in phase space.

For natural events, this transfer function represents the actual physical processes of particle interaction with detector material—ionisation in tracking chambers, showering in calorimeters, and signal processing in readout electronics. For simulated events, the detector simulation models these processes through programs like GEANT4 [461, 621,

655] that track particles through detailed detector geometries, applying the current best understanding of particle matter interactions.

I.A.6.v Operational framework.

In practice, these four dataset categories work in concert to enable precision measurements. Paired Generation–Simulation samples provide the response matrix linking particle and detector levels, while control regions in Data validate the simulation fidelity. Closure tests using independent Monte Carlo samples verify the unfolding procedure, and the final measurement applies these validated corrections to data to infer the underlying Truth.

This framework implicitly assumes that the detector response is sufficiently similar between data and simulation, an assumption that must be continuously validated through auxiliary measurements and control samples. The robustness of any measurement depends on understanding the limitations of this assumption and properly propagating the associated systematic uncertainties.

The subsequent chapters of this dissertation will repeatedly reference these four dataset categories in the development of sophisticated methods for extracting particle level information from detector level observations. This terminology provides the consistent language necessary for discussing the various approaches to the unfolding problem and their relative merits in different physics contexts.

I.A.7 A note on conventions and clarity.

This work strives to maintain a balance between mathematical rigour and physical insight. Where conventions differ between communities,¹ attempts are made to note both conventions.

¹For example, particle physicists and machine learning researchers.

I.B The Standard Model: Theoretical framework.

I.B.1	Fundamental particles and forces.	19
I.B.1.i	Fermions: The building blocks of matter.	20
I.B.1.i.a	Quarks.	20
I.B.1.i.b	Leptons.	21
I.B.1.ii	Bosons: Force carriers.	21
I.B.2	Theoretical framework and symmetries.	22
I.B.3	The Higgs mechanism and mass generation.	23
I.B.4	Limitations and beyond the Standard Model (BSM) physics.	23

The Standard Model of particle physics represents one of the most significant intellectual achievements in modern science. Developed throughout the latter half of the 20th century, it provides a quantum field theory framework that describes three of the four known fundamental forces—the electromagnetic, weak, and strong interactions—in addition to classifying all known elementary particles. The mathematical formulation of the Standard Model is based on gauge theory, specifically quantum chromodynamics (QCD) and the electroweak theory, underpinned by the gauge symmetry group $SU(3)_C \times SU(2)_L \times U(1)_Y$.

The predictive power of the Standard Model has been repeatedly validated through precision experiments across multiple energy scales, from low energy nuclear phenomena to the highest energy particle collisions achievable at modern accelerators. Its crowning achievement came with the discovery of the Higgs boson in 2012 at the Large Hadron Collider (LHC) [543, 550], confirming the mechanism through which elementary particles acquire mass. A schematic illustration of the standard model can be found in Figure I.3.

I.B.1 Fundamental particles and forces.

The Standard Model categorises elementary particles into two main families: fermions, which comprise matter, and bosons, which mediate forces between matter particles.




Figure I.3: A schematic illustration of the Standard Model [42].

I.B.1.i Fermions: The building blocks of matter.

Fermions, characterized by half-integer spin, obey the Pauli exclusion principle and satisfy Fermi–Dirac statistics. Fermions are further classified into quarks and leptons, each arranged in three generations of increasing mass.

I.B.1.i.a Quarks.

Quarks are spin = $\frac{1}{2}$ particles that are categorised into three generations as follows:

- Up (u) and down (d),

- Charm (c) and strange (s),
- Top (t) and bottom (b).

Quarks carry fractional electric charge and colour charge, and experience all fundamental forces. They are confined within hadrons—composite particles categorized as baryons² or mesons³.

I.B.1.i.b Leptons.

Like quarks, leptons too are spin $1/2$ particles that are categorised into three generations.

- Electron (e) and electron neutrino (ν_e),
- Muon (μ) and muon neutrino (ν_μ),
- Tau (τ) and tau neutrino (ν_τ).

Electrons, muons, and taus carry unit electric charge and interact through the electromagnetic and weak forces, while neutrinos are electrically neutral and interact only through the weak force, making them notoriously difficult to detect.

I.B.1.ii Bosons: Force carriers.

Bosons, with integer spin values, mediate the fundamental interactions. They are not subject to the Pauli exclusion principle and instead satisfy Bose–Einstein statistics. The Standard Model comprises the following bosons:

²three quark states, like protons and neutrons.

³quark–antiquark pairs.

The photon (γ) is a massless spin-1 boson that mediates the electromagnetic force.

$W^\pm$ **and** Z **bosons** are massive spin-1 bosons that mediate the weak force.

Gluons (g) are a set of eight massless spin-1 bosons that mediate the strong force.

The Higgs boson (H) is a massive spin-0 boson associated with the Higgs field that gives mass to elementary particles.

I.B.2 Theoretical framework and symmetries.

The Standard Model is constructed through principles of quantum field theory where particles are excitations of underlying quantum fields. Its mathematical structure is determined by local gauge invariance under the following specific symmetry transformations:

$U(1)_Y$ is associated with electroweak hypercharge and is the symmetry of electroweak theory,

$SU(2)_L$ describes the weak isospin, and acts on left-handed fermions, and

$SU(3)_C$ governs the strong interactions through colour charge in QCD.

Electroweak unification, demonstrated by Glashow [770], Weinberg [742], and Salam [740], demonstrates how the electromagnetic and weak forces emerge as different aspects of a single electroweak interaction, which undergoes spontaneous symmetry breaking at low energies.

I.B.3 The Higgs mechanism and mass generation.

The Higgs mechanism, proposed by Peter Higgs in the 1960s [765], addresses the theoretical inconsistency of massive gauge bosons in a gauge invariant theory. The mechanism introduces a scalar field—the Higgs field—that permeates space and spontaneously broke the electroweak symmetry when the universe cooled after the Big Bang.

This symmetry breaking generates masses for the W and Z bosons while leaving the photon massless, explaining the significant difference between the electromagnetic and weak forces at ordinary energies. Additionally, the Higgs field couples to fermions through Yukawa interactions [731], generating their masses with coupling strengths proportional to the particle masses.

The discovery of the Higgs boson at the LHC in 2012, with properties consistent with Standard Model predictions, provided crucial experimental validation of this mechanism and completed the Standard Model’s particle roster.

I.B.4 Limitations and beyond the Standard Model (BSM) physics.

Despite its remarkable success, the Standard Model has several well recognised limitations, including,

1. It does not incorporate gravity, the fourth fundamental force.
2. It fails to explain the observed matter–antimatter asymmetry in the universe.

3. It does not account for dark matter or dark energy, which together constitute about 95% of the universe's energy content.
4. It requires fine tuning of parameters, raising theoretical concerns like the hierarchy problem.
5. It does not explain neutrino masses, which must exist given observed neutrino oscillations.

These limitations motivate theoretical extensions and experimental searches for physics beyond the Standard Model, including supersymmetry, grand unified theories, and various dark matter candidates. Precision measurements at particle physics experiments provide one of the most powerful approaches to probe these potential extensions, making analysis techniques like those discussed in this dissertation essential for advancing our fundamental understanding of nature.

I.C Fundamental role of cross section measurements in particle physics.

Differential cross section measurements are the fundamental currency of scientific exchange in particle physics, serving as the primary bridge between theoretical predictions and experimental observations. These measurements quantify the probability density of specific particle interactions as a function of kinematic variables, providing the essential link between theoretical predictions and experimental observations. A cross section quantifies the probability of a specific particle interaction occurring and is typically expressed in units of area (barns).⁴ This seemingly simple concept forms the cornerstone of how we test and validate our understanding of fundamental physics.

The Standard Model makes precise predictions for cross sections that can be directly tested at particle physics experiments. Any statistically significant deviation between measured cross sections and theoretical predictions may signal the presence of new physics beyond the Standard Model [249].⁵

Cross sections are particularly powerful because they encode the underlying quantum field theory structure in a form that can be directly probed by experiment. For instance, measurements of jet production cross sections at different energy scales reveal the running of the strong coupling constant α_S [38], while precision electroweak cross section measurements constrain the properties of the Higgs boson and other

⁴1 barn = 10^{-24} cm².

⁵Such deviations might also signal errors in the theoretical framework used for predictions or in the experimental procedures used to measure the cross section.

fundamental particles [628]. In searches for physics beyond the Standard Model, differential cross section measurements can reveal subtle deviations that point to new particles or interactions, even when direct observation is beyond experimental reach.

These measurements also serve a crucial role in constraining effective field theories (EFTs) that parameterise potential new physics in a model independent way. By measuring differential distributions with high precision, experiments can place bounds on EFT coefficients, narrowing the space of viable theoretical extensions to the Standard Model [465].

For example, the ongoing precision program at the Large Hadron Collider (LHC) relies heavily on refined cross section measurements to extract maximum physical insight from collected data. In addition to driving comparisons with theoretical models, cross section measurements are also used at high energy physics experiments for MC tuning [15] and consistency checks [28] among other applications.

I.D Cross section measurements: From theory to experiment.

I.D.1	Theory.	28
I.D.2	Experimental measurement.	30
I.D.3	Applications in particle physics.	31
I.D.3.i	Theory validation.	31
I.D.3.ii	Monte Carlo tuning.	32
I.D.3.iii	Consistency checks.	33
I.D.3.iv	Luminosity determination.	34
I.D.3.v	Background estimation.	34
I.D.3.vi	Parton luminosity.	34
I.D.3.vii	Detector performance validation.	35
I.D.3.viii	Systematic uncertainty evaluation.	35
I.D.3.ix	Future planning.	36

I.D.1 Theory.

Classically, the cross section (σ) represents the effective area within which two particles must interact for a particular process to occur.

Definition 6 (Cross section). For collisions between discrete particles, the *cross section* is defined as the area transverse to their relative motion. If the particles were to interact via contact forces (e.g., hard spheres), the cross section corresponds to their geometric size. For long range forces however, the cross section is larger than the physical dimensions of the particles due to action-at-a-distance effects.

Definition 7 (Differential cross section). The *differential cross section* ($\frac{d\sigma}{d\Omega}$) provides additional granularity by describing how the probability of scattering depends on specific final state variables, such as scattering angle (θ) or energy transfer. It is defined as

$$\frac{d\sigma}{d\Omega} = \frac{\text{Number of events scattered into } d\Omega}{\text{Incident flux} \times \text{Target density}}. \quad (\text{I.11})$$

The total cross section can be recovered by integrating over solid angle:

$$\sigma = \int_{4\pi} \frac{d\sigma}{d\Omega} d\Omega. \quad (\text{I.12})$$

While the classical picture above is intuitive, scattering at HEP experiments is governed by quantum field theory (QFT). In this framework the probability for a process is encoded in a Lorentz invariant matrix element $\mathcal{M}$.

For a $2 \rightarrow n$ reaction with incoming 4-momenta $p_{1,2}$ and final state phase space $d\Phi_n$, the fully differential cross section is

$$d\sigma = \frac{(2\pi)^4 \delta^{(4)}(p_1 + p_2 - \sum_{i=1}^n p_i)}{4 \sqrt{(p_1 \cdot p_2)^2 - m_1^2 m_2^2}} |\mathcal{M}|^2 d\Phi_n, \quad (\text{I.13})$$

where the denominator is the flux factor and $d\Phi_n = \prod_{i=1}^n \frac{d^3 p_i}{(2\pi)^3 2E_i}$ is the Lorentz invariant phase space element.⁶ Equation (I.13) reduces to the classical area when $|\mathcal{M}|^2$ is replaced by a contact interaction and the final state integral collapses to a single kinematic configuration. Integrating Equation (I.13) over final state kinematics yields the total cross section, $\sigma = \int d\sigma$.

At tree level, $|\mathcal{M}|^2$ is computed from Feynman rules derived from the Lagrangian, while higher order corrections incorporate loops, parton showers, and non-perturbative effects such as hadronisation. For practical experimental predictions one folds $|\mathcal{M}|^2$ with parton distribution functions (PDFs) and convolves the result with detector response—precisely the forward process that the unfolding methods developed in this dissertation seek to invert.

Differential cross sections have a long history of providing valuable insights for probing fundamental properties of particles and interactions. Their use dates all the way back to Rutherford’s scattering experiments that revealed the existence of atomic nuclei by analysing angular distributions of scattered alpha particles [797].

⁶Standard derivations can be found in [127, 517, 522, 700].

I.D.2 Experimental measurement.

At modern colliders⁷ the two beams themselves act as both “projectile” and “target.” The basic experimental quantity is the instantaneous luminosity.

Definition 8 (Luminosity). The *luminosity* $\mathcal{L}(t)$ is defined such that the interaction rate for a process with cross section σ is $dN/dt = \mathcal{L}(t) \sigma$.

Time integrating over a data taking period $[t_0, t_f]$ yields the integrated luminosity.

Definition 9 (Integrated luminosity). The *integrated luminosity* $\mathcal{L}_{\text{int}}$ is defined as

$$\mathcal{L}_{\text{int}} = \int_{t_0}^{t_f} \mathcal{L}(t) dt, \quad \sigma = \frac{N_{\text{obs}} - N_{\text{bkg}}}{\mathcal{L}_{\text{int}} \epsilon A}. \quad (\text{I.14})$$

Here N_{obs} is the number of selected events, N_{bkg} an estimate of background contaminations, ϵ the detector and selection efficiency and A the geometric-plus-kinematic acceptance of the analysis.

For binned measurements one bins events in an observable X ⁸ and divides by the bin width.

$$\left. \frac{d\sigma}{dX} \right|_{X_i} = \frac{1}{\mathcal{L}_{\text{int}} \Delta X_i} \frac{N_i^{\text{obs}} - N_i^{\text{bkg}}}{\epsilon_i A_i}, \quad (\text{I.15})$$

with the index i denoting the i^{th} bin.

Luminosity determination is itself a precision measurement, usually performed with dedicated luminometers that exploit van der Meer scans or pileup counting techniques.

⁷This section focusses on collider experiments, but similar analyses can be applied to non-collider HEP experiments as well, such as fixed target experiments. [703] and chapters 31 to 38 of [249] provide a comprehensive and detailed exposition of experimental measurement in HEP. [4, 531] focus specifically on fixed target techniques and phenomenology. [1] compares collider and fixed target formalisms, reviews luminosity analogues, and details the role of detector simulations and unfolding in a fixed target context.

⁸For example, transverse momentum p_T or rapidity y .

The efficiency–acceptance term ϵA is obtained from full detector simulations and corrected in data using control samples and “tag–and–probe” methods.

Equations (I.14) and (I.15) thus link the theoretically calculated parton level cross sections (*vide* Equation (I.13)) to the raw observables recorded by the a detector, completing the chain from theory to experiment.

I.D.3 Applications in particle physics.

As mentioned above, cross section measurements serve as the fundamental currency of particle physics, translating abstract theoretical predictions into measurable experimental quantities. However, their applications extend far beyond simple theory validation into the operational heart of how experiments function, analyse data, and cross-validate results.

I.D.3.i Theory validation.

As discussed above, comparing measured cross sections with predictions from quantum field theory validates and tests theoretical models like quantum chromodynamics and electroweak theory, by encapsulating interaction probabilities in a measurable form. Deviations from expected cross sections may indicate new phenomena, such as supersymmetric particles or dark matter candidates.

Differential cross sections also provide constraints on effective field theories and parton distribution functions (PDFs), essential for understanding the internal structure of hadrons. Unfolded cross section measurements allow comparisons with theoretical models years after data collection, even if detector simulations are no longer available,

further enhancing their utility, and future proofing the data. Their determination requires careful design and analysis techniques to account for systematic uncertainties introduced by detector effects.

I.D.3.ii Monte Carlo tuning.

Definition 10 (Monte Carlo tuning). *Monte Carlo tuning* is the iterative process of adjusting simulation parameters to match measured cross sections, ensuring that detector simulations accurately reproduce real experimental data.

Cross section measurements are extensively used in Monte Carlo (MC) tuning.

Every particle physics analysis relies on sophisticated simulations that model everything from the initial parton interactions through hadronisation to the detector response. These simulations contain dozens of phenomenological parameters, such as the strong coupling constant at various scales and non-perturbative fragmentation functions.

Definition 11 (Parton distribution functions). *Parton distribution functions* (PDFs) are probability distributions that describe the probability of finding a parton (quark or gluon) in a hadron at a given momentum fraction x and scale Q^2 . PDFs are determined from global fits to a wide range of hard scattering processes, including deep inelastic scattering, Drell–Yan production, and jet production.

Measured cross sections provide the ground truth that anchors these simulations to reality.

Consider the following example: When CMS measures the inclusive jet cross section at a new centre-of-mass energy, that measurement immediately becomes a crucial input for tuning generators like PYTHIA or HERWIG. The differential distributions—whether in transverse momentum, rapidity, or invariant mass—reveal where the models succeed and where they fail. A discrepancy in the high- p_T tail might indicate a need to adjust the modelling of initialstate radiation; unexpected structure in angular distributions could point to missing higher order QCD effects.

I.D.3.iii Consistency checks.

Cross-validations between different experimental approaches, detector configurations, or analysis methods ensure measurement reliability and identify systematic biases. Cross sections serve as essential tools for consistency checks across multiple dimensions of experimental physics. Within a single experiment, measuring the same process through different decay channels provides a powerful systematic cross check. For instance, measuring the W boson production cross section through both electronic and muonic decays tests the understanding of lepton universality while simultaneously validating detector calibrations. Any significant deviation signals either new physics, errors in the phenomenological method used, or unaccounted systematic effects.

Between experiments, cross section measurements enable crucial cross-experiment validation. When CMS and ATLAS measure the same process with independent detectors and analysis chains, agreement within uncertainties validates both measurements. Disagreement, conversely, can reveal subtle systematic effects or push calculations to higher precision.

These applications cascade through every level of experimental operations.

I.D.3.iv Luminosity determination.

Luminosity determination becomes possible through processes with well-known theoretical cross sections. Van der Meer scans calibrate the absolute luminosity scale, but elastic scattering and other standard candle processes provide continuous monitoring. The uncertainty on integrated luminosity, typically 2 to 3%, directly impacts every cross section measurement, creating a web of interdependencies.

I.D.3.v Background estimation.

Background estimation in searches for new physics relies on measured cross sections of Standard Model processes. When searching for supersymmetric particles, the irreducible backgrounds from $W + \text{jets}$ or $t\bar{t}$ production must be understood at the percent level. Control regions enriched in backgrounds, combined with precise cross section measurements, enable data driven background estimates that would be impossible from simulation alone.

I.D.3.vi Parton luminosity.

Definition 12 (Parton luminosity). *Parton luminosity* is the effective luminosity for specific parton–parton interactions, calculated by convolving the total luminosity with parton distribution functions.

A particularly elegant application of cross sections emerges in parton luminosity calculations. Since protons are composite objects, the effective luminosity for producing

heavy particles depends on the convolution of PDFs with the partonic cross section. Measurements of Drell–Yan production at different invariant masses directly probe the quark and antiquark distributions, while inclusive jet production constrains the gluon PDF. This creates a self-consistent feedback loop, where better PDFs enable more precise predictions, which enable more sensitive measurements, which further constrain the PDFs.

I.D.3.vii Detector performance validation.

Detector performance validation represents another major application of cross section measurements. Measured cross sections for well understood processes serve as standard candles for monitoring detector stability over time. For example, slow drift in the measured $Z \rightarrow \mu\mu$ cross section might indicate degrading muon chamber performance long before it would be noticed in individual event displays. These measurements become part of the experiment’s data quality monitoring, flagging problems in real time.

I.D.3.viii Systematic uncertainty evaluation.

The role of cross section measurements in systematic uncertainty evaluation cannot be overstated. Every measurement must account for theoretical uncertainties in signal and background processes. By measuring auxiliary cross sections, for instance, $Z + \text{jets}$ production when studying $W + \text{jets}$, experiments can constrain these uncertainties using data rather than relying solely on theoretical estimates. This *in situ* constraint often reduces systematic uncertainties by factors of two or more.

I.D.3.ix Future planning.

Finally, cross sections enable physics program planning for future experiments. The measured production rates at current energies, extrapolated using theoretical calculations, determine required luminosities and detector capabilities for next-generation experiments. For example, the surprisingly large Higgs production cross section at the LHC, for instance, has already influenced design considerations for future electron-positron Higgs factories. In sum, cross sections certainly do function as the Rosetta Stone of particle physics, translating between the languages of theory, simulation, and experimental measurement while maintaining coherence across all three domains. However, they serve not merely as endpoints of analyses, but as the connective tissue that binds together theory, simulation, and experiment into a coherent whole. They simultaneously test theoretical understanding, calibrate experimental tools, and point the way toward new discoveries. This multiplicative utility explains why cross section measurements, even of well studied processes, remain at the heart of every particle physics experiment.

I.E Detector response in precision measurements.

The direct comparison between theoretical predictions and experimental measurements is complicated by detector effects. HEP detectors are technological marvels that capture the trajectories of charged particles, energy deposits in calorimeters, and timing and pattern-recognition information from tracking and particle-identification systems, but they introduce distortions that must be carefully accounted for to extract the true physical distributions of interest. Particle physics detectors represent some of humanity’s most sophisticated sensing apparatus—capturing particle trajectories with silicon sensors operating at liquid helium temperatures, measuring energy deposits in dense calorimeter crystals, and reconstructing vertices with sub-millimetre precision. Yet these technological marvels inevitably introduce systematic distortions that transform the pristine theoretical predictions into the messy reality of experimental data.

Consider the information degradation that occurs in every measurement. Finite resolution creates fundamental blurring, much like how a camera lens distorts an image. The detector’s discrete sensing elements can only measure particle energies, momenta, and positions to finite precision, creating an inherent convolution between the true physics distribution and the instrument response function. Geometric acceptance imposes hard boundaries on observable phase space. Particles scattered into the forward beam pipe or extreme backward angles simply vanish from the recorded dataset, creating holes in the measurement that no amount of statistics can fill.

Detection efficiency varies across the detector’s active volume, introducing a complex weighting function that depends on particle type, energy, and trajectory. A high-

energy muon might traverse the entire detector with near-perfect efficiency, while a low-energy hadron could be absorbed in the first layers of material. Particle misidentification compounds these challenges through cross-contamination between categories, such as when hadronic shower fluctuations cause a pion to masquerade as a kaon [529].

In this way every detector measurement embeds two irreducible probability relationships.

Detection incompleteness $\mathbb{P}(\text{measured} \mid \text{true}) < 1$ True events that fail to be recorded

Measurement impurity $\mathbb{P}(\text{true} \mid \text{measured}) < 1$ Recorded events that represent contamination

This demonstrates why detector corrections are fundamentally different from simple calibrations. Unlike adjusting a scale that consistently reads, say, 5% high, detector response involves dual information loss that operates asymmetrically.

The first inequality captures the selection bias where certain true configurations have zero probability of detection, creating null spaces in the measurement. The second inequality captures the contamination bias where every reconstruction category contains some fraction of misclassified events.

Background contamination represents a third problem—every measurement category contains some admixture of misclassified events. When hadrons interact in electromagnetic calorimeters, they can mimic electron signatures. When cosmic ray muons traverse the detector during a collision, they contribute to the muon count

despite having no connection to the physics of interest. This contamination creates what in signal processing is called the false positive rate.

The mathematical relationship between true particle level distributions and observed detector level measurements follows the convolution integral

$$p(x) = \int r(x | z) p(z) dz \quad (\text{I.16})$$

Where $p(x)$ is the detector level density, $p(z)$ is the particle level density and $r(x | z)$ serves as the response kernel, the conditional probability density that maps each possible true configuration z to the distribution of possible detector measurements x . This kernel encapsulates the entire cascade of finite resolution information degradation.

The response kernel is analogous to the optical transfer function in image processing. It describes how the “lens” of the detector blurs and distorts the perfect theoretical “image”. The response function inherently embeds both probabilistic asymmetries. Regions where $\int r(x | z) dx < 1$ reveal acceptance holes, i.e. true configurations that produce no detector signal whatsoever. Conversely, the convolution structure itself ensures that multiple truth distributions can yield identical detector observations, creating the degeneracy problem that makes direct inversion impossible.

A central challenge then, for particle physics, is inverting this response kernel to recover $p(z)$ from observed data $p(x)$. This inversion is mathematically ill posed precisely because of the information loss. Standard matrix inversion fails catastrophically, amplifying statistical noise into wild oscillations that bear no resemblance to the underlying physics.

The resolution requires sophisticated regularisation methods that impose additional constraints such as smoothness assumptions, positivity requirements, and prior

knowledge about the expected signal shape. These constraints transform the ill posed inverse problem into a well defined statistical inference challenge, though at the cost of introducing systematic uncertainties that must themselves be carefully validated.

The detector response problem exemplifies a universal pattern in experimental science: the tension between instrumental precision and information preservation. From astronomical imaging through medical diagnosis to HEP measurements, the fundamental trade-off between sensitivity and purity governs all attempts to extract signal from noise.

I.F Challenges at modern experiments.

Several challenges at the modern HEP experiments make cross section measurements particularly demanding.

High dimensional phase spaces Modern measurements often involve multiple correlated observables, creating high dimensional distributions that are difficult to analyse with traditional methods.

Limited statistics in extreme regions Rare processes or the tails of distributions often contain valuable physics information but suffer from limited statistics.

Complex detector effects Detectors have non-trivial response functions that can vary significantly across phase space, and are only known implicitly through precision simulations. Their explicit functional form is unknown.

Theoretical uncertainties Precision measurements are increasingly limited by theoretical uncertainties in both signal and background modelling.

Computational constraints Detailed simulation of detector response requires substantial computing resources, limiting the statistical precision of response modelling.

These challenges make the unfolding problem increasingly difficult, particularly as measurements probe more complex final states and differential distributions. For example, measurements of jet substructure, which probe the detailed radiation pattern within collimated sprays of particles, involve observables with complex correlations

and detector effects that vary based on jet energy, rapidity, and substructure properties themselves [336, 377, 443].

The need for unfolding arises from the fundamental requirement to present results in a detector independent form that can be directly compared with theory predictions or results from different experiments. Without this correction, theoretical interpretations would need to incorporate experiment specific detector simulations, significantly complicating scientific exchange and theoretical analysis, and inter-experiment comparisons would simply not be possible.

I.G Dissertation scope and physics impact.

This dissertation focuses on developing, analysing, and applying novel machine learning methods for cross section measurements in particle physics, with particular emphasis on unbinned approaches that overcome limitations of traditional techniques. The work spans the spectrum from improving binned methods with neural posterior estimation to completely unbinned approaches for both full distributions and statistical moments. All of the original research presented in this dissertation represents collaborative efforts with many brilliant scientists, to each of whom I am deeply grateful. The views I present here my personal perspective on these developments. Where I discuss or present specific results primarily generated by my collaborators, I explicitly indicate this attribution. The collaborative nature of this work has been one of its greatest strengths, and I am indebted to all those who contributed their expertise and insights to these projects.

The primary contributions of this dissertation include:

1. Development of NEURAL POSTERIOR UNFOLDING (NPU), enhancing binned approaches through normalising flows and amortised inference.
2. Introduction of MOMENT UNFOLDING, directly deconvolving distribution moments without binning.
3. Creation of REWEIGHTING ADVERSARIAL NETWORKS (RAN), a general framework for unbinned spectrum unfolding.

4. Analysis of event correlations in unfolded data and their impact on uncertainty estimation.
5. Investigation of symmetry discovery with SYMMETRYGAN and its connections to measurement constraints.

These methodological advances address fundamental challenges in experimental particle physics, potentially enhancing the precision and scope of measurements at present and future HEP experiments. These methods can have a wide range of applications in particle physics, including:

- Improved precision in jet substructure measurements, enabling better discrimination between different theoretical models of QCD radiation.
- Enhanced sensitivity to effective field theory parameters by directly deconvolving distribution moments.
- More robust uncertainty quantification in high dimensional measurements.
- Computational efficiency gains allowing for more detailed systematic studies.
- A rigorous statistical framework for incorporating detector response uncertainties in the unfolding process.

By bridging sophisticated machine learning techniques with the specific requirements of particle physics measurements, this work aims to advance the ability to extract fundamental physical insights from complex experimental data. The methods

developed here have applications beyond particle physics, potentially benefiting any field where deconvolution of instrumental effects is necessary for scientific inference.

Chapter II

Theoretical foundations.

II.A	Statistical formulation of the unfolding problem.	48
II.A.1	The detector response and forward problem.	49
II.A.2	The inverse problem: Unfolding.	49
II.A.3	Likelihood based formulation.	50
II.A.4	Regularisation techniques.	50
II.A.5	Challenges in high dimensional phase spaces.	51
II.B	Forward and inverse problems in HEP.	57
II.B.1	Mathematical formulation.	58
II.B.2	Challenges in inverse problems.	59
II.B.3	HEP specific considerations.	59
II.C	Historical development: From matrix inversion to modern approaches.	61
II.D	Traditional unfolding methods in experimental analyses.	63
II.D.1	Bin by bin correction.	64

II.D.2	Matrix Inversion	65
II.D.3	Bayesian methods.	65
II.D.4	Tikhonov regularisation.	68
II.D.5	Template fitting.	68
II.D.6	Regularised Poisson likelihood.	69
II.D.7	Summary.	69
II.E	Regularisation: Need, approaches, and limitations.	71
II.E.1	Necessity.	72
II.E.2	Limitations and practical challenges.	72
II.F	Unbinned methods: statistical considerations.	75
II.F.1	Principles and implementations.	76
II.F.2	Statistical considerations in unbinned regimes.	80
II.F.3	Limitations in complex phase spaces.	81
II.G	Evaluation metrics for unfolding.	83
II.G.1	Statistical metrics for evaluating point estimates.	84
II.G.2	Uncertainty quantification metrics.	88
II.G.3	Evaluation of correlation structure.	92
II.G.4	Method specific evaluation metrics.	96
II.G.5	Practical considerations.	97

II.A Statistical formulation of the unfolding problem.

II.A.1	The detector response and forward problem.	49
II.A.2	The inverse problem: Unfolding.	49
II.A.3	Likelihood based formulation.	50
II.A.4	Regularisation techniques.	50
II.A.5	Challenges in high dimensional phase spaces.	51
II.A.5.i	Binned methods.	51
II.A.5.ii	Unbinned methods.	55

Unfolding, also known as deconvolution, is the process of correcting detector distortions in experimental data to recover the true particle level distributions. This procedure is critical for comparing experimental results with theoretical predictions and for enabling detector independent analyses. The unfolding problem is inherently statistical and presents unique challenges due to its ill posed nature.

II.A.1 The detector response and forward problem.

The relationship between the particle level truth distribution $p(z)$ and the detector level measured distribution $p(x)$ is governed by the detector response function $r(x | z)$, which encapsulates the resolution effects.¹ Equation (I.16) describes the **forward problem**, where the true distribution $p(z)$ is mapped to the measured distribution $p(x)$. The detector response function $r(x | z)$ can often be estimated through detailed simulations.

II.A.2 The inverse problem: Unfolding.

The goal of unfolding is to invert the forward problem and estimate the Truth, $p(z)$ from Data $p(x)$. Mathematically, this requires solving

$$p(z) = \int r^{-1}(z|x) p(x) dx, \quad (\text{II.1})$$

where $r^{-1}(z | x)$ represents the inverse response kernel. However, this inversion is ill posed because small fluctuations in $p(x)$ can lead to large variations in $p(z)$ [634]. Regularisation techniques are therefore essential to stabilise the solution.

¹Efficiency and acceptance effects can also be incorporated if “empty” events are allowed.

II.A.3 Likelihood based formulation.

In practice, unfolding is performed using statistical inference methods. Given a set of measured data $\mathbf{X}_{i=1}^N$, the likelihood function for a proposed truth distribution $p(z; \theta)$, parametrised by θ , is

$$\mathcal{L}(\theta; \mathbf{X}) = \prod_{i=1}^N p(x_i; \theta), \quad (\text{II.2})$$

where

$$p(x; \theta) = \int r(x | z) p(z; \theta) dz. \quad (\text{II.3})$$

Maximising this likelihood yields an estimate of the parameters θ , which define the unfolded truth distribution. Regularisation can be incorporated into this framework by adding penalty terms to the likelihood or by constraining the parameter space.

Note. The likelihood function in Equation (II.2) represents the probability of the data $\mathbf{X}$ given the parameters θ for relative cross sections. If one were to use the likelihood function for absolute cross sections, the likelihood would also depend on the luminosity, and therefore include a poisson term for N .

II.A.4 Regularisation techniques.

Regularisation mitigates the instability of unfolding by imposing constraints on the solution. Common approaches include Tikhonov regularisation [636], in which one adds a penalty term proportional to the norm of the second derivative of $p(z)$, enforcing smoothness, and iterative methods which gradually refine estimates of $p(z)$, regularising by stopping before convergence. These techniques balance fidelity to the measured data (prior independence) with stability of the unfolded solution.

II.A.5 Challenges in high dimensional phase spaces.

II.A.5.i Binned methods.

A significant challenge that traditional unfolding methods face is the *curse of dimensionality*. As the number of measured observables increases, the statistical power required to populate discrete bins grows exponentially, quickly overwhelming even the largest datasets collected at modern experiments.

Consider a measurement involving just four kinematic variables, say, transverse momentum, pseudorapidity, azimuthal angle, and invariant mass. With a modest 20 bins per dimension, the resulting joint histogram requires 1.6×10^5 bins. Most of these bins will contain zero events, creating a sparse matrix that renders traditional unfolding techniques numerically unstable.

Theorem 1 (Scaling of relative uncertainty). Consider a d dimensional measurement space partitioned into histogram bins of widths $\Delta z_1, \dots, \Delta z_d$, and let N_{events} denote the total number of events uniformly distributed over the full measurement range. Then the relative uncertainty in the estimated differential event density in any given bin obeys

$$\text{Rel. uncertainty} = \frac{\sigma_Q}{Q} \propto \frac{1}{\sqrt{N_{\text{events}}} \prod_{i=1}^d \Delta z_i}, \quad (\text{II.4})$$

where Q is the differential quantity (events per unit hypervolume).

Proof of Theorem 1.

A single histogram bin occupies hypervolume

$$V = \prod_{i=1}^d \Delta z_i \quad (\text{II.5})$$

in the measurement space. Under the uniform-density assumption, the expected number of events in that bin scales as

$$N_{\text{bin}} \propto N_{\text{events}} V. \quad (\text{II.6})$$

By Poisson statistics the absolute uncertainty on the bin count is

$$\sigma_{N_{\text{bin}}} = \sqrt{N_{\text{bin}}}, \quad (\text{II.7})$$

and hence the relative uncertainty on N_{bin} is

$$\frac{\sigma_{N_{\text{bin}}}}{N_{\text{bin}}} = \frac{1}{\sqrt{N_{\text{bin}}}} \propto \frac{1}{\sqrt{N_{\text{events}} V}}. \quad (\text{II.8})$$

Since the differential measurement is $Q = N_{\text{bin}}/V$, its relative uncertainty coincides with that of the bin count.

$$\frac{\sigma_Q}{Q} = \frac{(\sigma_{N_{\text{bin}}}/V)}{(N_{\text{bin}}/V)} \quad (\text{II.9})$$

$$= \frac{\sigma_{N_{\text{bin}}}}{N_{\text{bin}}} \quad (\text{II.10})$$

$$\propto \frac{1}{\sqrt{N_{\text{events}}} \prod_{i=1}^d \Delta z_i}, \quad (\text{II.11})$$

which establishes the claimed scaling. $\square$

This relationship reveals why traditional binned methods become increasingly impractical as dimensionality increases: maintaining fixed statistical precision requires exponentially more data or exponentially coarser binning, both of which severely limit the measurement's resolving power. Beyond a point, the *response matrix*, $\mathbf{R}$, becomes not just ill conditioned but genuinely rank deficient, as sufficiently many bins have

not just few events in them, but rather no events in them at all. Hence, entire swaths of phase space remain unmeasurable regardless of accumulated statistics.

Binning artifacts compound these statistical challenges through systematic information loss. The process of discretising continuous distributions into finite bins introduces artificial boundaries onto the underlying physics. This discretisation becomes particularly pernicious when dealing with correlation structures that span multiple dimensions. Real particle interactions generate complex kinematic relationships—the angular distribution of decay products correlates with their energies, jet substructure variables depend on the overall jet momentum, and detector response functions couple seemingly independent observables. Binning destroys these correlations by requiring low dimensional projections.

For instance, in jet substructure measurements, the relationship between different substructure variables contains valuable information about the underlying physics that can be obscured by independent binning of each variable. Moreover, many theoretical predictions in particle physics are at the level of statistical moments or other distribution properties rather than full differential spectra. Traditional unfolding methods require first unfolding the full distribution and then calculating these properties, which can lead to reduced precision in the moment predictions.

The computational burden grows even faster than the statistical requirements. Matrix inversion algorithms scale as $\mathcal{O}(N^3)$ with bin number, meaning that the aforementioned four dimensional histogram with 1.6×10^5 bins requires approximately $\mathcal{O}(10^{16})$ floating point operations to invert. More fundamentally, the *null spaces* and *degeneracies* that plague low dimensional unfolding become dramatically worse

in higher dimensions, where multiple truth configurations can project to identical detector signatures along numerous measurement axes simultaneously, because high dimensional measurements create complex geometric relationships between true and observed phase spaces.

In addition to these limitations, traditional binned approaches suffer from an additional systematic weakness; the response matrix $\mathbf{R}$ depends on numerous variables that are typically marginalised over in order to bin the data. While conventional methods construct $\mathbf{R}$ based on a limited set of binned observables, the true detector response depends on a much richer set of event characteristics. These include additional kinematic variables not captured in the chosen binning scheme, event level properties such as particle multiplicity and missing energy, detector conditions like instantaneous luminosity and pileup activity, and correlations with other particles produced in the same collision.

Traditional unfolding methods are effectively forced to assume that $\mathbf{R}$ remains constant when averaged over these marginalised variables, but this assumption is patently incorrect when the detector response varies systematically across different regions of this extended phase space. For example, the energy resolution for jets may depend not only on the jet's transverse momentum and pseudorapidity, variables often chosen for binning, but also on the jet's substructure, the presence of nearby particles, and the overall event topology. By marginalising over these features, binned methods introduce systematic biases that can propagate through the unfolding procedure and distort the final measurements in ways that are difficult to quantify or correct.

II.A.5.ii Unbinned methods.

Unbinned approaches emerge as a natural response to these challenges, operating directly on individual events rather than aggregated histogram counts. These methods exploit the event level structure that binning destroys, preserving the full kinematic information and correlation patterns within each recorded data point. Instead of discretising phase space into predetermined categories, unbinned techniques allow the data itself to determine the relevant resolution scales and correlation structures.

Rather than explicitly modelling the response function $r(x | z)$, modern techniques like OMNIFOLD [309] employ iterative reweighting strategies that learn implicit mappings between particle level and detector level distributions. These methods sidestep the curse of dimensionality by avoiding explicit probability density estimation, instead focusing on weight optimisation that preserves marginal distributions while respecting the detector response.

The computational complexity shifts from matrix algebra to optimisation landscapes, that scale more favourably with dimensionality. While traditional unfolding requires inverting matrices that grow exponentially with dimension, unbinned methods typically employ gradient based optimisation that scales polynomially. This trade-off exchanges the well understood numerical properties of linear algebra for the more complex but ultimately more scalable challenges of machine learning optimisation.

It is important to acknowledge that this transition introduces new sensitivities to model and hyperparameter choices that traditional methods avoid. The implicit effect of the learning dynamics make it crucial to understand how optimisation biases influence the final results. These considerations establish the foundation for exploring

how modern machine learning approaches navigate these challenges while preserving the statistical rigour that particle physics demands.

II.B Forward and inverse problems in HEP.

II.B.1	Mathematical formulation.	58
II.B.2	Challenges in inverse problems.	59
II.B.3	HEP specific considerations.	59

The measurement process in high energy physics experiments inherently involves two complementary mathematical challenges: the forward problem of predicting detector responses from the theoretical understanding of particle level interactions, and the inverse problem of recovering true physics distributions from observed detector measurements. These twin challenges form the conceptual foundation for understanding detector effects and developing unfolding methodologies.

II.B.1 Mathematical formulation.

The relationship between particle level truth distributions and detector level observations is governed by the Fredholm integral equation of the first kind [798],

$$p(x) = \int r(x | z) p(z) dz + \epsilon(x), \quad (\text{II.12})$$

where $p(z)$ represents the true particle level distribution, $r(x | z)$ is the detector response kernel encoding resolution effects and acceptance, $p(x)$ is the observed detector level distribution, and $\epsilon(x)$ accounts for measurement noise [768]. This equation encapsulates the *forward problem* when predicting $p(x)$ given $p(z)$, and the *inverse problem* when estimating $p(z)$ from data $p(x)$.

For discrete histogram representations, this becomes a matrix equation,

$$\boldsymbol{\mu} = \mathbf{R}\boldsymbol{\nu} + \boldsymbol{\epsilon}, \quad (\text{II.13})$$

where $\boldsymbol{\nu}$ and $\boldsymbol{\mu}$ are vectors of true and observed bin counts, respectively, $\mathbf{R}$ is the smearing matrix containing conditional probabilities, $R_{ij} = \mathbb{P}(\text{observed bin } i | \text{true bin } j)$ [566], and $\boldsymbol{\epsilon}$ is the vector of measurement noise.

II.B.2 Challenges in inverse problems.

The inverse problem in HEP is fundamentally and intrinsically ill posed. The response kernel is non-injective, i.e. different true distributions can produce identical observed distributions after detector smearing [68, 681, 768]. and, the distributions are ill conditioned; small measurement errors ϵ amplify into large fluctuations in unfolded solutions due to small singular values in $\mathbf{R}$ [5, 514, 610]. These challenges, intrinsic to inverse problems, are compounded by the fact that modern analyses involve a large number of observables, making brute force phase space discretisation computationally prohibitive [167, 211, 330]. The curse of dimensionality is particularly pronounced in high energy physics, where the number of observables can be in the tens or hundreds. Addressing these challenges necessitates regularisation techniques that impose physical constraints on solutions, classic examples of which are Tikhonov regularisation or iteration cut-offs before convergence, described in Section II.C.

II.B.3 HEP specific considerations.

Three aspects particularly complicate unfolding in particle physics compared to other inverse problem domains. First, the response matrix $\mathbf{R}$ is estimated from from detailed Monte Carlo simulations² that encode complex, non-Gaussian systematic uncertainties through their modelling assumptions [54, 392, 449, 564].

Second, detector response models contain hundreds to thousands of correlated nuisance parameters—including jet energy scale, b -tagging efficiencies, and pile up

²E.g., PYTHIA [213] for hadronisation, GEANT4 [461] for detector physics.

effects—that must be profiled or marginalised alongside the unfolding procedure [151, 188, 325, 551].

Finally, the discrete nature of particle counting combined with highly variable event rates across phase space creates a challenging statistical landscape where Poisson uncertainties dominate in tails and signal regions, while Gaussian approximations may be valid elsewhere [329, 439, 589, 658].

A representative unfolding example is differential jet substructure measurements, where detector effects smear the true distribution of observables like jet mass or N –subjettiness. The forward problem involves simulating jets through hadronisation models and detector response, producing a migration matrix that relates true and reconstructed substructure observables. The inverse problem requires unfolding these observables from reconstructed jets while accounting for correlated uncertainties in jet energy scale, angular resolution, and pile up contamination [456].

Similar challenges arise in unfolding differential cross sections as functions of transverse momentum, rapidity, or invariant mass, where detector acceptance and resolution vary significantly across the measurement range [485].

II.C Historical development: From matrix inversion to modern approaches.

The problem of unfolding has a rich history in high energy physics, with methods evolving alongside computational capabilities and statistical sophistication. Early approaches relied primarily on simple correction factors applied to individual bins of histograms, appropriate only when detector effects were minimal.

As measurements became more precise, regularised matrix inversion techniques emerged as the standard approach. These methods discretise both the particle level and detector level distributions into bins, relating them through a response matrix R_{ij} that describes the probability for an event in particle level bin j to be observed in detector level bin i . In component form, Equation (II.13) can be written as

$$\mu_i = \sum_j R_{ij} \nu_j + \epsilon_i, \quad (\text{II.14})$$

where μ_i is the expected number of counts in detector level bin i , ν_j is the expected number of events in particle level bin j , and ϵ_i is the measurement noise in detector level bin i . Naively, one might attempt to solve this system by simply inverting the response matrix:

$$\nu_j = \sum_i R_{ji}^{-1} \mu_i \quad (\text{II.15})$$

However, this direct inversion leads to wildly oscillating solutions with large variances—a manifestation of the ill posed nature of the unfolding problem. Additionally, $\mathbf{R}$ need not, and often is not, a square matrix, and so *the* inverse is not even well defined.

To address this issue, a series of techniques were developed to invert Equation (II.13) by imposing additional constraints.

Iterative Bayesian unfolding [450, 754, 759] , also known as Lucy–Richardson deconvolution, uses Bayes’ theorem to iteratively update the estimate of the true distribution, with the number of iterations controlling regularisation strength.

SVD unfolding [693] applies singular value decomposition to the response matrix and suppresses contributions from small singular values that amplify statistical fluctuations.

TUnfold [557] formulates unfolding as a least squares problem with Tikhonov regularisation to penalise large second derivatives, preserving smoothness.

These methods, which will be discussed in more detail in Section II.D, have served the field well for decades, particularly for one dimensional measurements where binning is manageable. However, they all share the common limitation of requiring discretisation of the underlying distributions, which becomes increasingly problematic as measurements probe higher dimensional spaces and more complex observables.

II.D Traditional unfolding methods in experimental analyses.

II.D.1	Bin by bin correction.	64
II.D.2	Matrix Inversion	65
II.D.3	Bayesian methods.	65
II.D.3.i	Iterative Bayesian unfolding (IBU).	66
II.D.4	Tikhonov regularisation.	68
II.D.5	Template fitting.	68
II.D.6	Regularised Poisson likelihood.	69
II.D.7	Summary.	69

Traditional unfolding methods form the bedrock of detector corrections in high energy physics, balancing statistical rigour with computational practicality. This section provides an overview of established techniques, their mathematical foundations, implementation nuances, and limitations.

II.D.1 Bin by bin correction.

The simplest unfolding approach applies multiplicative correction factors to observed bin counts.

$$\hat{\nu}_j = \frac{\mu_j - b_j}{C_j}, \quad C_j = \frac{\nu_j^{\text{MC}}}{\mu_j^{\text{MC}}}, \quad (\text{II.16})$$

where μ_j is the observed count in bin j , b_j the estimated background, and C_j the correction factor derived from Monte Carlo (MC) simulations relating particle-level generated (ν_j^{MC}) and detector-level simulated (μ_j^{MC}) events [277].

This method has the advantage of being computationally trivial, with no bin-to-bin correlations. However, it fails to account for a non-diagonal response matrix ($\exists i \neq j \ R_{ij} \neq 0$). This is illustrated most dramatically by the observation that biases persist even with $C_j \rightarrow 1$ due to ignored cross-bin migrations [586].

Used primarily in early LHC analyses³, bin-by-bin correction remains viable only for coarse binnings with negligible migration ($\lesssim 5\%$ [566]) between adjacent bins.

³E.g., ATLAS [8, 561] jet cross-sections

II.D.2 Matrix Inversion

When $n_{\text{bins, Gen.}} = n_{\text{bins, Sim.}}$, the response matrix $\mathbf{R}$ is square. Formally one can write an unfolded solution as

$$\hat{\boldsymbol{\nu}} = \mathbf{R}^{-1} \boldsymbol{\mu}, \quad (\text{II.17})$$

and even propagate the covariance as

$$\Sigma_{\hat{\boldsymbol{\nu}}} = \mathbf{R}^{-1} \Sigma_{\boldsymbol{\mu}} (\mathbf{R}^{-1})^T. \quad (\text{II.18})$$

However, in practice, direct inversion is highly pathological. This pathology can be quantified by the condition number

$$\kappa(R^{-1}) = \frac{|\lambda_{\max}(R^{-1})|}{|\lambda_{\min}(R^{-1})|} \sim 10^3 - 10^6 \quad (\text{II.19})$$

where $\lambda_{\max}(R^{-1})$ and $\lambda_{\min}(R^{-1})$ are the largest and smallest eigenvalues of R^{-1} respectively. The condition number measures how much a perturbation in the measured counts $\delta\mu$ perturbs the predicted truth counts $\delta\nu$. The large condition number amplifies statistical fluctuations [498, 632]. Further, unphysical solutions such as negative bin count values can arise from noise dominated eigenvectors.

Methods have been suggested to control this variance, such as Truncated SVD [109], involving discard singular values $\sigma_i < \lambda_{\text{cut}}$ [75], and Wiener-SVD [454], a frequency domain filtering method to maximise signal-to-noise ratio [701]. Despite these proposals, matrix inversion's instability limits its utility.

II.D.3 Bayesian methods.

Bayesian methods regularise through prior distributions $p(z)$, yielding posterior estimates.

$$p(z \mid x) \propto \mathcal{L}(x \mid z) p(z) \quad (\text{II.20})$$

The initial guess $\boldsymbol{\nu}^{(0)}$ biases the solution. Common priors include Generation, $\boldsymbol{\nu}_{\text{MC}}$, a uniform distribution, and data driven backwards folding $\mathbf{R}^T \boldsymbol{\mu}$. The entropy maximisation prior, $p(z) \propto \exp(-\sum z_j \ln z_j)$ [539] and Gaussian processes for enforcing smoothness [392] have also been proposed as suitable priors for IBU. These provide natural uncertainty quantification but suffer from high computational cost, scaling poorly with dimensionality [615], sensitivity to prior misspecification, especially in low-statistics regions [623], and difficulty interpreting credible intervals as frequentist coverage [650].

II.D.3.i Iterative Bayesian unfolding (IBU).

Despite its name, IBU is not a true Bayesian method. While Bayesian approaches regularise through prior distributions $p(z)$ to yield posterior estimates via Bayes' theorem, IBU is an expectation maximisation (EM) algorithm that uses an initial guess rather than a prior distribution. IBU, also known as Lucy–Richardson deconvolution or D'Agostini iterative unfolding [587] is an expectation maximisation (EM) algorithm iteratively updates an initial guess $\boldsymbol{\nu}^{(0)}$ through the procedure described in Algorithm 1.

This method regularises via early stopping, by terminating at $K \sim 4 - 6$ iterations before noise amplification [696, 735].

While computationally efficient, this approach lacks objective stopping criteria, requiring heuristic cross-validation [666] and underestimates uncertainties due to

Algorithm 1 Iterative Bayesian unfolding (Richardson–Lucy deconvolution/D’Agostini iteration)

Input: Response matrix R_{ij} ,
 histogram counts of Data, μ_i ,
 initial estimate $\nu_j^{(0)}$,
 # iterations K ,
 # bins $N_{\text{Sim.}}$, $N_{\text{Gen.}}$.

Output: Unfolded (Truth) distribution $\nu_j^{(K)}$

```

1 for  $k \leftarrow 0$  to  $K - 1$  do
2   for  $j \leftarrow 1$  to  $N_{\text{Gen.}}$  do
3     
$$\nu_j^{(k+1)} \leftarrow \nu_j^{(k)} \frac{\sum_{i=1}^{N_{\text{Sim.}}} R_{ij} \mu_i}{\sum_{\ell=1}^{N_{\text{Gen.}}} R_{i\ell} \nu_\ell^{(k)}} \quad (\text{II.21})$$

4   // Normalisation to enforce  $\sum_j \nu_j = 1$ 
5   for  $j \leftarrow 1$  to  $N_{\text{Gen.}}$  do
6     
$$\nu_j^{(k+1)} \leftarrow \frac{\nu_j^{(k+1)}}{\sum_{m=1}^{N_{\text{Gen.}}} \nu_m^{(k+1)}}$$

7 return  $\left\{ \nu_j^{(K)} \right\}_{j=1}^{N_{\text{Gen.}}}$ 

```

ignored iteration dependent covariance [679]. IBU is dominant in LHC analyses⁴ because it balances simplicity, interpretability and computational efficiency with accuracy.

⁴E.g., differential jet substructure measurements [93].

II.D.4 Tikhonov regularisation.

Tikhonov regularisation is a penalised least-squares minimisation method

$$\hat{\boldsymbol{\nu}} = \arg \min_{\boldsymbol{\nu}} [||\boldsymbol{\mu} - \mathbf{R}\boldsymbol{\nu}||^2 + \lambda ||\mathbf{L}(\boldsymbol{\nu} - \boldsymbol{\nu}_0)||^2] \quad (\text{II.22})$$

where $\mathbf{L}$ is typically the discrete curvature operator⁵ and $\boldsymbol{\mu}_0$ a prior estimate [596] that anchors solutions to MC predictions. L-curve optimisation balances the residual norm against the solution norm to choose λ [568]. The choice of λ sets the bias variance trade-off. $\lambda \rightarrow 0$ represents the high variance, low bias limit and $\lambda \rightarrow \infty (\implies \hat{\boldsymbol{\nu}} \rightarrow \boldsymbol{\nu}_0)$ represents the low variance, high bias limit. This method is implemented through the TUnfold package, which also provides automated λ tuning via global correlation minimisation [557]. However Tikhonov regularisation struggles with non-differentiable features like threshold effects due to biased curvature penalties, and requires *ad hoc* λ selection via L-curve curvature maximisation.

Note. The RooUnfold package [562] provides implementations of bin-bin-bin corrections, matrix inversion, IBU, SVD, and TUnfold.

II.D.5 Template fitting.

Template fitting is a method suitable in cases where $N_{\text{Sim.}} \gg N_{\text{Gen.}}$. In this case, one can construct detector level templates for each particle level bin.

$$\mu_i = \sum_{j=1}^{N_{\text{Gen.}}} R_{ij} \nu_j + b_i \quad (\text{II.23})$$

⁵E.g., discrete second derivatives.

with χ^2 minimisation:

$$\chi^2 = \sum_{i=1}^{N_{\text{Sim.}}} \frac{(\mu_i - \sum_j R_{ij} \nu_j - b_i)^2}{\sigma_i^2} \quad (\text{II.24})$$

The solution then is over-constrained through the use of $N_{\text{Sim.}}/N_{\text{Gen.}} \sim 2 - 3$ for stability [216]. Nuisance parameters are systematically modelled in this approach via template morphing [479].

Template fitting requires dense detector level binning, which inflates statistical uncertainties. This method is commonly used in Higgs coupling measurements where broad mass resolutions necessitate wide truth bins.

II.D.6 Regularised Poisson likelihood.

For low statistics regions, [330] advocates minimising

$$-\ln \mathcal{L}(x | z) + \lambda S(z). \quad (\text{II.25})$$

$S(z)$ penalises non monotonicity in sharply falling spectra. Using cubic B-splines with entropy regularisation, this method avoids binning artifacts through continuous representations [330]. However, it requires careful basis function placement to prevent endpoint spikes [226] and demands specialised optimization protocols (e.g., cooling schedules for λ [372]).

II.D.7 Summary.

Table II.1 summarises the properties and limitations of the methods discussed above. These limitations motivate the use of machine learning in unfolding, a transition

Method	MC dependence	Uncertainty propagation
Bin-by-bin	Extreme	Underestimated
Matrix inversion	None	Exact but unstable
IBU	Moderate	Partial
Tikhonov	Moderate	Full
Template fit	Low	Full
Regularised Poisson	Moderate	Full

Table II.1: Comparing the MC dependence and uncertainty propagation of traditional unfolding methods. Note that Tikhonov regularisation, template fitting and regularised Poisson likelihood are all variations on a similar unfolding strategy in that they all perform a binned fit of the forward folded data. They differ from each other in the choice of the regularisation scheme applied.

explored in subsequent sections. However, traditional methods remain indispensable for validation and low dimensional precision measurements where interpretability is crucial.

II.E Regularisation: Need, approaches, and limitations.

II.E.1	Necessity.	72
II.E.2	Limitations and practical challenges.	72
II.E.2.i	Subjectivity–objectivity trade-off.	72
II.E.2.ii	High dimensional regimes.	73
II.E.2.ii.a	Spectrum dependent biases.	74

The inherent ill posedness of unfolding necessitates regularisation to stabilise solutions against statistical fluctuations while preserving physical meaning. This section systematically examines the theoretical justification for regularisation, surveys dominant methodologies, and critically evaluates their strengths and limitations in high energy physics applications.

II.E.1 Necessity.

As discussed earlier, inverse problems in HEP exhibit pathological characteristics that demand regularisation. Regularisation counteracts unphysical variance by introducing prior knowledge about $p(z)$, typically favouring smoothness or similarity to Monte Carlo predictions. However, this unavoidably discards information—regularised solutions cannot resolve features finer than the detector resolution or distinguish theories predicting distributions within the regularisation bias [27, 727].

II.E.2 Limitations and practical challenges.

II.E.2.i Subjectivity–objectivity trade-off.

All regularisation methods inject subjective choices—smoothness scales, prior distributions, stopping criteria and so on—that bias results. This trade-off reveals a deeper epistemological issue: regularisation transforms the question from “what does the data show?” to “what does the data show given our assumptions about smoothness?”

While [467, 584] argue for publishing unregularised results alongside regularised ones, this approach, though transparent, may be insufficient. The unregularised results often contain artifacts that obscure physical interpretation, and a well chosen regularisation scheme can eliminate solutions that are patently unphysical.

A more nuanced approach would involve explicitly testing regularisation assumptions against physical models where possible, and developing domain specific regularisation schemes that incorporate known physical constraints rather than generic smoothness priors. This would shift the subjectivity from mathematical convenience to physics informed choices, making the trade-offs more scientifically meaningful rather than purely computational.

II.E.2.ii High dimensional regimes.

Traditional methods fail catastrophically in $d \gtrsim 4$ phase spaces and struggle even before that for multiple reasons. Binned approaches require n^d histogram bins struggling to effectively sample an increasingly sparse phase space, are shown in Section II.A.5.i. Global smoothness assumptions become untenable for multi-scale features [264, 328, 693] straining regularisation methods that rely on them.

As the number of dimensions increases, the binning also increasingly distorts error propagation. Bayesian credible intervals can exhibit poor frequentist coverage, as shown Fig. 4 of [630] and as argued in [503, 660], and correlated systematic uncertainties⁶ introduce non-convex likelihoods [162, 422] that strain inference methods.

⁶E.g., jet energy scale.

II.E.2.ii.a Spectrum dependent biases.

Sharply falling spectra⁷ exacerbate regularisation artifacts [330]. Entropic priors overweight high- z regions, distorting tails [370, 594, 644, 669], finite sample sizes truncate the measurable phase space, creating cutoff-induced spikes [49, 556], and curvature penalties conflict with natural spectral shapes, requiring physics informed regularisation strategies [190, 310, 406, 475].

Recent advances aim to mitigate these limitations in various ways. For example, adversarial regularisation involves training discriminators to enforce physical consistency rather than explicit smoothness [389]. Differentiable unfolding methods embed the detector response in neural networks enabling gradient based λ optimisation [345].

However, no universal solution exists, and the choice of regularisation must align with analysis specific priorities. As detector granularity increases, developing dimension agnostic regularisation schemes remains an open challenge requiring collaboration between statisticians and physicists to design methods that are statistically sound and suitable for the physics being analysed.

⁷E.g., proton momentum in [734].

II.F Unbinned methods: statistical considerations.

II.F.1	Principles and implementations.	76
II.F.1.i	Reweighting methods.	76
II.F.1.ii	Generative modelling.	78
II.F.2	Statistical considerations in unbinned regimes.	80
II.F.3	Limitations in complex phase spaces.	81
II.F.3.i	Model misspecification.	81

The evolution from binned to unbinned unfolding methodologies represents a paradigm shift in high energy physics, driven by the need to preserve fine grained kinematic information while managing the statistical and computational complexities of high dimensional phase spaces. This section systematically analyses the theoretical foundations, practical challenges, and performance trade-offs that this transition entails.

II.F.1 Principles and implementations.

The unbinned approach to unfolding represents a fundamental shift from traditional histogram based methods. Rather than discretising data into bins, unbinned unfolding preserves the complete kinematic information in the data by operating directly on individual event tuples $\{(z_1, x_1), \dots, (z_N, x_N)\}$ where each z_i represents a particle level event and x_i its corresponding detector level measurement. This approach eliminates information loss from binning and naturally handles high dimensional phase spaces where binning becomes prohibitive.

Modern unbinned methods leverage machine learning techniques, particularly those designed for density (ratio) estimation. These approaches fall into two main categories: reweighting methods and direct generative modelling.

II.F.1.i Reweighting methods.

The central goal of one class of unbinned unfolding methods is to estimate a reweighting function $w(z)$ that transforms Generation to match the underlying Truth

that was forward folded into the Data. Formally, this function satisfies

$$p(z) = w(z) \cdot q(z), \quad (\text{II.26})$$

where $q(z)$ is the distribution of Generation.

The most established approach in this category is OMNIFOLD [309], which employs an iterative procedure inspired by the Iterative Bayesian Unfolding (IBU) algorithm. OMNIFOLD alternates between two steps that estimate likelihood ratios using binary classifiers.

In each iteration k , OMNIFOLD first trains a classifier to distinguish between Data and Simulation yielding the ratio

$$\nu^{(k)}(x) = \frac{p(x)}{q^{(k)}(x)}, \quad (\text{II.27})$$

where $q^{(k)}(x)$ represents the distribution from Simulation weighted by the current iteration's weights.

This detector level ratio is then propagated back to particle level through the Monte Carlo pairing. For each Generation event z_i with corresponding Simulation event x_i , the weight update follows:

$$w^{(k+1)}(z_i) = w^{(k)}(z_i) \cdot \nu^{(k)}(x_i). \quad (\text{II.28})$$

The process, in principle, continues until convergence, at which point the final weights $w(z)$ provide the desired transformation from Generation to Truth.

The connection to likelihood ratio estimation can be explicitly proven—for a binary cross entropy optimised classifier f , the likelihood ratio between the distributions it classifies is

$$LR = \frac{f}{1 - f}. \quad (\text{II.29})$$

Algorithm 2 details the OMNIFOLD procedure in pseudocode.

Algorithm 2 OMNIFOLD

Input: Data samples $\{x_j^{\text{Data}}\}$,
 Gen.–Sim. pairs $\{(z_i, x_i)\}$,
 Initial gen level weights $w^{(0)}(z_i)$,
 # iterations K

Output: Unfolded weights $w^{(K)}(z_i)$

```

1 for  $k \leftarrow 1$  to  $K$  do
2   Detector level step:
3   Train classifier  $C_{\text{det}}^{(k)}$  on  $\{x_j^{\text{Data}}, +1\} \cup \{x_i, -1, w^{(k-1)}(z_i)\}$ 
4   Compute detector ratio
      
$$\nu^{(k)}(x) = \frac{C_{\text{det}}^{(k)}(x)}{1 - C_{\text{det}}^{(k)}(x)} \quad \forall x \quad (\text{II.30})$$


```

II.F.1.ii Generative modelling.

An alternative approach employs generative models to directly learn the conditional distribution $p(z \mid x)$. Among these, Conditional Invertible Neural Networks (cINNs) [267] offer a particularly elegant framework.

cINNs learn a diffeomorphic mapping

$$g_\theta : Z \rightarrow X \quad (\text{II.32})$$

between particle and detector levels, parameterised by neural network weights θ . The invertibility constraint ensures that for any detector level observation x , one

Algorithm 2 (continued)

```

5 for  $k \leftarrow 1$  to  $K$  do
6   Pull-back weights:
7   for all pairs  $(z_i, x_i)$  do
8      $\omega_i^{(k)} \leftarrow \nu^{(k)}(x_i)$ 
9   Particle-level step:
10  Train classifier  $C_{\text{part}}^{(k)}$  on  $\{z_i, +1, \omega_i^{(k)}\} \cup \{z_i, -1, w^{(k-1)}(z_i)\}$ 
11  Compute particle ratio
12    
$$\rho^{(k)}(z) = \frac{C_{\text{part}}^{(k)}(z)}{1 - C_{\text{part}}^{(k)}(z)} \quad \forall z \quad (\text{II.31})$$

13  Update weights:
14  for all generator events  $z_i$  do
15     $w^{(k)}(z_i) \leftarrow w^{(k-1)}(z_i) \rho^{(k)}(z_i)$ 
16 return  $\{w^{(K)}(z_i)\}$ 

```

can directly sample the corresponding particle level distribution through the inverse mapping. The key advantage lies in the tractable Jacobian computation,

$$p(z \mid x) = p(x \mid z) \frac{p(z)}{p(x)} = \left| \det \frac{\partial g_\theta}{\partial z} \right|^{-1} p(g_\theta^{-1}(x) \mid x), \quad (\text{II.33})$$

enabling direct sampling from the posterior $p(z \mid x)$ without iterative procedures [311].

Generative models, due to the much greater expressiveness that is necessary for them to model the functional representations needed, are susceptible to training

challenges such as mode collapse⁸. Ensuring stable convergence of generative models can get increasingly difficult for high dimensional problems.

More recently, Schrödinger Bridge Unfolding [30] has emerged as a method that frames unfolding as an optimal transport problem. This approach minimises the Kullback–Leibler divergence between joint distributions while constraining the marginal to match observed data.

$$\inf_{p(z,x)} D_{\text{KL}}\left(p(z,x) \parallel p_{\text{MC}}(z,x)\right) \text{ subject to } p(x) = p_{\text{Data}}(x). \quad (\text{II.34})$$

This formulation provides theoretical guarantees on the uniqueness of the solution and offers improved stability compared to purely generative approaches.

II.F.2 Statistical considerations in unbinned regimes.

Neural networks implicitly regularise via inductive controls.⁹ However, this introduces model dependent smoothing scales requiring careful validation against closure tests [570].

Reweighting based methods can propagate one part of the overall uncertainty through event weights.

$$\Sigma[O] = \sum_{i=1}^N w_i^2 O(z_i)^2 - \left(\sum_{i=1}^N w_i O(z_i) \right)^2, \quad (\text{II.35})$$

where $\Sigma(O)$ is the covariance of observable $O(z)$ [555]. While avoiding binning induced correlations, unbinned inference requires reconceptualising what the unbinned

⁸Normalising flows however, due to the structural constraints that invertibility imposes, are resistant to mode collapse.

⁹E.g., convolutional layers enforce translational symmetry in jet images [57].

equivalent of the covariance matrix would be in order to appropriately account for correlations in the unfolded data.

II.F.3 Limitations in complex phase spaces.

II.F.3.i Model misspecification.

Generative models assume that

$$q(z) > 0 \implies p(z) > 0, \quad (\text{II.36})$$

meaning that wherever Generation assigns positive probability density, $q(z) > 0$, the Truth distribution also has positive probability density, $p(z) > 0$. This assumption ensures that the generative model does not learn to generate events in regions of phase space where the true physics has zero probability. It is fundamental to the validity of generative unfolding approaches such as conditional Invertible Neural Networks (cINNs) and Variational Autoencoders (VAEs) [77]. Generative models learn to map from detector level measurements to particle level distributions by generating samples from the learned posterior. If the Generation were to sample in kinematically forbidden regions or unphysical phase space, the unfolded results would include spurious events that contaminate the measurement. The assumption requires that the generative model be properly constrained to respect the physical boundaries of the true distribution.

Violations of this assumption can occur when generative models, trained on imperfect or limited data, learn to extrapolate beyond the true physical support and generate samples in unphysical regions. This is particularly problematic given the

tendency of neural networks to produce overconfident predictions in regions with sparse training data.

Conversely though, enforcing this support containment requirement has practical implications for unfolding performance. Since the generative model can only learn to assign meaningful probability density to regions of phase space adequately represented in the training data, kinematic regimes that are poorly sampled or entirely absent from the Monte Carlo generation are rendered effectively invisible to the unfolding procedure regardless of their importance in the truth distribution. This limitation becomes particularly concerning for new physics searches, where signatures may manifest in previously unexplored regions of phase space. Novel phenomena occurring outside the support of the generative model where cannot be properly unfolded and may remain undetected [605].

The assumption thus places stringent requirements on Monte Carlo coverage and highlights the importance of comprehensive phase space sampling in training data preparation for generative unfolding methods. Hybrid approaches combining discriminative and generative components show promise for anomaly detection [364, 389].

II.G Evaluation metrics for unfolding.

II.G.1	Statistical metrics for evaluating point estimates.	84
II.G.1.i	Residual based metrics.	84
II.G.1.ii	Distributional distance metrics.	86
II.G.1.ii.a	Kullback–Leibler divergernce.	86
II.G.2	Uncertainty quantification metrics.	88
II.G.2.i	Pull distributions.	88
II.G.2.ii	Coverage properties.	91
II.G.2.iii	Variance and bias decomposition.	91
II.G.3	Evaluation of correlation structure.	92
II.G.3.i	Covariance matrix assessment.	93
II.G.3.ii	Event–to–event correlation metrics.	93
II.G.4	Method specific evaluation metrics.	96
II.G.4.i	Iterative methods.	96
II.G.4.ii	Bayesian methods.	97
II.G.5	Practical considerations.	97

The evaluation of unfolding methods presents unique challenges due to the ill posed nature of the inverse problem. While the goal of unfolding is conceptually straightforward, to recover the true particle level distribution from detector level observations, quantifying the success of this recovery requires careful consideration when the downstream use of the data is unknown.

Task specific evaluation becomes more tractable when clear physics objectives are defined. For example, when unfolding is performed to measure a specific parameter such as α_S , evaluation metrics can be designed to directly assess how well the unfolding procedure enables accurate parameter extraction. This section focuses on metrics and approaches for evaluating unfolding performance, considering both traditional binned techniques and modern unbinned methods, for general purpose unfolding applications where the downstream use is not known. Metrics for assessing accuracy, precision, and uncertainty quantification are assessed and practical considerations for their application in high energy physics analyses are discussed.

II.G.1 Statistical metrics for evaluating point estimates.

II.G.1.i Residual based metrics.

The most intuitive approach to evaluating an unfolding method is to compare the unfolded distribution to the true distribution when it is known.¹⁰ Simple residual based metrics quantify the difference between the estimated and true distributions.

¹⁰E.g., in simulation studies.

Definition 13 (Residuals). For binned methods, the *bin by bin residual* is defined as

$$\delta_i = \hat{t}_i - t_i \quad (\text{II.37})$$

where $\hat{t}_i$ is the unfolded count in bin i and t_i is the true count.

Various summary statistics of these residuals can be computed, including

Mean squared error

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (\hat{t}_i - t_i)^2, \quad (\text{II.38})$$

Root mean squared error

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{t}_i - t_i)^2}, \quad (\text{II.39})$$

Mean absolute error

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |\hat{t}_i - t_i|. \quad (\text{II.40})$$

The MSE might be preferred because it is easier to do calculus with, the RMSE preferred because it is expressed in the units of the original data rather than squared units, but they are manifestly functionally equivalent. The choice between the MSE and MAE depends on the specific requirements of the analysis. MSE more heavily penalises large errors due to the squaring operation, making it more sensitive to outliers and extreme residuals. MAE treats all errors equally regardless of magnitude and proves more robust to outliers, making it preferable when the unfolding procedure should not be dominated by a few problematic bins or events. The square function though is more amenable to calculus than the absolute value function. In high energy

physics applications, RMSE is the most commonly used metric, while MAE may be provided as valuable complementary information.

While these metrics are straightforward, they have limitations in the context of unfolding. In particular, they treat all bins equally, even though certain regions of phase space may be physically more significant than others. Additionally, these metrics do not account for the correlations between bins introduced by the unfolding process.

For unbinned methods, where the output is a set of weights or a continuous probability density, these metrics must be adapted. One approach is to bin the unbinned unfolded distributions and then apply the above metrics, though this introduces the binning artifacts that the unbinned method was designed to avoid.

II.G.1.ii Distributional distance metrics.

Given the limitations of simple residual metrics, distributional distance measures provide a more comprehensive assessment of unfolding performance. These metrics compare the entire unfolded distribution to the true distribution.

II.G.1.ii.a Kullback–Leibler divergence.

The *Kullback–Leibler (KL) divergence* [779] measures the information lost when using the unfolded distribution to approximate the true distribution.

Definition 14 (KL divergence).

$$D_{\text{KL}}(p||q) = \int p(x) \ln \frac{p(x)}{q(x)} dx, \quad (\text{II.41})$$

where p is the true distribution and q is the unfolded distribution.

While theoretically sound, KL divergence can be numerically unstable when the support of the distributions differs.

The *Vincze–Le Cam (VLC) divergence* [728, 739] is symmetric alternative to KL divergence that is both bounded and highly convex [341].

Definition 15 (VLC Divergence).

$$\Delta_{VLC}(p, q) = \frac{1}{2} \int \frac{(p(\lambda) - q(\lambda))^2}{p(\lambda) + q(\lambda)} d\lambda. \quad (\text{II.42})$$

This metric is particularly useful for comparing unfolding methods as it provides a balanced assessment of differences across the entire distribution and has been used in comparative analyses of various unfolding approaches [289, 309].

The *Wasserstein metric* [762], also known as the *Earth mover’s distance*,” [682] provides a measure of the minimum “work” required to transform one distribution into another.

Definition 16 (Wasserstein α –metric).

$$W_\alpha(p, q) = \left(\inf_{\gamma \in \Gamma(p, q)} \iint |x - y|^\alpha d\gamma(x, y) \right)^{\frac{1}{\alpha}}, \quad (\text{II.43})$$

where $\Gamma(p, q)$ is the set of all joint distributions with marginals p and q .

This metric accounts for both the magnitude and location of discrepancies between distributions, and is hence helpful for comparing distributions with sharp cutoffs or low statistics regions.

Table II.2 summarises the various distributional metrics and their relative strengths for unfolding evaluation.

Table II.2: Distributional distance metrics for unfolding evaluation

	Symmetric	Bounded	Support sensitivity	Compute Cost
D_{KL}	No	No	High	Low
Δ_{VLC}	Yes	Yes	Medium	Low
W_2	Yes	No	Low	High

II.G.2 Uncertainty quantification metrics.

Beyond point estimates, properly evaluating unfolding methods requires assessing the accuracy of their uncertainty estimates. This is particularly important in particle physics, where uncertainties propagate to downstream analyses such as parameter fitting.

II.G.2.i Pull distributions.

Pull distributions offer a rigorous way to evaluate the calibration of reported uncertainties.

Definition 17 (Pull distribution.). For a given unfolded bin or parameter θ , the *pull* is defined as

$$\text{Pull } \theta = \frac{\hat{\theta} - \theta_{\text{true}}}{\sigma_{\hat{\theta}}} \quad (\text{II.44})$$

where $\hat{\theta}$ is the unfolded estimate, θ_{true} is the true value, and $\sigma_{\hat{\theta}}$ is the reported uncertainty.

For a well calibrated method, the pull distribution across many pseudo-experiments should follow a standard normal distribution, $\mathcal{N}(0, 1)$. This follows from the definition of random variables. Since the pull is defined as the ratio of the bias to the estimated

uncertainty, for a properly calibrated unfolding method, two conditions must be satisfied simultaneously. First, the method must be unbiased, meaning that across many pseudo experiments, the estimated values should equal the true values on average, resulting in zero mean for the numerator. Second, the reported uncertainties must accurately reflect the actual variability of the estimates, such that the denominator correctly represents the standard deviation of the numerator.

When both conditions are met, the pull becomes a standardised random variable. For sufficiently large sample sizes, many unfolding estimators approach normality due to the Central Limit Theorem, since they often involve weighted sums or averages of many events. When the underlying estimator is approximately normally distributed, if it satisfies the two conditions above that set its mean and standard deviation to 0 and 1 respectively, the population of sample pulls must be distributed as a standard normal. This is formalised in Theorem 2.

Theorem 2 (Asymptotic standard normality of pulls). Let $\hat{\theta}_n$ be an estimator of a scalar parameter θ based on n independent events, and let σ_n denote the reported (1σ) uncertainty for $\hat{\theta}_n$. Assume:

Asymptotic unbiasedness: $\mathbb{E}[\hat{\theta}_n] - \theta \rightarrow 0$ as $n \rightarrow \infty$.

Variance calibration: $\sigma_n^2 / \text{Var}(\hat{\theta}_n) \xrightarrow{p} 1$.

Asymptotic normality: $\sqrt{n}(\hat{\theta}_n - \theta) \Rightarrow \mathcal{N}(0, \tau^2)$ for some $\tau^2 > 0$.

For the *pull*, P_n defined according to Definition 17,

$$\lim_{n \rightarrow \infty} P_n \sim \mathcal{N}(0, 1). \quad (\text{II.45})$$

That is to say, across pseudo-experiments, the empirical distribution of pulls converges to the standard normal law.

Proof of Theorem 2. By asymptotic normality,

$$\lim_{n \rightarrow \infty} \frac{\hat{\theta}_n - \theta}{\sqrt{\text{Var}(\hat{\theta}_n)}} \sim \mathcal{N}(0, 1). \quad (\text{II.46})$$

By variance calibration,

$$\lim_{n \rightarrow \infty} \frac{\sqrt{\text{Var}(\hat{\theta}_n)}}{\sigma_n} = 1. \quad (\text{II.47})$$

Hence, by Slutsky's theorem [793],

$$\lim_{n \rightarrow \infty} P_n = \lim_{n \rightarrow \infty} \frac{\hat{\theta}_n - \theta}{\sqrt{\text{Var}(\hat{\theta}_n)}} \cdot \lim_{n \rightarrow \infty} \frac{\sqrt{\text{Var}(\hat{\theta}_n)}}{\sigma_n} \sim \mathcal{N}(0, 1) \cdot 1 = \mathcal{N}(0, 1) \quad (\text{II.48})$$

□

Theorem 2 provides a rigorous theoretical foundation for the intuitive notion that a pull distribution with non-zero mean indicates systematic bias in the unfolding procedure. Deviations from unit standard deviation indicate either overestimation or underestimation of uncertainties. The theorem also provides a diagnostic for the quality of the uncertainty estimates, as a pull distribution with non-zero mean or non-unit standard deviation indicates that the uncertainty estimates are not accurate.

In the context of binned unfolding, pull distributions can be computed for each bin, while for unbinned methods, they can be applied to derived quantities or parameters of

interest. For Bayesian methods, pulls can be calculated using the mean and standard deviation of the posterior distribution [553].

II.G.2.ii Coverage properties.

Related to pulls, but more direct, is the evaluation of coverage properties of confidence intervals or credible intervals. For a nominal 68% confidence interval, approximately 68% of intervals computed across many pseudo-experiments should contain the true value. Systematic deviations from nominal coverage indicate issues with the uncertainty estimation. Coverage can be assessed through closure tests. These involve generating multiple datasets from a known truth, applying the unfolding procedure, and checking the fraction of times the true value falls within the reported confidence intervals.

Coverage plots plot the actual coverage versus the nominal coverage across different confidence levels. The example in Figure II.1 shows expected coverage properties for two different unfolding methods, Tikhonov regularisation and IBU, as a function of the regularisation parameter.

II.G.2.iii Variance and bias decomposition.

The total error of an unfolding method can be decomposed into bias and variance components,

$$\text{MSE}(\hat{t}) = \text{Bias}^2(\hat{t}) + \text{Var}(\hat{t}), \quad (\text{II.49})$$

where $\text{Bias}(\hat{t}) = \mathbb{E}[\hat{t}] - t$ and $\text{Var}(\hat{t}) = \mathbb{E}[(\hat{t} - \mathbb{E}[\hat{t}])^2]$.




Figure II.1: Coverage of 95% confidence intervals with Tikhonov regularisation (left) and IBU (right). The error bars are given by the 95% Clopper–Pearson intervals and the nominal confidence level is shown by the dotted line. When the regularisation is strong (fewer iterations in the case of IBU), both methods undercover substantially. Image credit: [439]

This decomposition is particularly valuable for understanding the trade-offs inherent in regularised unfolding methods, where stronger regularisation typically reduces variance at the expense of increased bias. Different applications might prioritise minimising one component over the other, making this decomposition important for method selection and hyperparameter choices.

II.G.3 Evaluation of correlation structure.

Traditional evaluation metrics often focus on marginal distributions, overlooking an important aspect of unfolding: the correlation structure between different bins or events. Properly accounting for these correlations is crucial for downstream analyses.

II.G.3.i Covariance matrix assessment.

For binned methods, the full covariance matrix of the unfolded distribution provides information about bin-to-bin correlations.

Definition 18 (Covariance matrix.). The *covariance matrix* Σ is defined as

$$\Sigma_{ij} := \mathbb{E}[(\hat{t}_i - \mathbb{E}[\hat{t}_i])(\hat{t}_j - \mathbb{E}[\hat{t}_j])]. \quad (\text{II.50})$$

Definition 19 (Correlation matrix.). A useful, related measure is the *correlation matrix*, ρ , defined as the normalised covariance matrix,

$$\rho_{ij} = \frac{\Sigma_{ij}}{\sqrt{\Sigma_{ii}\Sigma_{jj}}}, \quad (\text{II.51})$$

where Σ is the covariance matrix.

Comparing the correlation structure of the unfolded distribution to that of the true distribution¹¹ can reveal systematic distortions introduced by the unfolding procedure.

II.G.3.ii Event-to-event correlation metrics.

For unbinned methods event-to-event correlations in the unfolded weights can significantly impact downstream inference. These correlations can be quantified by studying the weight correlation as a function of distance. For any pair of events, one can compute the correlation between their weights as a function of their distance in feature space.

¹¹When known, e.g., in simulation studies.

One can also estimate the reduction in statistical power due to correlated weights

$$N_{\text{eff}} = \frac{(\sum_i w_i)^2}{\sum_{i,j} w_i w_j \boldsymbol{\rho}(w_i, w_j)} . \quad (\text{II.52})$$

In this equation, N_{eff} represents the effective number of independent observations when event weights are correlated, derived from the variance properties of weighted sums. The numerator, $(\sum w_i^2)$ represents the square of the total weighted sample size.

The denominator can be decomposed into of two components, which together account for the increased variance introduced by weight correlations. The first component, $\sum w_i^2$, captures the variance contribution from individual event weights, similar to the standard effective sample size formula for independent weighted observations. The second component, $2 \sum_{i < j} w_i w_j \boldsymbol{\rho}(w_i, w_j)$, accounts for the covariance contributions between all pairs of events, where the factor of 2 arises because each pair appears only once in the sum over $i < j$.

When weights are uncorrelated, $\boldsymbol{\rho}(w_i, w_j) = 0$, and the formula reduces to the familiar $N_{\text{eff}} = (\sum w_i)^2 / \sum w_i^2$. However, upon unfolding methods, nearby events in phase space often receive similar weight corrections¹², leading to correlations that increase the denominator and reduce the effective sample size. This reduction quantifies the loss of statistical power compared to an ideal scenario with independent weights.

The formula emerges from considering the variance of weighted observables. For a weighted sum, $S = \sum w_i x_i$, the variance is

$$\sigma_S^2 = \sum_i \sum_j w_i w_j \Sigma_{ij}, \quad (\text{II.53})$$

¹²As will be discussed in greater detail in Chapter VII.

which is the denominator of the Equation (II.52). The effective sample size is then defined as the ratio of the squared expectation to the variance, providing a measure of statistical efficiency that accounts for both weight magnitudes and their correlations. Formally,

Theorem 3 (Effective sample size with correlated event weights). Let $w_1, \dots, w_N$ be event weights, and let $\rho(w_i, w_j)$ denote the correlation coefficient between the contributions of events i and j to a given observable. The *effective sample size* accounting for correlations defined in Equation (II.52) computes the ratio of the squared total weight to the variance of the weighted sum of N correlated unit-variance contributions.

Proof of Theorem 3. Consider the weighted sum $S := \sum_{i=1}^N w_i X_i$ with mean μ and variance 1 for all i . The variance of S is

$$\sigma_S^2 = \sum_{i=1}^N w_i^2 \sigma_i^2 + 2 \sum_{1 \leq i < j \leq N} w_i w_j \Sigma_{ij} \quad (\text{II.54})$$

$$= \sum_{i=1}^N w_i^2 + 2 \sum_{1 \leq i < j \leq N} w_i w_j \rho(w_i, w_j). \quad (\text{II.55})$$

The definition of effective sample size equates σ_S^2 to the variance of N_{eff} independent, equally weighted observations of variance 1.

$$\sigma_S^2 = \frac{\left(\sum_{i=1}^N w_i \right)^2}{N_{\text{eff}}}. \quad (\text{II.56})$$

Solving for N_{eff} yields Equation (II.52). $\square$

If $\rho(w_i, w_j) = 0$ for all $i \neq j$, the denominator reduces to $\sum_i w_i^2$ and N_{eff} simplifies to the familiar independent-weights expression $N_{\text{eff}} = (\sum_i w_i)^2 / \sum_i w_i^2$. In unfolding




Figure II.2: Average weight correlation between two events as a function of the absolute distance between the events in the observable for Gaussian data unfolded using OMNiFOLD. Figure created by Owen Long [44].

applications, nearby events in phase space often have $\rho(w_i, w_j) > 0$, modifying the denominator of Equation (II.52) and thus reducing N_{eff} , which quantifies the corresponding change of statistical power.

Figure II.2 illustrates how event correlations typically decay with distance, with the correlation length scale increasing with detector resolution effects.

II.G.4 Method specific evaluation metrics.

II.G.4.i Iterative methods.

For iterative methods like Iterative Bayesian Unfolding (IBU) or OMNiFOLD, convergence behaviour provides important diagnostic information. To study the convergence behaviour, one plots metric values (e.g., χ^2 or NLL) as a function of

iteration number. One then compares unfolded distributions at different iterations to assess stability and analyse how the bias–variance trade-off evolves with iteration number.

II.G.4.ii Bayesian methods.

For Bayesian unfolding methods such as Fully Bayesian Unfolding (FBU) [548] or NEURAL POSTERIOR UNFOLDING (NPU) [91], additional posterior specific metrics are relevant.

One can compare detector level data to detector level predictions generated from the posterior. One assesses convergence using standard MCMC based diagnostics like the Gelman–Rubin statistic [713] or effective sample size. The width of the posterior allows one to evaluate the posterior uncertainty in relation to the true frequentist variance.

II.G.5 Practical considerations.

In general, when comparing different unfolding methods, a structured evaluation framework ensures fair and comprehensive assessment. Such a framework should consider the following aspects:

Computational efficiency: Measuring training time, inference time, and memory requirements.

Dimensionality scaling: Assessing how performance metrics change as the dimensionality of the problem increases.

Prior dependence: Evaluating robustness to different initial simulations.

Regularisation parameter sensitivity: Comparing how performance varies with changes in regularisation strength.

In real experimental settings where the truth is unknown, evaluation presents additional challenges that require alternative pragmatic approaches. For example, when the true distribution is unavailable, data splitting techniques can provide useful validation. The two most commonly used techniques are cross validation, which involves splitting the detector level data, unfolding one portion, then forward folding it, and comparing predictions against the held out portion; and bootstrapping, where one generates multiple resampled datasets to assess the unfolding procedure.

Closure tests involve applying the full analysis chain (forward model followed by unfolding) to a known input distribution. While not a direct evaluation of performance on real data, closure tests provide confidence in the methodology. The simplest kinds of closure tests involve applying detector simulation to a known particle level distribution, then unfolding the resulting detector level distribution and comparing with the original input. This procedure can then be modified by using a different particle level input than the one used to train the unfolding method, testing robustness to prior misspecification.

Evaluating how unfolding methods propagate systematic uncertainties is crucial for real world applications. One can test the sensitivity of the method to systematic uncertainties by applying variations to the response matrix based on known systematic uncertainties and assessing the impact on unfolded distributions. For methods that

support nuisance parameter profiling, evaluating how effectively nuisance parameters are profiled out is a valuable test for the effectiveness of the method.

Rigorous evaluation of unfolding methods requires a multifaceted approach that considers accuracy, uncertainty quantification, and computational performance. The metrics and frameworks presented in this section provide a comprehensive foundation for assessing both traditional and machine learning-based unfolding techniques. For binned methods, established metrics like χ^2 and coverage tests remain valuable, while for unbinned approaches, distributional metrics like Wasserstein distance and VLC divergence offer more appropriate evaluation.

Regardless of the method used, uncertainty calibration through pull distributions and correlation structure assessment are important to validate any measurement. As unfolding methods continue to evolve, particularly with the advent of machine learning approaches, evaluation metrics must adapt accordingly. The framework presented here is designed to be extensible, accommodating new methods and application domains while maintaining rigour and comparability

Chapter III

Machine learning for unfolding.

III.A	The emergence of machine learning in particle physics.	103
III.A.1	A paradigm shift in data analysis.	104
III.A.2	Evolution of ML applications in HEP.	105
III.A.3	Machine learning approaches to unfolding.	107
III.A.4	Early successes and current challenges.	109
III.B	Introduction to neural networks.	111
III.B.1	Essential concepts.	112
III.B.2	Neural networks as universal approximators.	119
III.B.3	Activation functions, optimisation, and regularisation.	121
III.C	Supervised learning approaches for unfolding.	128
III.C.1	Mathematical framework.	129
III.C.2	Statistical interpretation.	130
III.C.3	Types of supervised learning.	131

III.C.4	Learning and generalisation.	131
III.C.5	Evaluation.	132
III.C.6	Unfolding as a supervised learning problem.	132
III.C.7	Regularisation strategies.	134
III.C.8	Advantages and challenges.	135
III.D	Deep Learning Architectures	137
III.D.1	Convolutional neural networks.	138
III.D.2	Recurrent neural networks.	139
III.D.3	Graph neural networks.	140
III.D.4	Transformer based architectures.	141
III.D.5	Physics informed neural networks.	141
III.E	Modern machine learning frameworks.	147
III.E.1	Generative models.	149
III.E.2	Discriminative Models	160
III.E.3	Adversarial models.	171
III.F	Comparative analysis of ML based unfolding methods.	183
III.F.1	Method selection framework.	184
III.F.2	Performance trade-offs.	185
III.F.3	Open challenges and research frontiers.	188
III.G	Case studies: ML based unfolding in HEP analyses.	190
III.G.1	OMNIFOLD.	191
III.G.2	NEURAL POSTERIOR UNFOLDING.	192
III.G.3	Invertible neural networks for unfolding.	194

III.G.4	MOMENT UNFOLDING	195
III.G.5	Reweighting Adversarial Networks (RANs)	196
III.G.6	Schrödinger Bridge Unfolding.	197
III.G.7	Unified Neural Folding GAN	199
III.H	ML upstream and downstream of unfolding.	201
III.H.1	Uncertainty quantification in ML based unfolding.	202
III.H.2	Hybrid approaches and integration with simulation.	203
III.H.3	Equivariant networks for physics informed unfolding.	204
III.H.4	Comparison with traditional methods.	205
III.H.5	Implementation considerations from real experiments.	206

III.A The emergence of machine learning in particle physics.

III.A.1	A paradigm shift in data analysis.	104
III.A.2	Evolution of ML applications in HEP.	105
III.A.2.i	Classification tasks.	105
III.A.2.ii	Regression and anomaly detection.	106
III.A.2.iii	Generative models.	107
III.A.3	Machine learning approaches to unfolding.	107
III.A.4	Early successes and current challenges.	109

III.A.1 A paradigm shift in data analysis.

The past decade has witnessed a remarkable transformation in how particle physicists analyse data, driven by the adoption and adaptation of machine learning (ML) techniques. This revolution has been catalysed by the confluence of three factors: the growing volume and complexity of data from modern HEP experiments, the substantial increase in available computing resources and the rapid advancement of ML algorithms and frameworks. The LHC alone produces petabytes of data annually, with individual collisions generating thousands of particles across multiple detector subsystems, creating a rich, high dimensional dataset that traditional analysis methods struggle to fully exploit [7].

Particle physics presents unique analytical challenges that align particularly well with the strengths of modern machine learning approaches. The field routinely deals with high dimensional feature spaces, complex non-linear correlations between observables, rare signal processes buried within overwhelming backgrounds, and some of the largest scientific datasets in existence [428]. These characteristics create an ideal test bed for advanced ML techniques, which excel at discovering patterns in precisely such environments.

The relationship between particle physics and machine learning is not entirely new. Neural networks were first applied to high energy physics problems in the late 1980s and 1990s [613]. However, these early applications were limited by computational resources and algorithmic capabilities. The contemporary renaissance began around 2012 to 2014, coinciding with the broader deep learning revolution across computer science [410]. This timing allowed particle physicists to leverage developments in

computer vision, reinforcement learning, and generative modelling, adapting these techniques to the unique requirements of high energy physics.

III.A.2 Evolution of ML applications in HEP.

III.A.2.i Classification tasks.

The first wave of modern ML applications in particle physics focused predominantly on classification tasks, particularly signal/background discrimination and particle identification. These problems are naturally framed as binary or multiclass classification, making them accessible entry points for machine learning techniques. Boosted decision trees initially dominated these applications, particularly in the analysis chains that led to the Higgs boson discovery [509]. However, deep neural networks quickly demonstrated superior performance by automatically discovering complex patterns in high dimensional data, often outperforming carefully hand crafted physics inspired variables [269, 480, 507]. For example, neural networks trained on low level detector information have been shown to match or exceed the performance of approaches using physics motivated high level features for tasks like quark gluon discrimination [380] and top quark tagging [445].

The ATLAS and CMS experiments have now incorporated deep learning based taggers for identifying hadronically decaying W/Z bosons [319], top quarks [348], and Higgs bosons [203], significantly enhancing their sensitivity to new physics.

While early applications leveraged convolutional neural networks to exploit spatial correlations in calorimeter deposits [403, 468] and recurrent or graph neural networks to handle variable-length, unordered collections of particles [343, 347], transformer

architectures and foundation models have emerged as strong contenders for many classification tasks [63, 143, 251]. Transformers excel at modelling complex dependencies through self attention mechanisms, naturally processing sets of particles by capturing relationships between a span of pairs of particles simultaneously [455]. The Particle Transformer (ParT) architecture [251], specifically designed for HEP applications, has demonstrated superior performance on jet classification tasks by incorporating physics motivated constraints while leveraging the power of self-attention to identify relevant particle interactions and correlations [144].

III.A.2.ii Regression and anomaly detection.

As the field matured, ML applications expanded beyond classification to include regression tasks, such as energy calibration [184], pileup mitigation [402], and particle momentum reconstruction [369]. These applications require models to predict continuous quantities rather than discrete classes, introducing additional complexity but extending ML based methods to a larger class of problems. Simultaneously, unsupervised and weakly supervised learning techniques emerged as powerful tools for anomaly detection, identifying potential new physics without explicit models [98]. These methods include autoencoders that flag events with high reconstruction loss, density estimation techniques that identify low probability regions of phase space, and weakly supervised classifiers that can distinguish data mixtures without event by event labels.

III.A.2.iii Generative models.

Perhaps the most significant recent development in ML for HEP has been the adoption of generative models, which learn to produce samples from complex probability distributions. This capability is particularly valuable in particle physics, where accurate simulation is crucial but computationally expensive.

Generative adversarial networks (GANs) [516], variational autoencoders (VAEs) [375], and normalising flows [376] have all been successfully applied to particle physics problems. These models can generate synthetic collision events[248], simulate detector responses[206, 224], and model complex differential distributions,[30] often accelerating these processes by orders of magnitude compared to traditional Monte Carlo techniques.

The development of these generative models opened new possibilities for addressing inverse problems in particle physics, including the unfolding problem at the centre of this dissertation. By learning complex, high dimensional probability distributions directly from data, these methods offer a natural framework for tackling unfolding without the limitations of binning.

III.A.3 Machine learning approaches to unfolding.

The application of machine learning to unfolding represents a particularly promising frontier. Traditional unfolding methods face significant challenges when dealing with high dimensional spaces, correlated variables, and complex detector effects. Machine learning approaches offer several potential advantages compared to traditional statistical methods.

Neural networks excel at capturing patterns in high dimensional spaces, helping to mitigate the curse of dimensionality that plagues binned methods. While traditional approaches become computationally prohibitive beyond a few dimensions, ML based methods can effectively unfold many variables simultaneously [289]. This capability enables jointly unfolding multidifferential measurements that would be impractical with conventional techniques.

Many ML based unfolding methods operate directly on continuous distributions, eliminating binning bias and the need to predetermine bin boundaries. This approach preserves fine grained information that might otherwise be lost through discretisation, and it enables the extraction of arbitrary derived observables from the unfolded distribution [77].

Furthermore, the architecture and training procedure of neural networks can provide natural regularisation without requiring explicit constraints. Rather than manually tuning regularisation parameters as in traditional methods, ML approaches incorporate regularisation through network depth, network width, dropout rates and early stopping [698]. This implicit regularisation can adapt more naturally to the varying complexity across different regions of phase space.

Additionally, some ML methods can learn the complete detector to particle transformation in a single training procedure, avoiding the separate steps of response matrix estimation and inversion required by traditional approaches. This integrated optimisation may reduce the compounding of uncertainties that can occur when errors propagate through multiple discrete stages of conventional unfolding pipelines.

Once trained, many ML models also enable rapid inference on new data without retraining. This amortised approach is particularly valuable for unfolding tasks that need to be repeated with slight variations, such as systematic uncertainty studies or detector condition changes [315].

III.A.4 Early successes and current challenges.

Machine learning approaches to unfolding, such as OMNIFOLD, have demonstrated promising results on jet substructure measurements, showing improved performance in high dimensional spaces compared to traditional methods. OMNIFOLD has been successfully applied in experimental analyses at H1 [21, 105, 170, 221], STAR [132, 200], T2K [54], ATLAS [41, 103], CMS [41], and LHCb [169], validating its practical utility.

However, ML based unfolding methods also face significant challenges. Unlike traditional techniques with decades of validation, ML based unfolding is relatively new and requires careful validation to ensure physical consistency. Key concerns include proper uncertainty quantification, interpretability of the results, sensitivity to the training data, and potential biases introduced by the neural network architecture or training process [294]. Additionally, the field is still developing consensus on best practices for hyperparameter selection, architecture design, and evaluation metrics specific to unbinned unfolding. The balance between flexibility and stability remains an active area of research, with new approaches continuing to emerge.

As machine learning continues to evolve, both within particle physics and in the broader scientific community, further innovations in unfolding techniques are expected.

Promising directions include physics informed neural networks that incorporate known conservation laws or symmetries [383], uncertainty aware models that provide more reliable error estimates [393], and techniques that combine the strengths of traditional and ML based approaches [261].

The rapid progress in this field demonstrates the synergistic relationship between particle physics and machine learning. Particle physics provides challenging problems and unique datasets that drive methodological innovations; machine learning offers powerful tools that enable new measurements and insights. This dissertation builds upon this foundation, exploring novel machine learning approaches to improve the precision and scope of differential cross section measurements through novel unfolding techniques.

III.B Introduction to neural networks.

III.B.1	Essential concepts.	112
III.B.1.i	Neural network fundamentals.	112
III.B.1.ii	Forward propagation.	115
III.B.1.iii	Training objectives and loss functions.	116
III.B.1.iv	Backpropagation and parameter learning.	117
III.B.1.v	Batching and training dynamics.	118
III.B.2	Neural networks as universal approximators.	119
III.B.2.i	Relevance to physics applications.	119
III.B.3	Activation functions, optimisation, and regularisation.	121
III.B.3.i	Activation functions.	121
III.B.3.ii	Optimisation techniques.	123
III.B.3.ii.a	Gradient descent.	123
III.B.3.ii.b	Momentum based methods.	124
III.B.3.ii.c	Adaptive methods.	124
III.B.3.ii.d	ADAM.	125
III.B.3.iii	Regularisation techniques.	126

Neural networks are the fundamental building blocks of modern machine learning applications. This section provides a rigorous mathematical framework for understanding neural networks.

III.B.1 Essential concepts.

III.B.1.i Neural network fundamentals.

At its core, a neural network is a parametric function approximator loosely inspired by biological neural systems.

Definition 20 (Neuron). The fundamental unit of a neural network, a *neuron* or *node*, computes a weighted sum of its inputs followed by a non-linear transformation:

$$y = \sigma(w^T x + b), \quad (\text{III.1})$$

where $x \in \mathbb{R}^n$ is the input vector, $y \in \mathbb{R}^m$ is the output vector, $w \in \mathbb{R}^{n \times m}$ is a weight matrix, $b \in \mathbb{R}^m$ is a bias term, and $\sigma(\cdot)$ is a (typically non-linear) activation function.

Definition 21 (Depth of a neuron). The *depth* of a neuron is defined as the length of the longest path between the input and the neuron.

Definition 22 (Layer of a neuron). The set of all neurons at depth ℓ is called *layer* ℓ of the network.

Definition 23 (Width of a layer). The *width* of a layer is the number of neurons in that layer.

This organization of neurons into layers, creates a hierarchical structure.




Figure III.1: Diagram illustrating the key structural parameters of a neural network: depth (number of layers), width (units per layer), and their relationship to input, hidden, and output layers [180].

Definition 24 (Depth of a neural network). The *depth*, $d = \ell_{\max}$, of a neural network is defined as the depth of its deepest layer.

A graphical overview of these concepts is depicted in Figure III.1.

Definition 25 (Hidden layer). Any layer of a neural network except the input layer (often labelled layer 0) and the output layer ($\ell_{\max}$) is called a hidden layer.

Definition 26 (Deep neural network). A *deep neural network* is a neural network with depth, $d > 2$. That is to say standard deep feed forward neural network consists of an input layer, *more than one* hidden layers, and an output layer.




Figure III.2: Comparison between a shallow neural network (left) with a single hidden layer and a deep neural network (right) with multiple hidden layers [497].

Figure III.2a depicts a shallow neural network comprising a single hidden layer, while Figure III.2b shows a deep neural network with multiple hidden layers. The shallow architecture can model only a limited set of simple, low level features, whereas the deep architecture composes successive non-linear transformations to extract increasingly abstract representations. Increasing depth enhances representational capacity of the network and its capacity for hierarchical feature learning, albeit at the cost of greater computational complexity and potential optimization challenges.

For a network with d layers, at each layer l the neural network computes

$$y^{(l)} = \sigma^{(l)}(W^{(l)}y^{(l-1)} + b^{(l)}) \quad (\text{III.2})$$

where $y^{(l)}$ represents the output of layer l , $W^{(l)}$ is the weight matrix, $b^{(l)}$ is the bias vector, and $\sigma^{(l)}$ is the activation function of later l . By convention, $y^{(0)} = x$ is the input, and $y^{(d)} = y$ is the output.

This architecture enables neural networks to approximate arbitrarily complex functions, as formalised in Theorem 4, which loosely states that a deep neural

network of sufficient depth and width comprising a finite number of neurons can approximate any continuous function on compact subsets of $\mathbb{R}^n$ given certain mild conditions on the activation function. Importantly, neural networks achieve superior approximation rates compared to traditional methods like polynomials, particularly in high dimensional settings. While polynomial approximation suffers from the curse of dimensionality, requiring $\mathcal{O}(\varepsilon^{-d})$ parameters to achieve accuracy ε in d dimensions, neural networks can approximate smooth functions with rates of $\mathcal{O}(n^{-r/d})$ where n is the number of parameters and r is a smoothness parameter, often with significantly better constants and the ability to exploit compositional structure in the target function [409, 440].

III.B.1.ii Forward propagation.

Definition 27 (Forward propagation). *Forward propagation* is the process of computing the network's output given an input.

From its input $y^{(0)} = x$ the network sequentially computes the output of each layer

$$z^{(1)} = W^{(1)}y^{(0)} + b^{(1)} \quad (\text{III.3})$$

$$y^{(1)} = \sigma^{(1)}(z^{(1)}) \quad (\text{III.4})$$

$$\vdots$$

$$z^{(d)} = W^{(d)}y^{(d-1)} + b^{(d)} \quad (\text{III.5})$$

$$y^{(d)} = \sigma^{(d)}(z^{(d)}) . \quad (\text{III.6})$$

where $z^{(l)}$ represents the pre-activation values at layer l . Figure III.3 depicts the computation in the forward pass. Successive layer activations are computed by




Figure III.3: Illustration of the forward propagation process in a neural network, showing how the input vector is transformed by each layer's weights and biases and then passed through non-linear activation functions to yield the final output [552].

applying the weight matrices and bias vectors to the input, then passing the result through each layer's non-linear activation function until the final output is produced.

III.B.1.iii Training objectives and loss functions.

In order to quantify the accuracy with which a neural network f models a desired target function g for any particular input x , a suitable measure of the error between $f(x; W)$ and $g(x)$ is required.

Definition 28 (Divergence). The *divergence* $\text{Div}(f(x; W), g(x))$ is a continuous valued proxy for the error between $f(x; W)$ and $g(x)$.

The goal of training can then be defined to be to learn the values of W that minimise the expected divergence.¹ The optimal value of W is then given by

$$\widehat{W} = \arg \min_W \int_X \text{Div} \left(f(x; W), g(x) \right) p_X(x) \, dx. \quad (\text{III.7})$$

In practice however, g is often not known for all $x \in X$, it is only known at some subset $\{x_i\} \subseteq X$.

Definition 29 (Loss function). The *loss function* L is defined as the unbiased empirical average divergence between the neural network output and the target output over all training instances.

$$L(W) = \frac{1}{|X|} \sum_{i=1}^{|X|} \text{Div} \left(f(x_i; W), g(x_i) \right). \quad (\text{III.8})$$

This definition entails that the expected value of the loss function is the expected divergence, which is precisely what it means for the loss function to be an unbiased estimate of the expected divergence. However, this does not guarantee that minimising the loss function will minimise the expected divergence.

III.B.1.iv Backpropagation and parameter learning.

The training of neural networks revolves around finding the optimal values for weights and biases that minimise the loss function L . Backpropagation is the algorithm used to efficiently compute gradients of the loss function with respect to each network parameter. The core idea of backpropagation relies on the chain rule of calculus. For

¹For notational convenience, from here on out, b will be collapsed into W as its 0th column.

a loss function L and parameters $\{W^{(l)}\}_{l=1}^d$ the network computes

$$\frac{\partial L}{\partial W^{(l)}} = \frac{\partial L}{\partial z^{(l)}} \frac{\partial z^{(l)}}{\partial W^{(l)}} = \delta^{(l)} (y^{(l-1)})^T, \quad (\text{III.9})$$

where $\delta^{(l)} = \frac{\partial \mathcal{L}}{\partial z^{(l)}}$ is the error term for layer l .

These error terms are computed recursively, starting from the output layer.

$$\delta^{(L)} = \frac{\partial \mathcal{L}}{\partial y^{(L)}} \odot \sigma'^{(L)}(z^{(L)}), \quad (\text{III.10})$$

$$\delta^{(l)} = ((W^{(l+1)})^T \delta^{(l+1)}) \odot \sigma'^{(l)}(z^{(l)}), \quad (\text{III.11})$$

where $\odot$ denotes the Hadamard (elementwise) product and σ' is the derivative of the activation function. These derivatives are thus propagated backwards through the network; beginning at the output layer, and propagated to the input layer.

III.B.1.v Batching and training dynamics.

Definition 30 (Batch). A *mini-batch*, often simply called a *batch*, is a subset of the training data used to compute the gradient of the loss function. In practice, neural networks are trained using mini-batch gradient descent, where parameter updates are performed using gradients computed on subsets (batches) of the training data.

For a batch of size B , the loss is

$$L_{\text{batch}} = \frac{1}{B} \sum_{i=1}^B L(x_i, y_i, W). \quad (\text{III.12})$$

Training proceeds in *epochs*, where each epoch represents one complete pass through the training dataset. To prevent overfitting, a separate validation dataset is typically used to monitor performance, and techniques such as early stopping may be applied when validation performance plateaus or degrades.

III.B.2 Neural networks as universal approximators.

The power of neural networks in physics applications stems from their ability to approximate arbitrary functions without requiring exponentially many parameters, unlike traditional methods such as polynomial approximation which suffer from the curse of dimensionality. This property is formalised in Theorem 4, which has several variations but essentially states that a feed-forward network of sufficient, with a single hidden layer can approximate any continuous function on compact subsets of $\mathbb{R}^n$ to arbitrary precision, given sufficient (finite) width.

Theorem 4 (Universal Approximation Theorem).

$$\begin{aligned} \forall (\sigma, g, \epsilon) \in C^0(\mathbb{R}) \times C^0([0, 1]^n) \times (0, \infty) \exists M \in (0, \infty) \\ \left[\forall x \in \mathbb{R} |\sigma(x)| < M \implies \exists N \in \mathbb{N} \exists (c, b, W) \in \mathbb{R}^N \times \mathbb{R}^N \times \mathbb{R}^{N \times n} \right. \\ \left. \{ \forall z \in [0, 1]^n \|c^T \sigma(Wz + b) - g(z)\| < \epsilon \} \vee \{ \exists A \in \mathbb{R} \forall x \in \mathbb{R} \sigma(x) = A \} \right]. \end{aligned} \quad (\text{III.13})$$

This theorem was first proved by Cybenko for sigmoid activation functions [721] and later extended by Hornik et. al to include a broader class of activation functions [722]. Modern versions of the theorem have relaxed various assumptions and expanded the results to deeper networks [94].

III.B.2.i Relevance to physics applications.

The universal approximation property is particularly significant in physics. Many physical systems are governed by complex, non-linear differential equations that lack closed form solutions. Neural networks can approximate these solutions without

explicitly knowing the underlying equations, hence serving as excellent implicit simulators. Additionally, physics problems often involve high dimensional spaces.² Neural networks excel at learning in such spaces, where traditional numerical methods become computationally intractable. The universal approximation property makes neural networks well suited for these tasks, where the mappings to be learned are highly complex. When theoretical descriptions are incomplete, neural networks can discover patterns in data that suggest new physical models or refinements to existing ones.

Neural networks as universal approximators are also helpful for simulation based inference. Modern physics often relies on computer simulations rather than closed-form likelihood functions. Neural networks are well suited to approximate these implicit models for fast posterior inference.

While the Theorem 4 guarantees the existence of a neural network capable of approximating any continuous function, it does not provide guidance on architecture design or training procedures. Nor does the theorem provide any guarantees about the rate of convergence of training procedures. In practice, deeper networks with multiple hidden layers have provide faster convergence, and greater parameter efficiency as discussed above. However, the benefits of increasing depth are bounded, and decay exponentially, and hence the choice of architecture can involve carefully balancing the increased computational cost of a deeper network with the gains it provides.

Extensions of the theorem show that depth can exponentially reduce the number of neurons required to approximate certain function classes, explaining the empirical

²E.g., the phase space for jet physics.

success of deep learning in physics applications. These theoretical insights have motivated the development of specialised architectures tailored to specific physics problems, as will be discussed in subsequent sections.

III.B.3 Activation functions, optimisation, and regularisation.

III.B.3.i Activation functions.

Activation functions introduce non-linearity into neural networks, enabling them to learn complex patterns. Since the composition of an arbitrary number of affine functions is itself an affine function, without the non-linearity introduced by the activation function, a neural network would be incapable of modelling any non-linear function irrespective of its number of layers or parameters.

Listed below are some of the most commonly used activation functions in high energy physics applications.

Sigmoid:

$$\sigma(z) = \frac{1}{1 + e^{-z}} \quad (\text{III.14})$$

An activation function that outputs values between 0 and 1, it is most often used in problems involving binary classification. The sigmoid activation function has a rich and important history as one of the earliest activation functions to be used in machine learning applications. However, it suffers from vanishing gradient problems, where training instances far from its inflection point do not provide much information to the system. Additionally, when many sigmoid

functions are composed in deep networks, the resulting function can become very sharp and step-like, leading to training instabilities and poor gradient flow.

Hyperbolic Tangent (tanh):

$$\tanh(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}} \quad (\text{III.15})$$

The hyperbolic tangent function is very similar to sigmoid, except that it outputs values between -1 and 1 . It is preferred in some applications because it is zero-centered, which can help in training dynamics, but it also suffers from the same vanishing gradient problems as the sigmoid.

Rectified Linear Unit (ReLU):

$$\text{ReLU}(z) = \max(0, z) \quad (\text{III.16})$$

Simple, and incredibly computationally efficient, ReLU has become one of the most frequently used activation functions both within HEP and in the AI/ML space more broadly. It addresses the vanishing gradient problem for positive inputs. However, since it is simply zero for half its support, it can suffer from the “dying ReLU” problem where units can become permanently inactive.

Leaky ReLU:

$$\text{LeakyReLU}(z) = \max(\alpha z, z) \quad (\text{III.17})$$

$\alpha < 1$ is a small positive constant. LeakyReLU is a modification of ReLU to allow small negative valued outputs [538], which helps mitigate the dying ReLU problem.

Some other variants of ReLU include Parametric ReLU (PReLU) [489] where α is learnable, Exponential Linear Unit (ELU) [463] which provides smooth transitions for negative inputs, and Gaussian Error Linear Unit (GELU) [183] which has become popular in modern transformer architectures.

Softmax:

$$\text{softmax}(z_i) = \frac{e^{z_i}}{\sum_j e^{z_j}} \quad (\text{III.18})$$

Softmax is most often used for multiclass classification output layers because it produces a probability distribution over classes. The softmax reduces to the sigmoid when there only two classes, rendering binary classification simply a special case of multiclass classification.

These are but a few examples of the most common activation functions. The choice of activation function significantly impacts network performance, convergence speed, and the occurrence of issues like vanishing or exploding gradients.

III.B.3.ii Optimisation techniques.

As discussed, neural network training is fundamentally an optimisation problem. Various algorithms have been developed to find optimal parameters efficiently.

III.B.3.ii.a Gradient descent.

Gradient descent is the oldest and simplest approach to optimisation.

It involves updating parameters in the direction of the negative gradient.

$$W_{t+1} = W_t - \eta \nabla_W L(W_t) , \quad (\text{III.19})$$

where η is the learning rate.

While conceptually simple, gradient descent can be slow to converge and may get trapped in suboptimal solutions such as local minima.

Definition 31 (Stochastic gradient descent (SGD)). *Stochastic gradient descent* is an adaptation of gradient descent that updates parameters using gradients computed on mini-batches, introducing noise that can help escape local minima.

$$W_{t+1} = W_t - \eta \nabla L_{\text{batch}}(W_t) . \quad (\text{III.20})$$

III.B.3.ii.b Momentum based methods.

A significant innovation in loss function optimisation was the invention of momentum based methods. These methods accelerate convergence by accumulating a “velocity” vector in directions of persistent reduction in the loss function.

$$v_{t+1} = \gamma v_t + \eta \nabla L(W_t) \quad (\text{III.21})$$

$$W_{t+1} = W_t - v_{t+1} , \quad (\text{III.22})$$

where γ referred to as the momentum coefficient. In HEP applications, a typical value of γ is 0.9.

III.B.3.ii.c Adaptive methods.

Definition 32 (Adaptive methods.). *Adaptive methods* are optimisation techniques that vary the learning rate η over the course of the training rather than keeping it constant, and allow different learning rates for different dimensions of the input, rather than the traditional scalar learning rate that was constant across dimensions.

One of the earliest adaptive methods proposed was ADAGRAD [572], which adapts learning rates per parameter based on historical gradients. ADAGRAD was however limited by its aggressive learning rate reduction that slowed down training considerably.

ADAGRAD was soon replaced in almost all applications by *RMSProp* [552], an unpublished algorithm proposed by Hinton in course lectures [448], which addresses ADAGRAD's aggressive learning rate reduction. Today, RMSProp is a widely used optimiser in ML applications, and is especially effective at stabilising training when gradients are sparse and/or the loss function is dominated by saddle points.

III.B.3.ii.d ADAM.

ADAM is an optimisation technique that is both momentum based and adaptive. It involves combining the gradient update steps of momentum based methods and RMSProp, and is currently the most widely used optimiser in ML applications [436]. ADAM's optimisation procedure is can be described as

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) \nabla_W L(W_t) , \quad (\text{III.23})$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) (\nabla_W L(W_t))^2 , \quad (\text{III.24})$$

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t} , \quad (\text{III.25})$$

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^t} , \quad (\text{III.26})$$

$$W_{t+1} = W_t - \frac{\eta}{\sqrt{\hat{v}_t} + \epsilon} \hat{m}_t . \quad (\text{III.27})$$

Typical values in HEP applications are $\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^8$.

Definition 33 (Learning rate scheduling). *Learning rate scheduling* is the process of dynamically adjusting the learning rate during training.

Some of the more common variants are step decay,³ exponential decay,⁴ and cosine annealing.⁵

III.B.3.iii Regularisation techniques.

To improve stabilise the training, prevent overfitting, and increase generalisation, various regularisation methods are employed.

L^p Regularisation involves adding penalty terms to the loss function to penalise large weights.

$$\mathcal{L}_{\text{reg}} = \mathcal{L} + \lambda_p \|W\|^p. \quad (\text{III.28})$$

The most common values of p are 1 and 2. For certain applications, a hybrid $L1 - L2$ regularisation may also be appropriate [471].

Dropout is a form of regularisation that involves randomly setting a fraction of neuron outputs to zero during training,

$$y_{\text{dropout}}^{(l)} = m \odot y^{(l)}, \quad (\text{III.29})$$

where m is a binary mask with entries drawn from a Bernoulli distribution with parameter $p \in [0, 1]$.

³Reducing learning rate by a fixed factor after a set number of epochs.

⁴ $\eta_t = \eta_0 e^{-kt}$.

⁵ $\eta_t = \eta_{\min} + \frac{1}{2}(\eta_{\max} - \eta_{\min})(1 + \cos(\frac{t\pi}{T}))$.

Batch normalisation is the process of normalising the inputs of a layer to have zero mean and unit variance.

$$\hat{y}^{(l)} = \frac{y^{(l)} - \mu_B}{\sqrt{\sigma_B^2 + \epsilon}} , \quad (\text{III.30})$$

$$z_{\text{BN}}^{(l+1)} = W^T \hat{y}^{(l)} + b , \quad (\text{III.31})$$

where μ_B and σ_B^2 are the batch mean and variance.

These regularisation techniques, combined with appropriate optimisation strategies and activation functions, form the foundation for effectively training neural networks for high energy physics applications. The principles outlined here will be applied in subsequent sections focusing on specific neural network architectures and their applications to unfolding problems.

III.C Supervised learning approaches for unfolding.

III.C.1	Mathematical framework.	129
III.C.2	Statistical interpretation.	130
III.C.3	Types of supervised learning.	131
III.C.4	Learning and generalisation.	131
III.C.5	Evaluation.	132
III.C.6	Unfolding as a supervised learning problem.	132
III.C.6.i	Classification based approaches	132
III.C.6.ii	Regression based calibration	133
III.C.6.ii.a	Deterministic regression	133
III.C.6.ii.b	Probabilistic regression.	134
III.C.7	Regularisation strategies.	134
III.C.8	Advantages and challenges.	135

Supervised learning provides a powerful framework for addressing the unfolding problem in particle physics. Unlike traditional matrix inversion methods, supervised learning approaches can leverage the flexibility and expressiveness of machine learning models to handle high dimensional data and complex detector responses without requiring explicit binning schemes.

III.C.1 Mathematical framework.

Before exploring specific applications to unfolding, it is essential to establish what supervised learning is from a mathematical and statistical perspective. Supervised learning represents one of the primary paradigms in machine learning where an algorithm learns a mapping from inputs to outputs using labelled training data.

Supervised learning involves using a dataset, $D = \{(x_i, y_i)\}_{i=1}^N$ consisting of N input–output pairs. Each input $x_i \in X$ is a feature vector, and each output $y_i \in Y$ is a label or target value. The objective is to learn a function $f : X \rightarrow Y$ that approximates the true relationship between inputs and outputs.

Definition 34 (Risk function.). For a space of functions f_θ parametrised by θ the *risk function* is defined as

$$R(f_\theta) = \mathbb{E}_{(x,y) \sim P(X,Y)}[\text{Div}(y, f_\theta(x))] , \quad (\text{III.32})$$

where Div is the divergence between the predicted output $f_\theta(x)$ and the true output y , and $P(X, Y)$ is the joint probability distribution of inputs and outputs.

Supervised learning aims to learn the parameters θ_* that minimise the expected value of the risk function.

Definition 35 (Empirical risk). The *empirical* risk as an approximation of the risk function defined as

$$\hat{R}(f_\theta) = \frac{1}{N} \sum_{i=1}^n L(y_i, f_\theta(x_i)) , \quad (\text{III.33})$$

where L is the loss function, the empirical expected divergence.

Since the true distribution $p(X, Y)$ is unknown, one can only minimise the empirical risk. This empirical risk minimisation is accomplished through the optimisation of the loss function using techniques described in Section III.B.3.ii.

III.C.2 Statistical interpretation.

From a statistical standpoint, supervised learning can be viewed as conditional density estimation. For classification problems, the network estimates $\mathbb{P}(Y | X)$, the probability of a particular class given the input features. For regression problems, it estimates the conditional expectation $\mathbb{E}[Y | X]$ or the full conditional distribution $p(Y | X)$.

The choice of divergence function directly relates to the statistical assumptions being made. The mean squared error,

$$\text{Div}_{MSE}(x) = (y - f_\theta(x))^2, \quad (\text{III.34})$$

corresponds to maximum likelihood estimation under Gaussian noise assumptions. The binary cross-entropy,

$$\text{Div}_{BCE}(x) = -y \ln f_\theta(x) - (1 - y) \ln(1 - f_\theta(x)), \quad (\text{III.35})$$

corresponds to maximum likelihood estimation for categorical distributions. Quantile divergences correspond to estimating specific quantiles of the conditional distribution

III.C.3 Types of supervised learning.

Supervised learning problems broadly fall into two categories:

Classification: When the output space Y is discrete (categorical), the task is to predict the class label of a given input. Common loss functions include the cross-entropy loss and the hinge loss [652].

Regression: When the output space Y is continuous, the task is to predict a real-valued output. Common loss functions include the mean squared error, the mean absolute error, and Huber loss [766].

For unfolding problems, both paradigms can be relevant. Classification approaches often estimate density ratios between distributions, while regression approaches directly estimate mappings between detector level and particle level observables.

III.C.4 Learning and generalisation.

A fundamental aspect of supervised learning is generalisation i.e. the ability of the model to perform well on unseen data. This requires balancing two competing concerns: underfitting, when the model is too simple to capture the underlying structure of the data, and overfitting, when the model captures noise in the training data rather than the underlying structure. Regularisation techniques help prevent overfitting by constraining the complexity of the model.

III.C.5 Evaluation.

Supervised learning models are typically evaluated by measuring their performance on a separate test set not used during training. Common evaluation metrics include accuracy, precision, recall, F1-scores, and ROC curves for classification problems; mean squared error, mean absolute error, and R-squared values for regression problems. In the context of unfolding, additional evaluation criteria can include physical consistency, preservation of known symmetries and robustness to statistical fluctuations.

III.C.6 Unfolding as a supervised learning problem.

The unfolding objective of recovering the mapping from detector level distributions to particle level truth distributions can be naturally framed as a supervised learning task where the model learns from paired Generation and Simulation events. The supervised learning approach to unfolding typically uses paired data (z, x) , where z represents particle level quantities and x represents detector level observations. The forward problem (detector simulation) maps $z \mapsto x$ through some response function, while unfolding attempts to estimate the inverse mapping.

III.C.6.i Classification based approaches

A prominent class of supervised learning methods for unfolding comprises methods that use binary classification as their foundation. These approaches leverage the ability of classifiers to estimate likelihood ratios between distributions.

Of this class, OMNIFOLD, described in Section II.F.1.i is perhaps the most widely adopted classification based unfolding method and has been applied to several experi-

mental measurements. It uses an iterative procedure inspired by Iterative Bayesian Unfolding (IBU), generalising it to the unbinned case. At each iteration, a pair of classifiers are trained in order to estimate likelihood ratios. OMNIFOLD effectively retains the full phase space information without requiring dimensionality reduction.

III.C.6.ii Regression based calibration

Calibration corrects individual detector level observables x_i on an event by event basis. Unfolding, in contrast, tackles the ill posed inverse problem of inferring the Truth *distribution* $p(z)$ from the smeared Data distribution $p(x)$. Regression models that remap X are therefore calibration tools rather than unfolding tools *per se*; they do not, by themselves, invert the response matrix or recover $p(z)$.

It is still worthwhile to mention these regression based methods at this point, because they can reduce the bias and variance of traditional unfolding by feeding better behaved inputs into response matrix construction, and because probabilistic regressors provide conditional densities that can be integrated into likelihood based or Bayesian unfolding frameworks.

III.C.6.ii.a Deterministic regression

Early neural network approaches attempted to learn a deterministic function f that predicts a point estimate. Although simple, such models tend to underestimate uncertainties, leading to biased predictions when used naively.

III.C.6.ii.b Probabilistic regression.

Modern methods replace point estimates with conditional density estimators that yield conditional probability density functions. Mixture Density Networks [58, 70, 71, 218, 272], Gaussian process regressors [185], and normalising flows [112, 212] have each been explored for this purpose, providing a principled propagation of detector smearing into unfolding stage uncertainties.

In practice, conditional density can be sampled to build an event wise response ensemble, or inserted as the likelihood term in Bayesian unfolding. Nevertheless, the global inverse problem remains one must still correct for resolution⁶ regression alone cannot address.

III.C.7 Regularisation strategies.

Supervised learning approaches to unfolding must address the inherently ill posed nature of the inverse problem. This requires effective regularisation strategies to prevent the amplification of statistical fluctuations.

Neural network architectures themselves provide a form of implicit regularisation. The choice of network depth, width, activation functions, and training protocols significantly impacts the solution's smoothness properties. Convolutional layers, for instance, can encode physical priors such as translational invariance, constraining the space of possible solutions.

⁶Efficiency and acceptance effects also need to be corrected for, but they are typically not encapsulated in the response matrix for unfolding.

In iterative approaches like OMNIFOLD, early stopping serves as a form of regularisation. By halting the iteration process before full convergence, the method prevents overfitting to statistical fluctuations in the data.

Domain specific physics knowledge can also be incorporated into the learning process through custom loss functions, model architectures, or constraints. For example, conservation laws, symmetries, or known theoretical behaviours can be encoded to guide the supervised learning process toward physically plausible solutions.

III.C.8 Advantages and challenges.

Supervised learning approaches to unfolding can operate directly on unbinned data, avoiding information loss and artifacts from binning. They scale better to high dimensional observables and can effectively capture complex non-linear detector responses. Additionally, they provide flexible regularisation through model architecture and training.

However, supervised approaches also face challenges. They require large amounts of paired simulated training data with accurate detector modelling, where both detector level and particle level information must be available for each event. Unless designed carefully for interpretability, the black box nature of neural networks can obscure the physical interpretation of the results. For unbinned unfolding methods, uncertainty quantification remains challenging, particularly in capturing the correlations between events when computing confidence intervals and validating the unfolding performance requires careful cross-checks and closure tests.

Despite these challenges, supervised learning approaches represent a significant advancement in unfolding methodology, enabling measurements that would be impractical with traditional binned techniques, in high dimensional phase spaces relevant to modern particle physics analyses.

III.D Deep Learning Architectures

III.D.1	Convolutional neural networks.	138
III.D.2	Recurrent neural networks.	139
III.D.3	Graph neural networks.	140
III.D.4	Transformer based architectures.	141
III.D.5	Physics informed neural networks.	141
III.D.5.i	Energy Flow Networks and Particle Flow Networks.	142
III.D.5.i.a	Relationship to other architectures. . . .	144
III.D.5.i.a.1	Deep sets.	144
III.D.5.i.a.2	GNNs and transformers	145

High energy physics (HEP) presents unique data analysis challenges that have inspired the adoption and adaptation of specialised deep learning architectures. These architectures are designed to handle the distinctive properties of particle physics data, including high dimensionality, sparsity, permutation invariance, and complex correlations. This section explores the most prominent deep learning architectures employed in HEP applications, with particular emphasis on their relevance to unfolding tasks.

Note. The architectures described in this section, namely CNNs, RNNs, GNNs, Transformers, EFNs and PFNs, are not used in the methods described in Chapters IV to VI and VIII. They are discussed here in order to provide a more comprehensive view of the array of methods employed in the field.

III.D.1 Convolutional neural networks.

Convolutional neural networks (CNNs) have found extensive application in HEP, particularly for analysing data with inherent spatial or geometric structure. Originally developed for image recognition tasks, CNNs are well-suited for handling detector data that can be represented in image-like formats.

A prominent application of CNNs in HEP is in the analysis of “jet images” — representations of energy deposits in calorimeters. Treating jet constituents as “pixels” in a coordinate system of pseudorapidity (η) and azimuthal angle (ϕ) creates image-like representations that CNNs can process effectively. For unfolding applications, CNNs can learn complex mappings between detector level jet images and their corresponding particle level representations. The translation invariance property of CNNs naturally

encodes physical symmetries, providing a form of implicit regularisation that is beneficial for the ill posed unfolding problem [304].

CNNs have also been successfully applied to model energy depositions in calorimeters [359]. These architectures can capture the spatial correlations in shower development patterns, enabling more accurate unfolding of particle energies and types from detector responses. Three dimensional CNNs have been employed to handle the full volumetric nature of calorimeter data, treating detector cells as voxels in a 3D space [147]. These approaches have shown promising results in reconstructing particle properties from complex shower patterns [6, 387].

III.D.2 Recurrent neural networks.

Recurrent neural networks (RNNs), particularly long short-term memory (LSTM) networks and gated recurrent units (GRUs), have been applied to sequential aspects of HEP data analysis. In many HEP applications, particles can be naturally ordered by properties such as transverse momentum or angular distance from a reference axis. RNNs are well suited to process these ordered sequences while maintaining information about dependencies between particles.

For unfolding tasks involving sequential data, RNNs can model the mapping between detector level sequences and their particle level counterparts, capturing complex ordered dependencies that might be lost in other architectures. RNNs are also applicable to time-dependent detector responses or beam conditions that evolve over time. By incorporating time information, these models could help unfold distributions that are affected by time-varying detector effects.

III.D.3 Graph neural networks.

Graph neural networks (GNNs) have emerged as one of the most promising architectural paradigms for HEP applications. These networks explicitly model relationships between particles or detector elements as graphs, with nodes representing individual entities and edges representing their interactions or relationships.

Particle interactions naturally form graph-like structures, where each particle can be considered a node with properties (features) such as momentum, energy, and charge. GNNs can process these “particle clouds” directly, without requiring conversion to image or sequence formats. The key advantage of GNNs for unfolding is their ability to preserve the permutation invariance of particle collections; the physical properties of a jet should not depend on the arbitrary order in which particles are processed. By operating directly on graphs, GNNs respect this physical constraint.

GNNs have also been applied to model the complex geometry of particle detectors, where detector elements can be represented as nodes and their physical or electronic connections as edges [256, 347]. This approach enables more accurate modelling of detector responses, which is crucial for reliable unfolding. A specialised class of GNNs called message passing neural networks (MPNNs) show particular promise in physics applications. [145] These networks iteratively update node representations by passing information along edges, mimicking the way physical interactions propagate through a system.

III.D.4 Transformer based architectures.

Transformer architectures, which have revolutionised natural language processing, are increasingly being applied to HEP problems due to their ability to model complex dependencies through self-attention mechanisms [141]. Transformers can naturally process sets of particles by using self-attention to capture relationships between pairs and groups of particles. This makes them particularly well suited for unfolding tasks involving collections of particles with complex interrelationships. The Particle Transformer (ParT) architecture,[251] specifically designed for HEP applications, incorporates physics motivated constraints and has demonstrated strong performance on jet classification tasks [20].

Similar principles can be applied to develop transformer based unfolding methods. OMNILEARN [62], a foundation model built on the transformer architecture, demonstrates the ability of such models to improve the accuracy, precision, or speed of a wide range of physics tasks.

The attention mechanisms in transformers serve to highlight which detector level features are most informative for reconstructing particle level properties. This interpretability is valuable for understanding the unfolding process and identifying potential biases or limitations.

III.D.5 Physics informed neural networks.

A growing trend in HEP applications is the development of physics informed neural networks.

Definition 36 (Physics informed neural network). A *physics informed neural network* is a neural network that explicitly incorporates domain knowledge about physical laws, conservation principles, and symmetries.

Networks that preserve known physical symmetries, such as Lorentz invariance or gauge symmetry, can provide more physically plausible unfolding results. Architectures such as Lorentz Group Equivariant Networks[313] ensure that the neural network respects these fundamental symmetries by construction. For problems where energy conservation is essential, specialised architectures can enforce this constraint by design. These networks ensure that the total energy is preserved between detector level and particle level representations, reducing the space of possible solutions and improving physical consistency. Beyond specialised architectures, physical constraints can be incorporated into the loss function or network design. For example, constraints on momentum conservation, charge conservation, or known physical boundaries of observables can guide the network toward physically plausible solutions.

III.D.5.i Energy Flow Networks and Particle Flow Networks.

Energy Flow Networks (EFNs) and Particle Flow Networks (PFNs) [379] represent specialised deep learning architectures developed specifically for jet physics in HEP applications. These networks leverage the fundamental principle that jets can be represented as collections of particles, each characterised by its energy or transverse momentum, and angular coordinates.

The theoretical motivation for EFNs stems from the importance of infrared and collinear (IRC) safety in jet physics. While traditional approaches like Energy Flow

Polynomials [404] provide a complete basis of handcrafted IRC-safe observables, EFNs offer a data driven alternative that learns optimal IRC-safe representations directly from the energy flow of jets. PFNs extend this framework by incorporating additional particle level information beyond geometric coordinates, trading theoretical constraints for increased flexibility.

Both architectures share a common structural feature that jet observables can be effectively decomposed into per-particle features and their aggregated relationships. This decomposition naturally aligns with how QCD radiation patterns manifest in jet structure. The networks implement this as follows:

1. **Per-particle feature extraction:** Each particle is processed independently through a shared neural network, Φ , referred to the particle embedding network, mapping individual particle features to a latent representation.
2. **Permutation invariant aggregation:** The individual particle embeddings are combined through a permutation invariant operation. For EFNs, this is specifically a weighted sum using particle energies, ensuring IRC safety. PFNs typically use unweighted summation.
3. **Global processing:** The aggregated representation is processed by another neural network F , called the latent space network to produce the final output.

A primary distinction between these architectures is their treatment of IRC safety. EFNs exclusively process geometric information ⁷ with energy-weighted aggregation, ensuring that the learned representations respect IRC safety by construction. This

⁷ (η, ϕ) coordinates.

design choice guarantees that EFNs can only learn observables that are theoretically well defined in QCD.

PFNs, conversely, can incorporate additional per-particle features such as particle type, charge, or identification probabilities. While this makes PFNs more expressive for certain tasks, they do not provide IRC safety guarantees.

For unfolding applications, this distinction has important implications. EFNs' inbuilt IRC safety ensures that unfolded distributions preserve fundamental theoretical properties of QCD, providing implicit regularisation aligned with physical principles. This could eliminate certain unphysical solutions and help constrain the unfolding problem. PFNs offer complementary advantages when IRC-unsafe features, such as particle identification or flavour tagging information might be relevant, or even central, to the unfolding task.

Both architectures naturally handle the variable multiplicity of jets without requiring padding or truncation, making them well suited for realistic jet data. In unfolding contexts, they can be employed as components within iterative approaches (like OMNIFOLD).

III.D.5.i.a Relationship to other architectures.

III.D.5.i.a.1 Deep sets.

From a broader machine learning perspective, EFNs represent specialised implementations of deep sets [199] for processing permutation invariant sets of varying sizes. This connection to established theoretical frameworks facilitates both the analysis of these architectures and the development of future innovations in physics aware neural network design.

III.D.5.i.a.2 GNNs and transformers

PFNs can also be interpreted as a special case of more general architectures like graph neural networks (GNNs) and transformers, which have gained prominence in particle physics applications. In the GNN interpretation, a jet can be represented as a fully connected graph where particles are nodes and their relationships are edges. While standard PFNs implicitly assume equal interactions between all particles through their summation pooling, GNNs explicitly model particle–particle interactions through message passing. This allows GNNs to learn which particle relationships are most relevant for a given task, potentially capturing substructure patterns that simple summation might miss.

The connection to transformers is even more direct. The self–attention mechanism can be viewed as learning dynamic, data–dependent weights for particle interactions. Recent architectures like the Particle Transformer [251] and LorentzNet [228] use this idea to learn complex inter–particle relationships while preserving desired invariances. These transformer based approaches effectively interpolate between the simplicity of PFNs and the expressiveness of full graph networks, learning to focus on the most relevant particle combinations for each jet. For unfolding applications, such architectures could potentially learn to identify and correct for detector specific correlations between particles, though care must be taken to ensure the learned representations remain physically meaningful and do not overfit to detector artifacts.

The deep learning architectures discussed in this section provide a rich toolkit for addressing the challenges of unfolding in HEP. By selecting and adapting architectures based on the specific properties of the data and the physics requirements of the analysis,

researchers can develop more accurate and physically consistent unfolding methods. The next section will explore advanced machine learning frameworks that build upon these architectures to provide even more powerful approaches to unfolding.

III.E Modern machine learning frameworks.

III.E.1	Generative models.	149
III.E.1.i	Variational Autoencoders (VAEs)	150
III.E.1.ii	Normalising flows.	153
III.E.1.ii.a	Mathematical formulation.	153
III.E.1.ii.b	Variations.	156
III.E.1.ii.b.1	Conditional normalising flows	156
III.E.1.ii.b.2	Other innovations.	157
III.E.1.ii.c	Applications.	158
III.E.1.ii.c.1	Normalising flows for unfolding.	158
III.E.1.ii.d	Limitations.	159
III.E.2	Discriminative Models	160
III.E.2.i	Training dynamics and optimisation.	162
III.E.2.ii	Applications and challenges.	162
III.E.2.iii	Traditional discriminative architectures.	163
III.E.2.iii.a	Support vector machines (SVMs)	164
III.E.2.iii.b	Boosted decision trees.	166
III.E.2.iv	Discriminative models and deep learning.	167
III.E.2.v	Graph neural networks.	168
III.E.2.v.a	Transformer models.	169
III.E.3	Adversarial models.	171
III.E.3.i	Generative adversarial networks (GANs).	172
III.E.3.i.a	Applications.	174

III.E.3.i.a.1	GANs for unfolding.	175
III.E.3.i.b	Limitations.	177
III.E.3.ii	Wasserstein GANs.	177
III.E.3.iii	Conditional adversarial networks.	179
III.E.3.iv	Other GAN based methods.	180
III.E.3.v	Adversarial autoencoders.	181
III.E.3.vi	Cycle consistent adversarial networks.	182

This section explores sophisticated machine learning frameworks that have demonstrated significant potential for addressing the unfolding problem in high energy physics. Three principal categories are examined: generative models, discriminative models, and adversarial frameworks, each offering distinct approaches to learning complex distributions and relationships in particle physics data.

III.E.1 Generative models.

Generative models represent a powerful class of machine learning techniques that aim to learn and characterise the underlying probability distribution of observed data. Unlike discriminative models that focus on decision boundaries between classes, generative models capture the full data generation process, enabling them to produce new samples that resemble the training distribution [187].

The defining characteristic of these models is their ability to generate new, synthetic data points that follow the same statistical patterns as the training data. This capability makes them particularly valuable for applications in particle physics where simulation of complex physical processes is essential [357, 362]. In the context of unfolding, generative models offer a natural framework for modelling the mapping between particle level and detector level distributions.

Generative models typically approach the learning task either by modelling the probability density function of the data explicitly, or by learning to generate samples from the target distribution without explicitly estimating the density. Training generative models involves maximising the likelihood of the data, though the specific approach depends on the model architecture. For normalising flows, for example,

which have tractable likelihoods due to the change of variables formula, training directly maximises the exact likelihood.

$$\hat{\theta} = \operatorname{argmax}_{\theta} \prod_{i=1}^N p(x_i \mid \theta) , \quad (\text{III.36})$$

where $p(x \mid \theta)$ is the conditional probability density function at x given θ and $\{x_i\}_{i=1}^N$ are the training data.

For Variational Autoencoders (VAEs) [536], on the other hand, the true likelihood is intractable due to the latent variable integration. Instead, VAEs maximise the Evidence Lower Bound (ELBO) [375], which provides a tractable lower bound on the log likelihood.

$$\ln p(x) \geq \mathbb{E}_{q(z|x)} [\ln p(x \mid z)] - D_{\text{KL}}[q(z \mid x) \parallel p(z)] , \quad (\text{III.37})$$

where the first term is the reconstruction loss and the second term is the KL divergence between the approximate posterior and the prior.

Both approaches ultimately aim to model the data distribution, but normalising flows offer exact likelihood computation at the cost of architectural constraints, while VAEs trade exact likelihood for greater flexibility in model design.

III.E.1.i Variational Autoencoders (VAEs)

Note. VAEs are not used in the methods described in Chapters IV to VI and VIII. They are discussed here in order to provide a more complete description of generative model architectures and their properties.

Variational Autoencoders represent a class of deep generative models that combine the principles of variational inference with neural network based autoencoders [161]. First introduced by Kingma and Welling in 2013 [536], VAEs provide a principled probabilistic framework for learning complex data distributions while enabling both generation of new samples and inference of latent representations. Fundamentally, VAEs extend traditional autoencoders by imposing a probabilistic structure on the latent space. While standard autoencoders learn deterministic encodings, VAEs encode inputs as probability distributions in the latent space. This probabilistic approach enables principled sampling and uncertainty quantification, which are critical requirements for unfolding applications.

The VAE architecture consists of two primary components, an encoder network $q_\phi(z | x)$ which maps inputs x to a latent space distribution, typically parametrised as a multivariate Gaussian:

$$q_\phi(z | x) = \mathcal{N}(z | \mu_\phi(x), \text{diag}(\sigma_\phi^2(x))) , \quad (\text{III.38})$$

where $\mu_\phi(x)$ and $\sigma_\phi^2(x)$ are parameters learned by the network, and a decoder network $p_\theta(x|z)$, which reconstructs inputs from latent samples, often parametrised as a Gaussian for continuous data,

$$p_\theta(x | z) = \mathcal{N}(x | \mu_\theta(z), \text{diag}(\sigma_\theta^2(z))), \quad (\text{III.39})$$

or a multivariate Bernoulli distribution [327] for categorical data.

This objective function minimised by VAEs, Equation (III.37), balances two competing terms, the *reconstruction term* $\mathbb{E}_{q_\phi(z|x)}[\ln p_\theta(x | z)]$ encourages accurate recon-

struction, and the KL divergence term $D_{KL}[q_\phi(z | x)||p(z)]$ that acts as a regulariser, encouraging the learned latent distribution to match a prior distribution.

The ELBO construction has deep connections to statistical mechanics and information theory. The KL divergence term bears striking resemblance to the free energy minimisation principle, where $p(z)$ serves as an analogue to the “ground state” or “vacuum state” distribution [546].

Once trained, VAEs enable principled sampling by first sampling from the latent distribution, $z \sim p(z) = \mathcal{N}(0, 1)$, and then generating a new observation by passing z through the decoder, $x \sim p_\theta(x | z)$. This generative capability is particularly valuable for modelling the detector response in particle physics, where the mapping from particle level to detector level observables involves complex transformations and uncertainties. Therefore in high energy physics, VAEs have found numerous applications including

Fast detector simulation where the VAE learns to map particle-level quantities to detector-level observables [116, 224],

Anomaly detection for beyond Standard Model physics searches [191],

Dimensionality reduction of for processing collider data [148], and

Unfolding detector effects [177].

For unfolding applications specifically, several features of VAEs make them naturally suited for the task. VAEs model the uncertainty in the inverse mapping from detector level to particle level distributions providing a direct probabilistic framework for

downstream analysis. One can regularise the problem by constraining the latent space to accord with the physics of the problem, which helps stabilise the training and reduce the variance. Once trained, VAEs provide fast amortised inference, enabling efficient processing of large collision datasets

Certain challenges posed by VAEs must be considered when selecting the appropriate model for a problem. VAEs with Gaussian latent spaces can have limited expressivity because the latent space may be too restrictive for complex physical distributions. Conversely, selecting a feature dense latent space (a) can deregularise the problem and (b) cause a loss of interpretability. VAEs also tend to produce smoothed-out samples, potentially losing fine details in particle distributions. Finally, balancing the reconstruction and KL terms requires careful tuning, and small changes in the hyperparameters can lead to significant changes in the learned function.

III.E.1.ii Normalising flows.

Normalising flows represent a powerful class of generative models that learn the exact likelihood computation through a series of invertible, differentiable transformations. Unlike VAEs, which approximate the likelihood, normalising flows directly learn a bijective mapping between the data distribution and a simple base distribution, typically a multivariate Gaussian.

III.E.1.ii.a Mathematical formulation.

The core principle behind normalising flows is the change of variables formula from probability theory.

Theorem 5 (Change of variables for pushforward densities). For all probability spaces $(\Omega, \mathcal{F}, \mathbb{P})$, for all random variables $Z : \Omega \rightarrow \mathbb{R}^d$ with measure law $\mu_Z \ll \lambda^d$ and density $p_Z = \frac{d\mu_Z}{d\lambda^d}$, for all open sets $U, V \subseteq \mathbb{R}^d$ and for all C^1 bijections with C^1 inverses (i.e. C^1 diffeomorphisms) $f : U \rightarrow V, X = f(Z)$ has measure law $\mu_X = f_{\#}\mu_Z$ *only if* $\mu_X \ll \lambda^d$ and its Radon–Nikodým derivative [10, 676] is, for almost every $x \in \mathbb{R}^d$,

$$\frac{d\mu_X}{d\lambda^d}(x) = \begin{cases} p_Z(f^{-1}(x)) |\det Df^{-1}(x)| & x \in V, \\ 0 & x \notin V. \end{cases} \quad (\text{III.40})$$

Equivalently, if $x = f(z)$ the for almost every $z \in U$,

$$\frac{d\mu_X}{d\lambda^d}(f(z)) = \frac{p_Z(z)}{|\det Df(z)|} \quad (\text{III.41})$$

Proof of Theorem 5. Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be an arbitrary Borel function. By the definition of the pushforward,

$$\int_{\mathbb{R}^d} g(x) d\mu_X(x) = \int_{\Omega} g \circ f \circ Z(\omega) d\mathbb{P}(\omega) = \int_{\mathbb{R}^d} g \circ f(z) p_Z(z) d\lambda^d(z). \quad (\text{III.42})$$

Apply the classical change-of-variables theorem for C^1 diffeomorphisms (equivalently, the area formula for diffeomorphisms) to the integrand $\psi(z) := g(f(z)) p_Z(z)$.

$$\int_U \psi(z) d\lambda^d(z) = \int_V \psi \circ f^{-1}(x) |\det Df^{-1}(x)| d\lambda^d(x). \quad (\text{III.43})$$

Since p_Z is zero almost everywhere outside U , one may restrict the domain of integration to U . This yields

$$\int_{\mathbb{R}^d} g \circ f(z) p_Z(z) d\lambda^d(z) = \int_V g(x) p_Z \circ f^{-1}(x) |\det Df^{-1}(x)| d\lambda^d(x). \quad (\text{III.44})$$

Combining Equations (III.42) and (III.44) yields

$$\int_{\mathbb{R}^d} g(x) d\mu_X(x) = \int_{\mathbb{R}^d} g(x) [\mathbb{1}_V(x) p_Z \circ f^{-1}(x) |\det Df^{-1}(x)|] d\lambda^d(x). \quad (\text{III.45})$$

Since the above equality holds for every bounded Borel function g , for almost every x the Radon–Nikodým theorem yields

$$\frac{d\mu_X}{d\lambda^d}(x) = \mathbb{1}_V(x) p_Z \circ f^{-1}(x) |\det Df^{-1}(x)|. \quad (\text{III.46})$$

as claimed. The equivalent form follows from

$$|\det Df^{-1}(f(z))| = \frac{1}{|\det Df(z)|}. \quad (\text{III.47})$$

□

In simple terms, the measure law of X is the pushforward measure $f_{\#}\mu_Z$, defined as

$$f_{\#}\mu_Z(B) = \mu_Z(f^{-1}(B)) \quad (\text{III.48})$$

for every Borel set $B \subseteq \mathbb{R}^d$.

Because $f : U \rightarrow V$ is a C^1 -diffeomorphism, $f_{\#}\mu_Z \ll \lambda^d$, and its Radon–Nikodým derivative satisfies, for λ^d -almost every $x \in V$,

$$\frac{d(f_{\#}\mu_Z)}{d\lambda^d}(x) = p_Z(f^{-1}(x)) |\det Df^{-1}(x)|, \quad \text{and} \quad \frac{d(f_{\#}\mu_Z)}{d\lambda^d}(x) = 0 \text{ for } x \notin V. \quad (\text{III.49})$$

The factor $|\det Df^{-1}(x)|$ is the Jacobian determinant of the inverse map and accounts for the local distortion of Lebesgue measure; its absolute value removes any dependence on orientation. Equivalently, for λ^d -almost every $z \in U$ with $x = f(z)$,

$$\frac{d(f_{\#}\mu_Z)}{d\lambda^d}(f(z)) = \frac{p_Z(z)}{|\det Df(z)|}. \quad (\text{III.50})$$

More generally, let $f = f_K \circ f_{K-1} \circ \dots \circ f_1$ be a composition of C^1 -diffeomorphisms $f_k : U_k \rightarrow U_{k+1}$ with $U_1 = U$ and $U_{K+1} = V$. Define $y_0 := z$ and $y_k := f_k(y_{k-1})$ for $k = 1, \dots, K$, so that $y_K = x$. Repeated application of the change-of-variables formula yields, for λ^d -almost every $x \in V$,

$$p_X(x) = p_Z(y_0) \prod_{k=1}^K \frac{1}{|\det Df_k(y_{k-1})|} \quad \text{with } y_0 = f^{-1}(x). \quad (\text{III.51})$$

Equivalently, writing $x_0 := x$ and $x_k := f_k^{-1}(x_{k-1})$ for $k = 1, \dots, K$ (so $x_K = z = f^{-1}(x)$),

$$p_X(x) = p_Z(x_K) \prod_{k=1}^K |\det Df_k^{-1}(x_{k-1})|. \quad (\text{III.52})$$

Normalising flows use precisely this structure to chain together multiple such transformations to create highly expressive mappings while maintaining invertibility.

III.E.1.ii.b Variations.

Since they were first proposed, several architectural innovations have expanded the expressivity and computational efficiency of normalising flows.

III.E.1.ii.b.1 Conditional normalising flows.

Conditional normalising flows are a particularly promising variant for unfolding, where the transformation depends on additional conditioning variables [390].

$$p_X(x \mid c) = p_Z \circ f^{-1}(x; c) \cdot \left| \det \left(\frac{\partial f^{-1}(x; c)}{\partial x} \right) \right|. \quad (\text{III.53})$$

In the unfolding context, the detector level observables can serve as the conditioning variables, with the flow mapping from a base distribution to the particle level distribution conditioned on detector measurements [158, 352].

This approach enables explicit modelling of the posterior distribution, $p(z \mid x)$, that unfolding aims to estimate, even when the spaces have different dimensions, effectively learning the inverse detector response while accounting for the inherent ambiguities and information loss in the measurement process. The conditioning mechanism preserves the exact likelihood computation that makes normalising flows so attractive for uncertainty quantification, while providing the flexibility needed to handle realistic experimental scenarios where dimensional consistency cannot be guaranteed. However, the choice of base distribution in normalising flows introduces an implicit prior that can significantly influence the unfolded results.

III.E.1.ii.b.2 Other innovations.

Coupling layers, such as NICE [487] and RealNVP [429], partition the input dimensions and apply transformations to one part conditioned on the other, ensuring tractable Jacobian determinants

$$y_1 = x_1 , \tag{III.54}$$

$$y_{d+1} = x_{d+1} \odot \exp(s(x_1)) + t(x_1) , \tag{III.55}$$

where s and t are scale and translation networks.

Autoregressive flows, such as IAF [437] and MAF [408] model dependencies between dimensions through an autoregressive structure

$$y_i = x_i \cdot \exp(s_i(x_1)) + t_i(x_1). \tag{III.56}$$

Continuous normalising flows [340] formulate the transformation as an ordinary differential equation (ODE), offering increased flexibility.

III.E.1.ii.c Applications.

Normalising flows have gained significant traction in particle physics applications. Some applications of normalising flows in HEP include

Simulation based inference: Using flows to perform likelihood free inference for physics parameters [332],

Fast detector simulation: Modelling the detector response through invertible mappings [189],

Phase space integration: Computing complex multidimensional integrals in perturbative calculations [262],

Unfolding: Inverting detector effects [149].

III.E.1.ii.c.1 Normalising flows for unfolding.

For unfolding specifically, normalising flows provide several compelling advantages. The exact likelihood computation provided by the model enables principled uncertainty quantification in the unfolded distributions, allowing for rigorous statistical inference and error propagation. The bijective nature of normalising flows aligns naturally with the physical process of detector effects and their inversion, as both involve transformations between probability distributions.

However, for standard normalising flows this direct correspondence is limited to cases where the measured and true distributions share the same dimensionality, a constraint that is frequently violated in realistic unfolding scenarios. In many experimental contexts, the detector and particle level quantities are not even the

same physical quantity, so standard normalising flows, being bijective transformations, cannot directly handle these dimensional differences.

Conditional normalising flows, discussed above in Section III.E.1.ii.b.1, elegantly address this limitation by treating the measured data as conditioning information rather than requiring a direct bijective mapping. Capable of capturing complex distributions common in particle physics, both traditional normalising flows and different variants of flow based models have been shown to be effective vehicles for unfolding, showing improved performance particularly for high dimensional and complex distributions [212].

III.E.1.ii.d Limitations.

Despite their theoretical elegance, normalising flows face several practical challenges in HEP applications. Each version of the model makes different trade-offs between expressivity, computational efficiency, and ease of training.

While normalising flows learn to transform from a simple base distribution (typically Gaussian) to the target distribution, this transformation is not unique—different choices of base distribution and flow architecture can lead to different learned mappings. In unfolding applications, this prior dependence manifests as a form of regularisation bias, where the specific parametrisation of the flow influences which solutions are preferred among the many that could explain the observed data. This issue becomes particularly acute in conjunction with the ill posedness unfolding problems where multiple true distributions could have produced the same detector level observations, making the implicit regularisation imposed by the flow architecture a critical but often unexamined aspect of the unfolding procedure.

The invertibility requirement limits the types of transformations that can be used, and the computational cost of evaluating Jacobian determinants for high dimensional data can impose practical limits on the dimensionality of the unfolding problem that flow based methods can reasonably be applied to. Furthermore, the prior dependence inherent in normalising flows presents a conceptual challenge for HEP applications. The base distribution and architectural choices embed implicit assumptions about the structure of the solution space, biasing the results in ways that are difficult to quantify or control. Unlike some of unfolding methods that have been discussed so far, where regularisation assumptions are explicit and interpretable, the prior induced by the flow architecture operates through complex, high dimensional transformations that resist straightforward interpretation. This opacity complicates the assessment of systematic uncertainties and the comparison of results obtained with different flow architectures or base distributions. Ongoing research continues to address these challenges through improved architectures and training procedures.

III.E.2 Discriminative Models

Note. Discriminative models are not directly used in the methods proposed in this dissertation, although they do serve as valuable components of the adversarial models that are used. They are detailed in this section in order to provide a more complete survey of ML architectures relevant to the study of cross section measurements.

Discriminative models, introduced in Sections III.A.2.i and III.A.2.ii, will be explored in greater depth in this section. Discriminative models parametrise the

conditional distribution, $p(y | x)$ using a function $f_\theta(x)$, where θ represents the model parameters.

$$p(y | x) = p(y | f_\theta(x)). \quad (\text{III.57})$$

For classification tasks with K classes, this typically takes the form of a categorical distribution,

$$p(y = k | x) = \frac{\exp f_\theta^k(x)}{\sum_{j=1}^K \exp f_\theta^j(x)}. \quad (\text{III.58})$$

The optimisation objective is to maximise the conditional likelihood,

$$\mathcal{L}(\theta) = \prod_{i=1}^N p(y_i | x_i; \theta), \quad (\text{III.59})$$

or equivalently, minimising the negative log-likelihood,

$$-\ln \mathcal{L}(\theta) = -\sum_{i=1}^N \ln p(y_i | x_i; \theta). \quad (\text{III.60})$$

This formulation provides a principled approach for both probabilistic classification and regression tasks that dominate HEP analyses.

Modern discriminative models in particle physics typically employ neural networks with various architectures tailored to specific data structures. The most common of these are fully connected networks, traditional deep neural networks with dense connections that transform inputs through a series of non-linear functions. Convolutional neural networks (CNNs) are useful for exploiting translational invariance in detector data through localised feature extraction.

$$y_{l,i,j}^k = \sigma \left(\sum_m \sum_{p,q} W_{p,q}^{k,m} y_{l-1,i+p,j+q}^m + b_l^k \right), \quad (\text{III.61})$$

where $y_{l,i,j}^k$ represents the output of the k -th feature map at position (i, j) in layer l .

Recurrent neural networks (RNNs) are most often used to capture sequential information in particle trajectories.

$$y_t = \sigma(W_x x_t + W_y y_{t-1} + b), \quad (\text{III.62})$$

where y_t is the hidden state at step t .

III.E.2.i Training dynamics and optimisation.

Training discriminative models involves gradient based optimisation with regularisation techniques to prevent overfitting. The cross-entropy loss described in Equation (III.35), derived from maximum likelihood is the most commonly used loss function to train discriminative models. In some cases, specialised losses are like focal loss are used address class imbalances in rare physics processes.

$$L_{\text{focal}} = - \sum_{i=1}^N (1 - p_i)^\gamma \ln(p_i), \quad (\text{III.63})$$

where p_i is the predicted probability of the correct class and γ is a focusing parameter.

III.E.2.ii Applications and challenges.

Discriminative models serve multiple critical functions in particle physics analyses.

Event classification: Separating signal from background events, with area under the ROC curve (AUC) values exceeding 0.99 in some analyses,

Particle identification: Distinguishing particle types based on detector signatures with significantly improved efficiency compared to cut-based methods,

Jet Tagging: Identifying jets originating from specific particles,

Unfolding: Estimating reweighting functions between detector level and particle level distributions, as implemented in OMNIFOLD.

Despite their successes, the use of discriminative models can pose a series of challenges. Neural network outputs often require calibration to provide accurate probability estimates for downstream statistical analysis. These models are also typically considerably more complex than their generative counterparts, and may identify features that lack clear physical interpretation. Quantifying model uncertainties remains challenging, particularly for ensemble approaches, when bootstrapping is computationally unfeasible.

The following section examines a few specific discriminative model architectures that have proven effective in particle physics applications. It begins with traditional approaches—Support Vector Machines and Boosted Decision Trees—which paved the way for ML adoption in HEP. It then discusses modern deep learning architectures, graph neural networks and transformers, which have pushed the boundaries of what is achievable in physics analysis. Each architecture offers unique advantages for different physics tasks, and understanding their relative strengths and limitations is crucial for selecting the appropriate model for a given analysis.

III.E.2.iii Traditional discriminative architectures.

Note. None of the architectures described in this section are used in the novel research presented in this dissertation. They are discussed here to provide a brief overview of historical and contemporaneous discriminative models used in the field.

Before the deep learning revolution, particle physics relied heavily on classical learning methods that demonstrated remarkable effectiveness for classification and regression tasks. These traditional discriminative architectures, while conceptually simpler than modern deep networks, established the foundation for machine learning adoption in high energy physics and continue to serve as reliable tools for many analyses. Two methods in particular, support vector machines and boosted decision trees, have been used extensively in physics analysis due to their interpretability, robustness, and strong performance on the structured, tabular data typical of particle physics experiments.

III.E.2.iii.a Support vector machines (SVMs)

SVMs find the maximum-margin hyperplane separating classes by solving the primary optimization problem

$$\min_{w,b} \frac{1}{2} \|w\|^2 \quad (\text{III.64})$$

$$\text{subject to } y_i (w^T x_i + b) \geq 1 . \quad (\text{III.65})$$

The dual formulation of this problem leads to the kernel trick, enabling non-linear decision boundaries,

$$\max_{\alpha} \sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j K(x_i, x_j) \quad (\text{III.66})$$

$$\text{subject to } \alpha_i \geq 0 \text{ and } \sum_{i=1}^N \alpha_i y_i = 0 , \quad (\text{III.67})$$

where $K(x_i, x_j)$ is a kernel function.

The choice of kernel dramatically affects SVM performance. The most common choices in particle physics problems are the

Linear kernel, $K(x_i, x_j) = x_i^T x_j$, which is efficient for high-dimensional but linearly separable data,

Radial basis functions (RBF), $K(x_i, x_j) = \exp(-\gamma \|x_i - x_j\|^2)$, which are capable of capturing nonlinear detector responses, and the

Polynomial kernel, $K(x_i, x_j) = (x_i^T x_j + c)^d$, which models polynomial relationships in feature space.

SVMs have been applied across various particle physics processes. Early applications in the Higgs discovery process demonstrated competitive performance and propelled these models forward in the discourse [424]. In analyses involving quark–gluon discrimination, SVMs with RBF kernels achieved discrimination power comparable to dedicated physics-inspired variables [540, 643]. Attempts to construct SVMs for unfolding focus on their potential to estimate the density ratio between detector level and particle level distributions

A persistent roadblock to their more widespread adoption has been the fact that traditional SVM implementations scale poorly with dataset size, $\mathcal{O}(n^2)$ to $\mathcal{O}(n^3)$. Unlike deep networks, SVMs also require careful feature selection, and standard SVMs don’t provide natural probability estimates, limiting their usefulness.

III.E.2.iii.b Boosted decision trees.

Historically one of the most popular models employed in jet physics analyses, boosted decision trees (BDTs) combine weak learners ⁸ into a powerful ensemble.

$$F(x) = \sum_{m=1}^M \beta_m h_m(x) , \quad (\text{III.68})$$

where h_m is the m^{th} tree and β_m is its corresponding weight.

Each tree recursively partitions the feature space according to information gain criteria,

$$\text{Gain} = \text{Impurity}(\text{parent}) - \sum_{j \in \{\text{children}\}} \frac{N_j}{N} \text{Impurity}(j) . \quad (\text{III.69})$$

Several boosting algorithms have either been developed specifically for HEP or been employed in HEP analyses.

AdaBoost [598, 697] iteratively reweights misclassified samples.

$$w_i^{(t+1)} = w_i^{(t)} \cdot e^{\alpha_t \cdot \mathbb{1}(y_i \neq h_t(x_i))} , \quad (\text{III.70})$$

where $\alpha_t = \frac{1}{2} \ln \left(\frac{1-\epsilon_t}{\epsilon_t} \right)$ and ϵ_t is the weighted error rate.

Gradient boosting [599] fits each tree to the negative gradient of the loss function with respect to the current prediction.

$$h_m = \underset{h}{\operatorname{argmin}} \sum_{i=1}^N L(y_i, F_{m-1}(x_i) + h(x_i)) . \quad (\text{III.71})$$

XGBoost [46, 462] incorporates second order derivatives and regularisation into the above schema.

$$L^{(t)} \approx \sum_{i=1}^n \left[g_i f_t(x_i) + \frac{1}{2} h_i f_t^2(x_i) \right] + \Omega(f_t) , \quad (\text{III.72})$$

⁸I.e. shallow decision trees.

where g_i and h_i are the first and second derivatives of the loss.

BDTs have been extensively deployed in particle physics analyses. The Toolkit for Multivariate Analysis (TVMA) [619] integrated BDTs into the ROOT framework, cementing their place as a staple in HEP analyses. BDTs were instrumental in the Higgs boson discovery by improving signal-to-background discrimination [399]. They can also be used in unfolding tasks, since they can model complex response matrices and estimate reweighting functions

Despite their success, BDTs present several challenges, causing them to be eclipsed more recently by deep learning methods. They struggle with highly correlated features common in detector data, and performance typically degrades in very high dimensional spaces. Tree based models also introduce discontinuities in the mapping function that can complicate uncertainty propagation.

III.E.2.iv Discriminative models and deep learning.

The transition from traditional learning to deep learning in particle physics has been driven by the need to handle increasingly complex data structures and extract more sophisticated feature representations. While classical methods excel with engineered features and tabular data, deep learning architectures can automatically learn complex representations from raw or minimally processed data. This capability has proven particularly valuable for particle physics applications involving complex topological relationships, variable length sequences, and high dimensional data. Modern deep learning approaches have thus expanded the analytical toolkit available to physicists, enabling new insights that were previously difficult or impossible to obtain.

III.E.2.v Graph neural networks.

GNNs, introduced in Section III.D.3, operate on graph-structured data, represented as $G = (V, E)$ with nodes $v_i \in V$ and edges $e_{ij} \in E$. The message passing framework updates node representations through the rule

$$y_i^{(l+1)} = \phi \left(y_i^{(l)}, \bigoplus_{j \in \mathfrak{N}(i)} \psi(y_i^{(l)}, y_j^{(l)}, e_{ij}) \right), \quad (\text{III.73})$$

where ϕ, ψ are learnable functions, $\mathfrak{N}(i)$ denotes neighbours of node i and $\bigoplus$ is the permutation invariant aggregation operator.

A few different specialized GNN variants have been developed with an eye towards specific applications in HEP and in physics more broadly. For example, interaction networks model physical interactions between particles with explicit edge functions. Dynamic graph CNNs can be used for an array of applications from construct graphs dynamically based on spatial proximity in detector space to encoding hierarchical quantum circuit representations, to studying the microscopic structure and emergent dynamics in glasses [56, 82, 192, 414].

As discussed in Section III.D.5.i.a.2, one possible interpretation of particle flow networks is as a special case of GNNs. PFNs incorporate physics informed constraints into the message-passing mechanism.

HEP specific objectives functions have also been developed to train GNNs for HEP applications. These include node level classification for identifying individual particles, graph level classification for categorising entire events, edge prediction for reconstructing particle interaction vertices, and energy regression for estimating particle energies and momenta.

Through these specialised architectures and objective functions, GNNs have shown promising performance across a range HEP use cases.

Jet Tagging: Representing jets as graphs of constituents for improved identification performance,

Event Reconstruction: Modelling entire collision events as interaction graphs,

Tracking: Reconstructing particle trajectories from detector hits,

Unfolding: Learning mappings between detector level and particle level observables while preserving physical or statistical invariants.

The most significant constraint on the use of GNNs in HEP is their computational complexity. Message passing operations scale with graph size and connectivity. While their ability to naturally handle variable sized inputs can be incredibly useful, implementing them requires careful batching strategies to maintain the differentiable structure of the network that backpropagation relies on.

III.E.2.v.a Transformer models.

Transformers process sequences through self-attention mechanisms,

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{Q K^T}{\sqrt{d_k}} \right) V, \quad (\text{III.74})$$

where Q, K, V are query, key, and value matrices derived from input embeddings. In HEP applications, self-attention enables modelling various relationships between particles, such as permutation invariance for the natural handling of varying particle multiplicities, long range dependencies, for capturing correlations between distant

detector regions, and interpretable attention maps for visualising learned physical interactions.

Transformer models, which revolutionised the field of natural language processing, are increasingly being adopted for complex HEP tasks. While transformers were originally designed to process sequential data where order not only matters, but also is expressly important to the learned representation, they can be adapted to be permutation invariant, a necessary property for particle physics applications where the ordering of particles or constituents is typically arbitrary and physically meaningless.

In jet classification tasks, transformers process jets as collections of constituents rather than sequences, learning representations that are invariant to the order in which particles are presented. This permutation invariance is achieved through modifications to the standard transformer architecture, such as removing positional encodings or using specialised attention mechanisms that treat input particles symmetrically. Transformers have been applied to event level classification tasks in particle physics, with applications including jet classification and event categorisation [29, 57, 293]. This capability has proven particularly valuable for event level analysis, where the transformer can capture complex relationships between particles throughout the entire collision while respecting the fundamental symmetries of the underlying physics.

However, despite their potential, some of the characteristics of transformer models have hindered their widespread adoption in HEP. Self-attention operations scale quadratically with sequence length, constraining the size of attention modules [455]. Transformer models typically require large training datasets [150], as a consequence of which, training runs can rarely exceed one epoch [125]. Hence, while transformers

can learn surface level relationships extremely effectively, they are limited in their ability to learn deeper structure that would require multiple passes over each data point to discover [34, 142, 144].

III.E.3 Adversarial models.

Adversarial models, which lie at the heart of several of the novel methods presented in this dissertation, leverage the competitive dynamics between two neural networks to achieve powerful generative and discriminative capabilities. These models have revolutionised high energy physics applications by enabling novel approaches to simulation, unfolding, and domain adaptation. Fundamentally different from traditional architectures, adversarial models employ a game theoretic framework that can capture complex, high dimensional distributions and transformations relevant to particle physics.

Adversarial models are typically constructed using two competing neural networks, a generator g and a discriminator d . The generator attempts to produce outputs that the discriminator cannot distinguish from real data, while the discriminator attempts to correctly classify inputs as either real or generated. This minimax game is formalised as the optimisation problem

$$\max_g \min_d L(d, g) = -\mathbb{E}_{x \sim p(x)}[\ln d(x)] - \mathbb{E}_{z \sim q(z)}[\ln (1 - d(g(z)))] , \quad (\text{III.75})$$

where $p(x)$ is the distribution of real data, $q(z)$ is a latent space distribution,⁹ $g(z)$ maps the latent space to the data space, and $d(x)$ outputs the probability that x came from the real data rather than the generator.

In the context of particle physics, this framework can be extended to incorporate domain specific constraints and physical laws, leading to variants such as

$$\begin{aligned} \max_g \min_d L(d, g) = & -\mathbb{E}_{x \sim p_{\text{data}}(x)}[\ln d(x)] \\ & -\mathbb{E}_{z \sim p_z(z)}[\ln(1 - d(g(z)))] + \lambda \mathcal{R}(g) , \end{aligned} \quad (\text{III.76})$$

where $\mathcal{R}(g)$ represents physics informed regularisation terms that enforce conservation laws, symmetries, or other physical constraints, weighted by hyperparameter λ .

The following sections will discuss some of the most common variants of adversarial models, and how they can be applied to the HEP analyses.

Note. While adversarial models are at the heart of multiple novel methods presented in subsequent chapters, not all architectures discussed in the following sections are directly used in the novel research proposed. The methods proposed in this dissertation instead take inspiration from extant models to invent entirely new algorithms. It is, however, important to better understand the insights provided by the existing landscape before proceeding to discuss novel contributions.

III.E.3.i Generative adversarial networks (GANs).

Generative adversarial networks (GANs) represent a fundamentally different approach to generative modelling in comparison to the generative models discussed

⁹Typically Gaussian noise, though in the variations presented in this dissertation, $q(z)$ will most often be the Generation distribution.

in Section III.E.1, employing an adversarial training framework rather than direct density estimation [279]. First introduced by Goodfellow et al. in 2014 [516], GANs have revolutionised generative modelling through their ability to produce remarkably realistic samples, particularly for high dimensional data like images [335, 360, 373, 391, 401, 486]. By training the two networks simultaneously in the aforementioned minimax game, the generator learns to produce increasingly realistic samples that can fool the discriminator, while the discriminator improves its ability to distinguish real from generated samples.

From a theoretical perspective, the GAN objective can be interpreted as minimising the Jensen–Shannon divergence, or a variant thereof, between the data distribution and the generator distribution [382, 545, 648, 661, 677, 717]. Training proceeds by alternating between optimising the two networks using opposing gradient information, *viz.* d is updated using gradient descent,

$$\theta_d \leftarrow \theta_d - \eta \nabla_{\theta_d} L(d, g), \quad (\text{III.77})$$

and g is updated using gradient ascent,

$$\theta_g \leftarrow \theta_g + \eta \nabla_{\theta_g} L(d, g). \quad (\text{III.78})$$

In equilibrium, the generator produces samples indistinguishable from the real data distribution, and the discriminator outputs a probability of 0.5 for all inputs. Formally the optimal discriminator computes

$$d(x) = \frac{p_{\text{data}}(x)}{p_{\text{data}}(x) + p_g(x)} \quad \text{for all } x, \quad (\text{III.79})$$

and the optimal generator perfectly captures the data distribution

$$p_g(x) = p_{\text{data}}(x). \quad (\text{III.80})$$

In practice, reaching this optimum can be challenging due to the min–max nature of the objective, leading to training instabilities like mode collapse and oscillations. For this reason numerous GAN variants have been proposed to address stability issues and enhance performance.

In addition to architectural modifications, there also exist techniques such as spectral normalisation, which can be applied to GAN-like methods across the board to potentially improve stability.

Definition 37 (Spectral normalisation). *Spectral normalisation* is the process of controlling the Lipschitz constant of a discriminator by normalising its weights,

$$W_{\text{SN}} = \frac{W}{\sigma(W)}, \quad (\text{III.81})$$

where $\sigma(W)$ is the spectral norm of weight matrix W , provides additional tools to avoid training instabilities in GANs [389, 405].

Spectral normalisation, which has long been studied for its effectiveness in regularising inverse problems [274, 326, 625], therefore is a natural fit for GAN based methods for unfolding.

III.E.3.i.a Applications.

These advances have made GANs more practical for scientific applications like those in particle physics.

In high energy physics, GANs have found diverse applications.

Fast detector simulation: GANs can approximate the mapping from particle-level to detector-level observables, significantly accelerating the simulation process [114, 160, 193, 198, 236, 407].

Event generation: GANs can generate collision events, potentially replacing or augmenting traditional Monte Carlo approaches [361, 365].

Anomaly detection: Identifying rare or anomalous collision events that deviate from Standard Model predictions [98, 204, 288, 364].

III.E.3.i.a.1 GANs for unfolding.

The unfolding problem in high energy physics presents a two level challenge that makes it particularly well suited for GAN based approaches. At its core, unfolding seeks to recover the true particle level distribution from the measured detector level observations [136, 422, 569, 592, 664]. This inverse problem is complicated by the fact that detector measurements are affected by numerous physical processes including finite resolution, acceptance limitations, detector inefficiencies, and backgrounds [557, 656]. The mapping from particle level to detector level involves complex non-linear transformations that traditional unfolding methods struggle to capture fully [562, 693].

GANs offer several fundamental advantages for addressing this two level structure.

1. The adversarial training framework naturally accommodates the forward and inverse mappings between detector and particle levels without requiring explicit likelihood functions [116, 312, 394].

2. The generator network can learn to transform detector level distributions to particle level distributions, while the discriminator ensures that the generated particle level events are indistinguishable from true simulated events [397]. This approach circumvents the need for regularised matrix inversion or other the kinds of regularisation schemes that plague traditional methods [254, 698], instead learning the unfolding transformation directly from data.
3. Perhaps most importantly, GANs excel at handling the dimensional mismatch between detector and particle levels, a common scenario in experimental physics. The detector level space often has different dimensionality than the particle level space due to several factors: detector granularity may lead to merged objects at detector level that are resolved at particle level, some particles may escape detection entirely, and detector noise can introduce spurious signals [289, 309]. Traditional unfolding methods might be forced to select unnatural binning strategies or use dimensionality reduction to handle these mismatches [557].

In contrast, GANs can seamlessly learn mappings between spaces of different dimensions through their neural network architectures [312]. The generator network can incorporate latent variables to account for missing information when mapping from lower to higher dimensions, while attention mechanisms and pooling operations handle dimension reduction in the opposite direction [366, 394]. This flexibility allows GANs to preserve more information during the unfolding process and avoid the information loss associated with binning in traditional methods [309].

III.E.3.i.b Limitations.

GANs are however infamous for their training instability. The adversarial objective can lead to convergence issues, particularly for the sparse, high dimensional distributions common in particle physics. Even when the two networks are well balanced, the min-max optimisation can lead to oscillatory behaviour rather than convergence.

Another particularly well studied failure mechanism is mode collapse, in which GANs may fail to capture the full diversity of the target distribution, focusing instead on a limited subset of modes. This occurs when the generator network is too powerful, and discriminator is unable to provide receive gradient information from it. Conversely, when the discriminator is too powerful and only provides minimal gradient information to the generation, the training stalls due to the vanishing gradients.

These challenges have motivated hybrid approaches that combine the strengths of GANs with other novel architectural design choices to use them more effectively for unfolding.

III.E.3.ii Wasserstein GANs.

The Wasserstein GAN (WGAN) represents a significant advancement in generative adversarial network training by addressing one of the fundamental problems that plagued earlier GAN architectures. One critical limitation of the JS divergence, which traditional GANs rely to measure the difference between the real data distribution and the generated data distribution, is that when the two distributions have little or no overlap,¹⁰ the JS divergence saturates at its maximum value and provides

¹⁰For example, early in training when the generator produces poor samples.

uninformative gradients [419, 420, 434]. This leads to the aforementioned problem of vanishing gradients, where the discriminator becomes too good and gives the generator no useful signal for improvement, causing training to stall [433, 470].

The Wasserstein distance, also known as the Earth Mover’s Distance, introduced in Section II.G.1.ii offers a more principled solution to this problem. Unlike JS divergence, the Wasserstein distance provides a smooth, continuous measure of the distance between distributions even when they don’t overlap. Intuitively, it measures the minimum “cost” of transforming one distribution into another, where cost is defined as the amount of “mass” that needs to be moved multiplied by the distance it needs to travel [602]. This property ensures that the Wasserstein distance always provides meaningful gradients that point in the direction of improvement, regardless of the difference in the distributions [420]. As a result, WGANs exhibit more stable training dynamics, with the loss function actually correlating with sample quality [201, 271, 363].¹¹ The improved gradient flow allows for more reliable convergence and reduces the likelihood of the common failure modes of mode collapse, where the generator learns to produce only a limited variety of samples, and vanishing gradients, where the discriminator no longer provides meaningful gradients to the generator, stalling training.

The optimisation for a WGAN can be formulated as

$$\max_g \min_{d \in D} \mathbb{E}_{z \sim q(z)} [d(g(z))] - \mathbb{E}_{x \sim p_{\text{data}}(x)} [d(x)], \quad (\text{III.82})$$

where D is the set of 1-Lipschitz functions.

¹¹Unlike traditional GANs where the loss might not, in fact, be a reliable measure of performance.

III.E.3.iii Conditional adversarial networks.

Conditional GANs (cGANs) extend the standard GAN by incorporating auxiliary conditioning information y into both generator and discriminator,

$$\begin{aligned} \max_g \min_d L(d, g) = & -\mathbb{E}_{(x, y) \sim p_{\text{data}}(x, y)} [\ln d(x | y)] \\ & -\mathbb{E}_{z \sim q(z), y \sim p_{\text{data}}(y)} [\ln(1 - d(g(z | y) | Y))] . \end{aligned} \quad (\text{III.83})$$

This formulation enables the generation of samples conditioned on specific physical parameters, detector configurations, or particle properties. Conditioning information can be incorporated into the GAN architecture in a few different ways. Most directly, conditioning information can be concatenated as additional input features

$$g(z, y) = g_\theta([z, y]) \quad d(x, y) = d_\phi([x, y]) . \quad (\text{III.84})$$

Conditional normalisation is an approach that involves modulating normalisation parameters based on conditioning

$$\gamma(y) = \text{NN}_\gamma(y), \quad \beta(y) = \text{NN}_\beta(y), \quad (\text{III.85})$$

$$\text{CN}(h, y) = \gamma(y) \cdot \frac{h - \mu(h)}{\sigma(h)} + \beta(y). \quad (\text{III.86})$$

A more recent approach is to incorporate conditioning information instead through feature-wise linear modulation (FiLM) layers [17, 72, 101, 242, 258, 446], that scale and shift feature maps.

$$\text{FiLM}(h_i, y) = \gamma_i(y) \cdot h_i + \beta_i(y) . \quad (\text{III.87})$$

Conditional GANs provide powerful tools for physics constrained generation. In HEP studies, they have been used for

Parametrised simulation: Generating detector responses across different operating conditions,

Conditional unfolding: Correcting detector effects with explicit energy dependence,

Systematic variation generation: Creating samples with varied systematic uncertainties,

Model parametrisation: Generating samples across theoretical model parameter spaces.

III.E.3.iv Other GAN based methods.

One innovation designed to improving stability during training is the Progressive GAN [401]. In this approach, one gradually increases model complexity during training, improving stability and quality.

Another class of GAN based methods implement *feature matching*, such as in the McGan [444], which encourages the generator to match statistics of intermediate discriminator activations,

$$\mathcal{L}_{\text{feature}} = \|\mathbb{E}_{x \sim p_{\text{data}}}[f(x)] - \mathbb{E}_{z \sim q}[f(G(z))]\|_2^2, \quad (\text{III.88})$$

where $f(x)$ represents activations from an intermediate layer of the discriminator.

III.E.3.v Adversarial autoencoders.

Adversarial Autoencoders (AAEs) combine autoencoder reconstruction with adversarial training.

$$L_{\text{AAE}} = L_{\text{recon}} + \lambda L_{\text{adv}} \quad (\text{III.89})$$

where $L_{\text{recon}} = \|x - \text{Dec}(\text{Enc}(x))\|^2$ measures reconstruction quality, and L_{adv} is the adversarial loss that forces the encoder's latent space distribution to match the sample distribution,

$$L_{\text{adv}} = \mathbb{E}_{z \sim q} [\ln d(z)] + \mathbb{E}_{x \sim p_{\text{data}}} [\ln (1 - d(\text{Enc}(x)))]. \quad (\text{III.90})$$

AAEs are trained in two phases per iteration.

1. **Reconstruction Phase:** Update encoder and decoder to minimise reconstruction error.
2. **Regularization Phase:**
 - Update discriminator to distinguish between samples from the prior and encoded data,
 - Update encoder to fool the discriminator.

AAEs have been used especially in applications involving dimensionality reduction, to learn physics preserving low dimensional representations; for anomaly detection, to identify unusual events through reconstruction errors or latent space deviations; and for simulation enhancement, to improving the fidelity of simplified simulations.

III.E.3.vi Cycle consistent adversarial networks.

CycleGANs enable unpaired domain translation through cycle consistency.

$$L_{\text{CycleGAN}} = L_{\text{GAN}}(G, D_Y, X, Y) + L_{\text{GAN}}(F, D_X, Y, X) + \lambda L_{\text{cyc}}(G, F) , \quad (\text{III.91})$$

where $G : X \rightarrow Y, F : Y \rightarrow X$ are mapping functions between domains, and the cycle consistency loss is

$$L_{\text{cyc}}(G, F) = \mathbb{E}_{x \sim p_{\text{data}}} [\|F(G(x)) - x\|_1] + \mathbb{E}_{y \sim p_{\text{data}}} [\|G(F(y)) - y\|_1] . \quad (\text{III.92})$$

CycleGANs employ two generators and two discriminators. Generator G : maps domain X to domain Y , generator F : maps domain Y back to domain X , discriminator D_X distinguishes real X from generated $F(Y)$, and discriminator D_Y distinguishes real Y from generated $G(X)$.

In HEP, CycleGANs can enable applications, such as

Detector level mapping: Mapping between simulated and real detector data without paired examples,

Cross experiment translation: Translating measurements between different experimental setups,

Systematic variation generation: Creating physically plausible variations for systematics studies.

Research addressing the challenges adversarial models pose continues to expand the applicability of adversarial models in particle physics, with recent advances incorporating invariant representations, equivariant architectures, and physics informed losses to enhance both performance and reliability.

III.F Comparative analysis of ML based unfolding methods.

III.F.1	Method selection framework.	184
III.F.2	Performance trade-offs.	185
III.F.2.i	Synthesis of experimental experience.	187
III.F.3	Open challenges and research frontiers.	188

The preceding sections have introduced a diverse landscape of machine learning approaches to the unfolding problem, from discriminative methods to generative models and adversarial frameworks. This section provides a comparative synthesis to guide method selection, analyse performance trade-offs, and identify emerging best practices from experimental implementations.

III.F.1 Method selection framework.

The choice of unfolding method depends critically on the interplay between measurement requirements, computational constraints, particularities of the data available and prior knowledge. This section seeks to present a framework for method selection based on key experimental considerations. For measurements where full posterior distributions are required, such as when propagating uncertainties through complex analyses or when non-Gaussian error structures are expected, Bayesian approaches like NEURAL POSTERIOR UNFOLDING [91] or conditional normalising flows [311] provide natural solutions. These methods excel when the measurement goal involves parameter estimation or when correlations with nuisance parameters must be marginalised.

Conversely, when computational efficiency is paramount, particularly for measurements that must be repeated with multiple systematic variations, reweighting based methods like OMNIFOLD [309] offer favourable scaling properties. The amortised nature of the trained classifiers enables rapid evaluation across different data configurations without retraining, a significant advantage for systematic uncertainty studies.

The dimensionality of the measurement phase space presents another decision point. While OMNIFOLD has demonstrated robust performance up on high dimensional data [309], generative models based on normalising flows can struggle with the exponentially growing complexity of high dimensional transformations [408]. For extreme dimensionality, hybrid approaches that combine dimensionality reduction with subsequent unfolding may be necessary.

The nature of the physics observable itself influences method selection. For observables defined through complex, non-linear transformations of detector quantities,¹² methods that can learn these transformations implicitly, such as graph neural networks combined with adversarial training [347], may outperform approaches requiring explicit feature engineering.

III.F.2 Performance trade-offs.

A quantitative comparison of ML based unfolding methods reveals fundamental trade-offs that must be balanced against analysis requirements. Computational scaling presents the most immediate practical constraint. Training times for different methods span orders of magnitude: OMNIFOLD typically requires 10 min to 30 min for training on typical jet physics datasets [308], while conditional normalising flows can require 2 h to 10 h for comparable problems [311]. However, inference times show the opposite pattern, with flow-based methods providing near instantaneous sampling once trained, while OMNIFOLD requires iteration through multiple classifier evaluations.

¹²Such as jet substructure variables computed from constituent particles.

Statistical efficiency, quantified through the effective sample size after unfolding, varies significantly across methods. Reweighting approaches inherently suffer from weight variance that reduces statistical power [41]. Generative approaches avoid this loss but introduce sampling variance that must be controlled through ensemble methods.

Model dependence and bias represents another consideration. While OMNIFOLD iteratively reduces dependence on the initial Monte Carlo distribution, it cannot eliminate dependence on the simulated detector response kernel [309]. Normalising flow approaches embed prior assumptions through their choice of base distribution and architectural constraints [408]. Empirical studies suggest that prior sensitivity is most pronounced in phase space regions with limited statistics, where regularisation effects dominate.

Uncertainty quantification capabilities also vary across methods. Bayesian approaches provide natural uncertainty estimates through posterior sampling, though the computational cost scales linearly with the number of posterior samples required. Frequentist reweighting methods require bootstrapping or asymptotics [570] introducing additional computational overhead. An alternative involves ensemble methods, training multiple networks with different initialisations, which can provide robust uncertainty estimates at moderate computational cost [294]. These approaches though, require careful validation to ensure proper uncertainty coverage, and do not provide much interpretability about the uncertainty sources.

III.F.2.i Synthesis of experimental experience.

The growing body of experimental applications provides crucial insights into practical implementation challenges and solutions. Analysis of recent measurements reveals common patterns and emerging best practices.

Closure test design has evolved significantly from early implementations. The H1 collaboration’s pioneering application of OMNIFOLD established the importance of stress tests using deliberately mismodelled simulations [221]. Subsequent analyses have standardised on a three tier validation hierarchy.

1. Closure with matched simulation,
2. Stressed closure with reweighted simulation to mimic common mismodeling patterns, and
3. Data driven validation using control regions.

Systematic uncertainty propagation remains an active area of development. The ATLAS collaboration’s approach of training separate unfolding models for each systematic variation [41] provides an estimate of coverage at substantial computational cost, but is quite *ad hoc* in the choice of models used, and at best gestures towards the coverage. Alternative approaches using single models with systematic uncertainties encoded as additional inputs show promise but require careful validation to ensure proper uncertainty coverage.

The integration of ML based unfolding into experimental workflows has revealed unexpected challenges. Version control for trained models, particularly when dealing with ensemble methods, requires infrastructure beyond traditional analysis frameworks.

The CMS collaboration has developed specialised containers that package trained models with their associated preprocessing pipelines, ensuring reproducibility across different computing environments [104]. Practical experience has also highlighted the importance of physics informed regularisation. Pure ML approaches can produce unphysical features in regions where training data is sparse. The incorporation of known conservation laws, either through architectural constraints or specialised loss terms, significantly improves the physical plausibility of unfolded results [141].

III.F.3 Open challenges and research frontiers.

Despite substantial progress, several fundamental challenges remain that define the research frontier in ML based unfolding. The treatment of correlations in unbinned unfolded data represents perhaps the most pressing theoretical challenge. While binned methods produce well defined covariance matrices, the unbinned output of ML methods¹³, lacks an analogous mathematical framework for correlation quantification. The study presented in Chapter VII suggests that the effective sample size reduction due to weight correlations can be substantial [44], but general methods for propagating these correlations through subsequent analyses remain underdeveloped.

The development of genuinely model independent unfolding methods remains elusive. All current approaches require some prior knowledge, whether encoded in the training simulation, architectural choices, or regularisation schemes. Work on domain adaptation techniques shows some preliminary promise for reducing model

¹³Typically event weights or generated samples.

dependence [389], but fundamental information theoretic limits suggest complete model independence may be impossible for a given acceptable level of variance.¹⁴

Scalability to extreme dimensionality presents both computational and statistical challenges. While ML methods handle moderate dimensionality better than traditional approaches, measurements involving hundreds of observables¹⁵ remain computationally prohibitive. Hierarchical approaches that unfold in stages, possibly combining different methods at each stage, might represent a possible research direction.

The standardisation of ML based unfolding methods across experiments remains an important practical challenge. Unlike traditional methods implemented in common frameworks like RooUnfold [562], most ML approaches currently lack standardised implementations¹⁶. This fragmentation complicates method comparison and slows adoption of improvements across the field.

Finally, the connection between unfolding and downstream physics inference requires deeper investigation. Most current methods optimise for accurate recovery of particle level distributions, but many analyses ultimately aim to constrain theoretical parameters. Direct optimisation for parameter sensitivity, rather than distribution accuracy, could provide more efficient approaches for specific physics goals [315]. The development of end-to-end differentiable analysis chains, from raw data through unfolding to parameter inference, represents an ambitious but potentially transformative research direction.

¹⁴*vide* The No Free Lunch Theorem in machine learning [688].

¹⁵Such as full jet constituent information.

¹⁶Although the recent implementation in [11] represents a very meaningful stride in this direction.

III.G Case studies: ML based unfolding in HEP analyses.

III.G.1	OMNIFOLD.	191
III.G.2	NEURAL POSTERIOR UNFOLDING.	192
III.G.3	Invertible neural networks for unfolding.	194
III.G.4	MOMENT UNFOLDING	195
III.G.5	Reweighting Adversarial Networks (RANs)	196
III.G.6	Schrödinger Bridge Unfolding.	197
III.G.7	Unified Neural Folding GAN	199

This section examines some new machine learning based unfolding methods and successful applications of machine learning methods in recent high energy physics analyses, demonstrating their practical utility and impact on physics measurements.

Note. Not all of the methods detailed in this section are directly relevant to or used in the novel research presented in this dissertation. They are included here in order to provide a relatively wide cross section of modern ML based approaches to unfolding.

III.G.1 OMNIFOLD.

The OMNIFOLD algorithm represents the first ML based unfolding method to achieve widespread adoption in experimental particle physics, establishing a new paradigm for unbinned measurements across multiple experiments. The H1 Collaboration at HERA pioneered its experimental implementation, initially applying OMNIFOLD to measure lepton–jet momentum balance and azimuthal correlations in deep inelastic scattering [221]. This groundbreaking analysis demonstrated the feasibility of ML based unfolding in a production environment. The collaboration subsequently extended the approach to more complex observables, including jet substructure measurements [21, 170], showcasing the method’s versatility across different physics applications.

Building on H1’s success, experiments at the LHC have applied OMNIFOLD for increasingly sophisticated measurements. The LHCb collaboration applied the technique to study jet fragmentation functions, unfolding the longitudinal momentum fraction of charged particles within jets [169]. This analysis highlighted OMNIFOLD’s

capability to handle the complex correlations inherent in fragmentation processes while maintaining model independence across a broad kinematic range.

The CMS [35] and ATLAS [41, 103, 131] collaborations have similarly adopted OMNIFOLD for jet substructure measurements, contributing to a growing body of cross-experiment validation. These independent implementations using different detector technologies and analysis frameworks provide crucial confirmation of the method’s robustness and general applicability.

Beyond particle physics experiments, the STAR collaboration at RHIC has successfully applied OMNIFOLD to heavy ion collision data [200], demonstrating its effectiveness in the challenging environment of ultrarelativistic nuclear collisions. The method’s ability to handle the large multiplicities and complex backgrounds characteristic of heavy ion physics represents a significant extension of its capabilities. These advances position OMNIFOLD not merely as an unfolding tool but as a framework for optimal information extraction from complex experimental data.

III.G.2 NEURAL POSTERIOR UNFOLDING.

NEURAL POSTERIOR UNFOLDING represents an innovative approach that enhances traditional binned unfolding through the integration of normalising flows with Bayesian inference frameworks. While NPU operates within the binned paradigm rather than an unbinned framework, it addresses several fundamental limitations of classical binned approaches that make it particularly valuable for cross section measurements.

The method’s major contribution lies in its transformation of the binned unfolding problem from point estimation to full posterior inference. Traditional binned methods

such as iterative Bayesian unfolding or singular value decomposition produce only point estimates with approximate uncertainty bands. In contrast, NPU leverages normalising flows to learn the complete posterior distribution over the unfolded bin contents, providing access to the full correlation structure and non-Gaussian features of the solution space [91]. This capability proves especially valuable when propagating uncertainties through complex analysis chains or when the unfolded results feed into likelihood fits for physics parameters.

Preliminary studies applying NPU to jet distributions in simulated LHC data have demonstrated several key advantages over conventional binned methods [91]. The normalising flow architecture provides natural regularisation through its limited expressivity, avoiding the *ad hoc* regularisation parameter tuning required by traditional methods. Furthermore, the learned posterior automatically adapts its regularisation strength across phase space, applying stronger smoothing where statistics are limited while preserving fine structure in well-populated regions.

The amortised inference capability of NPU addresses a critical computational bottleneck in systematic uncertainty evaluation. Once trained, the neural posterior estimator can instantly provide posterior distributions for new data without retraining, enabling efficient evaluation of systematic variations that would require complete re-execution of traditional unfolding methods. This computational efficiency becomes particularly important for analyses requiring hundreds of systematic variations or for data quality monitoring applications.

NPU's continued relevance in the era of unbinned methods stems from practical considerations. Many theoretical calculations, particularly fixed-order QCD predic-

tions, naturally provide binned predictions that can be directly compared with binned measurements. Legacy measurements spanning decades of particle physics research exist in binned form, and NPU enables improved reanalysis of historical data with modern statistical methods. Additionally, established practices in the field often favor binned representations for data preservation and sharing. Repositories like HEP-Data [442] and journal publications traditionally present results in tabulated form, making binned measurements more accessible for reinterpretation and combination with other results.

The familiarity of binned representations also facilitates communication with the broader physics community and enables straightforward comparison with the vast corpus of existing measurements. The method, in this way, serves as a valuable bridge between traditional and modern approaches, allowing experimentalists to gradually adopt machine learning techniques while maintaining familiar binned workflows. This transitional role proves particularly important for analyses where full unbinned methods face adoption challenges. Chapter IV provides a comprehensive treatment of NPU’s theoretical foundations and implementation details.

III.G.3 Invertible neural networks for unfolding.

Invertible neural networks and their conditional variants represent a theoretically motivated approach to unfolding that leverages the bijective structure of normalising flows. While these models learn deterministic bijective mappings between latent and data spaces, their integration with probabilistic frameworks requires careful interpretation.

The key insight of INN based unfolding lies in modelling the detector response as an invertible transformation. Bellagente et al. pioneered this approach for high energy physics applications, demonstrating that conditional normalising flows could learn the mapping between detector level and particle level distributions [312]. Their method conditions the flow on detector level observations to generate samples from the particle level distribution, effectively learning the inverse of the detector response function. The probabilistic interpretation of cINNs emerges not from the bijective mapping itself, but from the explicit modeling of the conditional distribution $p(z | x)$ through the change of variables formula laid out in Theorem 5 [95]. The learned transformation provides samples from this conditional distribution, though the quality of uncertainty estimates depends critically on how well the flow captures the true posterior structure.

This represents a different uncertainty quantification paradigm than traditional Bayesian approaches. The flow learns to generate samples that reflect the ambiguity in the inverse mapping rather than providing explicit credible intervals [267]. Experimental applications of cINNs to unfolding remain in the exploratory phase, with studies conducted on simulated data [95]. The computational efficiency of trained flows makes them attractive for applications requiring rapid inference, such as real-time data quality monitoring or iterative analysis optimisation.

III.G.4 MOMENT UNFOLDING

MOMENT UNFOLDING is a recently proposed method, focusing specifically on unfolding statistical moments of distributions rather than entire spectra [110]. This

method is particularly valuable when theoretical predictions are at the level of moments rather than differential distributions, as is often the case in QCD calculations, such as the energy dependence of the spectra of many quark and gluon jet observables, the strong coupling constant and leading power correction matrix elements [544, 600, 631, 641, 642, 654, 748, 750, 758].

MOMENT UNFOLDING employs a GAN-inspired structure with a generator that implements Boltzmann weighting factors to constrain the moments of the unfolded distribution. By focusing only on a small number of moments, this method provides natural regularisation and has demonstrated excellent precision in jet substructure applications where scaling behaviours and moments are of primary interest.

MOMENT UNFOLDING is also able to unfold statistical moments as a function of another observable (e.g., jet transverse momentum), enabling detailed studies of energy scale dependence without requiring full spectral unfolding. Feature specific methods like MOMENT UNFOLDING have the potential to offer precision and computation cost improvements over generalised unfolding methods for comparisons focused on individual features rather than the entire density. MOMENT UNFOLDING is discussed in further detail in Chapter V.

III.G.5 Reweighting Adversarial Networks (RANs)

Building on the conceptual framework introduced in MOMENT UNFOLDING, Reweighting Adversarial Networks (RANs) extend the methodology to unfold full spectra through an adversarial learning framework. A RAN implements particle level

reweighting functions guided by detector level classifiers, enabling non-iterative full spectral unfolding.

Unlike OMNIFOLD, which requires multiple iterations and separate classifiers at each step, a RAN employs a single adversarial training process to determine optimal weights. This approach offers computational advantages while maintaining performance comparable to iterative methods.

RANs have been used successfully on simulated data to unfold jet substructure observables in multidimensional spaces, where their non-iterative nature provides significant computational advantages. Implemented specifically as a Wasserstein GAN, a RAN is effective in scenarios with limited detector level overlap between Simulation and Data, a challenging regime for other reweighting based methods. Chapter VI delves into more detail about RANs.

III.G.6 Schrödinger Bridge Unfolding.

The application of Schrödinger bridge processes to unfolding represents a theoretically distinct approach that addresses specific limitations of existing ML methods [111, 303]. Unlike reweighting methods that can suffer from weight degeneracy or flow-based approaches that may struggle with multi-modal distributions, Schrödinger Bridge Unfolding frames the unfolding problem as finding the most likely stochastic evolution between the detector and particle level distributions.

This optimal transport formulation has some unique advantages. First, the entropic regularisation inherent in the Schrödinger Bridge framework naturally prevents overfitting without requiring *ad hoc* smoothness constraints. Second, the method

provides theoretical guarantees on the existence and uniqueness of solutions under mild conditions, offering more robust convergence properties than adversarial training approaches. Third, SBU excels at preserving multi-modal structure in distributions, a critical capability for measurements involving resonances or threshold effects where other regularisation approaches might artificially smooth important features [30].

Recent applications have demonstrated SBU’s effectiveness in challenging unfolding scenarios where traditional ML methods struggle [111]. For measurements involving rare processes with complex phase space boundaries, the optimal transport framework naturally respects the support of the distribution without requiring explicit constraints. This makes SBU particularly valuable for searches where signal regions occupy disconnected regions of phase space.

While the previously discussed methods approach unfolding through classification, direct generation, or adversarial training, the theoretical foundations of the unfolding problem connect deeply to optimal transport theory. The Schrödinger Bridge formulation represents the natural extension of this connection, providing a mathematically principled framework that unifies several desirable properties: entropic regularisation, theoretical guarantees, and the ability to interpolate between distributions in a physically meaningful way. This class of transport based methods may become increasingly important as the field seeks more theoretically grounded approaches to uncertainty quantification.

III.G.7 Unified Neural Folding GAN

The Jefferson Lab collaboration has recently proposed UNF-GAN, a generative adversarial network approach that seeks to address several limitations of existing ML based unfolding methods [157]. This method distinguishes itself through its unified treatment of both the forward folding and unfolding processes within a single adversarial framework, offering a more self consistent approach to modeling detector effects.

UNF-GAN employs a unique dual-generator architecture where one generator learns the forward mapping from particle level to detector level (approximating the detector response), while a second generator learns the reverse map for unfolding. The discriminator enforces consistency between both transformations, ensuring that the composition of forward and reverse mappings approximates the identity transformation. This bidirectional training strategy provides stronger constraints than methods that only model the unfolding direction, potentially reducing the solution space ambiguity inherent in ill posed inverse problems.

The method demonstrates particular strength in handling non-Gaussian detector resolutions and asymmetric migration effects. Traditional unfolding methods often assume approximately Gaussian smearing, but UNF-GAN's non-parametric approach can capture complex, multi-modal response functions that arise from threshold effects, detector non-linearities, or particle misidentification. Initial applications to simulations of electron scattering data at Jefferson Lab have shown improved resolution recovery compared to conventional methods, particularly in kinematic regions where detector acceptance varies rapidly.

A key innovation of UNF-GAN lies in its treatment of null spaces in the response matrix. By simultaneously learning both the forward and inverse mappings, the method can identify regions of phase space where information is irretrievably lost and provide appropriate uncertainty estimates rather than attempting to reconstruct information that does not exist in the data. This principled handling of information loss represents an advancement over methods that might produce spurious structure in under-constrained regions

III.H ML upstream and downstream of unfolding.

III.H.1	Uncertainty quantification in ML based unfolding.	202
III.H.2	Hybrid approaches and integration with simulation.	203
III.H.3	Equivariant networks for physics informed unfolding.	204
III.H.4	Comparison with traditional methods.	205
III.H.5	Implementation considerations from real experiments.	206

III.H.1 Uncertainty quantification in ML based unfolding.

Robust uncertainty quantification remains a central challenge in ML based unfolding, with the field developing multiple complementary approaches to address different sources of uncertainty. Bayesian neural networks have been explored as a framework for propagating parameter uncertainty through the unfolding process [55, 214, 231, 285, 321, 353, 393, 521, 585, 646]. These methods place prior distributions over network weights and use variational inference or Monte Carlo sampling to approximate the posterior weight distribution given the training data. The resulting ensemble of weight configurations provides a mechanism for quantifying uncertainty arising from finite training data and network parameter estimation. Vanilla BNNs primarily capture epistemic uncertainty [292] related to the network parameters themselves, rather than the systematic uncertainties arising from physics modelling choices or detector simulation accuracy, which must still be evaluated through traditional approaches such as varying the training samples or response matrices. However, extensions of BNNs that incorporate aleatoric uncertainties have been proposed in recent years [196, 395, 435].¹⁷

Ensemble methods have emerged as a practical alternative that balances computational efficiency with uncertainty coverage. Rather than the full Bayesian treatment, these approaches train multiple networks with different random initialisations, architectures, or training data subsets. The spread of predictions across the ensemble provides an empirical estimate of model uncertainty. Ensemble strategies for systematic uncertainty evaluation combine ensemble predictions with traditional approaches

¹⁷Though none in physics.

such as varying simulation parameters or detector conditions. The ensemble spread captures the sensitivity of the unfolding procedure to the specific optimisation path taken during training, which can be particularly important for methods involving non-convex optimisation landscapes.

For unbinned unfolding results, characterising uncertainties requires careful consideration of event level correlations. Standard statistical procedures that assume independence between events can significantly misestimate uncertainties when applied to correlated unfolded weights. The importance of accounting for these correlations in downstream analyses is highlighted in [44], with methods ranging from bootstrap resampling that preserves correlation structure to analytical approximations of the effective sample size. Chapter VII provides a detailed treatment of these correlation effects and their impact on statistical inference.

III.H.2 Hybrid approaches and integration with simulation.

Hybrid approaches that combine traditional and ML based unfolding have emerged as popular, pragmatic solutions in experimental settings. The ATLAS collaboration has explored methods that use ML techniques for response modeling while employing traditional matrix inversion for the actual unfolding step [259], leveraging the strengths of both approaches.

Integration of differentiable simulation into the unfolding pipeline represents another promising direction. By making detector simulations differentiable, these approaches enable end-to-end optimisation of the entire measurement chain [128, 176]. Recent work has also explored direct integration of theory predictions into ML based un-

folding [55], enabling simultaneous unfolding and parameter fitting. This unified approach can provide more direct constraints on physical parameters while properly accounting for all experimental effects.

III.H.3 Equivariant networks for physics informed unfolding.

Equivariant neural networks, which preserve symmetry transformations by design, have recently been applied to unfolding problems in particle physics. These architectures incorporate physical symmetries as inductive biases, naturally handling rotational, translational, or permutation symmetries common in physics applications.

Particle Flow Networks with equivariance properties have shown promise for unfolding jet constituent distributions [309], preserving permutation symmetry while capturing complex detector effects. These approaches leverage the rich structure of jets while respecting fundamental physical constraints. Equivariant graph neural networks have been explored for unfolding decay topologies at the LHCb experiment [43, 108, 179], where the underlying physical process exhibits various symmetries that can be exploited to improve unfolding precision and physical consistency.

These physics informed approaches highlight the broader trend toward integrating explicit physical knowledge into ML based unfolding methods, combining the flexibility of neural networks with the strong constraints provided by fundamental physical principles. Chapter VIII presents SYMMETRYGAN, a method that has been used to discover the symmetries which has subsequently been used to inform physics inspired networks.

III.H.4 Comparison with traditional methods.

An important aspect of ML for HEP research is the direct comparison of proposed ML based methods with traditional unfolding techniques in controlled experimental settings. The STAR collaboration applied both OMNIFOLD and traditional Singular Value Decomposition (SVD) methods to unfold jet measurements in heavy-ion collisions [200]. Analyses at ATLAS [41] and LHCb [169] have compared OMNIFOLD against IBU. Their results demonstrated that ML methods could achieve similar central values but with improved handling of boundary effects and systematic uncertainties.

Studies using simulated data can provide additional insights through controlled comparisons. OMNIFOLD has been compared against commonly used methods implemented in the RooUnfold package, including Iterative Bayesian Unfolding (IBU), SVD, and bin-by-bin correction factors in [270] on simulated top quark pair production data, and on neutrino cross section measurements [54]. A comprehensive survey by Huetsch et al. [55] evaluated multiple ML based unfolding methods alongside traditional approaches on standardised simulated datasets. This work emphasised that while traditional methods remain robust for simple, low dimensional problems, ML approaches offer unique advantages for modern analyses requiring unbinned, high dimensional unfolding.

These comparative studies consistently demonstrate that ML methods complement rather than replace traditional approaches. Traditional methods such as SVD and IBU remain valuable for their theoretical transparency, well understood uncertainty quantification, and computational efficiency in low dimensional problems. However,

ML methods excel in scenarios involving high dimensional phase spaces, complex detector responses, and the need for simultaneous unfolding of multiple correlated observables. The choice of method ultimately depends on the specific requirements of the analysis, with many experiments now employing both approaches as cross-checks to ensure robust results.

III.H.5 Implementation considerations from real experiments.

These case studies have revealed important practical considerations for implementing ML based unfolding in experimental analyses. One key lesson has been the importance of careful validation procedures. Experiments typically employ closure tests, linearity tests, and stress tests with modified simulation samples and in order to gain acceptance in the field, any proposed ML based unfolding method must be demonstrate robustness to these tests.

Another practical consideration is computational efficiency. While ML methods typically require significant resources for training, optimised implementations that leverage high performance computing resources can make these methods practical for large scale analyses [358].

Systematic uncertainty evaluation remains a continuing challenge. Recent analyses have developed methods for efficiently propagating systematic uncertainties through ML unfolding pipelines, often employing ensemble techniques or incorporating uncertainty sources directly into the training procedure [47, 54].

Chapter IV

NEURAL POSTERIOR UNFOLDING

IV.A	Motivation for improved binned unfolding	209
IV.A.1	Degeneracy and null spaces in unfolding.	210
IV.A.2	Regularisation limitations.	211
IV.B	Normalising flows for inverse problems.	213
IV.B.1	A Bayesian perspective on inverse problems.	214
IV.B.2	Conditional normalising flows for inverse problems.	215
IV.C	The NEURAL POSTERIOR UNFOLDING Algorithm	217
IV.C.1	Statistical foundation.	218
IV.C.2	Machine learning architecture.	219
IV.C.3	Addressing degeneracy with NPU.	221
IV.C.4	Implicit regularisation in NPU.	221
IV.D	Numerical results.	224
IV.D.1	2 bin degenerate response example.	225

IV.D.2	Gaussian example.	226
IV.D.3	Particle physics example.	232
IV.D.4	Summary of numerical results.	236
IV.E	Beyond binning: the path forward.	239

IV.A Motivation for improved binned unfolding

IV.A.1	Degeneracy and null spaces in unfolding.	210
IV.A.2	Regularisation limitations.	211

Definition 38 (Binned unfolding). A *binned unfolding method* is one in which both the particle level and detector level distributions are discretised into histograms, transforming the continuous convolution of Equation (I.16) into the matrix equation Equation (II.13).

IV.A.1 Degeneracy and null spaces in unfolding.

One of the fundamental challenges in unfolding collider data comes from degeneracies in the response matrix. Consider a case where two particle level bins α and β are indistinguishable at the detector level. This statement can be formalised as, there exists a detector level bin κ such that for all detector level bins ι :

$$R_{\iota\alpha} = R_{\iota\beta} = \delta_{\iota\kappa}, \quad (\text{IV.1})$$

where δ is the Kronecker delta.

In such scenarios, traditional methods like Iterative Bayesian Unfolding (IBU) produce catastrophically poor results, because the unfolded result becomes entirely dependent on the simulation, and learns no information from the data. The update rule for IBU takes the form

$$\forall n \in \mathbb{N} \quad \nu_{\alpha}^{(n)} = \left(\frac{\nu_{\alpha}^{(0)}}{\nu_{\alpha}^{(0)} + \nu_{\beta}^{(0)}} \right) \cdot \mu_{\kappa} \quad \nu_{\beta}^{(n)} = \left(\frac{\nu_{\beta}^{(0)}}{\nu_{\alpha}^{(0)} + \nu_{\beta}^{(0)}} \right) \cdot \mu_{\kappa}. \quad (\text{IV.2})$$

This means the relative contributions of bins α and β to the unfolded result are entirely determined by the MC, regardless of the observed data. Even worse, IBU (with, e.g., bootstraps) will report zero uncertainty on the ratio $\nu_{\alpha}/\nu_{\alpha}+\nu_{\beta}$ when in reality this ratio should have maximal uncertainty since the detector cannot distinguish between these contributions.

IV.A.2 Regularisation limitations.

To address the ill posedness detailed in Section II.C, traditional methods employ various forms of regularisation, examples of which are provided in Section II.D. However, all these approaches require a careful balance between fidelity to the data and conformity to the regularisation constraints. Traditional unfolding methods typically do not provide a natural framework for propagating uncertainties and can significantly underestimate the true uncertainty in degenerate regions. These limitations motivate the development of new unfolding methods that can

1. Naturally handle degeneracies and null spaces in the response matrix,
2. Provide proper uncertainty quantification, especially in poorly constrained regions,
3. Reduce dependence on arbitrary regularisation parameters, and
4. Offer computational efficiency for large scale problems.

NEURAL POSTERIOR UNFOLDING [91] addresses these challenges head on, by leveraging normalising flows for neural posterior estimation [469]. Unlike traditional methods that provide point estimates, NPU yields a full posterior distribution over the unfolded parameters, naturally capturing uncertainties in poorly constrained regions. Additionally, the regularisation in NPU emerges implicitly from the neural network architecture and training protocol, reducing the need for manual tuning of regularisation parameters.

The following sections introduce normalising flows and neural posterior estimation as the foundation for the NPU method, demonstrating how these modern machine learning techniques address the fundamental challenges of binned unfolding.

IV.B Normalising flows for inverse problems.

IV.B.1	A Bayesian perspective on inverse problems.	214
IV.B.2	Conditional normalising flows for inverse problems.	215

IV.B.1 A Bayesian perspective on inverse problems.

The unfolding problem can be naturally framed in Bayesian terms. Given detector level measurements x , we aim to reconstruct the posterior distribution over particle level quantities z ,

$$p(z | x) = \frac{p(x | z) p(z)}{p(x)}, \quad (\text{IV.3})$$

where $p(x | z) = r(x | z)$ is the detector response, $p(z)$ is the prior for particle level quantities, and $p(x)$ is the evidence.

Traditional methods typically focus on finding a point estimate that maximises this posterior, but a full Bayesian treatment would characterise the entire posterior distribution to properly account for uncertainties and degeneracies.

Normalising flows offer several key advantages for modelling the posterior distribution in inverse problems.

Complex distribution modelling: Unlike standard parametric approaches, normalising flows can represent multimodal, asymmetric posteriors that often arise in inverse problems with degeneracies where multiple particle level configurations can lead to similar detector level observations.

Amortised inference: Once trained, normalising flows enable rapid posterior sampling without requiring new optimisation runs for each new observation. This is particularly valuable when processing large datasets or when rerunning statistical analyses with different systematic variations.

End-to-end differentiability: The entire posterior estimation process can be optimised end to end, enabling gradient based learning approaches that are more efficient than traditional MCMC methods for high dimensional unfolding.

Implicit regularisation: Neural network architectures naturally introduce an inductive bias that serves as implicit regularisation, potentially reducing overfitting to statistical fluctuations compared to direct matrix inversion approaches without requiring explicit regularisation parameters.

IV.B.2 Conditional normalising flows for inverse problems.

In the context of unfolding, one is specifically interested in modelling the conditional distribution $p(z \mid x)$, the particle level distribution given detector level observations. Conditional normalising flows, described in Section III.E.1.ii.b.1 are an architecture almost perfectly designed for this task. A flow f_θ with learnable parameters θ can be trained to model the particle level distribution z ,

$$z = f_\theta(u \mid y), \tag{IV.4}$$

where u is drawn from a simple base distribution,¹ and x is the conditioning variable, detector level data.

This conditional structure allows the flow to learn the posterior distribution specifically tailored to each detector level observation, capturing the uncertainty and potential multimodality of the solution even in regions where the detector response is non-injective.

¹Typically a standard normal.

While Markov Chain Monte Carlo (MCMC) methods have traditionally been used for Bayesian inference in inverse problems [136, 427, 634, 635, 638]², normalising flows have some unique properties that make them especially well suited to this task.

1. No burn in period is required once the flow is trained.
2. Samples are generated independently rather than in a correlated chain.
3. The computational cost of generating samples does not increase with the dimensionality of the parameter space.
4. Due to the fully differentiable structure of the network, training can leverage GPU acceleration and parallelisation.

These advantages make normalising flows particularly attractive for inverse problems where MCMC methods may face mixing and convergence challenges or become computationally prohibitive to use.

²As they are in Fully Bayesian Unfolding [548].

IV.C The NEURAL POSTERIOR UNFOLDING

Algorithm

IV.C.1	Statistical foundation.	218
IV.C.2	Machine learning architecture.	219
IV.C.3	Addressing degeneracy with NPU.	221
IV.C.4	Implicit regularisation in NPU.	221

NEURAL POSTERIOR UNFOLDING (NPU) is a novel approach to unfolding that combines the rigour of Bayesian statistics with the flexibility and computational efficiency of modern machine learning. Rather than focusing solely on point estimates, NPU aims to characterise the full posterior distribution over particle level observables, providing comprehensive uncertainty quantification while addressing the key limitations of traditional unfolding methods.

IV.C.1 Statistical foundation.

At its core, NPU leverages normalising flows for neural posterior estimation in the context of binned unfolding. The fundamental insight is to directly estimate the conditional posterior distribution $p(z | x)$ through $\mathbb{P}(\boldsymbol{\nu} | \boldsymbol{\mu})$ the probability that the true particle level histogram counts are $\boldsymbol{\nu}$ given the measured detector level histogram counts are $\boldsymbol{\mu}$.

The relevant statistical quantities are

- A detector level histogram $\boldsymbol{\mu} = \{\mu_j\}_{j=1}^{N_D}$, where μ_j is the number of events in detector level bin j
- A particle level histogram $\boldsymbol{\nu} = \{\nu_i\}_{i=1}^{N_T}$, where ν_i is the number of events in particle level bin i
- A response matrix $\mathbf{R}$ with elements $R_{ji} = \mathbb{P}(\mu_j | \nu_i)$, the probability of an event in particle level bin i being measured in detector level bin j

The forward model relating these quantities is given by Equation (II.13). In this framework the posterior distribution of Equation (IV.3) is approximated by

$$\mathbb{P}(\boldsymbol{\nu} \mid \boldsymbol{\mu}) = \frac{\mathbb{P}(\boldsymbol{\mu} \mid \boldsymbol{\nu}) \mathbb{P}(\boldsymbol{\nu})}{\mathbb{P}(\boldsymbol{\mu})} , \quad (\text{IV.5})$$

where $\mathbb{P}(\boldsymbol{\mu} \mid \boldsymbol{\nu})$ is the likelihood, $\mathbb{P}(\boldsymbol{\nu})$ is the prior, and $\mathbb{P}(\boldsymbol{\mu})$ is the evidence. For Poisson distributed measurements, the likelihood takes the form

$$\mathbb{P}(\boldsymbol{\mu} \mid \boldsymbol{\nu}) = \prod_{j=1}^{N_D} \frac{(\sum_i R_{ji} \nu_i)^{\mu_j} e^{-\sum_i R_{ji} \nu_i}}{\mu_j!} . \quad (\text{IV.6})$$

NPU is conceptually similar to Fully Bayesian Unfolding (FBU), but replaces the traditional Markov Chain Monte Carlo (MCMC) sampling with normalising flows for neural posterior estimation, providing computational advantages and increased flexibility while maintaining the same full Bayesian treatment of the problem [500].

IV.C.2 Machine learning architecture.

NPU builds upon the neural posterior estimation framework, which uses conditional normalising flows to directly model the posterior distribution. The NPE framework comprises the base distribution, a simple distribution $p_u(u)$ that is easy to sample from, an invertible transformation, a learnable, invertible function $f_\theta(u \mid \boldsymbol{\mu})$ that transforms samples from the base distribution into samples from the approximate posterior, conditioned on the measured data $\boldsymbol{\mu}$, and a density transformation, encoding the the change of variables formula,

$$\mathbb{P}_\theta(\boldsymbol{\nu} \mid \boldsymbol{\mu}) = \mathbb{P}_u(f_\theta^{-1}(\boldsymbol{\nu} \mid \boldsymbol{\mu})) \cdot \left| \det \left\{ \nabla_{\boldsymbol{\nu}} f_\theta^{-1}(\boldsymbol{\nu} \mid \boldsymbol{\mu}) \right\} \right| . \quad (\text{IV.7})$$

The neural network parameters θ are optimised to maximise the likelihood of the true posterior, using pairs of simulated (ν_i, μ_j) examples. The training objective is to minimise the loss function

$$L(\theta) = -\mathbb{E}_{(\boldsymbol{\mu}, \boldsymbol{\nu}) \sim \mathbb{P}_{\text{sim}}(\boldsymbol{\nu}, \boldsymbol{\mu})} [\ln \mathbb{P}(\boldsymbol{\nu} \mid \boldsymbol{\mu})] \quad (\text{IV.8})$$

This approach allows the model to learn complex posterior distributions, including multimodal and highly correlated structures that often arise in unfolding problems [469].

The implementation uses a normalising flow based on the MADE (Masked Autoencoder for Distribution Estimation) architecture [408]. The model is composed of an invertible transformation network with three fully connected layers, each containing 50 to 100 nodes with Swish activation functions [447]. Conditional inputs incorporated through an auxiliary fully connected layer. The network is then trained for 1000 to 1500 epochs using the ADAM optimiser with a learning rate of 10^{-4} and a batch size of 10^4 . The loss function is the negative log likelihood described above in Equation (IV.8).

This architecture provides sufficient flexibility to model complex posterior distributions while remaining computationally tractable [376, 500]. After training the normalising flow, the unfolded response is obtained through a maximum likelihood estimation (MLE) step, where the negative log likelihood of the observed data conditioned on the learned model is minimised. This optimisation process also uses the same ADAM optimiser.

IV.C.3 Addressing degeneracy with NPU.

A major advantage of NPU over traditional unfolding methods is its ability to naturally handle degeneracies in the response matrix. Consider the two bin degenerate example presented earlier, where two particle level bins α and β are indistinguishable at the detector level. In traditional methods like IBU, the unfolded result becomes entirely dependent on the prior and the method incorrectly reports zero uncertainty on the ratio $\nu_\alpha/\nu_\alpha+\nu_\beta$.

In contrast, NPU naturally returns a broad posterior distribution that properly reflects the uncertainty. This is demonstrated in [91] with a two bin example where the response becomes increasingly degenerate, with correlation coefficient $\rho \rightarrow 1$. As the detector loses its ability to differentiate between truth bins, NPU produces credible intervals that encompass all values consistent with the total counts, appropriately reflecting the inherent uncertainty.

Simultaneously, while allowing for uncertainty in degenerate directions, NPU maintains constraints in well determined directions, such as the total sum $\nu_\alpha + \nu_\beta$. This behaviour emerges naturally from the Bayesian formulation, without requiring explicit identification of the degenerate subspaces. By modelling the full posterior, NPU provides a principled way to represent uncertainty in precisely those directions where the data provides little or no constraint.

IV.C.4 Implicit regularisation in NPU.

A significant advantage of NPU is its implicit regularisation, which arises naturally from the neural network architecture and training procedure rather than requiring

explicit regularisation terms or early stopping rules. Several mechanisms contribute to this implicit regularisation.

Network capacity: The finite capacity of the neural network limits the complexity of the posterior approximation, preventing overfitting to statistical noise in a manner similar to how traditional regularisation constrains solution complexity.

Smooth transformations: Normalising flows use smooth transformation functions, which naturally bias the solution toward smooth posterior distributions rather than ones with sharp, unphysical features.

Amortised inference: By learning a conditional distribution that applies across many possible measurements, NPU averages over many training examples, reducing sensitivity to outliers or statistical fluctuations in any single measurement.

Architectural inductive bias: The specific architecture of the normalising flow introduces an inductive bias. For instance, MADE architectures model dependencies between variables in a structured way that can align with physical correlations between bins.

Optimisation dynamics: The stochastic gradient descent training process itself provides regularisation through early stopping based on validation performance, preventing overfitting to the training data.

This implicit regularisation offers some advantages over the explicit regularisation used in traditional methods. It adapts automatically to the complexity of the problem

rather than requiring manual tuning of regularisation parameters. It can handle different degrees of regularisation for different regions of the solution space, providing stronger regularisation where constraints are weak and less regularisation where data provides strong constraints and it emerges from general principles rather than specific assumptions about the solution,³ potentially leading to less biased results. The effectiveness of this implicit regularisation can be observed empirically in cases where NPU produces unfolded distributions with excellent agreement with the truth, despite the ill posed nature of the problem.

The next section will demonstrate the application of NPU to concrete examples, including both controlled Gaussian distributions and realistic jet substructure observables, to illustrate its performance in practice.

³Such as smoothness or proximity to a prior.

IV.D Numerical results.

IV.D.1	2 bin degenerate response example.	225
IV.D.2	Gaussian example.	226
IV.D.2.i	Experimental setup.	227
IV.D.2.ii	Results and comparison.	229
IV.D.2.ii.a	Calibration and coverage.	229
IV.D.3	Particle physics example.	232
IV.D.3.i	Dataset and Observables	233
IV.D.3.ii	Results	234
IV.D.4	Summary of numerical results.	236

This section presents comprehensive empirical validation of NEURAL POSTERIOR UNFOLDING (NPU) through a series of increasingly complex examples. Consider first, a simple two-bin degenerate case to illustrate NPU’s advantages in handling unconstrained regions of phase space, before moving to a Gaussian example to systematically assess statistical performance, and finally studying the method’s efficacy on realistic jet substructure simulation from the LHC.

IV.D.1 2 bin degenerate response example.

The first example illustrates NPU’s behaviour by honing in degenerate scenarios where traditional methods struggle. Construct a two-bin setup where $\boldsymbol{\mu}, \boldsymbol{\nu} \in \mathbb{R}^2$ with a response matrix $\mathbf{R} \in \mathbb{R}^{2 \times 2}$. Fix the Truth values at $\nu_0 = \nu_1 = 5 \times 10^4$ and parametrise the response matrix using a correlation coefficient ρ and diagonal elements $\sigma_0 = \sigma_1 \equiv \sigma = 0.8$.

When $\rho = 0$, the response matrix is well conditioned with a diagonal structure, meaning that each detector bin primarily receives contributions from a single truth bin. In this scenario, both IBU, with uncertainty quantification via bootstrapping, and NPU yield unfolded distributions that align well with the truth, producing statistically consistent confidence regions. These results are demonstrated in Figure IV.1a, where both methods successfully recover the true distribution within their respective uncertainty bands.

However, as $\rho \rightarrow 1$, the response matrix becomes degenerate, i.e., $R_{i\alpha} = R_{i\beta}$, because the detector loses its ability to differentiate between the two Truth bins. This

degeneracy means that multiple Truth configurations can produce identical Data observations, making the unfolding problem underconstrained.


Under these conditions, IBU struggles with the ill posed nature of the problem. It provides only a point estimate based entirely on the prior distribution, while incorrectly estimating zero uncertainty for the ratio $\nu_\alpha/\nu_\alpha+\nu_\alpha$. This failure to capture the true uncertainty is an inescapable limitation of IBU when dealing with degenerate responses.

In stark contrast, NPU handles this degeneracy gracefully by returning a full posterior distribution with credible intervals that appropriately span all values consistent with the observed total counts. This behaviour accurately reflects the inherent uncertainty and ambiguity of the degenerate unfolding problem, as demonstrated in Figure IV.1b. The NPU approach thus provides a more honest and statistically sound treatment of cases where the data cannot uniquely determine the truth distribution.

This example highlights a fundamental advantage of NPU and other Bayesian methods; they naturally capture uncertainty in unconstrained regions of phase space, while traditional methods like IBU can dramatically underestimate uncertainties when degeneracies are present.

IV.D.2 Gaussian example.

The second test employs a Gaussian distribution with known truth, allowing one to systematically evaluate NPU's performance with varying levels of detector smearing.



(a) 2 bin example with non-degenerate response (b) 2 bin example with degenerate response

Figure IV.1: Comparison between NPU and IBU using a two bin example. The MLE estimate of NPU is shown with an orange star, the IBU prediction with a blue star, and the true value with a green star. The orange and blue shaded regions represent the 95% credible intervals of NPU and IBU, respectively. **(a)** Non degenerate, diagonal response matrix ($\rho = 0$): Both IBU and NPU yield unfolded distributions that agree well with the truth. **(b)** Degenerate response ($R_{i\alpha} = R_{i\beta}$): IBU relies solely on the prior, while NPU produces a full posterior with credible intervals that properly account for uncertainty in unconstrained regions. The NPU maximum likelihood estimate closely aligns with the true distribution, and the NPU posterior is compatible with the statistical uncertainty evaluated by IBU using bootstraps. Figure created by Jingjing Pan [91].

IV.D.2.i Experimental setup.

Generate two from identical one dimensional Gaussian distributions with mean $\mu = 0$ and standard deviation $\sigma = 1$, a Generation dataset with 10^6 events for training, and a Truth dataset with 10^5 events for validation.

To simulate detector effects, apply Gaussian smearing with parameter $\epsilon = 0.5$ to both particle level datasets. This produces corresponding detector level distributions,




Figure IV.2: Gaussian unfolding example setup. **(a)** The experimental design: Truth (blue) represents the target particle-level Gaussian distribution to be recovered, while Data (black crosses) shows the corresponding detector level observations after Gaussian smearing. The goal of unfolding is to recover the Truth distribution from the observed Data. **(b)** The response matrix $\mathbf{R}$ constructed from Gen./Sim. pairs, which encodes the detector smearing transformation and serves as an input for the unfolding algorithms. The matrix shows how probability mass spreads from each particle level bin (horizontal axis) to detector level bins (vertical axis). Figure created by Jingjing Pan [91].

Simulation, from Generation, and Data, from Truth, both with the same mean but with increased width $\sigma_{\text{detector}} = \sqrt{1^2 + 0.5^2} \approx 1.12$.

The response matrix $\mathbf{R}$ is estimated from (Gen., Sim.) pairs. This response matrix is then used to unfold the observed Data, with results compared against the known Truth. The experimental setup and resulting response matrix are shown in Figure IV.2.

IV.D.2.ii Results and comparison.

The unfolded distributions from NPU (maximum likelihood estimate), IBU, and FBU all accurately recover the truth distribution, as shown in Figure IV.3a. Beyond point estimates, NPU provides a complete characterisation of the uncertainty through its full posterior distribution, including the bin-to-bin correlations displayed in the corner plot in Figure IV.3b. This corner plot reveals the pairwise relationships between bins and demonstrates strong agreement between the posterior distributions from NPU and FBU, with both methods properly encompassing the true distribution.

IV.D.2.ii.a Calibration and coverage.

To rigorously assess the statistical calibration of NPU, evaluate pull distributions across 100 pseudo-experiments. For each bin i in pseudo-experiment j , compute the standardised residual:

$$\text{Pull}_{ij} = \frac{\mu_{ij}^{\text{method}} - \nu_i}{\sigma_{ij}^{\text{method}}} \quad (\text{IV.9})$$

where μ_{ij}^{method} and $\sigma_{ij}^{\text{method}}$ are the posterior mean and standard deviation for that bin, and ν_i is the true value.

Recall from Section II.G.2.i and Theorem 2 that for properly calibrated uncertainty estimates, these pulls should follow a standard normal distribution with zero mean and unit variance.

As demonstrated in Figure IV.4a, the pull distributions for both NPU and FBU with moderate smearing ($\epsilon = 0.5$) are indeed well centred at zero with unit width, confirming proper statistical calibration. One can further test the robustness of this





Figure IV.3: Gaussian unfolding results comparison. **(a)** Unfolded distributions: NPU MLE (orange down triangles), IBU (red circles), and FBU (green up triangles) compared with truth (blue shaded region). All methods successfully recover the true distribution. **(b)** Corner plot showing pairwise posterior correlations between the four central bins from NPU (orange contours) and FBU (green contours, averaged over MCMC samples), with truth values marked in blue. Both methods show consistent posterior distributions that properly encompass the true values. The green for FBU is averaged over many draws. Figure created by Jingjing Pan [91].

calibration by varying the smearing parameter across $\epsilon \in [0.3, 0.6]$, finding that both methods maintain proper coverage throughout this range. Figure IV.4b demonstrates the calibration as function of smearing parameters for various methods.

This highlights an important practical advantage of NPU is its computational efficiency for repeated unfolding tasks. In the benchmark with 100 pseudo experiments, FBU using 10,000 MCMC draws requires approximately 40 s per experiment, totalling about 67 min. NPU requires about 280 s for initial training but processes new datasets

Table IV.1: Computational cost comparison between Fully Bayesian Unfolding (FBU) and NEURAL POSTERIOR UNFOLDING (NPU) for repeated unfolding tasks. NPU’s amortised inference approach provides a substantial speed up despite the initial training overhead.

Method	Setup(s)	Per run(s)	Total(min)
FBU (1×10^4 MCMC draws)	–	40	67
NPU (amortised)	280	2	5
<i>Speed-up: 13× faster</i>			



(a) Pull distributions for $\epsilon = 0.5$

(b) Calibration across smearing parameters

Figure IV.4: Statistical calibration analysis from 100 pseudo-experiments. **(a)** Pull distributions for NPU (orange) and FBU (green) with moderate smearing ($\epsilon = 0.5$), compared against the ideal $\mathcal{N}(0, 1)$ reference (black line). Both methods show proper calibration with pulls centred at zero and unit variance. **(b)** Pull distribution statistics as functions of smearing parameter $\epsilon \in [0.3, 0.6]$. The mean and standard deviation remain consistent with the ideal values ($\mu = 0$, $\sigma = 1$) across the range, confirming robust uncertainty quantification for both methods. Figure created by Jingjing Pan [91].




Figure IV.4: (continued on next page)

in just seconds each, yielding a total time of about 5 min for all 100 experiments. This $13 \times$ speed-up stems from NPU's amortised inference approach, which eliminates repeated MCMC sampling for new datasets. Table IV.1 summarises the computational cost of NPU and FBU for 100 runs.

IV.D.3 Particle physics example.

This experiment demonstrates NPU's application to a realistic high energy physics scenario, focusing on jet substructure measurements at the Large Hadron Collider (LHC).




Figure IV.4: (Continued from previous page) Unfolding results for jet substructure observables comparing Neural Posterior Unfolding (NPU), Forward Bayesian Unfolding (FBU), and Iterative Bayesian Unfolding (IBU) methods. The response matrix is constructed using PYTHIA 8.243 (Tune 21) MC distributions, while HERWIG 7.1.5 provides both the Data to unfolding and Truth for validation. Truth distributions are shown as blue shaded regions, with unfolded results from NPU (maximum likelihood estimate as orange points), FBU (green points), and IBU (red points). Error bars represent posterior standard deviations from the respective methods. NPU successfully recovers the particle-level truth distributions across all observables. Figure created by Jingjing Pan [91].

IV.D.3.i Dataset and Observables

The dataset comprises simulated proton–proton collisions at $\sqrt{s} = 14$ TeV, following the setup in [309]. Two simulation generators are employed,

HERWIG 7.1.5 [421] serves as the data and truth distributions, and

PYTHIA 8.243 with Tune 21 [213] is used to construct the response matrices.

Detector effects are emulated using DELPHES 3.4.2 [513], a fast simulation of the CMS detector with particle flow reconstruction. Jets are clustered using the anti- k_T algorithm [608] with radius parameter $R = 0.4$, as implemented in FASTJET 3.3.2 [547]. To minimise acceptance effects, this analysis includes only the leading jets in events containing a Z boson with transverse momentum $p_T^Z > 200$ GeV. After selection, approximately 1.6×10^6 events from each simulation are retained.

This analysis focus on four key jet substructure observables that are widely used in LHC analyses.

Jet width (w): the transverse momentum weighted first radial moment of radiation within a jet,

Jet constituent multiplicity (M): the number of constituents in a jet,

N -subjettiness ratio ($\tau_{21} = \tau_2^{\beta=1}/\tau_1^{\beta=1}$): a measure of the compatibility of a jet with a two-prong substructure hypothesis relative to a one-prong hypothesis [67, 558, 581]

Groomed momentum fraction (z_g): the momentum shared between subjects after soft drop grooming [519, 533]

These observables span a diverse range of physical characteristics and detector sensitivities, providing a comprehensive test of NPU’s capabilities.

IV.D.3.ii Results

Figure IV.4 presents the unfolded distributions for each observable, comparing results from NPU, FBU, and IBU against the known truth. For the FBU implemen-

tation in this more complex scenario, one needs to increase the number of MCMC steps ten fold compared to the Gaussian example, using 100 000 tuning steps and 500 000 draws.

For all the observables, NPU accurately recovers the truth, with uncertainty bands that properly account for statistical uncertainties. The results demonstrate a few different characteristics of NPU. First, NPU effectively handles the complex detector effects present in realistic LHC simulations. Second, the method exhibits robustness against simulation modelling differences: the response matrix is constructed using PYTHIA events, while the “nature” distribution NPU was expected to unfold was derived from HERWIG, representing different underlying physics models and parton shower algorithms. Finally, full posterior information enables rigorous uncertainty quantification across the entire distribution. The corner plots for these distributions Figure IV.5 reveal the complex correlation structure between bins, information that is typically unavailable with traditional unfolding methods unless explicitly computed via bootstrapping or similar approaches.

These results highlight NPU’s value for practical cross section measurements at HEP experiments, where detector effects are complex and true distributions may differ significantly from simulation. In a complete experimental analysis, uncertainties in the response matrix itself could be incorporated either by repeating the unfolding procedure with systematically varied response matrices or by including these systematic uncertainties directly in the likelihood formulation.




Figure IV.5: (continued on next page)

IV.D.4 Summary of numerical results.

These numerical studies demonstrate that NPU provides appropriate uncertainty quantification in degenerate scenarios where traditional methods fail, while maintaining proper statistical coverage across varying degrees of detector smearing. NPU efficiently processes multiple datasets through amortised inference, and accurately recovers truth distributions in both targeted degenerate toy examples and realistic high energy physics scenarios. The method combines the Bayesian foundations of FBU with several architectural and computational advantages.

The normalising flow architecture can represent a wide range of posterior distributions, including multimodal, asymmetric, and strongly correlated distributions that may be challenging for traditional MCMC methods. By providing differentiable




Figure IV.5: (Continued from previous page) Corner plots showing posterior correlations for jet substructure observables. Each panel displays pairwise correlations between bins for one observable, comparing Neural Posterior Unfolding (NPU, orange contours) and Forward Bayesian Unfolding (FBU, green contours, averaged over Markov Chain Monte Carlo samples) posteriors, with truth values marked in blue. The plots reveal complex bin-to-bin correlation structures that are typically unavailable from traditional unfolding methods. Strong agreement between NPU and FBU posteriors confirms the reliability of both approaches, while the encompassing of truth values validates proper uncertainty quantification. These correlations provide crucial information for propagating uncertainties in downstream physics analyses. Figure created by Jingjing Pan [91].

access to the posterior density, NPU enables gradient based optimisation for finding maximum likelihood estimates or other derived quantities.

While FBU requires MCMC sampling for each new measurement, NPU's amortised inference approach front loads computational cost into the training phase. Once trained, inference with NPU requires only forward passes through the neural network,

making it particularly valuable for analysing large datasets, performing multiple analyses with different systematic variations, bootstrapping for uncertainty estimation, and studies requiring many unfolding runs NPU also scales more favourably to high dimensional problems compared to MCMC based methods, which often suffer from the curse of dimensionality and mixing problems in complex posterior landscapes.

However, several limitations should be kept in mind. The flexibility of normalising flows comes with increased model complexity, requiring careful architecture design and hyperparameter tuning. While less explicit than in traditional methods, NPU's results still depend on the distribution of the generated data, which implicitly defines a prior over the particle level histograms. Also, as with any Bayesian method, it is important to validate that the posterior credible intervals have the correct frequentist coverage properties for critical analyses.

Despite these considerations, NPU represents a significant advancement in unfolding methodology, combining the statistical rigour of Bayesian inference with the flexibility and computational efficiency of modern deep learning approaches making it a promising tool for cross section measurements in high energy physics.

IV.E Beyond binning: the path forward.

The NEURAL POSTERIOR UNFOLDING method presented in this chapter represents a significant advancement in binned unfolding approaches. By incorporating normalising flows for posterior estimation, NPU addresses several limitations of traditional methods while maintaining the established binning paradigm that has served particle physics for decades. However, binning itself introduces fundamental constraints that motivate the development of alternative, unbinned methodologies.

Binning inherently discards information about the precise location of events within each bin, reducing statistical precision. The idea that the choice of binning is subjective and can introduce biases in the unfolded results is oft expressed in statistical treatments of the unfolding problem [27, 113, 277, 344, 453, 477, 485, 512, 570, 586, 595, 666, 693]. This discretisation is particularly problematic for observables with sharp features or resonances that may be obscured by bin boundaries. The restrictions the curse of dimensionality imposes on binned methods’ applicability to multivariate problems has already been discussed in Section II.A.5.i.

Binned methods require separate unfolding procedures for each differential distribution of interest. This approach is inefficient when multiple observables derived from the same dataset need to be unfolded, as is common in comprehensive physics analyses.

Recent years have seen the rapid development of unbinned unfolding methods powered by advances in deep learning. These innovations have enabled significant progress towards solving the challenges posed by binned methods. Although machine learning methods are typically a few orders of magnitude slower to train than

traditional methods like IBU, modern ML techniques can leverage efficient data representations and parallelised computation for faster inference, reducing the gap particularly for high dimensional problems.

The next chapters will explore some unbinned methods in detail, examining their theoretical foundations, practical implementations and performance characteristics on realistic physics examples. These methods build upon the insights from binned approaches like NPU while overcoming the intrinsic limitations of binning to enable new capabilities for precision measurements in high energy physics. As experimental datasets grow larger and theoretical predictions become more precise, unbinned methods will play an increasingly important role in extracting the full physics potential from experiments. The transition from binned to unbinned unfolding represents not just a technical evolution but a paradigm shift in how the field approaches the measurement and interpretation of cross section measurements.

Chapter V

MOMENT UNFOLDING:

Direct deconvolution of distribution moments.

V.A	Why moments: physics context and QCD calculations.	243
V.A.1	The theoretical significance of moments in physics.	244
V.A.2	Moments in QCD calculations.	245
V.A.3	Experimental significance of moments in QCD.	247
V.A.4	Applications in jet physics.	248
V.A.5	The case for direct MOMENT UNFOLDING	250
V.B	A GAN like method to unfold moments	253
V.B.1	Boltzmann inspired reweighting	254
V.B.2	Adversarial training framework.	255

V.B.3	Mathematical formalism.	256
V.B.4	Theoretical properties.	258
V.B.5	Analytic closure.	259
V.B.6	Comparison with traditional GAN architectures.	265
V.B.7	Practical considerations.	266
V.B.8	Extension to differential measurements.	266
V.B.9	Relationship to statistical mechanics.	267
V.C	Machine learning implementation.	269
V.C.1	Training Procedure	271
V.C.2	Gradient updates.	273
V.C.3	Implementation details.	273
V.C.4	Extensions to multiple observables.	274
V.C.5	Uncertainty estimation.	275
V.C.6	Validation procedures.	276
V.C.7	Code availability and reproducibility.	276
V.D	Case studies.	278
V.D.1	Gaussian experiments.	279
V.D.2	Jet substructure in particle physics.	282
V.D.3	Momentum dependent unfolding.	285
V.E	Comparison with alternative methods.	294
V.F	Conclusion.	298
V.F.1	Towards unfolding distributions.	299

V.A Why moments: physics context and QCD calculations.

V.A.1	The theoretical significance of moments in physics.	244
V.A.2	Moments in QCD calculations.	245
V.A.2.i	DGLAP evolution equations.	245
V.A.2.ii	Operator product expansion.	246
V.A.2.iii	Event shape moments.	247
V.A.3	Experimental significance of moments in QCD.	247
V.A.4	Applications in jet physics.	248
V.A.4.i	Moments in physics beyond the Standard Model. .	250
V.A.5	The case for direct MOMENT UNFOLDING	250
V.A.5.i	The case for unfolding unbinned moments.	250

Statistical moments of probability distributions serve as powerful tools in physics, providing a concise and theoretically meaningful way to characterise complex phenomena. In the context of particle physics, and particularly in quantum chromodynamics (QCD), working at the level of moments rather than distributions offers significant advantages for both theoretical calculations and experimental measurements. This section explores the fundamental importance of moments in physics, their theoretical foundations, and their specific applications in QCD.

V.A.1 The theoretical significance of moments in physics.

Statistical moments represent a systematic way to characterise probability distributions, with each successive moment providing additional information about the shape and properties of the distribution.

Definition 39 (Statistical moment.). The k^{th} moment of random variable Z a probability density $p(z)$ is defined as

$$\mathbb{E}[Z^k] = \int_{-\infty}^{\infty} z^k p(z) \, dz . \quad (\text{V.1})$$

Each moment has a direct physical interpretation—the first moment corresponds to the mean, the second moment to the variance, and higher moments characterise asymmetry (skewness) and peakedness (kurtosis) and so on. Moments therefore provide much more concrete insight into the properties of an observable than distributions. While the complete probability distribution contains all available information, working with moments is often much more tractable.

In many physical theories, moments can be calculated analytically even when full distributions cannot. For example, in QCD, there exist relatively simple evolution equations for how moments change as a function of energy scale. If certain moments exhibit universal behaviour across different physical systems, this too reveals universality in the fundamental underlying principles. Hence, for many physical systems, including those governed by QCD, moments provide a natural language for connecting theoretical predictions with experimental measurements, particularly when examining how observables scale with energy.

V.A.2 Moments in QCD calculations.

In quantum chromodynamics, moments play a particularly crucial role in understanding the scaling behaviour of parton distributions and fragmentation functions. The DGLAP equations form the theoretical framework for parton distribution function evolution, and simply considerably when expressed at the level of moments.

V.A.2.i DGLAP evolution equations.

The Dokshitzer–Gribov–Lipatov–Altarelli–Parisi (DGLAP) evolution equations [542, 748, 750, 758] govern how parton distribution functions evolve with energy scale. While these equations are integro–differential equations in their most general form, they transform into ordinary differential equations when expressed in terms of moments, significantly simplifying their solution [413, 518, 541, 604, 618, 651, 659, 733, 755–757].

Definition 40 (Moments of PDFs.). For a parton distribution function $f(x, Q^2)$, where x is the momentum fraction and Q^2 is the energy scale, the k^{th} moment is defined as

$$M_k(Q^2) = \int_0^1 x^{k-1} f(x, Q^2) dx. \quad (\text{V.2})$$

The DGLAP equations for these moments take the form

$$\frac{dM_k(Q^2)}{d \ln Q^2} = \frac{\alpha_s(Q^2)}{2\pi} P_k M_k(Q^2). \quad (\text{V.3})$$

where P_k are the moments of the splitting functions [139, 526, 573, 580, 678]. This formulation converts the complex integro-differential evolution equations into a set of simple ordinary differential equations, one for each moment independently.

V.A.2.ii Operator product expansion.

The operator product expansion (OPE) in QCD naturally expresses deep inelastic scattering structure functions in terms of moments [624]. In this formalism, the moments of structure functions are related to matrix elements of local operators, which have clear physical interpretations and scaling properties. For example, the moments of the structure function $F_2(x, Q^2)$ can be expressed as

$$M_k(Q^2) = \int_0^1 x^{k-2} F_2(x, Q^2) dx = \sum_i C_k^i(Q^2) \langle O_k^i \rangle, \quad (\text{V.4})$$

where C_k^i are Wilson coefficients calculable in perturbation theory, and $\langle O_k^i \rangle$ are matrix elements of local operators [745].

V.A.2.iii Event shape moments.

Definition 41 (Event shape observables.). *Event shape observables*, are observables which characterise the geometrical properties of energy flow in collisions, are often studied through their moments.

Moments of event shapes like thrust, broadening, and jet mass distributions provide sensitive tests of perturbative QCD and allow precise extractions of the strong coupling constant α_S [432, 709]. The moments of these event shapes can be calculated as

$$\langle \varepsilon^n \rangle = \int_0^{\varepsilon_{\max}} \varepsilon^n \frac{1}{\sigma} \frac{d\sigma}{d\varepsilon} d\varepsilon, \quad (\text{V.5})$$

where ε is the event shape variable, and $\frac{1}{\sigma} \frac{d\sigma}{d\varepsilon}$ is its normalised differential cross section. These moments can be predicted in perturbative QCD, typically expressed as power series in α_S :

$$\langle \varepsilon^n \rangle = A_n \left(\frac{\alpha_S}{2\pi} \right) + B_n \left(\frac{\alpha_S}{2\pi} \right)^2 + \mathcal{O}(\alpha_S^3), \quad (\text{V.6})$$

plus non-perturbative corrections that scale as inverse powers of the centre of mass energy.

V.A.3 Experimental significance of moments in QCD.

From an experimental perspective, moments offer several advantages that make them particularly valuable for QCD studies.

Some of the most precise determinations of the strong coupling constant α_S come from measurements of moments. For example, moments of event shapes measured at LEP have provided determinations of α_S with uncertainties at the few percent

level [640, 645]. Studies by TRISTAN [710, 723], TASSO [730, 732, 744], OPAL [637, 685, 690], DELPHI [653, 707], L3 [649, 662, 663, 671, 674, 683, 684, 689, 719], and ALEPH [687, 695, 712] collaborations have all used moment based analyses to extract α_S values that contribute significantly to the world average.

The scaling behaviour of moments with energy provides a direct and robust window into the running of QCD couplings. By measuring how moments of distributions change with energy scale, experiments can observe the logarithmic scaling violations that are a hallmark of asymptotic freedom in QCD [466, 718, 751, 756].

Many theoretical QCD calculations are more tractable and more precise for moments than for full distributions. For certain observables, perturbative calculations may exist to next to next to leading order (NNLO) or beyond for moments, while full distribution calculations may only be available at lower orders [431, 601, 709, 714, 715, 737].

In practice, measuring moments directly can reduce certain experimental uncertainties, particularly those related to detector resolution and acceptance effects in regions of phase space with sparse statistics.

V.A.4 Applications in jet physics.

In the context of jet physics specifically—a central arena for testing QCD—moments of jet substructure observables provide particularly powerful probes of QCD dynamics [106, 533]. Several notable applications deserve special mention.

The energy scale dependence of jet properties is best studied through their moments. The moments of jet substructure observables such as jet mass, width, and multiplicity

exhibit characteristic scaling with jet energy that can be predicted by perturbative QCD [533]. For example, the first moment of the jet mass scales as

$$\langle m^2 \rangle \propto \alpha_S(p_T) p_T^2 \ln \frac{R}{R_0} , \quad (\text{V.7})$$

where p_T is the jet transverse momentum, R is the jet radius, and R_0 is the reference scale [9, 747]. Higher moments have different scaling behaviours, providing multiple handles on the underlying QCD dynamics.

The moments of jet substructure variables are also used in the classification of jets. This is because jet moments differ characteristically between quark-initiated and gluon-initiated jets, making them valuable for jet flavour tagging. For instance, gluon jets typically have higher multiplicity and broader mass distributions than quark jets, reflected in the moments of these distributions [164, 253, 563].

Jet substructure moments also provide a window to study non-perturbative effects and hadronisation. The interplay between perturbative and non-perturbative effects in QCD is particularly evident in moments of jet observables. Lower moments often capture the perturbative, large angle radiation physics, while higher moments become increasingly sensitive to non perturbative effects like hadronisation [672, 691].

In the context of Soft-Collinear Effective Theory (SCET), moments of jet substructure observables provide direct constraints on the parameters of the effective theory, offering a bridge between first principles QCD and phenomenological models of jet formation [510, 588, 626].

V.A.4.i Moments in physics beyond the Standard Model.

While moments are crucial for precision QCD studies, their utility extends to searches for BSM physics. Certain moments of distributions can exhibit enhanced sensitivity to new physics scales. For example, higher moments of mass distributions can be more sensitive to heavy particle contributions than the full distribution average [163, 679].

Moments can therefore provide a model independent way to parametrise deviations from Standard Model predictions, similar to the role of effective field theory coefficients. By measuring a series of moments, experiments can place constraints on a wide class of possible new physics scenarios without committing to specific models [241, 336]. Hence, changes in the moment structure of distributions can serve as an early indicator of anomalous physics, even before the full nature of the anomaly is understood [124, 241, 297].

V.A.5 The case for direct MOMENT UNFOLDING**V.A.5.i The case for unfolding unbinned moments.**

Traditionally, there was no straightforward method to directly unfold moments. Instead, moments were computed after first measuring and unfolding the full distribution. This typically involved binning the data, unfolding the binned distribution, and then computing moments from the unfolded histogram. The limitation of this approach though, is the binning itself, not the moments *qua* moments.

The use of discrete bins introduces a bias in moment calculations. For a variable z binned into n_{bins} with centers $z_{\text{bin},i}$ and counts N_i , the binned approximation of the k^{th} moment is

$$\mathbb{E}_{\text{binned}} [z^k] = \frac{1}{N} \sum_{i=1}^{n_{\text{bins}}} N_i z_{\text{bin},i}^k. \quad (\text{V.8})$$

This approximation introduces a bias relative to the true moment, which grows with the order of the moment computed, and is particularly large in regions where the distribution varies rapidly within bins.

These limitations motivate the development of unbinned approaches that can directly unfold moments without first discretising the data, which is precisely the focus of the MOMENT UNFOLDING method described in this chapter. Given the theoretical significance of moments in QCD and their experimental advantages, there is a strong motivation for developing methods that can directly unfold moments from detector level data without first unfolding the entire distribution. Such methods would eliminate binning biases inherent in traditional approaches, while potentially improving precision and reducing computational cost by focusing directly on the quantities of interest. They would enable moment measurements in higher dimensional phase spaces where binned approaches become impractical to provide more direct comparisons with theoretical predictions that are formulated at the level of moments.

The remainder of this chapter presents the MOMENT UNFOLDING method, a novel approach that addresses precisely this need by directly deconvolving moments of distributions using machine learning techniques inspired by the Boltzmann distribution and implemented through a GAN like architecture. This method offers a significant over extant methods in moment measurements, enabling more stringent tests of

QCD predictions and potentially uncovering subtle effects that might be obscured in traditional analyses.

V.B A GAN like method to unfold moments

V.B.1	Boltzmann inspired reweighting	254
V.B.2	Adversarial training framework.	255
V.B.3	Mathematical formalism.	256
V.B.4	Theoretical properties.	258
V.B.4.i	Connection to the maximum entropy principle. . .	259
V.B.5	Analytic closure.	259
V.B.5.i	Moment learning with perfect detector response. .	261
V.B.5.ii	Universal but imperfect detector response.	262
V.B.5.ii.a	On closure and its obstructions.	264
V.B.6	Comparison with traditional GAN architectures.	265
V.B.7	Practical considerations.	266
V.B.8	Extension to differential measurements.	266
V.B.9	Relationship to statistical mechanics.	267

The direct extraction of distribution moments without first reconstructing the entire distribution presents a unique challenge in unfolding methodology. Traditional approaches typically reconstruct the full differential cross section before computing moments, introducing unnecessary computational complexity. The MOMENT UNFOLDING technique presented in this chapter takes an entirely different approach by directly targeting the moments themselves.

The basic idea underlying MOMENT UNFOLDING is the connection between statistical moments and the Boltzmann distribution from statistical mechanics [501, 746]. In thermodynamics, the Boltzmann distribution represents the probability distribution that maximises entropy while satisfying certain constraints on average quantities (the internal energy). This principle can be adapted for the unfolding problem, one formulation of which seeks to find a reweighting of Generation events that accurately reproduces the moments of the true particle level distribution.

V.B.1 Boltzmann inspired reweighting

The central mathematical construction of MOMENT UNFOLDING is a reweighting function inspired by the Boltzmann factor,

$$g(z; \boldsymbol{\beta}) = \frac{1}{P(\boldsymbol{\beta})} \exp \left[- \sum_{a=1}^n \beta_a z^a \right], \quad (\text{V.9})$$

where z is the particle level observable whose moments are to be unfolded, β_a are parameters to be determined,¹ P is a normalisation factor similar to the partition function, and n is the number of moments to be unfold simultaneously.

¹Analogous to Lagrange multipliers in used by Boltzmann to derive his eponymous distribution.

This form is particularly powerful because it directly connects to the maximum entropy principle in statistical mechanics. Just as the Boltzmann distribution maximises entropy subject to constraints on average energy, this reweighting function maximises the binary cross entropy loss of the discriminator’s classification output between the reweighted distribution and the prior distribution while enforcing constraints on the moments. When applied to Generation with probability density $q(z)$, this reweighting function produces a modified distribution

$$\tilde{q}(z) = g(z) \cdot q(z) . \quad (\text{V.10})$$

The training objective, then, is to determine the optimal values of the parameters β_a such that the moments of reweighted Generation, $\tilde{q}(z)$, match those of Truth, $p(z)$.

V.B.2 Adversarial training framework.

Given the theoretical foundation provided by the Boltzmann form, one now needs a practical method to determine the optimal parameters. The task can very naturally be framed as a two level optimisation problem.

1. The parameters β_a control the reweighting at the particle level, and
2. The quality of the reweighting must be assessed at the detector level

This task is highly non trivial because the detector response function, which maps particle level to detector level observables, is typically stochastic and known only implicitly through simulation.

To address this, MOMENT UNFOLDING employs an adversarial training architecture inspired by Generative Adversarial Networks (GANs), but with significant modifications tailored to the unfolding problem. Unlike traditional GANs that generate new samples, MOMENT UNFOLDING’s “generator” does not generate samples. Instead, $g(z; \beta)$ generates weights that reweight existing Generation events according to the Boltzmann form. The discriminator is a neural network classifier $d(x)$ that attempts to distinguish between Data and reweighted Simulation at detector level.

The system is then trained adversarially, with the generator attempting to fool the discriminator while the discriminator attempts to correctly classify events. This setup is illustrated in Figure V.1, showing the flow of information and the adversarial training process. In this framework, the optimal parameters β_a are those that make the reweighted Simulation indistinguishable from the Data. Once these parameters are found, the moments of the reweighted Generation are the unfolded moments.

V.B.3 Mathematical formalism.

The MOMENT UNFOLDING approach can be formalised as follows. Let $p(x)$ be the probability density function of the Data distribution and $q(z, x)$ be the joint distribution of particle level and detector level variables in the MC, Generation and Simulation. Then $q(z)$ and $q(x)$ are the marginal distributions of Generation and Simulation respectively. Let $r(x|z)$ be the detector response function. The assumption of a universal detector response can be represented as as

$$p(x|z) = q(x|z) = r(x|z). \quad (\text{V.11})$$

MOMENT UNFOLDING




Figure V.1: A schematic diagram of the MOMENT UNFOLDING algorithm. Like a GAN, g is the generator and d is the discriminator, but unlike a GAN g does not generate samples, it instead generates reweighting factors given by Equation (V.9). The reweighted Simulation dataset inherits its weight from the matching Generation dataset. The detector emulations are only run once, since a new simulated dataset is created via importance weights and not by changing the features themselves. The unfolded moments are validated against Truth distributions.

The goal is to find parameters β_a such that the moments of the reweighted distribution $g(z) \cdot q(z)$ match the true moments of the underlying Truth, $p(z)$.

The probability density function of the reweighted Simulation is

$$\tilde{q}(x) = \int g(z) q(z, x) dz. \quad (\text{V.12})$$

The adversarial training process optimizes a weighted binary cross entropy loss function,

$$L[g, d] = -\mathbb{E}_{X \sim p}[\ln d(X)] - \mathbb{E}_{X \sim \tilde{q}}[(1 - d(X))] \quad (\text{V.13})$$

This loss function encourages the discriminator to output high values for real data and low values for reweighted simulation, while the generator is trained to reweight the simulation to make it indistinguishable from real data.

Once the optimal parameters β_a are determined, the unfolded moments can be computed directly from the reweighted Generation

$$\mathbb{E}_{\text{unfolded}} [Z^k] = \frac{\sum_{z \in \text{Gen.}} g(z; \beta) \cdot z^k}{\sum_{z \in \text{Gen.}} g(z; \beta)} . \quad (\text{V.14})$$

This approach differs fundamentally from traditional unfolding methods that first reconstruct the full distribution and then compute moments. By directly targeting the moments themselves, MOMENT UNFOLDING is able to avoid both the binning artifacts and dimensionality challenges associated with binned methods and the computational complexity associated with unbinned full distribution unfolding.

V.B.4 Theoretical properties.

The MOMENT UNFOLDING method possesses several important theoretical properties that justify its application to physics problems. Under certain conditions, detailed in Section V.B.5, the MOMENT UNFOLDING method provably converges to the correct moments.

The parameter n (the number of moments being unfolded) serves as an implicit regularisation parameter. By limiting the number of moments, one can constrain the complexity of the reweighting function, helping to stabilise the unfolding process against statistical fluctuations. This creates a natural bias–variance trade-off: including more moments allows for a more flexible reweighting function but increases

sensitivity to statistical fluctuations, including fewer moments leads to stable unfolding but the generator may not have sufficient flexibility to capture interesting features of the data. In practice, the optimal choice of n depends on the physics application and available statistics.

V.B.4.i Connection to the maximum entropy principle.

The Boltzmann form of the reweighting function has a deep connection to the maximum entropy principle in statistical mechanics. The reweighting function maximises the entropy of the reweighted distribution relative to the simulated distribution, subject to constraints on the moments. This provides a principled physical and informational theoretic basis for the method: among all possible reweighting functions that reproduce the target moments, the Boltzmann form is the one that introduces the least additional information beyond what is required by the moment constraints.

V.B.5 Analytic closure.

Let p and q be probability densities on $\mathbb{R}$ with respect to Lebesgue measure such that $q(z) > 0$ for all z in the support of p . Fix an integer $n \geq 1$ and assume the first n moments of p exist and are finite. Define

$$\mathbb{E}_\rho[Z^a] := \int z^a \rho(z) \, dz \tag{V.15}$$

for the a^{th} moment under a density ρ . All integrals below are assumed finite and all variational steps are justified by dominated convergence.

Theorem 6 (Entropy projection with moment constraints). Under the assumptions above, the minimiser of $D_{\text{KL}}(\ell \| q)$ over densities all ℓ subject to $\mathbb{E}_\ell[Z^a] = \mathbb{E}_p[Z^a]$, $a = 0, 1, \dots, n$ (where $a = 0$ enforces normalisation) exists in the exponential family

$$\ell^*(z) = q(z) \exp\left(-1 - \sum_{a=0}^n \beta_a z^a\right), \quad (\text{V.16})$$

for some Lagrange multipliers $\{\beta_a\}_{a=0}^n$.

Proof of Theorem 6. Form the Lagrangian

$$\mathcal{L}(\ell, \beta) = D_{\text{KL}}(\ell \| q) + \sum_{a=0}^n \beta_a \left(\int z^a \ell(z) \, dz - \mathbb{E}_p[Z^a] \right). \quad (\text{V.17})$$

A standard variational calculation gives the Gâteaux derivative

$$\frac{\delta \mathcal{L}}{\delta \ell(z)} = \ln \frac{\ell(z)}{q(z)} + 1 + \sum_{a=0}^n \beta_a z^a. \quad (\text{V.18})$$

Setting this to zero yields Equation (V.16). Feasibility (i.e., choice of β) enforces the constraints. $\square$

Write $\ell^*(z) = g(z) q(z)$. Then

$$g(z) = \frac{1}{P} \exp\left(-\sum_{a=1}^n \beta_a z^a\right), \quad P = \int q(z) \exp\left(-\sum_{a=1}^n \beta_a z^a\right) \, dz, \quad (\text{V.19})$$

where β_0 has been absorbed into the normalisation P .

The remaining multipliers $\{\beta_a\}_{a=1}^n$ are fixed by the moment constraints and need not admit a closed form. Note that Equation (V.16) and everything that follows rely on the asymmetry of $D_{\text{KL}}(\cdot \| \cdot)$; replacing $D_{\text{KL}}(\ell \| q)$ with $D_{\text{KL}}(q \| \ell)$ changes the solution class.

V.B.5.i Moment learning with perfect detector response.

Let g be as in Equation (V.19) and define the reweighted density $\tilde{q}(z) := q(z)g(z)$.

Definition 42 (MLC loss). Consider the asymptotic *maximum likelihood classifier* (MLC) loss

$$L(d, g) = - \int p(z) \ln d(z) dz - \int \tilde{q}(z)(1 - d(z)) dz, \quad (\text{V.20})$$

where $d : \mathbb{R} \rightarrow (0, 1]$ is a discriminator.

Theorem 7 (Optimal discriminator and reduced objective). For fixed g , the unique maximiser of $d \mapsto -L(d, g)$ is $d^*(z) = \frac{p(z)}{\tilde{q}(z)}$. At d^* ,

$$\tilde{L}(g) = L(d^*, g) = -D_{\text{KL}}(p \parallel \tilde{q}). \quad (\text{V.21})$$

Proof of Theorem 7. Varying d in Equation (V.20) gives

$$\frac{\delta L}{\delta d(z)} = - \frac{p(z)}{d(z) + \tilde{q}(z)}. \quad (\text{V.22})$$

Hence $d^*(z) = p(z)/\tilde{q}(z)$. Substituting and using $\int p = \int \tilde{q} = 1$ gives Equation (V.21). $\square$

Theorem 8 (Closure under perfect response). Any stationary point of $g \mapsto \tilde{L}(g)$ of the form Equation (V.19) satisfies

$$\forall a \in \{1, \dots, n\} \quad \mathbb{E}_{\tilde{q}}[Z^a] = \mathbb{E}_p[Z^a]. \quad (\text{V.23})$$

Proof of Theorem 8. From Equation (V.21),

$$\tilde{L}(g) = - \int p(z) \ln \tilde{q}(z) dz + \text{const.} \quad (\text{V.24})$$

Hence, the functional derivative with respect to g is

$$\frac{\delta \tilde{L}}{\delta g(z)} = \frac{p(z)}{g(z)}. \quad (\text{V.25})$$

For the parametrisation Equation (V.19),

$$\forall a \in \{1, \dots, n\} \quad \frac{\partial g(z)}{\partial \beta_a} = -g(z)(z^a - \mathbb{E}_{\tilde{q}}[Z^a]). \quad (\text{V.26})$$

Since $\frac{\partial \ln P}{\partial \beta_a} = -\mathbb{E}_{\tilde{q}}[Z^a]$, thus

$$\frac{\partial \tilde{L}}{\partial \beta_a} = \int \frac{\delta \tilde{L}}{\delta g(z)} \frac{\partial g(z)}{\partial \beta_a} dz \quad (\text{V.27})$$

$$= - \int p(z) (z^a - \mathbb{E}_{\tilde{q}}[Z^a]) dz \quad (\text{V.28})$$

$$= -(\mathbb{E}_p[Z^a] - \mathbb{E}_{\tilde{q}}[Z^a]). \quad (\text{V.29})$$

At any stationary point, this vanishes for $a = 1, \dots, n$, proving Equation (V.23). $\square$

V.B.5.ii Universal but imperfect detector response.

For g as in Equation (V.19), define

$$\tilde{q}(x) = \int r(x | z) q(z) g(z) dz. \quad (\text{V.30})$$

Theorem 9 (Reduced objective at detector level). For the discriminator acting on x , the discriminator optimised MLC loss is

$$\tilde{L}(g) = -D_{\text{KL}}(p(x) \| \tilde{q}(x)). \quad (\text{V.31})$$

Proof of Theorem 9. Identical to Theorem 7 after replacing z by x and using Equation (V.30). $\square$

Theorem 10 (Moment matching). Under the assumptions above, any stationary point of $g \mapsto \tilde{L}(g)$ in the class Equation (V.19) satisfies

$$\forall a \in \{1, \dots, n\} \quad \mathbb{E}_{\tilde{q}}[Z^a] = \mathbb{E}_{p_{\text{mod}}}[Z^a], \quad (\text{V.32})$$

where the modified particle level density p_{mod} is

$$p_{\text{mod}}(z) = \int \frac{p(x) r(x | z) \tilde{q}(z)}{\tilde{q}(x)} dx. \quad (\text{V.33})$$

Proof of Theorem 10. By Equation (V.31) and the chain rule for functional derivatives,

$$\frac{\partial \tilde{L}}{\partial \beta_a} = \iint \frac{\delta \tilde{L}}{\delta \tilde{q}(x)} \frac{\delta \tilde{q}(x)}{\delta g(z)} \frac{\partial g(z)}{\partial \beta_a} dz dx. \quad (\text{V.34})$$

$$\frac{\delta \tilde{L}}{\delta \tilde{q}(x)} = \frac{p(x)}{\tilde{q}(x)} \quad (\text{V.35})$$

and

$$\frac{\delta \tilde{q}(x)}{\delta g(z)} = r(x | z) q(z). \quad (\text{V.36})$$

Using Equation (V.26) and $\tilde{q}(z) = q(z) g(z)$ gives

$$\frac{\partial \tilde{L}}{\partial \beta_a} = - \int \left(z^a - \mathbb{E}_{\tilde{q}}[Z^a] \right) \left(\int \frac{p(x)}{\tilde{q}(x)} r(x | z) dx \right) \tilde{q}(z) dz \quad (\text{V.37})$$

$$= - \left(\mathbb{E}_{p_{\text{mod}}}[Z^a] - \mathbb{E}_{\tilde{q}}[Z^a] \right), \quad (\text{V.38})$$

where p_{mod} is as in Equation (V.33). Stationarity enforces Equation (V.32). $\square$

Theorem 11 (Transfer operator representation). Let

$$\tilde{q}(z | x) := \frac{r(x | z) \tilde{q}(z)}{\tilde{q}(x)}. \quad (\text{V.39})$$

Define the transfer kernel

$$f(z | z') := \int \tilde{q}(z | x) r(x | z') dx. \quad (\text{V.40})$$

Then

$$p_{\text{mod}}(z) = \int f(z | z') p(z') dz'. \quad (\text{V.41})$$

Moreover, if $r(x | z) = \delta(x - z)$ (perfect detector), then $f(z | z') = \delta(z - z')$ and $p_{\text{mod}} = p$.

Proof of Theorem 11. Substitute $p(x) = \int r(x | z') p(z') dz'$ into Equation (V.33) and exchange integrals

$$p_{\text{mod}}(z) = \iint \frac{r(x | z) \tilde{q}(z)}{\tilde{q}(x)} r(x | z') dx p(z') dz' \quad (\text{V.42})$$

$$= \int \left(\int \tilde{q}(z | x) r(x | z') dx \right) p(z') dz'. \quad (\text{V.43})$$

If $r(x | z) = \delta(x - z)$, then $\tilde{q}(z | x) = \delta(z - x)$ and Equation (V.40) gives $f(z | z') = \delta(z - z')$. $\square$

V.B.5.ii.a On closure and its obstructions.

Equation (V.23) shows exact closure in the ideal detector limit. With a universal but imperfect detector, Equation (V.32) guarantees closure to the modified law p_{mod} , not necessarily to p . Deviations $p_{\text{mod}} \neq p$ grow when (i) the forward response r induces large distortions so that the inverse kernel $\tilde{q}(\cdot | x)$ in Equation (V.39) depends

sensitively on g , or (ii) q has poor overlap with p , forcing g to be large in regions where r is ill conditioned. These limitations are generic to unfolding and not specific to the MLC objective or the MOMENT UNFOLDING algorithm.

V.B.6 Comparison with traditional GAN architectures.

While MOMENT UNFOLDING draws inspiration from GANs, it differs from traditional GAN architectures, in that the so called “generator” does not generate data; it generates weights. Traditional GANs generate new samples by mapping random noise through a neural network. In contrast, MOMENT UNFOLDING reweights existing simulated events, preserving the physics correlations already present in the MC and focusing only on correcting the distribution’s moments.

Furthermore, unlike conventional GANs where the generator network is incredibly flexible with hundreds of millions of trainable parameters [239, 335, 360], the generator for MOMENT UNFOLDING has an extremely specific parametric form constrained by the Boltzmann expression, with at most a handful of trainable parameters. The Boltzmann constraint massively reduces the effective degrees of freedom in the optimisation process and encodes prior knowledge about the problem structure. The parameters β_a in the MOMENT UNFOLDING approach have clear physical interpretations as Lagrange multipliers enforcing moment constraints. This interpretability contrasts with the typically opaque nature of neural network weights in traditional GANs.

Finally, traditional GANs operate at a single level, with the generator and discriminator acting on the same “type” of data. MOMENT UNFOLDING operates across two levels, particle level and detector level, which might not even have the same

dimensionality, which is the source of the additional complexity but is essential for the unfolding problem.

V.B.7 Practical considerations.

While theoretically elegant, the MOMENT UNFOLDING method has practical considerations and limitations that must be understood in order to apply the method correctly and effectively.

Like all reweighting based methods, MOMENT UNFOLDING requires that the Generation distribution has overlapping support with the Truth distribution. Regions of phase space not covered by the MC cannot be properly unfolded. With finite statistics, extremely large weights can arise, in regions of phase space with sparse MC coverage as the generator tries and fails to reweight a zero density into a non-zero density. Various regularisation techniques must be employed to mitigate these effects if they arise in a given analysis.

Like other unfolding methods, MOMENT UNFOLDING too involves an *ad hoc* choice that must be imposed externally. The optimal number of moments to unfold depends on the physics application, available statistics, and the complexity of the true distribution. Too few moments may miss important features of the distribution, while too many may lead to overfitting and numerical instabilities.

V.B.8 Extension to differential measurements.

An extension of the MOMENT UNFOLDING approach is to study moments as a function of another observable, such as transverse momentum. This is achieved by

making the parameters β_a functions of the conditional variable,

$$g(z; \boldsymbol{\beta}(p_T)) = \frac{1}{P(\boldsymbol{\beta}(p_T))} \exp \left[- \sum_{a=1}^n \beta_a(p_T) z^a \right]. \quad (\text{V.44})$$

This allows the study of how moments evolve with energy scale or other experimental conditions, providing direct comparisons with theoretical predictions of scaling behaviours in QCD.

The parameters $\beta_a(p_T)$ can be modelled as polynomials, splines, or neural networks, depending on the expected complexity of their dependence on p_T . This extension significantly enhances the physics applications of the method, enabling studies of how the moments of one observable scale with another.

V.B.9 Relationship to statistical mechanics.

The mathematical form at the heart of MOMENT UNFOLDING has deep connections to statistical mechanics and information theory, providing a solid theoretical foundation for the method. In statistical mechanics, the Boltzmann distribution arises as the solution to the maximum entropy problem. The maximum entropy problem is the problem of finding the probability distribution that maximises entropy subject to constraints on expectation values (typically of energy). The solution takes the form

$$p(s) = \frac{1}{Z} e^{-\beta E(s)}, \quad (\text{V.45})$$

where $E(s)$ is the energy of state s , β is the inverse temperature, and Z is the partition function.

The MOMENT UNFOLDING algorithm draws inspiration from this idea to include multiple constraints on different moments, with each β_a serving as a Lagrange

multiplier for the constraint on the a^{th} moment. This connection to maximum entropy principles provides a theoretical justification for the specific form of the reweighting function and suggests that it is, in some sense, the most natural choice for the moment unfolding problem. The success of this approach in practical applications, as will be demonstrated in subsequent sections, confirms the value of this theoretical grounding and highlights the power of cross disciplinary approaches in developing new data analysis techniques for particle physics.

V.C Machine learning implementation.

V.C.1	Training Procedure	271
V.C.2	Gradient updates.	273
V.C.3	Implementation details.	273
V.C.4	Extensions to multiple observables.	274
V.C.5	Uncertainty estimation.	275
V.C.6	Validation procedures.	276
V.C.7	Code availability and reproducibility.	276

The theoretical framework introduced in the previous section provides the foundation for MOMENT UNFOLDING, but implementing this approach in practice requires careful architectural design and optimisation strategies. This section details the machine learning implementation that enables robust moment extraction through a GAN-like architecture.

The generator in MOMENT UNFOLDING differs fundamentally from traditional GAN generators. Instead of producing new samples through a complex neural network, the generator implements the Boltzmann-inspired weighting function, Equation (V.9). This function is applied to existing Generation events to reweight them, producing a weighted distribution that matches the moments of the true distribution. The parameters β_a are the trainable parameters of the generator, and P is a normalisation constant estimated at the batch level,

$$\hat{P} = \sum_{z \in \text{batch}} \exp \left[- \sum_{a=1}^n \beta_a z^a \right]. \quad (\text{V.46})$$

For differential moment measurements, where the unfolded information is the dependence of moments on another variable like transverse momentum, the β_a parameters become functions of p_T , as described in Equation (V.44). The experiments described below parametrise $\beta_a(p_T)$ as linear functions,

$$\beta_a(p_T) = \beta_a^{(0)} + \beta_a^{(1)} p_T. \quad (\text{V.47})$$

This linear parametrisation is motivated by empirical observations that the ratios of spectra between different simulations often exhibit approximately linear dependence on kinematic variables.

The discriminator network $d(x)$ takes detector level features as input and outputs a probability that a given event came from the real data rather than the reweighted simulation. This implementation uses a dense neural network with three fully connected hidden layers with 50 nodes per layer. ReLU [786] activation functions are applied to intermediate layers, and a sigmoid activation function is applied to the output layer.

This relatively simple architecture provides sufficient capacity to distinguish between data and simulation distributions. The training objective is formulated as a weighted version of the maximum likelihood classifier (MLC) loss [411],

$$L[g, d] = - \sum_{x \in \text{Data}} \ln d(x) - \sum_{(z, x) \in (\text{Gen.}, \text{Sim.})} g(z)(1 - d(x)) , \quad (\text{V.48})$$

where the first sum is over events from Data and the second sum is over matched pairs of Generation and Simulation events, weighted by the generator function $g(z)$.

This loss function differs from the standard Binary Cross Entropy (BCE) loss often used in GANs, providing better convergence properties for our specific application. While the BCE loss function could also be used (and indeed has been verified to give similar empirical performance), the MLC loss function provides certain theoretical guarantees about the convergence of the method to the correct moments, demonstrated in Section V.B.5 [110].

V.C.1 Training Procedure

The training procedure for MOMENT UNFOLDING follows the standard adversarial approach, but with specific adaptations for the reweighting framework.

The generator parameters β_a are initialised to small random values near zero, corresponding to minimal reweighting.² The discriminator weights are initialised using standard neural network initialisation techniques[490].

Each training step involves sampling batches of events from Data, Generation and Simulation. The training proceeds by alternating between optimising the discriminator and the generator. First, with the generator parameters fixed, the discriminator weights are updated to minimise the loss function. Then with the discriminator weights fixed, the generator parameters are updated to maximise the loss function. This alternating optimisation continues for a fixed number of epochs, or until convergence criteria are met.

The ADAM optimizer [436] is employed, with an initial learning rate of 5×10^{-4} for both the discriminator and generator. The learning rate is reduced by a factor of 0.5 every 20 epochs to ensure stable convergence.

To ensure stable training and prevent overfitting, a few regularisation techniques are employed. Gradient values are clipped to the range $[-1, 1]$ to prevent large updates that could destabilise training, particularly important for the generator parameters which can be sensitive to statistical fluctuations. Batch normalisation [491] is applied after each hidden layer in the discriminator to stabilise training and reduce sensitivity to weight initialisation and learning rates.

²I.e., starting close to the original MC.

V.C.2 Gradient updates.

For the discriminator, gradient computation follows standard backpropagation through the neural network. For the generator, gradients are computed with respect to the parameters β_a by differentiating the loss function.

$$\frac{\partial L}{\partial \beta_a} = - \sum_{(z,x) \in (\text{Gen.}, \text{Sim.})} \frac{\partial g(z)}{\partial \beta_a} (1 - d(x)), \quad (\text{V.49})$$

where

$$\frac{\partial g(z)}{\partial \beta_a} = -g(z) \left[z^a - \frac{1}{P} \frac{\partial P}{\partial \beta_a} \right] = -g(z) (z^a - \mathbb{E}_g [z^a]). \quad (\text{V.50})$$

Here, $\mathbb{E}_g [z^a]$ represents the a^{th} moment computed from the reweighted distribution. This formulation provides a direct connection between the gradient updates and the moment matching objective of the method.

V.C.3 Implementation details.

The method is implemented using KERAS [483] with the TENSORFLOW2 backend [460], leveraging their automatic differentiation capabilities for efficient gradient computation. NUMPY [333] is used for data manipulation, and MATPLOTLIB [617] is employed for visualisation.

The computational requirements of MOMENT UNFOLDING are remarkably light compared to many deep learning applications. Training typically takes less than five minutes per observable on an NVIDIA RTX6000 GPU for the case studies presented in this dissertation. Even the potentially computationally intensive calculation of the partition function P , which requires summation over all events in a batch is implemented efficiently using vectorised operations in TENSORFLOW.

Table V.1: Summary of training hyperparameters used in the implementation of the MOMENT UNFOLDING algorithm. These values control the architecture, optimisation behaviour, and regularisation of the training process.

Hyperparameter	Value	Description
Discriminator layers	3	Number of hidden layers in discriminator
Discriminator nodes	50	Number of nodes per hidden layer
Activation function	ReLU	Activation function for hidden layers
Output activation	Sigmoid	Activation function for the output layer
Optimizer	ADAM	Optimization algorithm used for training
Learning rate	5×10^{-4}	Initial learning rate for the optimizer
Batch size	10 000	Number of training events in each batch
Training epochs	50 to 100	Number of full passes through data
Gradient clip value	1.0	Maximum allowed gradient norm

Table V.1 summarises the key hyperparameters used in the implementation. These hyperparameter values were determined through systematic experimentation and provide a good balance between training stability and model performance across a range of physics scenarios.

V.C.4 Extensions to multiple observables.

While the basic formulation of MOMENT UNFOLDING applies to moments of a single observable, the method can be extended to handle multiple observables simultaneously. For multiple observables $z_1, z_2, \dots, z_d$, one can define joint moments as $\mathbb{E} \left[z_1^{k_1} z_2^{k_2} \dots z_d^{k_d} \right]$. The reweighting function then takes the form,

$$g(z_1, z_2, \dots, z_d) = \frac{1}{P} \exp \left[- \sum_{k_1, k_2, \dots, k_d} \beta_{k_1, k_2, \dots, k_d} z_1^{k_1} z_2^{k_2} \dots z_d^{k_d} \right] \quad (\text{V.51})$$

In practice, including all possible joint moments can quickly become intractable as the number of observables increases. To address this, one can employ strategies such as

- Limiting the total order of moments (e.g., $k_1 + k_2 + \dots + k_d \leq K_{max}$),
- Including only specific joint moments known to be physically relevant, and
- Using partially factorised forms that capture the most important correlations while maintaining computational feasibility.

V.C.5 Uncertainty estimation.

Accurate uncertainty estimation is crucial for meaningful physics measurements. In *MOMENT UNFOLDING*, uncertainties on the extracted moments are obtained through bootstrap resampling. Statistical uncertainties are estimated by generating N bootstrap replicas of the data by resampling with replacement, applying *MOMENT UNFOLDING* to each replica, and computing the standard deviation of the resulting moment values across all replicas. This procedure accounts for statistical uncertainties in both the data and the simulation.

Systematic uncertainties related to theoretical modelling, detector effects, and other sources can be evaluated by varying the underlying simulation according to systematic uncertainty sources, applying *MOMENT UNFOLDING* to each variation, and taking the envelope of the resulting moment values as the systematic uncertainty.

Such an approach would allow for comprehensive uncertainty quantification that accounts for all relevant sources of uncertainty in the measurement.

V.C.6 Validation procedures.

Several validation procedures are employed to ensure the reliability of the *MOMENT UNFOLDING* method. Closure tests verify that the method can correctly recover known moments from simulated data. The method’s internal consistency is verified by checking that the moments computed from the reweighted Generation match the target moments, the reweighted Simulation reproduces the features of the Data, and that the discriminator is unable to distinguish between Simulation and Data.

Results from *MOMENT UNFOLDING* are also compared with those from traditional unfolding methods like Iterative Bayesian Unfolding followed by moment calculation. Agreement between different methods provides additional confidence in the results, while disagreement may indicate methodological issues that require further investigation.

V.C.7 Code availability and reproducibility.

The implementation of *MOMENT UNFOLDING* is available as open-source Python code, facilitating reproducibility and adoption by the scientific community. This open-source approach promotes transparency, enables independent verification of results, and facilitates improvements and extensions to the method by the broader scientific community. The full code repository can be found at [45], and the specific datasets used for the studies in this dissertation are archived on Zenodo [3].

By combining a theoretically motivated approach with practical machine learning implementation details, *MOMENT UNFOLDING* provides a robust and efficient method for extracting moments directly from data. The next sections will demonstrate the

application of this method to Gaussian data and simulated jet substructure data, showcasing its performance on realistic physics problems.

V.D Case studies.

V.D.1	Gaussian experiments.	279
V.D.1.i	Experimental Setup	279
V.D.1.ii	Results	279
V.D.2	Jet substructure in particle physics.	282
V.D.2.i	Datasets.	282
V.D.2.ii	Observables	282
V.D.2.iii	Results.	284
V.D.3	Momentum dependent unfolding.	285
V.D.3.i	Inclusive distributions.	289
V.D.3.ii	Differential analysis.	289

V.D.1 Gaussian experiments.

Before applying *MOMENT UNFOLDING* to jet physics scenarios, it is essential to validate the method in a controlled environment where the ground truth is precisely known. To this end, the approach is first demonstrated using a one dimensional Gaussian example, providing a clear illustration of the method’s capabilities while establishing a benchmark for its performance.

V.D.1.i Experimental Setup

For this controlled study, generate datasets from Gaussian distributions with known parameters. The Truth dataset comprises 10 000 events drawn from a Gaussian distribution with mean $\mu_{\text{Truth}} = 0.0$ and variance $\sigma_{\text{Truth}}^2 = 1.0$. The Generation dataset comprises 100 000 events drawn from a Gaussian distribution with mean $\mu_{\text{Gen.}} = -0.5$ and variance $\sigma_{\text{Gen.}}^2 = 1.0$. Both distributions are then subjected to detector effects modelled as additive Gaussian noise with resolution parameter $\sigma_{\text{det}} = 0.5$. This setup mimics the typical scenario in particle physics where the MC differs from the true distribution, and both are observed through an imperfect detector.

Figure V.2 illustrates the particle level and detector level distributions, showing how the detector resolution affects both the Truth and Generation datasets.

V.D.1.ii Results

Figure V.3 shows the results of applying *MOMENT UNFOLDING* to the Gaussian example. The left panel displays the particle level distributions, the Truth distribution (blue), the Generation (orange), and the reweighted Generation (black dashed line).




Figure V.2: Effect of detector resolution on Truth and Generation datasets. The Truth dataset comprises 10 000 events drawn from a Gaussian distribution with mean $\mu_{\text{Truth}} = 0.0$ and variance $\sigma_{\text{Truth}}^2 = 1.0$. The Generation dataset comprises 100 000 events drawn from a Gaussian distribution with mean $\mu_{\text{Gen.}} = -0.5$ and variance $\sigma_{\text{Gen.}}^2 = 1.0$. Both distributions are subjected to detector effects modelled as additive Gaussian noise with resolution parameter $\sigma_{\text{det}} = 0.5$. **Left:** Particle level distributions showing the underlying Truth (blue) and Generation (orange) samples before detector effects. **Right:** Detector level distributions after applying detector resolution, demonstrating the broadening effect while maintaining the relative offset between the distributions. This setup mimics the typical scenario in particle physics where the Monte Carlo simulation differs from the true distribution, and both are observed through an imperfect detector.

Even visually, the close overlap between the Truth histogram and the reweighted Generation demonstrates the success of the MOMENT UNFOLDING procedure.

The right panel of Figure V.3 provides a more quantitative perspective by showing the discriminator optimised loss landscape as a function of the Boltzmann parameters β_1 and β_2 . This landscape visualisation allows us to verify that the maximum of the loss function (red star) aligns with the parameter values that produce the correct moments, providing empirical confirmation of the method's theoretical properties.




Figure V.3: (a) Distributions from the Gaussian example of particle level truth, generation, and reweighted generation i.e. MOMENT UNFOLDING. The agreement between the truth and reweighted samples demonstrates the qualitative performance of MOMENT UNFOLDING. (b) The weighted MLC loss from Equation (V.48) for fixed g but optimised over d , found by scanning over β_1 and β_2 . The correct value is indicated by a red star. Indicated in shaded blue is the 1σ bootstrapped interval for MOMENT UNFOLDING's prediction of β_a . The lower panel in each plot shows the ratio to truth.

These results confirm that MOMENT UNFOLDING successfully recovers the true moments of normal distributions. Since normal distributions are fully characterised by their first two moments, the reweighted distribution closely matches the entire truth distribution in this Gaussian case even though the method explicitly targets only the first two moments. We should not expect this to be the case for more complex distributions with non-trivial higher moments.

This controlled study confirms that MOMENT UNFOLDING performs as expected in a well understood scenario, providing confidence to extend the method to more complex physics environments.

V.D.2 Jet substructure in particle physics.

Having validated MOMENT UNFOLDING in a controlled environment, this section applies the method to a realistic high energy physics scenario: measuring moments of jet substructure observables in simulated proton–proton collisions at the Large Hadron Collider (LHC).

V.D.2.i Datasets.

The simulated samples used for this study follow the setup from [309]. Proton–proton collisions at $\sqrt{s} = 14$ TeV are simulated using PYTHIA 8.243 with Tune 26 (standing in for MC) and HERWIG 7.1.5 (standing in for nature). The DELPHES 3.4.2 fast simulation of the CMS detector is used to model detector effects, configured with particle flow reconstruction. Jets are clustered using the anti- k_T algorithm with radius parameter $R = 0.4$, implemented in FASTJET 3.3.2. To reduce acceptance effects, only leading jets in events with a Z boson having transverse momentum $p_T^Z > 200$ GeV are included. This setup tests MOMENT UNFOLDING in a realistic environment where the Truth (HERWIG particle level) and Generation (PYTHIA particle level) are different generators with distinct physics models, and both are observed through an imperfect detector to produce Data (HERWIG + DELPHES) and Simulation (PYTHIA + DELPHES).

V.D.2.ii Observables

Four key jet substructure observables are unfolded, each providing different insights into the internal structure of jets.

Jet mass (m): The invariant mass of the jet, calculated as

$$m = \sqrt{\sum_k E_k^2 - \sum_k \mathbf{p}_k^2}, \quad (\text{V.52})$$

where the sum runs over all constituents of the jet, and E_k and $\mathbf{p}_k$ are their energies and 3-momenta.

Jet charge (q): An infrared-safe but not collinear-safe observable that measures the jet's electric charge [186, 459].

$$q = \frac{1}{\sum_k \sqrt{p_{T,k}}} \sum_k q_k \sqrt{p_{T,k}}, \quad (\text{V.53})$$

where q_k is the electric charge of constituent k .

On account of being non collinear safe, jet charge is expected to be more susceptible to certain kinds of detector distortions. The first moment of the jet charge distribution as a function of jet p_T has been calculated in [334, 337, 338, 537, 559] and measured by ATLAS [459] and CMS [349, 452].

Jet width (w): A measure of the radial distribution of radiation within the jet.

$$w = \frac{1}{p_{T,\text{jet}}} \sum_k p_{T,k} \Delta R_k, \quad (\text{V.54})$$

where ΔR_k is the angular distance between constituent k and the jet axis.

Groomed momentum fraction (z_g) [495]: The momentum sharing between sub-jets after Soft Drop grooming [519, 533] with $z_{\text{cut}} = 0.1$ and $\beta = 0$.

$$z_g = \frac{p_{T,\text{subleading}}}{p_{T,\text{leading}} + p_{T,\text{subleading}}}, \quad (\text{V.55})$$

where $p_{T,\text{leading}}$ and $p_{T,\text{subleading}}$ refer to transverse momentum of the leading and subleading prongs respectively [219].

Table V.2: Moments of jet observables at particle level. First and second moments of m , q , w and z_g are shown for truth (HERWIG), generation (PYTHIA) and MOMENT UNFOLDING. Uncertainties in the truth and generation columns are estimated via bootstrap resampling; uncertainties in the unfolding column combine in quadrature the generation bootstrap uncertainty with the empirical 1σ spread from repeated unfolding on the same dataset.

		Truth	Generation	MOMENT UNFOLDING
$\mathbb{E}[M]$	$(\times 10^1)$	2.182 ± 0.030	2.064 ± 0.043	2.173 ± 0.047
$\mathbb{E}[M^2]$	$(\times 10^2)$	6.049 ± 0.222	5.360 ± 0.350	6.115 ± 0.364
$\mathbb{E}[Q]$	$(\times 10^{-2})$	1.006 ± 0.037	1.582 ± 0.038	1.090 ± 0.040
$\mathbb{E}[Q^2]$	$(\times 10^{-2})$	1.216 ± 0.082	1.508 ± 0.074	1.207 ± 0.074
$\mathbb{E}[W]$	$(\times 10^{-1})$	1.498 ± 0.025	1.231 ± 0.029	1.499 ± 0.029
$\mathbb{E}[W^2]$	$(\times 10^{-2})$	3.370 ± 0.113	2.421 ± 0.128	3.374 ± 0.128
$\mathbb{E}[Z_g]$	$(\times 10^{-1})$	2.334 ± 0.029	2.457 ± 0.030	2.353 ± 0.059
$\mathbb{E}[Z_g^2]$	$(\times 10^{-2})$	6.789 ± 0.166	7.425 ± 0.165	6.767 ± 0.330

These observables span a diverse range of jet properties, including both infrared and collinear safe observables (m and w), a Sudakov safe observable, z_g , and q , which is only infrared safe.

V.D.2.iii Results.

Table V.2 presents the first and second moments of each jet observable extracted using MOMENT UNFOLDING, compared with the true values (from HERWIG) and those from the generation (PYTHIA) in tabular form. The uncertainties in the Truth and Generation columns are computed by bootstrapping the respective datasets. For the MOMENT UNFOLDING column, the uncertainty includes the bootstrap uncertainty from the Generation and the empirical uncertainty from the unfolding process added

in quadrature. The empirical uncertainty from the unfolding process is computed by repeating *MOMENT UNFOLDING* on the same data and computing the 1σ spread

These results demonstrate that *MOMENT UNFOLDING* successfully recovers the true moments from the data for all observables within statistical uncertainties. The method effectively corrects for both the differences between the *PYTHIA* and *HERWIG* physics models and the detector distortions introduced by *DELPHES*.

Figure V.4 shows the particle level distributions for each observable, comparing the truth, generation, and reweighted generation after applying *MOMENT UNFOLDING*. While the method explicitly targets only the first two moments, it produces reasonable agreement across the entire distribution, particularly for the regions that contribute most significantly to the moments.

Figure V.5 displays the discriminator optimised loss landscapes for each observable, confirming that the maximum of the loss function aligns with the parameter values that produce the correct moments. The blue ellipses represent the 1σ confidence intervals for the *MOMENT UNFOLDING* parameters, and they all contain the true parameter values (red dots), indicating proper uncertainty estimation.

V.D.3 Momentum dependent unfolding.

One of the most powerful features of *MOMENT UNFOLDING* is its ability to extract moments as a function of another observable, such as jet transverse momentum. When implemented with momentum dependent parameters in the Boltzmann weight function, the generator described in Equation (V.44) unfolds moments conditional on the jet p_T .




Figure V.4: Distributions of (a) jet mass, (b) jet charge, (c) jet width, and (d) groomed momentum fraction at particle level for leading jets in $Z + \text{jet}$ events at $\sqrt{s} = 14$ TeV. The plots compare Truth, Generation, and Generation weighted by MOMENT UNFOLDING. MOMENT UNFOLDING successfully corrects Generation to match Truth across all four jet substructure observables. The lower panel in each plot shows the ratio to Truth.

This ability to unfold moments as functions of p_T is a necessary feature for any modern unfolding method. Many QCD phenomena exhibit strong scale dependence, including parton shower evolution, hadronisation effects, and detector response. Understanding the momentum dependence is necessary to account for these variations.




Figure V.5: The discriminator optimised maximum likelihood classifier (MLC) loss landscapes as a function of the MOMENT UNFOLDING parameters β_1 and β_2 for jet substructure observables. The true parameter values that correctly map between the generation and truth distributions are indicated by red dots, while the 1σ confidence intervals from MOMENT UNFOLDING’s parameter estimation are shown as blue ellipses. (Continued on next page.)

Jets at different p_T values populate different regions of the substructure phase space. A momentum dependent reweighting is therefore important for better coverage of the phase space. The momentum dependent approach effectively leverages the statistical power of the entire dataset while allowing for local adaptations in different p_T regions, resulting in more precise measurements overall.

The implementation of momentum dependent parameters, $\beta_a(p_T)$ introduces additional complexity that must be carefully managed. For this data, a linear parametrisation of $\beta_a(p_T)$ provides a good balance between flexibility and stability for the jet substructure observables studied.³ More complex parametrisations might be

³MOMENT UNFOLDING can be run as a binned method by parametrising $\beta_a(p_T)$ as a piecewise constant function.




Figure V.5: (Continued from previous page.) Loss landscapes for ((a), (b) on previous page) jet mass and charge, and (c) jet width and (d) groomed momentum fraction. The alignment of the loss function maxima with the true parameter values and the containment of these values within the uncertainty ellipses demonstrates both the accuracy of the method and the reliability of its uncertainty estimation.

warranted for observables with a strongly non-linear p_T dependence. However, increasing the number of parameters in the momentum dependent approach must be approached with caution, since that, in some ways, lifts the regularisation imposed by the MOMENT UNFOLDING method. If the training remains unstable despite appropriately constraining the functional form of $\beta_a(p_T)$ regularisation in the neural network training may become necessary. Some successful approaches to regularise such neural network trainings can be found in Section VI.C.2.

V.D.3.i Inclusive distributions.

Figure V.6 presents the inclusive distributions of jet mass, jet charge, jet width, and groomed momentum fraction after applying momentum dependent *MOMENT UNFOLDING*. Compared to the non-momentum dependent results shown in Figure V.4, these distributions demonstrate markedly better agreement with the truth. This improvement occurs because the momentum dependent approach effectively implements a more refined correction, allowing the reweighting function to adapt to the changing detector response and physics differences across the p_T spectrum. Parametrising β_a as a function of p_T accounts for the fact that both detector effects and physics modelling discrepancies can vary significantly with jet p_T .

The inclusive distributions obtained through momentum conditioned unfolding also reveal subtle physical features that might be obscured in the non-conditioned approach. For example, in the jet mass distribution, the shape of the low mass region shows better agreement with truth when using momentum conditioning, correctly capturing the interplay between soft radiation and resolution effects that vary with jet p_T .

V.D.3.ii Differential analysis.

When the momentum dependent unfolding approach is used for a differential analysis of moments as a function of p_T , it provides a comprehensive understanding of scaling behaviours. This dual perspective offers particularly powerful constraints on theoretical models, as it simultaneously tests the overall distribution shapes and






Figure V.6: Inclusive distributions of (a) jet mass, (b) jet charge, (c) jet width, and (d) groomed momentum fraction obtained using momentum dependent MOMENT UNFOLDING. The plots compare Truth, Generation, and Generation weighted by momentum conditioned MOMENT UNFOLDING. Compared to the momentum independent unfolding shown in Figure V.4, the momentum dependent approach achieves notably better agreement with the Truth. The lower panel in each plot shows the ratio to truth.

their evolution with energy scale. Figure V.7 shows the first and second moments of each observable as a function of p_T .

The results demonstrate that MOMENT UNFOLDING successfully recovers the p_T dependent moments of all the observables, capturing the variations in these quantities with jet p_T . This capability is particularly valuable for testing theoretical predictions of how jet properties scale with energy, a fundamental aspect of QCD phenomenology.




(a) Mean




(b) Variance

Figure V.7: Mean and variance of jet mass as a function of jet transverse momentum. The plots compare Truth (HERWIG), Generation (PYTHIA), and reweighted Generation, weighted with momentum dependent MOMENT UNFOLDING results. (Continued on next page.)




(c) Mean




(d) Variance

Figure V.7: (Continued from previous page.) Mean and variance of jet charge as a function of jet transverse momentum. (Continued on next page.)




(e) Mean




(f) Variance

Figure V.7: (Continued from previous page.) Mean and variance of jet width as a function of jet transverse momentum. For visual clarity, error bars show only statistical uncertainties on the truth distribution. (Continued on next page.)



(g) Mean



(h) Variance

Figure V.7: (Continued from previous page.) Mean and variance of groomed momentum fraction as a function of jet transverse momentum. The MOMENT UNFOLDING method successfully reproduces both the mean and variance across the full p_T range, with results consistent with truth within uncertainties.

These momentum dependent results enable the analysis of the scaling behaviour of each observable. The jet mass shows expected direct scaling with p_T , consistent with theoretical expectations from perturbative QCD [9, 172, 220, 275, 278, 351, 356, 377, 504], the variance of the jet charge decreases weakly with increasing p_T , reflecting the increased role of gluon radiation at higher energies [337, 452, 459], the mean jet width decreases with p_T , consistent with the collimation of high energy jets due to the Lorentz boost [476, 527], and the groomed momentum fraction shows minimal p_T dependence, a feature expected from the scale invariant nature of the splitting functions in QCD [22, 305, 322].

These results not only validate the *MOMENT UNFOLDING* method but also demonstrate its utility for extracting physically meaningful insights about jet formation and evolution. These scaling behaviours, when applied to real collider data, could serve as stringent tests of perturbative QCD and provide constraints on Monte Carlo tuning parameters.





Figure V.8: Comparison of different unfolding methods for extracting momentum dependent moments of jet mass as a function of jet transverse momentum. Methods compared are MOMENT UNFOLDING, OmniFold, IBU, and IBU with bin-by-bin corrections. The lower panel shows the ratio to truth. (Continued on next page.)

V.E Comparison with alternative methods.


To assess the performance of MOMENT UNFOLDING relative to existing approaches, this section compares the results from MOMENT UNFOLDING with those obtained from two alternative methods,

OMNIFOLD: A state of the art unbinned unfolding method that reconstructs the full distribution, from which moments can be computed, and

Iterative Bayesian unfolding (IBU): A traditional binned unfolding method followed by moment calculation.




(c) Mean




(d) Variance

Figure V.8: (Continued from previous page.) Mean and variance of jet charge as a function of jet transverse momentum. (Continued on next page.)



(e) Mean



(f) Variance

Figure V.8: (Continued from previous page.) Mean and variance of jet width as a function of jet transverse momentum. (Continued on next page.)




Figure V.8: (Continued from previous page.) Mean and variance of groomed momentum fraction as a function of jet transverse momentum. MOMENT UNFOLDING achieves precision comparable to OmniFold while significantly outperforming IBU, particularly at high p_T where binning effects become more pronounced.

For a comprehensive comparison, this study also includes a variation on IBU that applies bin by bin corrections to correct for binning effects in the moment calculation. Figure V.8 shows the moments of each observable as a function of jet p_T for all methods, comparing them to the true values. The lower panels show the ratio to truth, highlighting differences between methods.

This comparison highlights some of MOMENT UNFOLDING's most important properties. The MOMENT UNFOLDING algorithm achieves comparable or better precision than OMNIFOLD for most observables, despite, or perhaps due to, its more targeted approach. This demonstrates the advantage of directly targeting the moments without having to unfold the entire distribution. Both unbinned methods (MOMENT

UNFOLDING and OMNIFOLD) significantly outperform the binned approaches (IBU and IBU+Correction) in precision, highlighting the benefits of avoiding binning artifacts. The bin-by-bin correction improves the performance of IBU but still does not match the precision of the unbinned methods, indicating that the fundamental limitations of binning cannot be fully mitigated by *post hoc* corrections.

However, MOMENT UNFOLDING requires significantly less computational resources than OMNIFOLD due to its non-iterative nature. On average, MOMENT UNFOLDING can be run in about 1% of the time required for OMNIFOLD, while achieving comparable or better precision. This comparison highlights the practical advantages of MOMENT UNFOLDING in scenarios where computational resources are limited or where multiple unfolding tasks need to be performed, such as in systematic uncertainty studies. That said, IBU is about 10^4 times faster than MOMENT UNFOLDING, so in analyses where computation is at a high premium, and the binning artifacts introduced by IBU are acceptable, IBU continues to offer a low-tech, light unfolding procedure.

These results establish MOMENT UNFOLDING as a powerful tool for precision measurements of statistical moments in high energy physics, offering advantages in both statistical precision and computational efficiency compared to existing methods.

V.F Conclusion.

MOMENT UNFOLDING represents a significant advancement in the statistical toolkit for particle physics measurements, providing a novel approach to directly extract moments of distributions without the intermediate step of reconstructing the full spectrum. This methodology addresses several challenges in traditional unfolding techniques while offering unique advantages for theoretical comparison and interpretation. An intrinsically unbinned method, *MOMENT UNFOLDING* eliminates binning artifacts and dimensionality limitations that affect traditional methods. By using a physically motivated reweighting function based on the Boltzmann distribution, it provides a principled approach to the unfolding problem, with clear connections to the maximum entropy principle from statistical mechanics.

The adversarial training framework effectively addresses the two level nature of the unfolding problem, allowing corrections to particle level simulation to be optimised based on detector level comparisons between measured and simulated distributions. By extending this basic architecture to unfold moments conditional on momentum, *MOMENT UNFOLDING* can be used for detailed studies of how moments scale with energy, providing powerful constraints on theoretical models.

As a non-iterative method that is designed specifically to unfold moments rather than full spectra, *MOMENT UNFOLDING* offers significant computational advantages over iterative approaches, particularly for systematic uncertainty studies that require multiple unfolding runs. The application to jet substructure measurements in Section V.D.2 demonstrates the practical utility of *MOMENT UNFOLDING* for real physics analyses. The method successfully recovers both inclusive and differential

moments of a set of jet observables with diverse properties and behaviours, providing enhanced precision compared to traditional binned approaches and comparable or better accuracy than full distribution unbinned methods. These capabilities make MOMENT UNFOLDING particularly valuable for precision tests of QCD, where theoretical predictions are often most precise at the level of moments rather than full distributions.

Although MOMENT UNFOLDING excels at extracting moments, it does not guarantee optimal reconstruction of the full distribution, particularly when higher moments are significant. The choice of how many moments to unfold represents a regularisation parameter that must be carefully selected based on the physics application and available statistics. The basic formulation discussed above is designed for moments of a single observable, though extensions to joint moments are possible.

V.F.1 Towards unfolding distributions.

While MOMENT UNFOLDING offers significant advantages for moment specific measurements, many analyses do require the full unfolded distribution rather than just its moments. This raises the question: can the Boltzmann inspired approach be extended to unfold entire distributions? In a sense, this should be relatively straightforward. If the degree of the polynomial in the exponent of the Boltzmann factor governs the number of moments unfolded, by simply replacing the polynomial with a Taylor series,

$$g(z; \boldsymbol{\beta}) = \frac{1}{P} \exp \left[- \sum_{a=1}^{\infty} \beta_a z^a \right], \quad (\text{V.56})$$

one should be able to, in principle, constrain all moments simultaneously, thereby unfolding the entire distribution.

Even if unfolding infinitely many moments were not feasible, by Taylor’s theorem [802] in conjunction with the Hausdorff moment uniqueness theorem [794] and the Stone–Weierstrass approximation theorem [132, 800], for any arbitrary precision ϵ , there exists a natural number n such that the Hausdorff reconstruction from n or more moments approximates the probability density with precision better than ϵ . So even if it were, for some reason, not possible to unfold infinitely many moments, as long as it were possible to unfold arbitrarily large finite numbers of moments, that would suffice as a solution to this problem as well.

However, since the degree of the polynomial is precisely the regularising parameter for the inverse problem in the *MOMENT UNFOLDING* approach, by increasing it (to some large number or infinity) this approach lifts the regularisation central to the algorithm. The resulting ill posed inverse problem requires significant regularisation.

These challenges motivate the development of the Reweighting Adversarial Networks (RANs), which will be introduced in Chapter VI. A RAN builds upon the foundations laid by *MOMENT UNFOLDING*, extending the approach to full distribution unfolding while maintaining its unbinned, non-iterative, adversarial nature. By incorporating more flexible neural network parametrisations and Wasserstein metrics for distribution comparison, RAN addresses the limitations of *MOMENT UNFOLDING* while preserving its core strengths. This evolution from moment specific to full distribution unfolding represents a natural progression in developing a comprehensive framework for unbinned cross section measurements in particle physics.

Chapter VI

Reweighting Adversarial Networks.

VI.A	The need for full spectral measurements.	303
VI.A.1	Motivation: beyond moments and binned spectra.	304
VI.A.2	Use cases.	308
VI.A.3	Challenges in unfolding full distributions.	312
VI.A.4	Unbinned unfolding approaches.	316
VI.A.5	Addressing challenges.	320
VI.B	From moments to complete differential cross section spectra.	327
VI.B.1	Moments and the full probability distribution.	328
VI.B.2	Unfolding with moment constraints: classical and modern approaches.	331
VI.B.3	Benefits and challenges of MOMENT UNFOLDING.	336
VI.B.4	Extending MOMENT UNFOLDING.	339
VI.C	Methodology and regularisation.	345

VI.C.1	Historical approaches to regularisation.	346
VI.C.2	Regularisation in a RAN.	348
VI.D	Machine learning implementation.	362
VI.D.1	Neural network architecture.	363
VI.D.2	Adversarial training procedure.	366
VI.D.3	Datasets and preprocessing.	370
VI.D.4	Training, monitoring and validation.	374
VI.E	Results	379
VI.E.1	Gaussian model with smearing.	380
VI.E.2	Jet substructure unfolding results.	382

VI.A The need for full spectral measurements.

VI.A.1	Motivation: beyond moments and binned spectra.	304
VI.A.2	Use cases.	308
VI.A.3	Challenges in unfolding full distributions.	312
VI.A.4	Unbinned unfolding approaches.	316
VI.A.5	Addressing challenges.	320
VI.A.5.i	Subtracting backgrounds with RANs.	324

VI.A.1 Motivation: beyond moments and binned spectra.

Measurements in high energy physics traditionally report differential cross sections in a binned format, or even just a few summary statistics (moments) of a distribution. This approach has produced numerous important results in the past [26, 159, 307, 417, 679], but it fundamentally limits the information available for theoretical interpretation. A small set of moments offers only a coarse summary of a probability distribution and can hide critical features. In fact, infinitely many different distributions can share the same first few moments. Two observables with identical mean and variance may have drastically different tails or multi-modal structures that only a full distribution reveal. Thus, relying solely on low order moments risks missing new physics signals or subtle QCD effects that manifest as shape differences rather than overall normalisation changes.

Binning an observable into histograms can be a more detailed approach than quoting a small number of moments, but it too imposes significant limitations. These have been discussed at length in previous chapters, but this section recapitulates them as a reminder, and connects them to the limitations of moment focussed methods as well.

Finite bin widths smear out fine structures and impose an arbitrary discretisation on inherently continuous spectra. Moreover, once data are binned, it becomes impossible to reconstruct or analyse the distribution for any arbitrary transformation of that observable without returning to the original unbinned data. For example, if a cross section is unfolded in bins of an angle θ , one cannot later obtain the distribution in $\cos \theta$ or $\ln \theta$ except by repeating the unfolding from scratch.

In contrast, an unbinned, “full spectral,” measurement preserves maximum information, allowing *a posteriori* reprocessing such as deriving moments, binning into different intervals, or studying functions of the measured observable. This flexibility is especially valuable when comparing with various theoretical models, each of which may suggest different variables or summary statistics to highlight.

Another motivation for capturing the *full differential spectrum* is that many theoretical predictions in quantum chromodynamics (QCD) and beyond are sensitive to the detailed shape of distributions. New physics might appear as excess events in the tails of kinematic distributions, or subtle distortions across a spectrum rather than an overall rate change. Precision QCD studies often rely on the *scaling patterns* of entire distributions with energy or other parameters, not just their averages. By measuring the full spectrum, experimental results can be directly fed into global fits or theory calculations that integrate over the entire phase space. Hence, to fully exploit the data and enable the broadest possible comparisons to theory, it is imperative to go beyond a few moments or fixed bin histograms and aim to unfold the continuous differential cross section itself.

Information loss. Low order moments condense an entire distribution into a small handful of numbers. This sacrifices information about higher order fluctuations and tail behavior. Many distinct distributions can reproduce the same set of moments, so important differences¹ remain hidden if only moments are reported. Even in cases where theory predicts moments more readily than spectra², relying

¹E.g., a long tail vs. a sharp cutoff.

²As in some QCD calculations of energy scaling of moments.

exclusively on moments means discarding data that could otherwise constrain models.

Binned histograms similarly integrate the underlying distribution over each bin, blurring details smaller than the bin width. Fine binning could, in principle, mitigate this, but the detector resolution and finite statistics limit one’s ability to use bins that are too fine. Too fine bins lead to instability due to bin migration effects and large uncertainties [36, 87, 90, 120, 423, 458].

Binning artifacts and biases. Any choice of bin boundaries is inherently arbitrary.

Two experiments measuring the same observable might choose different binning schemes, complicating direct comparisons. Small shifts in bin edges can redistribute events and lead to apparent differences that are purely due to binning choices rather than physics.

Furthermore, when binning multivariate data³, the necessary discretisation in each dimension can introduce artificial discontinuities [14, 23, 115, 134, 159, 515]. These artifacts hinder comparisons of unfolded results with theoretical predictions or with results from other experiments that used different binning schema [317].

Additionally, binning can bias downstream analysis; for instance, extracting moments from a binned distribution requires assuming a shape within each bin.⁴ If the true distribution varies non-linearly inside the bin, the extracted moment is biased by the bin size and by the shape assumption, such that even

³Or examining an observable’s distribution in bins of another quantity.

⁴Often a constant or linear interpolation.

the sign of the error cannot be known *a priori* [110]. An unbinned measurement eliminates this intermediate step and the bias that can arise from it.

Limited dimensionality. Perhaps most importantly, binned unfolding severely constrains the number of observables one can unfold simultaneously. Each additional dimension (feature) requires exponentially more bins to maintain a given resolution. In practice, traditional unfolding is often performed in just one or two dimensions at a time because higher dimensional histograms would be too sparsely populated. This precludes measuring cross sections differential in many variables at once.

It also means that if one is interested in an observable O that is a complicated function of several kinematic quantities, one either must unfold those quantities jointly with fine binning (which is infeasible) or settle for unfolding a projection of O in one dimension (which loses information). In contrast, unbinned approaches can handle high dimensional data naturally, since they do not require constructing an d dimensional grid of bins. Unfolding the full phase space ⁵ is only realistic with an unbinned strategy.

Table VI.1 summarises these differences between unbinned moment based unfolding, binned density unfolding, and unbinned full spectrum unfolding methods. By construction, a full spectral measurement retains maximal information and avoids discretisation issues, at the cost of a more challenging unfolding procedure. These challenges were a major deterrent historically, explaining why experiments long favoured simplified results despite their limitations.

⁵I.e. a fully differential cross section in all relevant kinematic variables.

Table VI.1: Comparison of moment-only, binned spectrum, and unbinned spectrum unfolding approaches. Reporting only low-order moments loses most distribution information. Binned differential cross sections retain shape information but suffer from discretization and limited dimensionality. Full spectral measurements preserve the complete distribution, enabling maximal reusability and detailed theory comparisons, but require regularizing a much more ill posed training.

Aspect	Moments only	Binned spectrum	Full spectrum
Information	Low	Moderate	High
Reuse	High	Limited	High
Dimensions	Many	1 to 2	Many
Stability	Easy	Moderate	Hard

VI.A.2 Use cases.

Some of the most compelling motivations for unbinned unfolding of probability density functions come from specific classes of observables and analyses in high energy physics that demand detailed distributions rather than summary measures. A few examples are:

Jet substructure observables: Modern studies of jet physics often examine intricate internal properties of jets, for example, the distribution of jet mass, angularity, N -subjettiness ratios like τ_{21} , energy correlation functions, and so on. These observables have rich distributions shaped by QCD radiation and hadronisation inside the jet. For instance, the jet mass spectrum in hadronic collisions exhibits a steeply falling shape with resonance peaks or grooming induced features in certain regions; capturing these details is essential for validating parton shower models and grooming techniques.

If only the average jet mass or a few quantiles were reported, one would lose the structural information of how often jets are heavy versus light, or how substructure sculpts the distribution. Similarly, distributions of τ_{21} ⁶ contain information about the fraction of jets with two subjects versus one, which is crucial for signal (e.g. boosted W) vs background (QCD jet) discrimination. Capturing the entire τ_{21} spectrum allows experimentalists and theorists to identify where in the distribution their models agree or fail, rather than just comparing an efficiency at a fixed cut.

Unfolding these jet substructure distributions in an unbinned way provides a high resolution view of QCD dynamics and is increasingly necessary as theory tools⁷ improve to the point of predicting differential shapes [32, 78, 209]. In fact, recent measurements have demonstrated the power of full phase space unfolding for jets, using multivariate ML techniques to correct detector effects and obtain particle level jet observable spectra without binning [40, 64, 105, 210, 478]. These cases underscore that jet physics benefits enormously from preserving the full shape information.

Charge and flavour sensitive observables: Some observables aim to distinguish particle charge or flavour inside complex final states, and their distributions can be especially telling. An example is the jet charge distribution, the electric charge of a jet computed from the momentum weighted charges of its constituent

⁶A ratio used to tag two prong substructure.

⁷Like analytic resummation or first principles simulation.

particles. This quantity is used to tag whether a jet originated from a quark of a given electric charge.⁸

The jet charge is a continuous valued observable that can take positive and negative values, and experiments measure its probability distribution for jets of various momenta. The full shape of the jet charge distribution contains information about the underlying quark/gluon mixture and fragmentation processes; for instance, a broader distribution indicates a mix of high charge (quark origin) and zero charge (gluon origin) jets, whereas a narrow distribution peaked near zero suggests predominantly gluon jets. Only by unfolding the entire jet charge spectrum (including its dependence on jet p_T) can one provide data precise enough to tune fragmentation models and compare to theoretical calculations of charge transport in jets [24, 121, 207, 208, 336, 349, 452].

Another example is the distribution of identified particle multiplicities.⁹ The multiplicity distribution is discrete but often measured as a histogram. It is a particularly challenging observable to unfold because it is strongly sensitive to soft QCD and hadronisation effects [137, 154, 155, 482, 484, 524, 578, 590]. Two Monte Carlo generators might predict the same average multiplicity but differ in the width or tail of the multiplicity distribution, which affects extreme cases.¹⁰ Such observables are often *infrared unsafe* theoretically,¹¹ meaning theory must rely on phenomenological models.

⁸As opposed to a gluon, which has none.

⁹E.g., number of charged hadrons in an event or in a jet.

¹⁰Like very high multiplicity events.

¹¹Because soft particle emission has no cutoff, perturbative predictions diverge for the distribution shape.

The only way to constrain and improve those models is for experiments to provide the unfolded full distributions of these IR unsafe observables at particle level. In summary, any analysis where internal structure, charge assignments, or other detailed event properties matter will benefit from (or even require) full spectral unfolding rather than a few summary numbers.

Multidifferential and high dimensional measurements: The ultimate form of ‘full’ spectral measurement is unfolding in multiple kinematic dimensions simultaneously, effectively measuring a multi differential cross section over a high dimensional phase space. While this is extremely challenging, certain physics questions demand correlating several observables.

For example, consider measuring an observable O as a function of another variable (or collection of variables) $\mathbf{Q}$.¹² With traditional methods, one would perform a two dimensional unfolding, choosing a projection Q_j of $\mathbf{Q}$ if needed (bins in O vs bins in Q_j) in order to then extract moments or other features as a function of Q [152, 266, 493, 560, 579, 702]. Binning in two (or more) dimensions quickly suffers from sparse data in many bins and complicated systematic uncertainties.

An unbinned approach, by contrast, could in principle unfold the joint $(O, \mathbf{Q})$ distribution without requiring an explicit grid, allowing arbitrary slicing and analysis after the fact. This is especially relevant in the era of high luminosity experiments, where huge datasets invite more differential measurements. Rather than publish dozens of separate one dimensional spectra (each in a single

¹²Say, the distribution of jet charge q in bins of jet p_T or event energy Q .

kinematic region), one could publish a multi dimensional unfolded distribution that coherently captures all correlations. Such a result would be far more powerful for global interpretations, albeit significantly more complex to obtain. Unbinned density unfolding techniques are a step toward this ambitious goal, enabling higher dimensional unfolding than previously feasible with manageable uncertainties.

VI.A.3 Challenges in unfolding full distributions.

The push for unbinned, high resolution unfolded spectra comes with substantial challenges. Unfolding a full distribution, especially in multiple dimensions, is a markedly more ill posed problem than unfolding a small set of summary statistics or coarse bins. These challenges are statistical, physical, and computational.

At its core, unfolding requires inverting the detector response—a many to many mapping where a given particle level distribution can produce a range of detector level outcomes due to resolution and inefficiencies. Infinitely many particle level spectra are, in principle, consistent, within uncertainties, with a given set of detector level data, especially if the data are treated in fine detail. The problem of small fluctuations or statistical noise in the detector level histogram causing huge oscillations in the unfolded solution is exacerbated when the response “matrix” is essentially continuous, since here one is trying to reconstruct a function rather than a finite vector, which has infinitely many degrees of freedom.

Without any constraints or regularisation, unfolding is mathematically underdetermined; one must introduce additional information to obtain a physical solution. This

additional information can be statistical regularisation,¹³ or a strong prior¹⁴ to guide the result. Either way, the unfolded distribution will depend to some degree on these assumptions, which is a point of concern.

All unfolding methods require some initial model or ansatz for the true distribution, explicitly or implicitly. Traditional regularised unfolding (e.g. Bayesian iterative methods or Tikhonov regularisation) starts from a initial guess for the Truth distribution and updates it in light of the data. If the ansatz is significantly wrong in a region where data have low sensitivity, the unfolded result may inherit that wrong shape (bias) because there isn't enough information in the data to correct it. For binned methods with strong regularisation, the unfolded spectrum can end up looking very much like the ansatz except in regions where the data clearly indicate otherwise. When unfolding a full spectrum, managing bias can be even trickier than when unfolding moments.

With many bins or continuous degrees of freedom, there are more opportunities for the prior assumptions to creep in unless the method explicitly works to mitigate this. Modern ML based unbinned unfolding techniques like iterative reweighting (used by OMNIFOLD) attempt to reduce prior bias by gradually adjusting the prior to fit the data, but they still require that the initial simulation populate all regions of phase space that data might cover. In other words, if the Truth has support outside the domain of the Generation, no unfolding can recover that—a condition that applies to all known methods. This need for sufficient overlap in the support of the initial assumption and Truth distribution is a fundamental mathematical limitation: full

¹³E.g., penalizing roughness in the solution.

¹⁴Incorporated, e.g., initial guess of the spectrum.

spectral unfolding demands that the Generation model be flexible and broad enough to encompass the Truth, at least roughly, or else certain features will be missed entirely. Methods such as OMNIFOLD that involve classification and reweighting at detector level in addition to particle level, additionally require that Simulation populate the entire phase space of Data.

Unfolding more finely, or in more dimensions, inherently means extracting more parameters.¹⁵ from the same finite dataset. This trade-off means individual elements of the unfolded spectrum will have larger statistical uncertainties than if the data were aggregated into a few bins or moments. For example, unfolding 100 bins will yield each bin count with larger uncertainty than if one had combined them into 1 bin, since the the bin count of 1 bin is simply the total count. If not carefully handled, an unbinned unfolding will overfit statistical fluctuations in the observed data.¹⁶ Effective regularisation is essential to prevent noise amplification.

Additionally, the detector response kernel must be well modelled. Any mismodeling (systematic error) will imprint itself on the unfolded result in complex and unpredictable ways. When only a few numbers are extracted, one can sometimes correct for known detector biases by simple scale factors; but when unfolding an entire distribution, any mismodeling of the response shape can distort the unfolded spectrum non-uniformly. This places high demands on detector simulation fidelity and on methods to incorporate systematic uncertainties (e.g. variations of the response model) into the unfolding procedure. Fully Bayesian approaches or profile likelihood methods can propagate uncertainties to the unfolded spectrum, but doing so in high

¹⁵In effect, more bins worth of information.

¹⁶Referred to as *sculpting* in the machine learning literature [171, 263].

dimensions is computationally intensive. In short, the richer the information one unfolds, the more careful one must be to quantify the reliability of each feature of the spectrum.

Traditional unfolding algorithms scale poorly as the number of bins grows. Even if one had infinite amounts of data and so unpopulated bins were not a concern, these traditional algorithms like IBU that are relatively fast for tens of bins become very slow for hundreds of bins, and practically unusable for thousands of bins or continuous data, as each iteration must refine a fine grained distribution. Unbinned algorithms, on the other hand, typically rely on machine learning or Monte Carlo sampling. They may involve training complex models on large datasets or performing high dimensional optimisations. For instance, iterative reweighting methods train a classifier multiple times,¹⁷ and generative approaches might require training a high capacity generative model on millions of events. These computations require substantial computing resources (CPU and GPU), careful hyperparameter tuning, and can suffer from convergence issues. Ensuring that an unbinned unfolding converges to a stable solution without excessive computation is a highly non-trivial challenge. Thus practical feasibility of unfolding full spectra depends on advances in algorithms and computing, precisely the advances that modern machine learning based methods aim to provide.

Despite the challenges that unfolding full spectra presents, the drive to measure full spectra is strong because of the scientific payoff. The next subsection discusses

¹⁷Each iteration is a full supervised learning task on the dataset.

how recent machine learning methods, including the RAN architecture, tackle these issues.

VI.A.4 Unbinned unfolding approaches.

Note. This section focusses specifically on the regularisation methods that recent ML based unbinned unfolding approaches have successfully employed. For a broader conceptual overview of unbinned unfolding see Section II.F and for a survey of ML based unbinned unfolding methods, see Section III.G.

Recent years have seen rapid development of machine learning techniques for unfolding, which seek to overcome the limitations of traditional methods and make full spectral measurements possible in practice. Broadly, these approaches fall into two categories, those based on *reweighting* an existing MC sample to better agree with nature, and those based on *generating* new events¹⁸ to reproduce the Truth distribution. Both categories strive to avoid fixed histogram binning and instead work with unbinned data, using the power of ML to handle high dimensional inputs and complex detector responses.

One of the pioneering ML based methods is *OMNIFOLD*, which introduced an unbinned, multivariate unfolding technique using iterative reweighting. OMNIFOLD uses classifier neural networks to reweight a Monte Carlo sample.

Essentially, it trains classifiers to distinguish between Data and Simulation, and then uses the classifier output to assign weights to Simulation events such that the weighted Simulation better matches the Data. This procedure is iterated at both

¹⁸Typically with generative models.

detector and particle level in alternation. The method produces a set of weights for Generation events that yield an unfolded distribution. OMNIFOLD demonstrated, for the first time, that one can simultaneously unfold many observables¹⁹ without binning, given a sufficiently flexible classifier and a robust iterative scheme. It has been successfully applied in multiple experimental analyses.

The key advantages of OMNIFOLD are that it naturally accounts for high dimensional correlations²⁰ and it actively mitigates prior dependence by iterative refinement, effectively performing a form of expectation maximisation to find a self consistent unfolded result that does not overly rely on the initial generation distribution.

A noted limitation of OMNIFOLD is its computational cost: each iteration requires training two classifiers to convergence, and in realistic cases one might need 5 to 30 [154] iterations, amounting to training a large number of networks. The iterative nature can also complicate uncertainty evaluation and hyperparameter tuning.²¹ Still, OMNIFOLD set the stage for practical unbinned unfolding and proved the feasibility of full spectral measurements on real experimental data.

In parallel, other methods have explored using explicit generative models to unfold distributions. For example, VAEs have been employed to learn a mapping from random noise to particle level distributions such that, when those events are passed through a detector simulation, the output matches the observed data distribution. Similarly, normalising flows and other invertible neural networks have been used to model the probability density function of true observables and deform it until its

¹⁹And even the entire event record.

²⁰Since the classifier can use any or all features.

²¹E.g., one must subjectively choose when to stop iterating to avoid instability.

convolution with the detector response matches data. These approaches are extremely powerful in principle: a sufficiently flexible generator could capture the full true distribution without needing a starting simulation.

In practice, however, training such generative models is challenging. Pure generative unfolding would require vast amounts of data to constrain the high dimensional space and can suffer from mode dropping.²² Additionally, learning a full generative model from scratch means the method must learn *both* the underlying physics distribution *and* compensate for detector effects at the same time. This is a high dimensional optimisation that can be unstable or require careful conditioning. Some recent works have attempted to combine generative modelling with explicit usage of the known detector response to guide the training,²³ but these are still at the experimental stage.

Between pure reweighting (which uses an existing simulation as a starting point) and pure generation (which starts from random noise) lies an attractive compromise: using machine learning to *reweight* an initial MC in a single pass, by directly comparing its predictions to data. The **Reweighting Adversarial Network (RAN)** method is built on this approach. Other examples include optimal transport inspired methods that train a single neural network to learn a transport map by minimising some distance between weighted simulation and data distributions. Such approaches leverage the fact that modern simulations are a good but imperfect approximation of reality; rather than throw them away and learn from scratch, it may be easier to learn small corrections (reweightings) to the simulation. In this sense, RAN and

²²The generator failing to produce some less common features of the distribution if not properly incentivised.

²³For instance, by using differentiable detectors or gradient based deconvolution.

similar methods treat unfolding as a density ratio estimation problem: find a function $g(z)$ such that the weighted MC distribution $g(z) q(z)$ matches the true distribution $p(z)$. If Z denotes particle level kinematics, this $g(z)$ effectively encapsulates how the Monte Carlo needs to be modified to agree with nature.

The challenge is that $p(z)$ cannot be directly—one only has access to its smeared version at detector level. Therefore, a RAN sets up a two level training objective to optimise $g(z)$ based on how well the weighted events, after going through the detector simulation, agree with the detector level data. One might think that such an approach would be computationally unfeasible due to having to repeatedly rerun expensive detector simulation on reweighted Generation, and this would indeed be the case if the generator were modifying the features themselves. The intrinsic pairing of (Gen., Sim.) events, combined with the fact that the generator only modifies the importance weights of events rather than the events themselves means that the detector emulation only needs to be run once to create the original Simulation dataset, after which, Simulation can simply inherit its weights from the matching Generation event.

Table VI.2 provides a high level comparison of these methods. All these methods aim to enable unbinned, high dimensional unfolding; they chiefly differ in how they incorporate prior knowledge, whether they iterate, and how computationally intensive they are. Notably, most current ML based methods still require that the starting Simulation has support in the regions of interest of the Data. They also share common

Table VI.2: Comparison of key characteristics across different unfolding methods. “Unbinned” indicates whether the method can process continuous data without binning into fixed histograms. “Iterative” refers to whether multiple training iterations are required (distinct from neural network training epochs). “Uses simulation” indicates whether the method reweights existing Monte Carlo events (Yes) or generates new samples from scratch (No). The Reweighting Adversarial Network (RAN) method introduced in this chapter combines the advantages of unbinned processing and non-iterative training while leveraging prior simulations, positioning it between the computational efficiency of traditional methods and the flexibility of modern machine learning approaches.

Method	Unbinned	Iterative	Uses sim.
Traditional (IBU, TUnfold)	No	Sometimes	Yes
Discriminative (OMNIFOLD)	Yes	Yes	Yes
Generative (cINNs/VAEs)	Yes	No	No
Adversarial reweighting (RAN)	Yes	No	Yes

challenges of training stability and overfitting prevention, which are addressed through techniques like regularisation or specific loss functions.²⁴

VI.A.5 Addressing challenges.

The Reweighting Adversarial Network (RAN) approach described in this chapter is designed to address the aforementioned challenges head on, enabling full spectral unfolding with improved stability and efficiency. This section outlines how the methodology of RANs addresses the needs and difficulties detailed above.²⁵

²⁴For instance, OMNIFOLD uses early stopping of iterations, while adversarial methods use regularised neural network architectural constraints.

²⁵A complete description of the method follows in later sections, so the focus here is on conceptual understanding rather than implementation details.

A RAN explicitly targets the entire distribution rather than a fixed set of moments. In fact, it can be viewed as an extension of the MOMENT UNFOLDING method described in Chapter V to infinitely many moments. By using a flexible neural network to parametrise the weighting function, a RAN does not impose a rigid functional form limited to a few moment constraints. This means it has the capacity to adjust the simulation such that not only the first few moments, but rather *all* features of the distribution²⁶ are brought into agreement with the observed data.

Framing the problem in this fashion ensures that when a RAN converges, the theoretical guarantees of Section V.B.5 result in a faithful unfolded spectrum, from which moments or any other summary statistic can be derived. Importantly, because it handles the full spectrum at once, RANs avoid the bias of selecting certain observables upfront—it lets the data inform the entire shape. Hence the RAN method is aligned with the basic motivation of full spectral measurements: nothing gets thrown away or averaged out prematurely.

By construction, the RAN algorithm is an unbinned method. It operates on individual events, using distance measures defined on samples rather than comparing binned counts. The adversarial training framework means that a discriminator network looks at distributions of features (at detector level) and tries to tell apart weighted simulation from real data. If it finds any differences in the distributions, no matter how localised, the reweighting network is encouraged to adjust weights (at particle level) to eliminate that discrepancy. This is effectively a fine grained comparison across the full phase space, without ever projecting data into predetermined bins.

²⁶In principle, every differential element.

As a result, RANs do not suffer from the bin alignment issues or interpolation biases that plague binned unfolding. Different experiments using RANs on the same observable should, in principle, get comparable results without worrying that one used binning scheme A and the other binning scheme B, since neither would use bins at all. The output of a RAN can be presented as a smooth distribution or as a weighted event sample at particle level, which can then be binned or analysed downstream as needed. This flexibility maximises the utility of the measurement.

One of the core design goals of the RAN method is to eliminate the need for multiple iterative reweighting cycles (as in OMNIFOLD). Instead of training sequential classifiers for each iteration, RAN uses a single coupled training procedure where the particle level reweighting function and the detector level discriminator are learned together.²⁷ This yields a one shot solution for the weights after convergence. In practice, this means a RAN involves training one pair of neural networks with feedback from each other in one integrated run.

The computational cost is roughly equivalent to training a single GAN rather than a dozen separate classifiers. As demonstrated in Section VI.E, this can lead to a significant reduction in runtime and resource usage for large scale unfolding tasks. Furthermore, by using an optimal transport based loss (inspired by the Wasserstein GAN) and techniques like spectral normalisation in the discriminator, a RAN achieves stable training dynamics.

The Wasserstein metric provides a smooth loss function that correlates well with distribution difference, avoiding the chaotic or oscillatory behaviour that vanilla

²⁷Analogous to a generator and critic in a Wasserstein GAN.

GANs can exhibit, described in Section III.E.3.i. Spectral normalisation bounds the discriminator’s gradient, effectively regularising the learning and preventing mode collapse or instability. These choices were crucial technical innovations needed to extend MOMENT UNFOLDING to a full spectrum unfolding method. The end result is that a RAN converges reliably to a solution where the weighted generation agrees with truth, all in a single training loop. This addresses the computational challenge by trading an iterative series of simpler trainings for one more complex adversarial training, which has been shown to be tractable and efficient.

Like any reweighting based approach, a RAN does rely on the distribution of MCMC samples and will fail if the Truth distribution lies entirely the support of the Generation. However, a RAN upon convergence, provably does not rely on the exact Generation shape within the supported region i.e. it is unbiased, and converges to the MLE. The use of an adversarial loss means that a RAN is effectively solving a constrained optimisation to find the weight function that makes the simulation statistically indistinguishable from data by reweighting only at particle level.

The RAN method’s connection to the Boltzmann entropy inspired MOMENT UNFOLDING provides a theoretical understanding of how it converges the “maximum entropy” solution consistent with the constraints (the data). This is desirable because the maximum entropy solution is the least biased one given the information at available. In essence, a RAN inherits MOMENT UNFOLDING’s theoretical guarantees and compensates for the larger freedom of full spectra by adding proper regularisation in training, thereby taming the ill posedness.

An important practical aspect is that a RAN yields physically reasonable unfolded distributions, such as non negative weights and normalised total cross sections. Parametrising the weight function as $g(z) = \frac{1}{P} \exp(-f(z))$ (one convenient choice) ensures $g(z) \geq 0$ for all events, so the unfolded cross section is positive definite. The training objective can be set up to include a normalisation constraint or simply allow the weights to float the total normalisation to match the data yield. This avoids unphysical outcomes like negative weight factors or mismatched totals that can occur in some matrix inversion methods without extra constraints.

Moreover, because a RAN operates at the level of reweighting events, the resulting weighted event sample can be validated easily. One can always forward fold the weighted sample through the detector emulation to check that it indeed reproduces all features of the observed data. This built in consistency check is a powerful advantage of the reweighting paradigm in general. It is straightforward to verify that the full spectral measurement is successful, by confirming that no residual differences remain between data and weighted simulation across all distributions of interest.

VI.A.5.i Subtracting backgrounds with RANs.

It is worth noting that the RAN framework could potentially be extended to perform background subtraction simultaneously with unfolding by modifying the weight parametrisation to allow negative weights. While this extension has not been implemented or tested in the current work, allowing negative weights could potentially enable joint background subtraction and unfolding.

This approach might be especially valuable for heavy ion collisions where background separation is notoriously challenging due to the high multiplicity environment and complex underlying events [123]. Several studies have explored related joint unfolding and background estimation approaches, including work on machine learning based background subtraction methods that can reduce fluctuations below the statistical limit [123, 174, 368]. The difficult interplay between background subtraction and unfolding has been examined in detail in [528], which analysed how different background subtraction methods affect jet observables in heavy ion collisions. Recent research has also improved background subtraction techniques specifically for jet substructure measurements [247], demonstrating that proper handling of medium response is crucial for meaningful comparison with experimental data. A theoretical framework of transport theory deconvolution with background contributions could provide mathematical guidance for such extensions. Development of these capabilities would require careful validation and is left for future work.

In summary, a RAN is designed to meet the need for unfolding multivariate probability densities by addressing the key challenges that have historically prevented such measurements. It leverages modern ML²⁸ to unfold complete distributions without binning, while incorporating solutions to the statistical and computational pitfalls.²⁹ This section has established *why* a method like RAN is needed—because it enables the unfolding of maximal information from experimental data, the full differential cross section, in a way that is both theoretically and practically sound, thus

²⁸Specifically adversarial training and neural network flexibility.

²⁹Regularisation via WGAN, one pass training, *et cetera*.

empowering deeper insights into particle physics phenomena. The following sections will describe the RAN methodology in detail and demonstrate its performance.

VI.B From moments to complete differential cross section spectra.

VI.B.1	Moments and the full probability distribution.	328
VI.B.1.i	The maximum entropy principle and exponential families.	330
VI.B.2	Unfolding with moment constraints: classical and modern approaches.	331
VI.B.2.i	Linear regularised unfolding (Tikhonov and SVD).	332
VI.B.2.ii	Iterative Bayesian unfolding.	333
VI.B.2.iii	Entropy based methods.	334
VI.B.2.iv	Unbinned methods through the lens of moments.	335
VI.B.3	Benefits and challenges of MOMENT UNFOLDING.	336
VI.B.4	Extending MOMENT UNFOLDING.	339

Section VI.A discussed the motivation for unfolding complete differential cross sections—retrieving the full distribution of an observable at particle level from distorted detector level data. This section now builds a rigorous framework linking distribution moments to the full probability density, and details how leveraging moment constraints can facilitate unbinned unfolding of densities.

Treating moments as fundamental constraints, as opposed to bin by bin values, establishes a pathway from a limited set of summary statistics to a complete unfolded spectrum. This section first develops the theoretical foundation connecting moments and probability distributions,³⁰ then surveys classical and modern unfolding methods that rely on moment constraints.³¹ This is followed by a discussion of the advantages and pitfalls of moment based unfolding, notably issues relating ill posedness and non uniqueness, and the regularisation strategies to mitigate them. Finally, the MOMENT UNFOLDING method introduced in Chapter V will naturally be extended toward unfolding the entire distribution (all moments) via a GAN-like approach inspired by the Boltzmann distribution, setting the stage for unfolding a full, unbinned, differential cross section measurement.

VI.B.1 Moments and the full probability distribution.

Statistical moments provide a powerful, coarse grained description of a probability distribution. Moments about the mean (central moments) characterise the shape (variance, skewness, kurtosis, *et cetera*) of the distribution.

³⁰Including moment generating functions and maximum entropy arguments.

³¹Such as regularisation and entropy based techniques.

In principle, an infinite sequence of moments completely characterises a distribution. This is referred to as the moment representation of a distribution.

Definition 43 (Moment generating function (MGF)). The *moment generating function (MGF)* of a random variable Z , $M_Z(t) = \mathbb{E}[e^{tZ}]$, is defined as

$$M_Z(t) = \sum_{n=0}^{\infty} \frac{t^n}{n!} \mathbb{E}[Z^n]. \quad (\text{VI.1})$$

The MGF encapsulates the entire moment sequence via its Taylor expansion, and serves as a proxy for the full distribution.

Theorem 12 (Moment representation of a random variable [604, 722, 794, 800, 802]).

Let λ denote Lebesgue measure on $\mathbb{R}$ and $\mathcal{P}(\mathbb{R})$ the set of Borel probability measures on $\mathbb{R}$. Then

$$\begin{aligned} \forall m : \mathbb{N} \rightarrow \mathbb{R} \left[m(0) = 1 \implies \forall \mu \in \mathcal{P}(\mathbb{R}) \right. \\ \left. \left\{ \mu \ll \lambda \wedge \left(\forall n \in \mathbb{N} \int_{\mathbb{R}} z^n d\mu(z) = m(n) \right) \right. \right. \\ \left. \wedge \left(\exists \text{ open } I \subseteq \mathbb{R} \wedge 0 \in I \forall t \in I \int_{\mathbb{R}} e^{tz} d\mu(z) < \infty \right) \right. \\ \left. \implies \forall (\nu, n) \in \mathcal{P}(\mathbb{R}) \times \mathbb{N} \int_{\mathbb{R}} z^n d\nu(z) = m(n) \implies \nu = \mu \right\} \Big]. \end{aligned} \quad (\text{VI.2})$$

Equivalently, writing $d\mu = f dz$ and $d\nu = g dz$ when they exist,

$$\begin{aligned} \forall n \in \mathbb{N} \int_{\mathbb{R}} z^n f(z) dz = \int_{\mathbb{R}} z^n g(z) dz \wedge \exists \text{ open } I \ni 0 \forall t \in I \\ \int_{\mathbb{R}} e^{tz} f(z) dz < \infty \implies f = g \text{ almost everywhere} \end{aligned} \quad (\text{VI.3})$$

In fact, knowledge of $M_Z(t)$ (or the characteristic function) allows one to reconstruct the probability density $p(z)$ fully. In short, an infinite set of moment constraints is mathematically equivalent to knowing the complete distribution.

In practice, however, one can only estimate a finite number of moments from data with finite precision. A finite moment set underdetermines the distribution, because infinitely many distinct densities can share the same first n moments. This non-uniqueness under finite information is closely related to the ambiguity of solutions when attempting to invert detector effects: the data usually provide limited ‘moments’ of the true distribution. For example, binned event counts are themselves integrals of the true spectrum over bin ranges and therefore a kind of ‘moment’. As n increases, the moment constraints become more informative and the space of consistent solutions shrinks, but noise and uncertainties also grow. In the limit of $n \rightarrow \infty$, the true distribution would be recovered, but this limit cannot be reached exactly with finite and noisy data.

VI.B.1.i The maximum entropy principle and exponential families.

Given a few known moments, a common approach to approximate the underlying distribution is to apply the principle of maximum entropy [704]. This principle dictates that, among all distributions satisfying the known moment constraints, we should prefer the one with the largest entropy.³² Imposing constraints on expectations $\mathbb{E}[f_a(Z)] = c_a$ ³³ leads to a unique maximum entropy solution in the exponential family.

³²I.e. minimal additional assumptions or information should be added to the prediction beyond what the constraints impose.

³³For some set of functions $f_a(Z)$ defining the moments of interest.

In particular, one finds a probability density of the form

$$p^*(z; \boldsymbol{\beta}) = \frac{1}{P(\boldsymbol{\beta})} \exp \left[- \sum_{a=1}^{\infty} \beta_a f_a(z) \right], \quad (\text{VI.4})$$

where β_a are Lagrange multipliers adjusted to enforce the desired $\mathbb{E}_{p^*}[f_a(Z)] = c_a$, and $P(\boldsymbol{\beta}) = \int \exp(-\sum_a \beta_a f_a(z)) dz$ is the normalisation factor (partition function).

This is analogous to the Boltzmann distribution in statistical mechanics, which maximises entropy given a fixed average energy. Indeed, if $f_a(z) = z^a$ (the monomials), the above $p^*(z)$ reduces to the Boltzmann ansatz with coefficients β_a related to the distribution's moments. This exponential family is central to the moment based spectrum unfolding method. It provides a flexible yet principled parametrisation of the true distribution in terms of a set of parameters λ_a or equivalently a set of moments. Crucially, if the list of moment constraints is extended and refined, $p^*(z)$ can approximate the true distribution arbitrarily well (approaching the actual $p_Z(z)$ as the number of moments grows). Hence the full differential cross section can be unfolded to arbitrary precision in the limit of sufficiently many moment constraints.

VI.B.2 Unfolding with moment constraints: classical and modern approaches.

Many unfolding methods, both classical and modern, can be interpreted as using moment constraints or related regularisation assumptions to tackle the ill posed inversion of detector effects. In a binned setting, each bin count can be seen as a moment.³⁴ Unfolding those bin counts with minimal noise amplification often requires

³⁴As an integral of the continuous distribution against a top hat basis function.

additional constraints such as limiting the number of iterations. Although not usually thought of in this fashion, this can be equivalently reformulated as effectively limiting the space of possible moment values of these distributions convolved with a top hat function.

VI.B.2.i Linear regularised unfolding (Tikhonov and SVD).

One class of unfolding methods formulates the problem as the linear system of equations Equation (II.13). Solving for ν_j directly through matrix inversion or unregularised maximum likelihood is notoriously unstable. Although this has been discussed in Chapter II, one is now equipped to reformulate the problem in the language of moments. High frequency fluctuations in ν , equivalent to variations in high order ‘moments’, can fit statistical noise in μ . Tikhonov regularisation addresses this by adding a penalty on undesirable solutions, usually favouring smoothness [767]. For example, one penalises the squared second derivative of the unfolded spectrum or deviations from a prior guess. This effectively constrains the higher order moments of the solution, suppressing oscillatory components that are poorly determined by data. The result is a bias toward smooth moment behaviour, trading a bit of bias for a dramatic reduction in variance [665].

Similarly, the singular value decomposition (SVD) unfolding method truncates small singular values of the response matrix [706]. This is equivalent to discarding combinations of moments that cannot be determined well.³⁵ By keeping only the

³⁵Those along directions of the solution space corresponding to tiny eigenvalues would otherwise blow up with noise.

dominant modes, essentially the lowest frequency or moments with largest norm, SVD unfolding ensures stability at the cost of not fully utilising high frequency information.

Both Tikhonov regularisation and SVD thereby regularise the moment space, implicitly limiting the effective number of moments that contribute to the unfolded solution.

VI.B.2.ii Iterative Bayesian unfolding.

IBU's approach to unfolding can be reinterpreted as iteratively matching moments of the true distribution to the data. It starts with an initial guess for the true distribution, *Generation*, which corresponds to some initial moment estimates and iteratively updates the expectations to better agree with the observed data. At each iteration, the method reweights the Monte Carlo events by the ratio of data to simulation in each detector bin, effectively adjusting the candidate true distribution's moments to reduce the discrepancy in those bins.

This procedure can be interpreted as ensuring that the predicted detector level counts, moments in a top-hat basis for each bin, match the observed counts, one step at a time. The algorithm converges to a solution that maximises the likelihood in the limit of many iterations, but in practice one stops after a finite number of iterations to avoid over-fitting statistical fluctuations.

The regularisation can, then, be interpreted as limiting the effective degrees of freedom (moments) that are fitted, similar to Tikhonov or SVD albeit via a different mechanism.

VI.B.2.iii Entropy based methods.

Another class of unfolding techniques explicitly incorporates entropy or information criteria to regularise the solution. The maximum entropy (MaxEnt) unfolding approach chooses the unfolded distribution that maximises entropy subject to reproducing the observed detector data, usually via a likelihood term [729]. In practice this might mean maximising the entropy

$$S = - \sum_j \mu_j \ln(\mu_j) , \quad (\text{VI.5})$$

minus a term for agreement with data.

The solution tends to be as smooth and featureless as possible (high entropy) unless the data significantly demand a structure. Entropy regularisation thus penalises any extraneous moment fluctuations that are not required by the data, biasing toward a Gaussian distribution in the absence of strong evidence for shape.

Schmelling's method of reduced cross-entropy [706] is a related technique, effectively combining likelihood maximisation with an entropy prior. By treating the deviation from an ansatz in terms of Kullback–Leibler divergence (relative entropy), one can incorporate prior knowledge while still preferring the least structured solution beyond that prior. These methods make the connection to moment constraints explicit in that they treat the unfolding problem as one of satisfying certain expectation values the data constraints, while maximising uncertainty elsewhere. As discussed, this yields an exponential family solution. In fact, the MaxEnt solution can be written in the form $p(z) \propto \exp(-\sum_a \beta_a f_a(z))$ where $f_a(z)$ are the functions defining the data

constraints.³⁶ This is formally identical to the moment constrained maximum entropy distribution discussed above; the only difference is the nature of the constraints.³⁷

VI.B.2.iv Unbinned methods through the lens of moments.

With advances in computation, unbinned unfolding techniques have emerged that avoid histogramming data altogether. These methods often use machine learning to compare distributions without bins. While OMNIFOLD, for example, does not explicitly constrain low order moments or use an analytic parametrisation, it implicitly matches all features of the distribution that the classifier can discern—effectively attempting to equate the full set of moments by the end of the procedure. In each iteration, the classifier focuses on the current differences between data and weighted simulation, which are often in some *mode* of the distribution; iterating allows successively finer differences, higher order moments or localised shape features to be corrected. This is conceptually similar to performing an increasingly detailed moment matching, guided by the ML classifier as a flexible test statistic.

Other approaches employ adversarial neural networks or optimal transport metrics to directly learn the unfolded distribution in one go, rather than iteratively, using a generator model for the true distribution and a discriminator to enforce that the generator’s folded output looks like the data. Such setups, inspired by Generative Adversarial Networks (GANs), take advantage of the same idea that ultimately all

³⁶For example, indicator functions for each detector bin to ensure those predicted counts match the observed.

³⁷Data-driven constraints rather than actual physical moments of the distribution.

moments need to match for the generated distribution to be indistinguishable from the data.

ML based methods are moving toward treating unfolding as a fitting problem, using flexible function approximators and powerful statistical distances (e.g. Wasserstein distance) to overcome binning and to handle large dimensional data. These methods still must address ill posedness³⁸, but they offer a way to directly target the entire set of distributional degrees of freedom rather than pre-selecting a fixed set of basis functions.

VI.B.3 Benefits and challenges of MOMENT UNFOLDING.

Focusing on moments as the quantities to unfold offers several clear benefits. First, moments are often directly related to physics predictions. Many theoretical models provide predictions for means, variances, or other moment-like observables, especially in QCD where sum rules and scaling laws involve moments of distributions. Unfolding at the level of moments thus yields results that can be compared to theory with minimal further processing.

Second, by compressing the data into a few numbers, moment based unfolding can greatly reduce statistical fluctuations and noise sensitivity. A small set of global features (e.g. the first few moments) can usually be determined more precisely than an entire binned spectrum. The variance of estimators is lower since one has effectively integrated over many events to get each moment. This was one motivation for the

³⁸E.g., through network architecture choices, training tricks, or implicit regularisation like early stopping.

development of the MOMENT UNFOLDING method; measuring moments directly can be more precise and less sensitive to binning choices.

Third, MOMENT UNFOLDING can simplify high dimensional problems. In multi-differential measurements with many kinematic variables, fully binning the data becomes impractical due to sparsity (curse of dimensionality), but one might still meaningfully measure lower dimensional moments, such as the average of some observable as a function of another variable, avoiding the need to populate multi dimensional histograms. In summary, moment based approaches concentrate on the most salient features of the distribution, potentially yielding robust, computationally efficient unfolding results when only those features are of interest.

However, there are important limitations and challenges inherent to moment based unfolding. By construction, focusing on a limited set of moments discards information contained in the distribution beyond those moments. Two very different underlying distributions can share the same few moments. Thus unfolding only those moments provides an incomplete picture. This non-uniqueness means that moment based results must be interpreted carefully. They answer only the questions they explicitly ask. For instance, if only the first two moments (mean and variance) are unfolded, any differences residing in the non-Gaussian shape like skewness, tails, or multimodal structure will go undetected.

In this way MOMENT UNFOLDING shifts the ill posedness to the choice of moments. One must assume or hope that the chosen moment set is sufficient to capture the physically relevant differences. If an unexpected feature lies outside this span, it will be missed. This is connected to the concept of model dependence—choosing an

insufficient set of moments imposes a choice of model on the unfolding result, thereby biasing it.

A few different options exist to regularise the problem. For example, if unfolding a large number of moments, one might impose a smooth fall off in the moment values, since very high order moments tend to be increasingly sensitive to rare tails. In the adversarial MOMENT UNFOLDING approach, the number of moments n chosen to unfold acts as a regularisation knob: a small n strongly regularises, since it ignores any structure beyond the n th moment, while a larger n allows more detailed structure at the cost of higher variance and potential instability. This is analogous to the bias-variance trade-off in classic unfolding: using fewer moments yields a biased but low variance estimate; using more moments approaches an unbiased full result but with higher variance and higher risk of over fitting data fluctuations.

Another limitation of any reweighting based unfolding, moment based or otherwise, is the requirement of support overlap. The method can only correct distributions within the support of the MC distribution. If the true distribution has events in regions where the simulation has zero or negligible events, no amount of reweighting or moment adjustment can cover that gap. One would need to generate new events, a different approach entirely, or rely on extrapolation. Thus, MOMENT UNFOLDING, like other reweighting methods, assumes that the Generation's phase space is broad enough to contain the Truth or that any deficiencies are corrected by other procedures. It is worth noting that as one increases the number of moments or attempts to unfold the full distribution, the demand on simulation support becomes stricter: essentially

all features of the Truth distribution must be present in the Generation to be recovered by reweighting.

VI.B.4 Extending MOMENT UNFOLDING.

The Boltzmann ansatz can be viewed as a minimal deformation of the Simulation needed to match Data. $\beta_a = 0$ for all a corresponds to $g(z) = 1$. This implies no reweighting, i.e. using the raw Generation as is. As one turns on β_a s, one adjusts the weight of each MC event based on its z value. The exponential form guarantees that no arbitrary structure is introduced beyond that needed to fulfil the moment constraints i.e. it is the maximum entropy solution consistent with those constraints.

Another perspective is to see $g(z)$ as the Radon–Nikodym derivative between the true distribution $p(z)$ and the generation $q(z)$, restricted to the family of functions parametrised by β_a [790, 796]. If the true distribution $p(z)$ lies within this family for the chosen basis set, then there exists some β_* such that $p(z) = g(z; \beta_*) q(z)$. Even if $p(z)$ is outside this family, we expect that with a sufficiently rich set of basis functions, one can approximate $p(z)$ in the moment sense. In summary, Equation (VI.4) translates the physics question into a set of parameters β_a to be determined.

Once the optimal β_a parameters are found, the unfolded moments are obtained immediately by computing the weighted averages in the reweighted particle level sample. The unfolded k^{th} moment of Z is simply

$$\mathbb{E}_{\text{unfolded}} [Z^k] = \int z^k g(z; \beta) q(z) \, dz, \quad (\text{VI.6})$$

or in practice, the ratio of weighted sums over all generation events,

$$\mathbb{E}[Z^k] = \sum_{z \in \text{Gen.}} z^k g(z; \boldsymbol{\beta}) . \quad (\text{VI.7})$$

By construction, these unfolded moments will agree with the true moments of $p(z)$ if the procedure succeeds in finding β_a such that the reweighted simulation matches the data at detector level. g incorporates the normalisation factor $P(\boldsymbol{\beta})$ in the g . This ensures that $g(z; \boldsymbol{\beta})$ does not arbitrarily change the total number of events. In the language of statistical mechanics,

$$P(\boldsymbol{\beta}) = \int g(z; \boldsymbol{\beta}) q(z) \, dz \quad (\text{VI.8})$$

is the partition function ensuring probability conservation. It guarantees that the reweighted distribution properly normalises to the total cross section or total event count expected.

There is no explicit formula linking β_a to a simple goodness of fit measure at detector level, because the detector response $r(x | z)$ is typically a complex stochastic mapping. Instead, the machine learning training implicitly evaluates and drives this optimisation.

The MOMENT UNFOLDING method was demonstrated to yield accurate and statistically stable moment estimates in realistic scenarios. By avoiding any intermediate binning, it circumvents bin size biases and fully utilises the fine grained shape information in the data. Moreover, it is remarkably efficient. Focusing on a handful of parameters β_a is much simpler than attempting to learn a full function as a normal generative algorithm would. In essence, it reduces the unfolding task to a parameter fitting problem under the hood, albeit a high dimensional, simulation informed and

adversarially solved one. This focus provided a built in regularisation through the number of moments unfolded. The results in Section V.D showed that even with n as small as 2 one can capture key shape characteristics of distributions.

However, ultimately one may wish to unfold the entire distribution without having to pre-select moments. This is where the transition to extending the MOMENT UNFOLDING concept to all moments leads to the Reweighting Adversarial Networks (RAN) framework.

While unfolding a fixed set of moments is valuable, the ultimate goal in many analyses is to recover the complete differential cross section, effectively, to determine the true distribution $p(z)$ itself within resolution limits. The MOMENT UNFOLDING framework provides a natural stepping stone to this goal. Conceptually, one can imagine increasing the number of moments n in the weight function Equation (V.9) to capture finer and finer detail of the distribution. In the limit $n \rightarrow \infty$, the weight function $g(z, \boldsymbol{\beta})$ could represent an arbitrarily complex reweighting, capable of morphing the simulation into the true distribution. In practice, however, trying to include a very large number of moment parameters directly poses challenges.

The normalisation factor $P(\boldsymbol{\beta})$ becomes increasingly complicated to compute or differentiate when $\boldsymbol{\beta}$ is high dimensional, and the adversarial training becomes increasingly unstable due to the number degrees of freedom to correctly learn. In the initial MOMENT UNFOLDING studies, keeping n small was important for stable training—it acted as a strong regulariser and simplified the learning problem.

To extend the method toward full distributions, the Reweighting Adversarial Networks (RAN) approach was developed. RANs can be seen as MOMENT UNFOLDING

taken to the continuum limit, where instead of a few predetermined moments, all aspects of the distribution are learned. Practically, a RAN forgoes the explicit regularisation of $g(z, \boldsymbol{\beta})$ with a fixed set of β_a multiplying known basis functions. Instead, it employs a more general function approximator such as a neural network or some other high capacity ansatz to represent $g(z, \boldsymbol{\beta})$, which can adjust weights flexibly for different regions of phase space.

The adversarial setup remains, but now the generator’s parameters are not directly interpretable as specific moments;³⁹ they are a more granular description of the weight function. The discriminator in a RAN is tasked with comparing the fully continuous shapes of $\tilde{q}(x)$ and $p(x)$, in theory, enforcing an infinite number of constraints, since matching two distributions means matching an infinite set of moments or test statistics. In this way, a RAN builds on the moment based approach by removing the artificial limit on the number of moments: the method ‘learns to unfold all the moments of a distribution’.

One important innovation in the development of RANs was the use of the Wasserstein GAN framework [201], which provides a stable way to train the adversarial system even when the generator distribution and true distribution initially differ significantly.

This brings up a subtle issue. When extending to complete distributions, one must contend with the fact that $g(z)$ estimates the partition function from batch data. In a naive GAN, if $g(z)$ significantly changes normalisation, the discriminator could easily detect a total rate difference, and the generator would then simply rescale

³⁹Although the moment interpretation can be recovered by computing the Taylor coefficients of the log generator.

its prediction to match the event counts, potentially neglecting shape differences. By using an optimal transport loss and carefully constraining $g(z)$ to keep the total weight near unity, a RAN avoids trivial solutions and focuses on shape corrections. In other words, if a RAN does not incorporate the normalisation into its architectural design, then the architecture can lead to misestimations of the normalisation, and the RAN will have to learn the correct normalisation, at the expense of learning other features.

Although this was also true of MOMENT UNFOLDING, there is a quantitative difference. When n was small, a few β primarily modify the shape of a distribution within a highly restrictive scheme. Misestimations of the partition function become increasingly relevant when $g(z)$ becomes a free form function.

By allowing a more expressive reweighting function, a RAN is able to unfold the entire distribution in an unbinned, non-iterative manner. This stands in contrast to OMNIFOLD, which, while also ultimately recovering the full distribution, does so via multiple iterative reweighting steps. A RAN achieves a similar end point in one training loop by leveraging an adversarial paradigm. However, this power comes with technical challenges: the loss landscape is more complex and the risk of failure to converge is higher when the models essentially have as many ‘parameters’ as there are events. The implementation of a RAN thus requires additional care in training⁴⁰ to ensure convergence to a physically sensible solution that generalises and does not chase statistical fluctuations. The success of RANs in toy examples and simulated

⁴⁰E.g., gradient penalty terms for WGAN, spectral normalisation, *et cetera*.

jet data applications demonstrates that it is indeed possible to go from moments to complete spectra.

In this way, the progression from MOMENT UNFOLDING to full distribution unfolding can be seen as a continuum. At one end is a highly stable method focusing on a few moments; at the other end, a method that aims to extract the entire distribution, which requires stronger techniques to regularise. MOMENT UNFOLDING paves the way to the more ambitious goal of unbinned differential cross section measurement via its extension to RANs. This connection between moments and spectra ensures that the insights and constraints from one domain like theoretical expectations for certain moments can be seamlessly integrated into the process of obtaining the full spectrum. By casting the unfolding problem in the language of moment constraints and exponential families, one can gain both intuitive and quantitative control over the unfolding, whether one stops at a few moments or pushes onward to recover the complete distribution. This framework not only forms a bridge between techniques but also provides a clear rationale for why unfolding even a handful of moments is a stepping stone to unfolding distributions. The moments are the fingerprints of the distribution, and once one learns to reliably recover those fingerprints, one is better equipped to handle the the full differential cross section.

VI.C Methodology and regularisation.

VI.C.1	Historical approaches to regularisation.	346
VI.C.2	Regularisation in a RAN.	348
VI.C.2.i	Wasserstein GAN	349
VI.C.2.ii	Regularising the critic: Lipschitz constraints	353
VI.C.2.iii	Regularising the generator: initialisation and activation constraints.	354
VI.C.2.iv	Regularising with Generation	360

A variety of regularisation strategies have been developed historically and in contemporary machine learning approaches to ensure stable, physically meaningful unfolding solutions. This section begins by reviewing these strategies, from classical techniques to modern ML based methods, before detailing how the Reweighting Adversarial Network (RAN) method is designed to address regularisation. The focus of this section will be on the theoretical and conceptual aspects of the RAN approach, rather than on implementation details.

VI.C.1 Historical approaches to regularisation.

Classical unfolding methods introduce explicit regularisation to tame the ill posed nature of the problem. Several such examples are discussed in Sections II.D, III.B.3.iii and III.C.7. All of these classical techniques acknowledge the need for external constraints to obtain a stable inversion of the detector response.

In recent years, unbinned and high dimensional unfolding methods based on machine learning have emerged, necessitating new ways to incorporate regularisation. Many ML based approaches avoid explicit binning but must still combat the same amplification of statistical noise. Often, the regularisation in ML approaches is *implicit*, coming from network architectures and training protocols.

For example, deep neural networks have a finite capacity and tend to learn simpler patterns first, a form of Occam’s razor that can act as an ansatz. Choices of network depth, width, and activation functions impose smoothness biases on the learned mapping. In convolutional networks, weight sharing and locality encode translational invariance, effectively constraining the space of possible unfolding results to those

that respect certain symmetries. Many ML unfolding methods also apply standard regularisation techniques from predictive modelling, such as weight decay and L^2 penalties [107], dropout [525], or batch normalisation [491], to prevent overfitting. These generic measures help ensure the learned model does not simply reproduce statistical fluctuations in the training sample.

In adaptive or iterative ML unfolding algorithms, an explicit analogue of early stopping is used. The OMNIFOLD method, for instance, performs successive reweighting iterations with neural classifiers. By limiting the number of iterations and monitoring when the updates become consistent with statistical uncertainty, OMNIFOLD effectively regularises the solution. If iterated to completion, an algorithm like OMNIFOLD would overfit the data. This is analogous to IBU; halting earlier mitigates this risk.

Some neural network based unfolding methods also permit incorporation of physics informed constraints as a form of regularisation. For example, one can modify the loss function to penalise unphysical deviations or enforce known conservation laws and symmetries [52, 133, 175, 222, 230, 237, 302, 384]. Recent studies have integrated exact symmetry constraints into normalising flow or invertible network architectures used for unfolding, so that the learned particle level distribution automatically preserves quantities like momentum or charge [156, 311]. This reduces the solution space *a priori* to physically plausible ones, which is a powerful regularisation aligning with domain knowledge.

Beyond reweighting approaches, generative models learn to generate events from scratch so that, when passed through a detector simulation, they reproduce the observed data [318, 394]. Examples include normalising flows and variational au-

toencoders trained to invert detector effects. These models face even greater risk of instability due to their high flexibility. Regularisation here can take the form of strong prior distributions, for instance, seeding a flow with a known prior and only allowing limited deviation [117, 234], or carefully tuned network constraints. In practice, the success of ML generative unfolding has been limited by the need for large training samples and the difficulty of covering the full phase space without introducing spurious modes [32, 117]. Indeed, most existing ML unfolding methods still assume that the Simulation sample has sufficient overlap with the Data in all regions of interest. If the Simulation has zero or very little support in some region that the Data populates, no amount of reweighting or network fitting can compensate; as discussed, this can be both a feature or a bug.

VI.C.2 Regularisation in a RAN.

The Reweighting Adversarial Network approach is expressly designed to address regularisation challenges while performing unbinned unfolding in one stage. RAN can be viewed conceptually as an extension of the *Moment Unfolding* method, which unfolded only a fixed set of distribution moments and thereby enjoyed a strong built in regularisation by restricting to a few summary statistics. RAN lifts that restriction by attempting to unfold the entire distribution by extending to the ‘infinite moment’ limit, thus requiring new regularisation techniques to keep the problem well defined. Unlike OMNIFOLD, which iteratively trains many separate classifiers, RAN solves for all weights in a single training to convergence by jointly optimising g and d . This provides a significant computational advantage, but it demands a very stable training

procedure because the full flexibility of $g(z; \boldsymbol{\beta})$ is unleashed at once. To ensure stable and meaningful solutions, RAN incorporates several regularisation driven design choices in its methodology.

VI.C.2.i Wasserstein GAN

First and foremost, RAN employs a Wasserstein GAN (WGAN) framework for the adversarial training, in place of a standard Jensen–Shannon divergence based GAN. In a vanilla GAN, the discriminator $d(x)$ is trained to output a binary classification and the generator is trained to fool it, typically using a binary cross entropy or similar loss. Such losses correspond to minimising a divergence like the Kullback–Leibler divergence or the Jensen–Shannon divergence; if the real and model distributions do not overlap, the gradient of these losses vanishes, leading to unstable or mode collapsed solutions.

By contrast, a RAN leverages the Wasserstein distance as the measure of difference between distributions. The WGAN critic $d(x)$ does not output a class label but rather a real valued score, which should be high for real data and low for (in a RAN’s case, reweighted) simulation. The generator seeks to maximise this critic score on its output, thereby minimising the Wasserstein distance between the weighted Simulation and the Data distributions. The objective function adopted in training a RAN can be written as

$$L[g, d] = \mathbb{E}_{x \sim p} [d(x)] - \mathbb{E}_{(z, x) \sim q} [g(z; \boldsymbol{\beta}) d(x)] + \lambda \mathbb{E} [(\|\nabla_{\hat{x}} d(\hat{x})\| - 1)^2] , \quad (\text{VI.9})$$

where $x \sim p$ denotes a detector level event from the Data, and $(z, x) \sim q$ indicates a paired MC event with Generation features z and corresponding Simulation observation x . The MC is generative, so $z \rightarrow x$ are always paired via the detector model.

The first two terms estimate the Wasserstein-1 distance between the Data and reweighted Simulation distributions. The last term is a *gradient penalty* [205] with strength λ , enforcing the requirement that $d(x)$ be 1-Lipschitz continuous, as required for the Wasserstein distance theory [438]. This penalty drives $\|\nabla d\|$ toward unity on random points $\hat{x}$ interpolated between Data and Simulation samples, a technique introduced in [434] to stabilize WGAN training.

The minimax objective on the loss Equation (VI.9) with respect to g and d corresponds to the following adversarial optimisation:

1. d tries to maximise L_{WGAN} by making the Data–Simulation discrepancy as large as possible by assigning higher scores to real Data than to weighted Simulation samples.
2. g tries to minimise it to make the reweighted Simulation look as similar as possible to Data by adjusting weights $g(z; \beta)$.

At equilibrium, the weighted simulation cannot be distinguished from data by any 1-Lipschitz function, meaning the distributions are empirically matched in the optimal transport sense. This setup is schematically illustrated in Figure VI.1

The choice of a WGAN loss is fundamentally motivated by regularisation considerations. Because the Wasserstein distance evaluates the ‘distance’ between distributions in terms of an optimal transport cost, it provides smooth, meaningful gradients even

Reweighting Adversarial Networks




Figure VI.1: The training setup of a RAN. g generates weights to reweight the Generation. The Simulation dataset inherits its weights from the matching Generation events. The detector emulation is only run once, since a reweighted Simulation is created via importance weights and not by changing the features themselves. The critic d learns to classify events between Data and reweighted Simulation. When the discriminator is maximally confounded, the reweighting is successful.

when the support of the model and data distributions do not overlap perfectly [419, 420]. In other words, if the simulation is deficient in some region of phase space, a standard GAN discriminator would drive $g(z; \beta)$ to extreme values or simply saturate, providing no gradient in an attempt to cover the deficit, often leading to unstable oscillations in $g(z; \beta)$, with very large weights assigned to a few events. The WGAN critic, however, will assign a large positive $d(x)$ to real Data regions with no

Simulation equivalent, but the gradient of its linear output will encourage nearby Simulation events to increase in weight in a more controlled fashion, effectively telling the generator *which direction* to move probability mass to reduce the discrepancy. This greatly mitigates mode collapse and high variance weight solutions [298, 324].

In the context of an ‘infinite moment’ unfolding unconstrained by moment selection, using the WGAN framework is very helpful to avoiding the generator instabilities that occur with a naive GAN approach. Empirically, replacing the binary cross entropy GAN loss with the WGAN loss stabilises the training and yields a smoother weight distribution $g(z; \boldsymbol{\beta})$, especially in sparse regions of the detector feature space.

An additional benefit is that a RAN can, in fact, handle cases with minimal detector level overlap between Simulation and Data. Since the Wasserstein metric remains finite as long as the support of the weighted simulation can be transported to cover data, even if initially disjoint, RAN does not strictly require the Simulation and Data distributions to overlap at all. Of course, overlap at particle level is still required. RAN cannot invent Truth events outside the Generation’s support. Nonetheless, the ability to tolerate detector level mismatches is a major advantage over methods that rely on probability ratios, which diverge when supports do not overlap. In essence, the use of WGAN in RAN provides a strong form of regularisation by design: it selects a smoother notion of distribution difference that avoids infinite gradients and excessive weight variance.

VI.C.2.ii Regularising the critic: Lipschitz constraints

The critic network $d(x)$ in a RAN is subject to the 1–Lipschitz constraint required by the Wasserstein theory. This is enforced through two complementary regularisation techniques. The first, already noted, is the **gradient penalty** term in Equation (VI.9). Rather than clipping weights⁴¹, RANs adopt the improved strategy of penalising the norm of the critic’s gradient on randomly sampled interpolations between Data and Simulation to be close to 1. This smoothly encourages $d(x)$ to remain within the space of 1–Lipschitz functions without a sharp constraint on its parameters.

The second technique used to regularise the RAN’s training is **spectral normalisation** applied to the critic’s layers [405]. Spectral normalisation rescales the weight matrix of each layer such that its largest singular value is 1. Doing so at every training step effectively bounds the Lipschitz constant of each linear component of $d(x)$, ensuring that $d(x)$ cannot change faster than a certain rate with respect to its input [173, 195, 246, 385, 389]. Even small violations of the Lipschitz condition can lead to training instabilities, since the critic might exploit them to push L_{WGAN} lower, causing g to react with wildly large weights, making it essential that the Lipschitz constraint be observed strictly and carefully.

Spectral normalisation prevents the critic from assigning arbitrarily large scores to individual samples. This has an important regularising effect on the generator’s behaviour. If $d(x)$ is bounded and smooth, the incentives for $g(z)$ to produce extremely large weight factors are reduced, because no single event can ever receive a disproportionately large score advantage. In effect, spectral normalisation places

⁴¹The original WGAN approach, [420] which can impede learning by reducing capacity.

an upper bound on the influence of any individual Data—Simulation discrepancy on the loss. The combined use of gradient penalties and spectral normalisation in RANs was found to be effective for training stability—they keep the critic in check so that the adversarial training remains in a regime where gradients are informative and the generator’s updates remain moderate. With these in place, a RAN substantially avoids the divergence and mode collapse issues that plague naive adversarial training [273, 296, 354, 433].

VI.C.2.iii Regularising the generator: initialisation and activation constraints.

On the generator side, a RAN incorporates two important measures to regularise the solution. First, $g(z; \beta)$ is initialised close to the identity function before beginning adversarial training. In the context of unfolding, the ‘identity’ reweighting corresponds to assigning weight 1 to every MC event, i.e. initially assuming the Generation is the Truth and no reweighting is needed.. This initialisation reflects the belief that the generator only needs to make relatively small, smooth adjustments to the MC, which is a reasonable assumption in many cases where the Monte Carlo is approximate reality reasonably well.

By starting close to the identity mapping, a RAN avoids a situation where early in training the critic finds large differences and drives g to extreme compensations. Large fluctuations in weights early on can kick off a feedback loop of instability. The generator generates large weights in response to the critic’s output, the critic then sees a very irregular simulated distribution and reacts pathologically, and so on. The

initialisation acts as a strong regulariser by biasing g toward the null reweighting unless the data indicate otherwise.

This idea is analogous to using a Bayesian prior equal to the MC and regularising toward it at the start; only genuine discrepancies pull g away from 1. Indeed, from the perspective of classical regularisation, a RAN’s weight initialisation is equivalent to choosing the MC as a prior and initially penalising deviations from it. This leads to a much smoother evolution of the adversarial game—initially the critic cannot easily distinguish Data from reweighted Simulation because $g \approx 1$ still leaves them broadly similar. Over the course of the training, d learns gradually to distinguish the two, and g then adapts gradually in turn to make them more and more similar. This procedure significantly reduces the risk of the training falling into bad local minima or diverging in the first few epochs. In machine learning terminology, initialising close to the identity provides a ‘cold start’ regularisation that keeps the RAN in the perturbative regime where it performs best, rather than having to learn an unpredictable transformation from scratch [69, 217, 235, 265, 371].

The second generator-side regularisation is controlling the asymptotic behaviour of the generator output parametrisation for the weights. It is essential that $g(z)$ produce positive weights because events can only be reweighted by positive factors. It is also essential that the generator have the capacity to yield a wide range of values because some events might need weight $\gg 1$ if data are more abundant there, others weight < 1 if data are scarce.

The most straightforward choice is to define $g(z) = \exp[f_\theta(z)]$ where $f_\theta(z)$ is the output of a neural network with no constraints, and the exponential ensures positivity

and unboundedness. This choice is, in a sense, very well motivated by the Boltzmann distribution that inspired MOMENT UNFOLDING. In practice however, one finds that this choice, while mathematically valid, leads to numerical instability. The exponential is a very rapidly growing function, so any relatively large output of the network $f_\theta(z)$ would translate into an astronomically large weight, drastically skewing the training. Moreover, the gradient of $\exp[f_\theta(z)]$ is proportional to the output itself, so once a weight becomes huge, its gradient is huge too, often causing large and unstable oscillations.

To remedy this, a custom activation was designed, that tames the asymptotic growth of $g(z)$ at large network outputs while preserving the desired mathematical properties. In the final layer of the $g(z)$ network therefore one is advised to replace the $\exp$ with a composite function

$$g(z) = \frac{1}{P} \ln(1 + \text{softplus}(z)), \quad (\text{VI.10})$$

where $\text{softplus}(x) = \ln(1 + e^x)$ is a smooth approximation to a ReLU which grows linearly for large x , but is smooth and non-zero for all real-valued inputs [575, 675]. Wrapping softplus in $\ln(1 + \cdot)$ further slows the growth for large inputs: as $x \rightarrow \infty$, $\text{softplus}(x) \sim x$, so $\ln(1 + \text{softplus}(x)) \sim \ln(1 + x) \sim \ln x$.

Thus for very large raw network output x , this activation grows only logarithmically, instead of exponentially or even linearly.

Figure VI.2 illustrates the behaviour of the custom activation function compared to standard alternatives. The logarithmic scale reveals the dramatically different asymptotic behaviours: while the exponential function (orange dots) grows rapidly, the sigmoid function (violet dotted line) exhibits the opposite problem—it saturates

at 1, making it impossible to assign large weights even when the data requires them. RANs' custom activation (blue solid line) navigates between these extremes, growing exponentially near zero but only logarithmically for large inputs.

In the region near zero, shown in detail in the inset, the smooth activations exhibit similar behaviour, ensuring that small network outputs still produce meaningful weight variations. The ReLU (green dash-dot line), while computationally efficient, introduces a non-differentiable point at zero and null gradients for negative inputs that can complicate gradient based optimisation. The sigmoid, softplus (red dashed line) and RANs' custom activation all maintain smoothness throughout their domain, with the sigmoid showing its characteristic S-shaped curve that begins to plateau even for moderate positive values.

The key advantage of the log-softplus activation becomes apparent when considering the full range of possible inputs. Where the exponential would produce numerically unstable weights in the millions, ReLU and softplus would exhibit linear growth and the sigmoid would cap all weights at 1 regardless of how strongly the data demands larger values. The $\ln(1 + \text{softplus})$ activation provides unbounded but controlled growth. For large positive inputs, this activation produces weights of order 1—large enough to be meaningful but small enough to maintain numerical stability. This design acts as an adaptive regulariser, automatically preventing the numerical instabilities of exponential or even linear parametrisations while avoiding the expressivity limitations of bounded functions.

The key advantage of this activation is worth re-iterating. For even modest positive inputs, where the natural choice of the exponential would produce weights

in the millions or billions,⁴² a RAN's $\ln(1 + \text{softplus})$. activation yields much more manageable weights.⁴³ This compression of the output range prevents the loss function dynamics from experiencing numerical instability while preserving the relative ordering of weights. The activation thus acts as an adaptive regulariser, automatically preventing the numerical instabilities that plague naive exponential parametrisations without requiring explicit weight clipping or normalisation.

This adjusted activation guarantees that

1. $g(z) > 0$ for any input, by construction of softplus and $\log$,
2. $g(z)$ can still produce arbitrarily large values in principle. This activation function has no finite upper bound, unlike a sigmoid which saturates. This ensures that a RAN does not limit the solution space unduly,
3. Extremely large weights are disfavoured.

In practice, this means the network would have to invest a lot of capacity to achieve a huge weight on a single event, which is only worthwhile if doing so reduces the Wasserstein distance significantly. Typically, it will be more effective to increase weights more evenly on a group of events covering a region of phase space than to make one weight enormous. This activation function thus serves as an internal regulariser, curbing the tendency of the solution to form spikes or outlier weights. Since this is a soft constraint, large weights are not forbidden, and can be realised if the data demand them. The use of this modified activation proves to be crucial

⁴² $\exp(13) \sim 10^6$, $\exp(20) \sim 10^9$

⁴³ $\ln(1 + \text{softplus}(13)) \sim 2.5$, $\ln(1 + \text{softplus}(20)) \sim 3$.




Figure VI.2: Comparison of activation functions considered for parametrising the generator. The main plot shows five functions on a log scale: log-softplus (blue solid), exp (orange dots), ReLU (green dash-dot), softplus (red dashed), and sigmoid (violet waved). The inset has linear scaling. Exp exhibits rapid growth for moderately large inputs. The sigmoid saturates at 1, preventing large weights when needed. ReLU exhibits linear growth, but returns null gradients for negative inputs. Softplus also exhibits linear growth, but is a smooth function with non-zero gradients. log-softplus preserves smoothness, positivity, and positive gradients, with exp-like behaviour near zero to logarithmic growth for large z .

for achieving stable training in RAN; with the standard exponential, one observes frequent instances of a single event's weight blowing up and derailing the fit, whereas the log-softplus activation yields more balanced weight distributions.

VI.C.2.iv Regularising with Generation

Finally, it is worth noting an intrinsic form of regularisation in a RAN that the architecture includes, ‘for free’. By construction, the unfolded result is expressed as a reweighted version of the initial Generation sample. This means the unfolded distribution cannot introduce structures that were not present (even in latent form) in the MC. In effect, the Generation provides a prior support and shape for the solution. If the true underlying distribution has features outside the Generation’s support, RANs, like any reweighting method, cannot recover them. But conversely, RANs will not produce spurious artifacts that violate known physics encoded in the MC.

This is similar in spirit to a Bayesian prior or a template fit: one starts with a template (the Generation) and only deform it as necessary to fit the data. The closer the Generation is to reality, the less adjustment is needed. The design choice of ‘modulating a known density rather than learning a new one from scratch’ is motivated by the desire to leverage this built-in form regularisation. This allows a RAN’s solutions to be more statistically robust than those from unconstrained generative models, precisely because $g(z)$ operates within the scaffolding of the Monte Carlo sample. This comes at the cost of some bias if the Monte Carlo is poor, which no method can escape without additional external input, but is a conscious regularisation choice favouring stability over unrestricted flexibility.

In summary, the RAN methodology interweaves modern ML techniques with classical regularisation principles to achieve a stable unfolding. The use of the Wasserstein distance provides a more flexible way of comparing distributions that avoids the brittleness of binary classification metrics. Imposing Lipschitz continuity

on the critic via gradient penalties and spectral normalisation keeps the adversarial game well behaved and prevents the discriminator from overemphasising statistical noise. On the generator side, starting from the Generation and limiting the capacity for extreme weights through initialisation and activation choices anchor the solution close to expected physics and discourage overfitting. Together, these innovations allow RANs to unfold full distributions in one shot, non-iteratively, without the severe instabilities that one might fear from an unconstrained GAN approach.

In effect, the RAN algorithm attains a balance. It is flexible enough to accurately fit complex, high dimensional data, yet sufficiently regularised to suppress unphysical oscillations and variance. The result is an unfolding method that achieves competitive performance to iterative methods like OMNIFOLD, while operating in a single training pass and maintaining controlled behaviour even in challenging regimes⁴⁴. The following sections will demonstrate these points quantitatively with examples, but the methodological foundation laid out here is key to understanding why RAN performs as well as it does. It marries the strengths of adversarial learning with careful regularisation choices, yielding a novel and powerful approach to this classic problem.

⁴⁴Such as limited detector overlap or regions with low statistics.

VI.D Machine learning implementation.

VI.D.1	Neural network architecture.	363
VI.D.2	Adversarial training procedure.	366
VI.D.3	Datasets and preprocessing.	370
VI.D.3.i	Gaussian example.	370
VI.D.3.ii	Jet substructure dataset.	372
VI.D.4	Training, monitoring and validation.	374

VI.D.1 Neural network architecture.

A Reweighting Adversarial Network (RAN) consists of two components: a *generator* network that assigns event-wise weights, and a *critic* network that evaluates the discrepancy between weighted simulation and data. RAN implements both as small fully connected, dense, neural networks using the TENSORFLOW 2/KERAS framework for prototyping [460, 483].⁴⁵ The architecture was kept intentionally simple to ensure stable training and avoid over fitting, with identical layer widths for the generator and critic.

The generator $g(z; \beta)$ takes as input the *particle level* feature vector $z \in \mathbb{R}^{N_T}$ of a generated event, where N_T is the dimension of the Generation, and the Truth. For example, $N_T = 1$ for one dimensional toy data or $N_T = 6$ for the multi-observable jet dataset described below.

The generator outputs a scalar weight $w = g(z)$ that reweights that event in order to correct the Simulation towards Data. The generator is a feed forward multilayer perceptron with three hidden layers of 50 nodes each. Each hidden layer uses a Leaky Rectified Linear Unit (Leaky ReLU) activation [538], defined in Equation (III.17), with slope parameter $\alpha = 0.3$.

Batch normalisation and a dropout of 20% is applied after each dense layer to promote stable training and reduce overfitting. The final layer of the generator is a single linear neuron (no activation) producing a raw scalar t . This output is then transformed to a positive weight. Instead of a direct exponential mapping the

⁴⁵The design is straightforward to reproduce in PYTORCH using analogous layers and normalisation techniques.

stabilised transformation,

$$w = \ln \left(1 + \text{softplus}(t) \right), \quad (\text{VI.11})$$

is applied to t .

Finally, the weights w_i in each mini-batch are normalised so that their average is unity. This batch level normalisation means the total weight of the simulated sample remains consistent, preserving overall event counts, by approximating the ‘partition function’. The generator network thus defines a differentiable reweighting function $g(z)$ that can smoothly adjust the contribution of each simulated event during training.

The critic $d(x; \theta)$ is a discriminator-like network that takes as input the detector level feature vector $x \in \mathbb{R}^{N_D}$ of an event, where N_D is the dimension of Data and Simulation, and outputs a real valued score. This network is tasked with distinguishing the weighted Simulation from the Data in the detector space, and its output is used as an approximate distance measure between the two distributions. Like the generator, the critic is implemented as a fully connected network with three hidden layers of 50 nodes each, using Leaky ReLU activations and interleaved batch normalisation and 20% dropout. The output layer is a single linear node producing the critic score $d(x)$.

To satisfy the requirements of the Wasserstein GAN framework, the critic must be a 1-Lipschitz function. As discussed, spectral normalisation scales the layer weights such that their largest singular value is 1, effectively controlling the Lipschitz constant of d without weight clipping. This technique, combined with dropout, greatly improved training stability by preventing the critic from becoming overly




Figure VI.3: Architecture of a RAN. **Left:** The generator processes particle level features through three hidden layers (50 nodes each) with LeakyReLU activation, batch normalisation, and 20% dropout. The output layer produces event weights using the log-softplus activation. **Right:** The critic evaluates detector level features through spectrally normalised dense layers with LeakyReLU activation, layer normalisation, and dropout, outputting a scalar score. The networks are trained adversarially using a Wasserstein loss with gradient penalty, where the critic learns to distinguish weighted Simulation from Data while the generator learns to reweight Generation to induce weights on Simulation minimise this distinction.

sensitive to outlier events. Additionally, a mild output clipping⁴⁶ is applied to the critic scores, capping $d(x)$ within a reasonable range to eliminate spurious large values that could arise from statistical fluctuations, which in this context would correspond to unphysical negative probabilities or overly large separation scores [223, 420]. A schematic of architecture implemented is shown in Figure VI.3.

⁴⁶Not weight clipping.

The balanced capacity of the generator and critic, each with $\mathcal{O}(10^4)$ trainable parameters, was found sufficient to learn the required reweighting functions without overfitting, given the complexity of the tasks considered.

VI.D.2 Adversarial training procedure.

Training a RAN involves a minimax game between the generator and critic, similar in spirit to a Generative Adversarial Network (GAN). However, unlike a standard GAN that generates new events, RAN's generator only reweights existing events. RAN adopts the Wasserstein GAN (WGAN) approach for improved stability. The critic's objective is to maximise the statistical distance, in the Wasserstein-1 sense, between the weighted simulation and the true data distributions at the detector level, while the generator's objective is to produce weights for Generation such that the induced weights on Simulation minimise this distance. Formally, let $p(x)$ be the true data distribution in detector space and $q(z, x)$ be the joint distribution of generation and simulation. At each training step, the critic d is trained to minimise the loss function

$$L_d = -\mathbb{E}_{x \sim p}[d(x)] + \mathbb{E}_{(z, x) \sim q}[g(z) d(x)], \quad (\text{VI.12})$$

where (z, x) denotes a corresponding Generation and Simulation event pair. $-L_d$ is an empirical estimate of the Wasserstein-1 distance between Data and reweighted Simulation. The generator, on the other hand, is trained to maximise this same quantity (equivalently, to minimise the negative critic loss $-L_d$), thereby pushing the weighted Simulation to resemble the Data.

Unlike the binary cross-entropy loss in a traditional GAN, the Wasserstein formulation provides a smooth, continuous loss landscape even when the two distributions do not overlap perfectly. This is helpful for unfolding problems because even if detector effects cause $p(x)$ and the original $q(x)$ to have non-overlapping support in x , the WGAN critic can still provide informative gradients to the generator. This choice eliminates the severe mode collapse or large weight oscillations that occur when using the BCE loss in extant classifier based methods. The WGAN's stronger theoretical guarantees, combined with the property that RAN only classifies at detector level, and only reweights at particle level, ensure the training remains well behaved even with minimal overlap [434]

To further enforce the 1-Lipschitz condition required by the Wasserstein theory, one can include a gradient penalty term in the critic's loss. This penalty,

$$L_{GP} = \lambda \mathbb{E}_{\hat{x}} [(\|\nabla_{\hat{x}} d(\hat{x})\|_2 - 1)^2], \quad (\text{VI.13})$$

is computed on randomly sampled points $\hat{x}$ uniformly interpolated between Data and reweighted Simulation. λ usually set between five and ten. The gradient penalty, together with spectral normalisation, acts as a robust regulariser against critic over-training.

Training proceeds by alternating between critic and generator updates in each iteration, following the typical WGAN-GP strategy. The studies presented below used an update ratio of 3:2 for the critic:generator, meaning that the critic is updated slightly more frequently to keep it near optimality as the generator evolves. Pseudocode for one training cycle is outlined in Algorithm 3.

Algorithm 3 WGAN GP Training Step

Input: Data $\sim p(x)$; Generation, Simulation pairs $\sim q(z, x)$,
 critic d_θ , generator g_β ,
 gradient penalty weight λ ,
 critic steps n_d , generator steps n_g , batch size B .

Output: Trained parameters θ and β

```

for iteration = 1, 2, ... do
    // Critic updates
    for i in 1 ...  $n_d$  do
        Sample  $\{x_i^{\text{data}}\}_{i=1}^B \sim p(x)$ 
        Sample  $\{(z_j, x_j)\}_{j=1}^B \sim q(z, x)$ 
        Compute  $w_j = g(z_j; \beta)$ 
        Compute
        
$$L_d = \frac{1}{B} \sum_{i=1}^B d_\theta(x_i^{\text{data}}) - \frac{1}{B} \sum_{j=1}^B w_j d_\theta(x_j^{\text{sim}}) + \lambda L_{\text{GP}}(\theta) \quad (\text{VI.14})$$

         $\theta \leftarrow \theta + \eta_d \nabla_\theta L_d$ 

```

These steps constitute one training iteration, and they are repeated until convergence criteria are met. Both networks are optimised using the RMSProp optimiser, which was found to outperform ADAM in the WGAN framework, consistent with previous studies [400]. The learning rate for both generator and critic are set in the range $\eta \sim 1 \times 10^{-4}$ to 5×10^{-4} , with the exact value tuned by dataset. A need for explicit learning rate decay schedules beyond those automatically implemented by KERAS, was not observed as training typically converged to a stable solution before

Algorithm 3 (continued)

for iteration = 1, 2, ... **do***// Generator updates***for** j in $1 \dots n_g$ **do**Sample fresh $\{x_i^{\text{data}}\}_{i=1}^B \sim p(x)$ Sample $\{(z_j, x_j)\}_{j=1}^B \sim q(z, x)$ Compute $w_j = g(z_j; \beta)$

Compute

$$L_g = -\frac{1}{B} \sum_{j=1}^B w_j d_{\theta}(x_j^{\text{sim}}) + \frac{1}{B} \sum_{i=1}^B d_{\theta}(x_i^{\text{data}}) \quad (\text{VI.15})$$

 $\beta \leftarrow \beta - \eta_g \nabla_{\beta} L_g$

any signs of stalling. A mini-batch size of $B = 256$ was used for all experiments, providing a good balance between gradient estimate stability and computational efficiency.

The studies presented in this chapter do not involve an exhaustive hyperparameter search; instead, they rely on standard GAN/WGAN settings from the literature and minor tuning. In practice, the performance of the RAN method is found to be robust against moderate changes in these settings: for instance, altering the learning rate by a factor of two or using 40 or 60 nodes per layer, instead of 50, had no significant impact on the final unfolded distributions. This insensitivity suggests that the RAN method does not require fine-tuned hyperparameters to achieve reliable results, an important consideration for reproducibility.

Throughout training, key indicators should be continuously monitored to ensure the optimisation remains stable. In particular, one should track the Wasserstein critic loss L_d and the gradient penalty magnitude. A sharp increase in the gradient penalty or a sudden large oscillation in L_d can indicate that the critic is violating the Lipschitz constraint or that the generator is producing pathological weights, a potential onset of mode collapse. If such behaviour is observed, one can intervene by pausing training and lowering the learning rate or increasing regularisation before resuming from the last stable checkpoint.

In the experiments presented, the use of spectral normalisation and gradient penalty together largely prevented such divergences. Here, it is important to re-emphasise that RAN’s one shot training, as opposed to an iterative approach, simplifies the overall optimisation schedule. Once the adversarial equilibrium is reached, one immediately obtains the final reweighting function without needing multiple retraining cycles.

VI.D.3 Datasets and preprocessing.

The aforementioned architecture and training procedure was applied to two representative unfolding scenarios: a controlled 1–dimensional Gaussian example and a realistic high dimensional jet physics example. These two cases allow one to validate the RAN implementation on both simple and complex distributions.

VI.D.3.i Gaussian example.

The first dataset is a synthetic scenario with known analytical truth, designed to illustrate RAN’s behavior under varying degrees of detector smearing. Generate two

sets of events from normal distributions at the particle level, a Truth distribution $Z_T \sim \mathcal{N}(\mu_T, 1)$ and a Generation distribution $Z_G \sim \mathcal{N}(\mu_G, 1)$, with means $\mu_G = 0$ and $\mu_T = -1$. Sample 10^4 events from the Truth distribution and 10^5 events from the generation distribution.

In lieu of detector effects, each particle level value z is transformed by a simple deterministic response: $x = S \cdot z$, where S is a scalar ‘smearing’ factor which is constant for any given individual unfolding experiment. For example, when $S = 1$, the Data equals the Truth with no smearing, while $S > 1$ stretches the distribution, reducing the overlap between the Data ($X_D = S \cdot Z_T$) and Simulation ($X_S = S \cdot Z_G$). Data is unfolded over a range of S values representing increasingly degraded detector responses.

Since this is a one dimensional setup, $N_T = N_D = 1$ i.e, each event is described by a single number. Apart from the deterministic scaling, no additional noise or inefficiency is introduced, so the challenge lies purely in the shift of distributions and, more importantly the lack of detector level support overlap as S grows large. No additional preprocessing of the input features is necessary since they are already bounded and standardised aside from mean shifts.

During training, at each iteration draw a random batch of $B = 256$ values from the 10^4 available smeared Data points and another batch of 256 from the 10^5 smeared Simulation points. Because the Simulation sample is larger, cycle through it randomly shuffled such that multiple data epochs correspond to one pass through the simulation. It is beneficial to shuffle and reshuffle both samples frequently during training to avoid any periodic artifacts given the fixed deterministic smearing.

The generator is tasked with learning event weights $g(z)$ that correct the Generation distribution centred at -1 to match the true distribution centred at 0 , using only classification information from the X domain where in extreme cases the distributions barely overlap.

This toy example is especially useful for validating the implementation because the optimal reweighting function is known. In the limit of infinite statistics, the optimal weight would be proportional to the ratio of Truth to Generation densities at z , which in this case is

$$g_* = \exp \left[-\frac{1}{2}(z - \mu_T)^2 + \frac{1}{2}(z - \mu_G)^2 \right] \quad (\text{VI.16})$$

$$= \exp \left(z + \frac{1}{2} \right). \quad (\text{VI.17})$$

This example also allows one to verify that a RAN's adversarial training indeed converges to the correct solution, by comparing the unfolded distribution to the analytic expectation.

VI.D.3.ii Jet substructure dataset.

The second dataset studied is a realistic example taken from high energy physics, involving the unfolding of jet substructure observables. This experiment follows the setup of the OMNIFOLD study described in [309] to enable direct comparisons.

Proton–proton collision events at $\sqrt{s} = 14$ TeV are generated using two different Monte Carlo simulators, PYTHIA8.243 [502, 612] with Tune 26 [506] provides the initial Generation sample, and HERWIG 7.1.5 [421, 606] is used to generate an independent sample treated as the Truth distribution. Each event in both samples contains a

high- p_T Z boson with $p_T^Z > 200$ GeV decaying leptonically, produced back to back with a hadronic jet. The presence of a Z boson trigger ensures a well defined event sample and reduces acceptance differences. The final state particles in each event are passed through a fast detector simulation using DELPHES 3.4.2 [513] with a CMS detector configuration including particle flow reconstruction [451, 496]. Jets are then reconstructed in both the particle level and detector level event records using the anti- k_T clustering algorithm [607] with a radius parameter, $R = 0.4$, as implemented in FASTJET3.3.2 [547]

This dataset focuses on the hardest jet in each event, the jet that balances the Z boson's recoil, to define its observables. For each jet, this study considers six different substructure observables that characterise its internal pattern of particles. These include classic measures like the jet mass m and jet width w , as well as more specialised variables such as the N -subjettiness ratio τ_{21} , the groomed momentum fraction z_g from the Soft Drop algorithm, and a logarithmic version of the dimensionless jet mass $\ln \rho$. Precise definitions of these observables follow those in [309].

In total, each jet is described by a feature vector of $N_T = N_D = 6$ numbers (since the same set of observables is defined at particle level and, after DELPHES simulation, at detector level). These six features form the input to the RAN.

It is helpful to whiten (z-score) these inputs to minimise finite precision and overflow errors. The data is prepared such that the particle level Truth sample from HERWIG contains on the order of 10^6 jets, and the particle level Generation sample from PYTHIA, a similar order of magnitude. The detector level representations for

each sample are obtained via the same DELPHES procedure, yielding paired (z, x) examples for PYTHIA and independent x examples for HERWIG.

During training, the DELPHES+HERWIG jets are treated as the Data distribution $p(x)$ and the DELPHES+PYTHIA jets as the initial Simulation, $q(x)$. At each iteration, mini-batches of size $B = 256$ jets are drawn from each distribution. Given the relatively large size of the samples, one need not worry about cycling through the data too many times and can simply train for a fixed number of epochs determined by early stopping. One must ensure that each jet in the training sample is used many times over the course of training, which is important for the generator to refine the weights on all regions of phase space. Once training is complete, the learned weights $w = g(z)$ can be applied directly to the Generation events to produce an unfolded distribution, which can then be compared to the Truth, i.e. HERWIG particle level distribution.

It is important to emphasise that in this experiment the ‘Data’ is fully simulated, HERWIG+DELPHES rather than experimental observations or measurements. This is a common practice in validation studies, allowing one to quantitatively evaluate the accuracy of the unfolding by comparing to known truth. However, as is the case with any proposed method, the gold standard of its quality will be demonstrated through future work applying it to real experimental data.

VI.D.4 Training, monitoring and validation.

To ensure reliable performance and prevent overtraining, one should employ a rigorous monitoring and validation regimen during RAN training.

First, one should partition each dataset into a training set used to update network parameters, a smaller validation set, used for performance monitoring during the training, and a testing set. The ratios used were training : validation : testing = 70 : 15 : 15. The validation data is not seen by the networks during gradient updates, but one computes the same losses on it to check how well the model generalises. Specifically, after every fixed number of training epochs, one evaluates the critic loss L_d on the validation batch and tracks its evolution.

If one observes the validation loss starting to increase while the training loss continues to decrease, this is a clear sign of overfitting i.e. the critic is learning spurious differences that do not generalise, or the generator is over adjusting to peculiarities of the training sample. In such cases, one should invoke an *early stopping* criterion. In the implementation described, early stopping is triggered if the validation Wasserstein distance, approximated by $-L_f$ on the validation set, has not improved for 10 consecutive evaluation intervals. Once triggered, one stops training and rolls back to the generator state that had the best validation performance. This strategy ensures that one did not overtrain the critic to exploit statistical noise, which would manifest as unstable weights when applied to hitherto unseen validation data.

In addition to the WGAN loss, one can track several other high level divergence metrics between the reweighted Generation and Truth distributions. While these metrics are not used to train the model, since truth is not known in real applications, they are extremely useful to diagnose convergence and compare to alternative methods. One such metric is the **Wasserstein-1 distance** itself between the weighted and true distributions. Although the critic provides an estimate of this distance at detector

level via the loss, one can also compute it at particle level independently for simpler cases.

For the 1D Gaussian example, the Wasserstein distance has a closed form which reduces to the difference in means when both distributions are Gaussian with equal width, and for the jet observables one can calculate one dimensional Wasserstein distances for each observable's distribution individually. One can also monitor the Vincze–Le Cam (VLC) divergence, introduced in Section II.G.1.ii, between the unfolded reweighted Generation and the Truth. Lower values indicate closer agreement between distributions and $\Delta_{VLC} = 0$ if and only if $p = q$. This metric is particularly sensitive to differences in the tails of distributions and provides complementary check to the Wasserstein distance.

During training, one could periodically apply the current set of weights $g(z)$ to a large sample of simulation events and then evaluate Δ_{VLC} and the Wasserstein distance between this weighted sample and the truth sample using the known Truth events in validation studies. A decreasing trend in both metrics as training progresses signals that RAN is converging towards the correct reweighting.

For instance, in the Gaussian example with moderate smearing, one observes the VLC divergence between the weighted Generation and Truth decrease rapidly in the first few dozen epochs and then flatten out at a low value, indicating that the generator had found an effective weighting scheme. Similar behaviour is seen in the jet case: the combined six dimensional weight optimisation gradually improves the agreement in all observables. By the end of training, the unfolded distributions for all six jet observables are nearly indistinguishable from the truth, as quantified by final

Wasserstein distances on the order of 10^{-2} and VLC divergence on the order of 10^{-3} . Tables VI.3 and VI.4 are included in Section VI.E comparing these divergence metrics for RAN versus other methods; the RAN implementation achieves competitive or superior scores, validating the training procedure described.⁴⁷

The monitoring process also guards against the possibility of divergence or weight blow up. If at any point the training appears to destabilise,⁴⁸ one can intervene by stopping early and/or introducing additional regularisation. In practice, due to the regularisation implemented (spectral norm, gradient penalty, dropout), such pathological behaviour was rare.

One symptom one might encounter in some trials is the appearance of a few very large weights in sparsely populated regions of phase space. This is a known issue in unfolding problems. When there are regions that have little MC density but non-negligible data density, any reweighting based method will assign large weights to the few simulation events in that region, and a RAN is no exception. Nonetheless, the spectral normalisation on the critic help limit how fast and how large those weights could grow. If needed, one could impose an explicit cap on weights or an additional L_p norm term in the loss to penalise overly large weights, but this was not found to be necessary for these benchmarks.

In summary, the ML implementation of a 3RAN described above proves to be robust and reproducible across the different datasets. By carefully specifying the architecture with sufficient complexity to learn the needed reweighting, but regularised

⁴⁷See [309] for comparable benchmarks involving iterative methods.

⁴⁸E.g., the critic loss oscillates wildly or the weights $\{g(z)\}$ start to become extremely large for a small subset of events.

to avoid overfitting, by using a stable training objective in the Wasserstein divergence with gradient penalty and optimiser, RMSProp, and by employing diligent monitoring and early stopping criteria, one ensures that the training converges reliably to a good solution.

The design choices were informed by both theoretical considerations such as the advantage of WGAN for non-overlapping support and practical experimentation like the choice of 3 critic updates per 2 generator updates for efficiency. The resulting RAN model can be straightforwardly re-implemented in common deep learning frameworks and the reported hyperparameters can serve as a baseline for future applications. All training scripts and data will be made available open-source to facilitate verification and reuse.

VI.E Results

VI.E.1	Gaussian model with smearing.	380
VI.E.2	Jet substructure unfolding results.	382
VI.E.2.i	Overall performance.	384
VI.E.2.i.a	Stability and efficiency.	389
VI.E.2.ii	Generality of results.	391

This section evaluates Reweighting Adversarial Networks on the two case studies described in the preceding section. In each case, RANs’ performance is compared to two baseline methods, the OMNIFOLD algorithm and IBU. Quantitative divergence metrics are reported for each method that measure the algorithms’ ability to reproduce the Truth distributions, as well as the methods’ stability to statistical challenges.

VI.E.1 Gaussian model with smearing.

This study represents a targeted unfolding scenario to stress test RANs in a regime of limited detector level overlap. As the smearing factor increases above unity, the Data and Simulation become increasingly separated, with less and less overlapping support, even though the underlying particle level distributions remain fixed. This controlled setup allows an examination how each unfolding method copes with a progressively worsening mismatch between the Data and Simulation. Notably, methods like OMNIFOLD rely reweighting at the detector level, which can fail if the detector level supports do not overlap well. In contrast, RANs perform reweighting at particle level and use an optimal transport based loss to compare detector level outputs. One therefore expects a RAN to be more robust in scenarios with poor detector level overlap, provided the particle level distributions still cover each other.

To quantify performance the Vincze—Le Cam (VLC) divergence Δ_{VLC} between the unfolded outcome and the true distribution is used as a metric of accuracy. Figure VI.4 summarises the unfolding accuracy as a function of the smearing factor, with Δ_{VLC} plotted for RANs and OMNIFOLD.

At minimal smearing, both methods perform excellently, and OMNIFOLD achieves a slightly smaller divergence to truth than the RAN. For example, at smearing $S = 1$, Δ_{VLC} for OMNIFOLD is 0.057×10^{-3} , compared to the RAN's 0.122×10^{-3} . This confirms that when the detector effects are small and distributions have good support, the problem reduces to a straightforward reweighting and OMNIFOLD's direct classifier based approach is very effective. Indeed, it finds a nearly perfect weighting in this regime.

However, as the smearing increases and detector level overlap deteriorates, OMNIFOLD's performance degrades. Beyond a smearing factor of about 5, the OMNIFOLD solution's divergence from truth grows significantly. By the highest smearing values tested, where the Simulation and Data distributions are almost disjoint, OMNIFOLD's unfolded result significantly deviates from the true distribution, nearly two orders of magnitude worse Δ_{VLC} compared to the no smearing case.

These observations are consistent with expectations. OMNIFOLD tries to reweight at detector level, and when it cannot find common phase space between Data and Simulation, it struggles to assign meaningful weights. In stark contrast, RANs maintain stable performance across the entire range of smearings. As shown in Figure VI.4, RANs' Δ_{VLC} remains approximately constant, with only slight variation within statistical error, as smearing increases. Even at the largest smearing factors, the RAN's divergence with Truth is essentially the same as it was at zero smearing, indicating that the RAN successfully unfolds the distribution despite separation in detector level distributions. This behaviour can be attributed to the RAN's methodological property of only reweighting at particle level, combined with its use of

an optimal transport loss at detector level, which, unlike a classifier, does not require overlapping support to provide a meaningful gradient—the adversarial training can adjust weights to minimise the Wasserstein distance between the smeared Simulation and Data, even if those distributions are partially disjoint.

The error bars in Figure VI.4, representing 1σ uncertainties computed by bootstrapping the experiment, confirm the statistical stability of each method. RANs’ results are consistent across bootstrap resamples while OMNIFOLD’s uncertainty grows with smearing. In summary, this Gaussian study demonstrates that a RAN is robust to limited detector level support, recovering the true distribution reliably in a single pass, whereas OMNIFOLD requires sufficient detector level overlap to perform optimally. This advantage of RANs becomes valuable in scenarios where detector effects can be large.

VI.E.2 Jet substructure unfolding results.

This study assesses RANs on a realistic high energy physics unfolding task involving correcting jet substructure observables from detector level Data (simulated) to particle level Truth. This example serves as a stringent test of RANs’ ability to handle multi-dimensional, non-Gaussian distributions a range of complex features and long tails. The observables chosen span a range of distribution shapes—from broad distributions to sharply peaked ones and those with kinematic cut-offs—providing a comprehensive test bed. The RAN method is applied to learn a single weighting function $g(z)$ that assigns a single weight to each six dimensional Generation event based on its features




Figure VI.4: Comparison of unfolding performance under increasing detector smearing. The Δ_{VLC} between unfolded and true distributions is shown as a function of the smearing factor for a RAN (black circles) and OmniFold (red triangles). While OmniFold achieves slightly better accuracy at minimal smearing, its performance degrades as smearing increases, reflecting its reliance on detector level overlap. In contrast, the RAN maintains stable performance across all smearing levels, demonstrating robustness to limited detector level support. Generation (yellow diamonds) is shown for reference. Error bars indicate 1σ bootstrap confidence intervals.

to achieve this alignment between Data and reweighted Simulation, as described in Sections VI.C and VI.D.

For comparison OMNIFOLD and the classical IBU are also applied to the same task. Since IBU is a binned algorithm, one must perform separate one dimensional IBU unfolds for each observable using a fine histogram binning, to benchmark its performance on each one dimensional projection of the data.

VI.E.2.i Overall performance.

Figure VI.5 presents the unfolded particle level distributions for each of the six jet observables for the RAN, IBU, and OMNIFOLD, compared to the Truth distributions. Each panel shows the truth spectrum (HERWIG, treated as ‘Truth’) in filled blue, the generation in filled orange, the unfolded result from the RAN as a black solid line, the unfolded result from OMNIFOLD as a red dashed line and the result from IBU as a green dotted line. The ratio sub-panels below each panel display the unfolded/Truth ratio for each method, indicating how well each method approximates Truth across the range of each observable. Both the RAN and OMNIFOLD achieve good agreement with truth across all observables, a non-trivial accomplishment given the complexity of the distributions, and the RAN consistently provides a close match, particularly in regions that are challenging to unfold.

For example, in the jet mass distribution in Figure VI.5’s, top-left panel, the Generation initially underestimates the probability in the high mass tail and misplaces the peak relative to Truth. After unfolding, the RAN’s distribution closely follows the truth curve: it accurately reproduces the peak around $m \approx 20$ GeV and the long tail up to high masses. OMNIFOLD also significantly improves the agreement, but its ratio plot reveals a slight residual bias in the tail, with deviations of order 5–10% from unity in the highest mass bins, whereas the RAN’s ratio is within a few percent of unity throughout.




Figure VI.5: Particle level distributions of six jet substructure observables after unfolding. Each panel shows Truth as the blue filled histogram, Generation as the orange filled histogram, and unfolded results from the RAN (black solid), OMniFold (red dashed), and IBU (green dotted). The observables are jet mass m and soft-drop mass M (left column), constituent multiplicity and jet width τ_1 (middle column) and N -subjettiness ratio τ_{21} and groomed momentum fraction $\ln \rho$ (right column). The sub-panel under each panel shows the ratio of each unfolded result to truth, demonstrating that the RAN achieves close agreement across all observables, particularly in challenging regions such as distribution tails. Error bars represent the 1σ interval computed by training 20 independent runs for each method on the same dataset with different random initialisation.

A similar pattern is seen for the N -subjettiness ratio τ_{21} . The Generation initially produces a broader τ_{21} distribution than Truth (HERWIG), indicating an overestimation of two-prong substructure. RAN’s unfolding narrows this distribution to align almost exactly with the truth shape, correcting both the peak and the tail of the τ_{21} spectrum; OMNIFOLD moves in the right direction but leaves a noticeable difference. The OMNIFOLD unfolded τ_{21} distribution remains slightly too broad, with the ratio to truth dipping below 1.0 in the peak region and rising above 1.0 at higher τ_{21} .

For observables like the jet width w and logarithmic groomed mass $\ln \rho$, which have sharply peaked distributions near zero and long tails, RAN again demonstrates strong performance. It captures the steep drop off and the tail behaviour with high fidelity, where OMNIFOLD shows small mismodeling in the intermediate region for w and tail for $\ln \rho$.

The groomed momentum fraction z_g is a particularly challenging variable because of its hard cut-off at $z_g = 0.5$. Generative or reweighting methods often struggle to reproduce such cut-off behaviour precisely. The RAN handles this boundary reasonably well. The unfolded z_g distribution from the RAN matches the truth both in the low z_g region and near the cut-off, correctly recovering the falling slope as $z_g \rightarrow 0.5^-$. OMNIFOLD’s result for z_g is also quite close to truth with slight undercorrection near the cutoff.

Finally, the constituent multiplicity M , number of particles in the jet, is a particularly interesting case. It is a discrete distribution with a long tail. The RAN succeeds in reducing the discrepancy at both the low- M and high- M extremes, yielding an unfolded distribution that tracks the truth across the full range. OMNIFOLD

Method	m	M	w	τ_{21}	z_g	$\ln \rho$
RAN	3.35	1.19	0.59	0.15	0.80	0.20
OmniFold	8.02	4.43	0.81	0.78	6.32	4.15
IBU	7.40	16.07	5.28	1.52	19.18	10.31
Generation	11.71	25.39	26.95	5.50	12.85	61.54
Data	33.42	71.01	181.65	60.27	97.40	378.43

Table VI.3: Wasserstein-1 distances (W_1) between true distributions for six jet substructure observables and unfolded distributions produced by RAN, OmniFold, and IBU (performed on one dimensional projections onto each observable). Values are scaled by 10^{-2} for m , M , and τ_{21} ; by 10^{-3} for w , z_g , and $\ln \rho$. RAN achieves the smallest distance for all observables, with particularly dramatic improvements for challenging observables. Baseline distances for uncorrected Generation and Data are also shown for reference.

improves the multiplicity distribution in the bulk region but exhibits larger fluctuations in the very high-multiplicity tail partly due to limited statistics there and the difficulty for the classifier to learn in sparse regions. In summary, qualitatively, the RAN’s unfolded spectra are virtually indistinguishable from the truth for all six observables, within statistical uncertainties,

The ratio plots in Figure VI.5 highlight RAN’s performance. The RAN/Truth ratio is close to 1.0 across the phase space, while the. To make these comparisons quantitative, the Wasserstein–1 distances (W_1) and VLC divergences between each unfolded distribution and the truth distribution for all methods are computed. Tables VI.3 and VI.4 summarise these metrics for RAN, OMNIFOLD, and IBU on each jet observable. IBU is performed on one dimensional projections onto each jet

Method	m	M	w	τ_{21}	z_g	$\ln \rho$
RAN	2.31	5.83	1.29	0.92	0.80	1.27
OmniFold	5.06	6.32	1.33	7.91	6.84	2.34
IBU	5.11	16.07	7.06	17.51	11.90	3.53
Generation	34.92	12.89	21.86	20.33	29.17	18.86
Data	52.24	136.57	22.55	25.82	14.90	18.26

Table VI.4: Vincze–Le Cam divergence, Δ_{VLC} between true distributions for six jet substructure observables and unfolded distributions produced by RAN, OMNIFOLD and IBU (performed one observable at a time). Values are scaled by 10^{-3} for m , M , and τ_{21} ; by 10^{-4} for w , z_g , and $\ln \rho$. RAN achieves the smallest divergence for all observables. Baseline divergences for uncorrected Generation and Data are also shown for reference.

observable. Lower values indicate better agreement with Truth. The tables also include, for reference, the baseline distances for the Generation vs. Truth and Data vs. Truth distributions with no unfolding. RAN achieves the smallest divergence from Truth under both metrics.

One also observes that IBU, a binned, per-observable method, generally underperforms the unbinned methods here. This is likely due to binning and statistical issues: the need to choose finite bin widths leads to information loss and large uncertainties in sparse regions. For M , for example, the tail probabilities in high bins were not well estimated, leading to inflated divergence. In contrast, a RAN and OMNIFOLD operate on unbinned data and can leverage the full event information, yielding superior precision.

It is also noteworthy that the RAN’s advantage is achieved while producing a single set of event level weights that simultaneously correct all six distributions. In other words, RAN, like OMNIFOLD, inherently unfolds the joint multi-dimensional distribution of these observables without splitting the problem into one dimensional projections. The fact that a single model can balance all observables and still attain good per observable accuracy is a strong testament to the effectiveness of the adversarial reweighting approach.

VI.E.2.i.a Stability and efficiency.

Stability and efficiency, in addition to accuracy, are important considerations for unfolding methods. The stability of a RAN’s training can be probed by performing multiple independent training runs and by bootstrapping subsets of the jet dataset. The variations in the resulting unfolded distributions, and in the W_1 /VLC metrics are found to be small, on the order of a few percent, indicating that the RAN solution is robust to statistical fluctuations and the stochastic nature of training.

Likewise, OMNIFOLD’s iterative procedure, when run on the same data with different initialisations, converged to comparable results. Though minor differences in the weights per iteration can occur, the final distributions were consistent within uncertainties, lending confidence to the RAN’s unfolded output. IBU, being an analytic iterative method, showed negligible run-to-run variation for fixed binning.

OMNIFOLD requires a subjective choice of a stopping criterion or number of iterations. Too few iterations can under correct, too many can overfit the statistical noise. IBU similarly depends on a subjective choice of the number of iterations. In the OMNIFOLD implementation employed in the studies in this section, five iterations

are used, which is the standard approach advocated in [309], which was sufficient for good performance; using more iterations does not notably improve the agreement and can introduce instability. A RAN avoids this ambiguity entirely, since it is trained in one shot to convergence.⁴⁹ In practice, one finds that RANs train to convergence steadily without signs of overtraining, monitored via the validation loss, and reach a stable solution within a few dozen epochs.

In terms of computational cost, a RAN’s single pass adversarial training using the Wasserstein GAN-like setup is comparable to the cost of a one OMNIFOLD iteration. But since OMNIFOLD requires two classifier trainings per iteration, and OmniFold was trained for five iterations in these studies, the overall time needed to train a RAN was roughly a fifth of that for OMNIFOLD on these problems. IBU is extremely fast for one dimensional histograms, but extending IBU to many dimensions would be combinatorially expensive in terms of computational cost, and practically impossible due to a lack of enough data to fill the bins of high dimensional histograms, whereas RANs and OMNIFOLD scale gracefully to high dimensional data. Thus, a RAN offers an attractive trade-off: it achieves equal or better accuracy than OMNIFOLD on these benchmark tasks, while being simpler to use, since it has no iterative loop to manage, and potentially more efficient.

⁴⁹To be clear, a RAN, like any other machine learning based algorithm involves a host of subjective choices, including but not limited to the specific architecture chosen to create the neural networks, to the large array of hyperparameters involved in any neural network training. The point being made here is that iterative ML based methods also require all of these subjective choices, and require the subjective choice of the number of iterations in addition to the other choices.

VI.E.2.ii Generality of results.

Combined with the findings from the Gaussian experiments, these jet experiments suggest that a RAN’s advantages are not specific to a particular distribution but rather reflect general features of the algorithm. A RAN’s ability to handle non-overlapping detector effects, demonstrated in the Gaussian toy study, implies it could be helpful in experimental scenarios with poor detector resolutions or acceptance gaps, where traditional unfolding might struggle. The jet study confirms that RANs scale to complex, realistic tasks, providing high fidelity unfolding for multiple correlated observables in a single training. One expects that these properties would carry over to other unfolding applications. For example, measurements of other final states or higher dimensional distributions,⁵⁰ should similarly benefit from a RAN’s unbinned, non-iterative algorithm. Moreover, RAN’s inherently multivariate nature means that, unlike IBU, it can preserve and unfold correlations between observables, an important aspect for modern high dimensional analyses in particle physics.

Certainly, caution is warranted when extrapolating beyond the tested cases. Real experimental data may involve additional complexities in ways not captured by generator differences alone, or require careful treatment of systematic uncertainties in the unfolding procedure. Although the RAN performs well on the closure tests described in this section, it remains the case that RANs have only been validated so far using one simulator’s output as ‘pseudo-data’ and another’s as MC, rather than on real experimental measurements.

⁵⁰Such as multivariate phase space distributions.

In summary, the results presented in this section demonstrate that RANs achieve state of the art unfolding performance on both simple and complex tasks. They matches or exceed the accuracy of established methods like OMNIFOLD and IBU, while offering improved stability in difficult scenarios and simplified unfolding workflows. These characteristics make RANs a promising tool for future precision measurements and exploratory studies in which unfolding high dimensional data without bias is critical.

Chapter VII

Unbinned inference on correlated data.

VII.A	Introduction.	395
VII.B	Statistical independence in HEP.	396
VII.B.1	Poisson point processes for event counts.	397
VII.B.2	Independent and identically distributed (i.i.d.) events.	398
VII.B.3	Likelihood factorisation and conditional independence.	399
VII.C	Violation of statistical independence.	401
VII.C.1	Detector effects.	403
VII.C.2	Pileup.	404
VII.C.3	Unfolding and data processing.	406
VII.C.4	Global constraints.	408
VII.D	Consequences for inference.	412

<i>Chapter VII. Unbinned inference on correlated data.</i>	394
VII.E Formalism.	416
VII.F Uncertainty quantification.	426
VII.F.1 Analytic approaches: asymptotic theory and its limitations.	427
VII.F.2 Godambe information: the sandwich estimator)	429
VII.F.3 Wilks' theorem violations and Bartlett corrections	432
VII.F.4 Numerical resampling approaches: bootstrapping and toy Monte Carlo.	435
VII.G Case studies.	449
VII.G.1 Setup.	450
VII.G.2 Fully binned baseline.	450
VII.G.3 Correlation diagnostics after unbinned unfolding.	460
VII.G.4 Unbinned unfolding: binned and unbinned inference.	466
VII.G.5 Extension to higher dimensions.	475
VII.H Conclusions and outlook.	482

VII.A Introduction.

The preceding chapters developed a hierarchy of unfolding techniques, from classical binned approaches in Section II.D through NEURAL POSTERIOR UNFOLDING in Chapter IV and MOMENT UNFOLDING in Chapter V to RAN, a fully unbinned unfolding algorithm that operates directly on event level information in Chapter VI. Once the distributions are unfolded, parameter inference is performed on the unfolded data to summarise the the event information into a small number of observables whose differential cross sections are compared with theory by computing best fit parameters and confidence intervals. While performing inference, especially on unbinned data, an implicit assumption of *statistical independence* of individual events is often made. This independence guarantees that the joint likelihood factorises and allows the log-likelihood in Equation (VII.4) to be written as a sum over events.

In practice, however, the assumption of independence is violated whenever events are processed through an unfolding procedure. Deconvolution couples phase space regions and induces highly non-trivial, often long range, correlations between formerly independent data points. Ignoring those correlations can bias parameter estimates and/or lead to misestimation of confidence intervals.

VII.B Statistical independence in HEP.

VII.B.1 Poisson point processes for event counts.	397
VII.B.2 Independent and identically distributed (i.i.d.) events. . . .	398
VII.B.3 Likelihood factorisation and conditional independence. . . .	399

Definition 44 (Event independence in HEP). In high energy physics cross section measurements, standard statistical treatments for unbinned inference rely on a set of core assumptions collectively referred to as *event independence*. These assumptions posit that interaction events are generated and observed as independent trials of a stochastic process, which greatly simplifies the construction of likelihood functions and inference procedures.

This section will outline the key independence assumptions commonly made in unbinned analyses, and subsequently examine their applicability or limitations in experimental contexts.

VII.B.1 Poisson point processes for event counts.

A fundamental premise is that the total number of events N observed in an experiment follows a Poisson distribution. If $\mu(\theta)$ denotes the expected event yield for parameters θ ,¹ the probability of observing N events is

$$\mathbb{P}(N \mid \theta) = \frac{e^{-\mu(\theta)} \mu(\theta)^N}{N!}. \quad (\text{VII.1})$$

This reflects the physical picture of interaction events occurring as Poisson point processes in time, with each event occurrence independent of the last [494, 679]. The Poisson assumption is built into the extended likelihood formalism, ensuring that normalisation (total event count) is appropriately handled in parameter inference for cross sections [412, 430]. It provides a way to incorporate both the shape of distributions and the overall event yield into the likelihood. In practice, this means that

¹E.g., proportional to the integrated luminosity times the cross section.

in repeated identical experiments one would expect the observed N to fluctuate about $\mu(\theta)$ according to Poisson statistics, and it justifies including a factor $\exp(-\mu)\mu^N/N!$ in the likelihood function of a single dataset.

VII.B.2 Independent and identically distributed (i.i.d.) events.

Definition 45 (i.i.d.). *Independence and identical distribution* (i.i.d) is the assumption that each event can be treated as an independent draw from the same underlying probability density function $p(x \mid \theta)$, where x denotes the measured observables (kinematic variables, detector signals, *et cetera*) for that event. In other words, conditional on the physics model parameters θ , all events are statistically independent and governed by an identical distribution [263, 295, 299, 679].

This implies that one event's occurrence or properties do not influence any other event or equivalently, there are no intrinsic correlations or memory between events. The identical distribution assumption further requires that the experimental conditions remain stable so that each interaction is sampled from the same probability density function $p(x \mid \theta)$. This condition is typically enforced by construction, by dividing data-taking into consistent periods with fixed detector configuration and calibrations.

Combined with the Poisson law for N , the i.i.d. assumption forms the basis of the usual HEP data generation model, *viz.* that the data are viewed as a Poisson sample of independent events from $p(x \mid \theta)$. This assumption underlies virtually all unbinned

analysis techniques in HEP, from simple maximum likelihood fits to modern machine learning based inference methods [61, 74, 299, 316, 473].

VII.B.3 Likelihood factorisation and conditional independence.

Because events are modelled as independent draws, the joint probability density for N events factorises into a product of single event densities. For a given dataset $\{x_1, x_2, \dots, x_N\}$, one can write

$$p(x_1, \dots, x_N | \theta) = \prod_{i=1}^N p(x_i | \theta), \quad (\text{VII.2})$$

which in turn leads to a factorised likelihood function.

In the common case where N itself is treated as Poisson distributed, the full extended likelihood is

$$\mathcal{L}(\theta) = e^{-\mu(\theta)} \frac{\mu(\theta)^N}{N!} \prod_{i=1}^N p(x_i | \theta), \quad (\text{VII.3})$$

assuming all events are independent of each other [136, 140, 386, 412, 596, 614, 670]. The independence assumption allows the log-likelihood to be written as a sum over events,

$$\ln \mathcal{L}(\theta) = -\mu(\theta) + \sum_{i=1}^N \ln p(x_i | \theta), \quad (\text{VII.4})$$

up to the constant term $\ln(N!)$.²

Crucially, the factorisation in Equation (VII.4) holds only under the assumption that events are statistically independent. This property of likelihood factorisability

²In cases where N is fixed and not of interest, the $-\mu + \frac{\mu^N}{N!}$ portion is often dropped, yielding the simplified unbinned likelihood $\prod_i p(x_i | \theta)$.

enormously simplifies inference: it enables the use of well known asymptotic statistical results like Wilks' theorem on the sum of per event log-likelihood contributions, and it permits analytical derivations of estimators and uncertainties [570]. Indeed, widely used formulae for the variance of estimators and for test statistics³ presume an underlying model where $\ln \mathcal{L}$ sums over independent events.

Note that independence is typically understood as *conditional* independence given the model parameters θ and any nuisance parameters. For example, if events come from multiple sources, one usually models each source as an independent Poisson process; all events are then independent samples overall, with an additional mixture component in $p(x | \theta)$. Likewise, any uncertainty in calibration constants or other global parameters is introduced via nuisance parameters rather than treating events as correlated. The correlations induced by shared nuisance parameters are handled at the parameter level rather than by coupling events in the probability model. Under fixed values of all such parameters, events remain independent.

This conditional independence perspective justifies how multiple contributions are combined in likelihoods and how one accounts for systematic effects without introducing inter-event correlations.

³Traditionally used examples can be found in [263, 441]. More recent proposals can be found in [53, 166].

VII.C Violation of statistical independence.

VII.C.1	Detector effects.	403
VII.C.2	Pileup.	404
VII.C.3	Unfolding and data processing.	406
VII.C.4	Global constraints.	408

These statistical independence assumptions have served as the backbone of cross section measurements and other inference tasks in HEP for decades. They are reasonably well motivated by the physics of particle interactions. In particle physics experiments, each fundamental interaction (for example, a $p - p$ collision) is localised in spacetime and causally independent from other interaction. Detectors are typically designed and operated to measure each interaction event in isolation.

By treating events as independent, analyses can take full advantage of powerful likelihood-based inference techniques and modern machine learning methods that operate on event-by-event data [65, 79, 159]. Unbinned machine learning approaches to parameter estimation, for instance, explicitly rely on the factorised likelihood in Equation (VII.4) to avoid information loss from binning [112, 159, 210]. The independence assumption is also what allows one to rigorously define an asymptotic χ^2 or log-likelihood ratio statistic that sums over events and, in the large- N limit, follows known distributions for hypothesis testing [112, 166, 240]. In short, the usual paradigm treats each recorded event as an independent piece of evidence about the underlying physics parameters.

However, these assumptions are *idealisations* that can be violated in a variety of scenarios, thereby challenging the simplistic i.i.d. picture. It is essential to recognise and mitigate these violations because applying independent-event methods in their on correlated data can lead to biased estimates or misestimated uncertainties.

VII.C.1 Detector effects.

While the i.i.d. assumption requires a stable underlying distribution $p(x | \theta)$ for all events, in reality the detector and experimental conditions can evolve or fluctuate, effectively making events drawn from slightly different distributions. For example, changes in detector response due to calibration drifts, aging of instrumentation, or varying operational settings can cause the probability distribution of observables to shift over time. If unaccounted for, this means an event recorded early in the run is not identically distributed to an event recorded later. In addition, the finite resolution and inefficiencies of the detector are typically incorporated into $p(x | \theta)$ via a detector response model; the residuals due to imperfections in the modelling can introduce correlations.

Another detector induced effect is the presence of noise bursts or environmental backgrounds like cosmic rays or electronics noise that can affect multiple events in a correlated way. For instance, a temporary malfunction in a subdetector could bias a whole set of consecutive events in a similar manner. Experimentalists usually handle these issues by segmenting data gathering periods and calibrating each segment separately, or by introducing nuisance parameters to account for time dependent efficiency changes.

When treated properly, one can restore the approximation of identical distributions; but if such variations are neglected, the assumption of identically distributed independent events breaks down. In essence, the detector can introduce a slight correlation structure or at least a heterogeneity among events, violating the ideal independence

assumption, but experiments are typically cognisant of these effects and process them appropriately to restore near i.i.d conditions.

These effects often manifest as systematic uncertainties in the analysis. For example, although a calibration uncertainty correlates the predicted rates of all events, it can be handled by propagating those uncertainties appropriately rather than by treating events as correlated. Still, it is a reminder that truly independent, identical event samples exist only in the limit of a perfectly stable detector and complete modelling of its response.

VII.C.2 Pileup.

Definition 46 (Pileup). At particle physics experiments multiple events can occur in the same same detector read-out window and be recorded together. This phenomenon is known as *pileup*. Pileup entails that what is recorded as a single ‘event’ may in fact be a composite of several independent interactions superimposed in the detector readout.

High luminosity running conditions can lead to dozens of overlapping interactions per event. In an ideal scenario, these would be uncorrelated collisions, and indeed they are physically independent interactions; however, the merging of their signals confounds the independence at the level of reconstructed events.⁴

⁴In nucleus–nucleus (heavy ion) experiments the main contaminant is not separate overlapping interactions but the extremely high multiplicity *underlying event* produced in the *same* ion–ion interaction. Typical central Pb–Pb collisions at the LHC yield $\mathcal{O}(10^3)$ charged particles within $|\eta| < 1$, orders of magnitude above $p - p$ pileup levels [549, 622].

Sophisticated event reconstruction algorithms attempt to separate pileup interactions. For example, at collider experiments, one can identify multiple distinct collision vertices and attribute detector hits to different vertices. Despite these efforts, residual entanglement remains: additional tracks or energy deposits from pileup interactions can slip into the reconstruction of the hard scatter event of interest, affecting measured quantities like jets, missing energy, or lepton isolation. As a result, the observables x_i for a given recorded event can include contributions from other simultaneous events, blurring the notion that each x_i is drawn from the single event distribution $p(x | \theta)$.

In extreme cases, one interaction's presence can influence whether another interaction is recorded at all, due to trigger bandwidth or detector occupancy limits, introducing a form of inter-event dependence. Pileup therefore violates the assumption that each event is an independent trial of the same distribution. Instead one is truly measuring conglomerate events with extra particles, and the probability of certain features (e.g. multiple soft jets, high detector occupancy) is higher than what a single interaction $p(x | \theta)$ would predict.

To mitigate pileup, experiments apply corrections or weights to event observables such as subtracting average energy from soft interactions [350, 426, 481], or they incorporate pileup explicitly into the modelling by extending $p(x | \theta)$ to include pileup contributions. When pileup is included in the Monte Carlo model and properly tuned, the effective independence can be partially restored by enlarging the event description (i.e. treat the hard collision and pileup as parts of one compound event).

Nonetheless, any imperfections in pileup mitigation mean that events in a high pileup dataset are not as independent as assumed. This is especially pertinent for

precision measurements at future high luminosity colliders, where average pileup multiplicities will be extremely high [102, 244, 290, 339, 402]. Overlapping events can also occur in other experimental contexts. Neutrino experiments might have cosmic ray overlays, astrophysical observations might have coincident signals, *et cetera*—all such scenarios require care if one is to maintain the statistical interpretation of independent events.

VII.C.3 Unfolding and data processing.

A common workflow in HEP is to first *unfold* the measured data, correcting for detector effects to infer particle level distributions, and then perform inference or fits on these unfolded results. A large class of unfolding algorithms, including modern unbinned ones, assign weights or probabilities to events in order to transform detector level data to a particle level prediction. However, the unfolding procedure itself induces statistical correlations among the resulting events.

In an unbinned unfolding, the output events, often weighted MC, are no longer independent: they are coupled by the global constraints of the unfolding such as conserving overall normalisation and matching distributions. For example, if one unfolded event’s weight fluctuates upward, typically some other event’s weight must fluctuate downward to compensate and preserve the total yield or certain distribution constraints. Thus, the set of unfolded events has an empirical global covariance structure because knowledge of one event’s outcome gives information about others.

If one were to naïvely apply an unbinned likelihood analysis assuming those unfolded events are independent, the likelihood factorisation would be incorrect, and one might

misestimate uncertainties or bias the fit. This situation has been explicitly observed. The work presented in this chapter [44] shows that treating unfolded events as independent can misestimate the confidence intervals, with asymptotic error formulae significantly misestimating the true uncertainty.

This is a prominent example where the conditional independence assumption of events (conditional on θ) does not hold because the unfolded events carry correlations from the statistical inversion procedure. In practice, experimental analyses that use unfolded data either avoid unbinned fits on the raw unfolded events, instead aggregating unfolded results into bins with an associated covariance matrix that captures the induced correlations [54] or simply ignore the covariance structure of unbinned unfolded data [32, 210]. For instance, if an unfolded spectrum is used to extract a physics parameter, data is provided in set bins alongside a covariance matrix Σ such that a binned χ^2 fit can be performed,

$$\chi_{\text{full}}^2(\theta) = (D - \mathbb{P}(\theta))^T \Sigma^{-1} (D - \mathbb{P}(\theta)), \quad (\text{VII.5})$$

where D is a vector of normalised bin counts of the unfolded data and $\mathbb{P}(\theta)$ the corresponding model prediction [33, 104, 181, 276].

This χ^2 , or the equivalent binned likelihood, correctly accounts for correlations via Σ . By contrast, a fully unbinned likelihood $\prod_i p(x_i|\theta)$ on the unfolded events x_i would ignore the Σ information and treat events as independent, which is formally unjustified. The advent of machine learning has enabled high dimensional unbinned unfolding, increasing the urgency to address these correlations [32, 48, 268].

It would therefore be prudent to avoid analytic uncertainty formulae in such cases and rely on bootstraps to evaluate uncertainties. Future research efforts could

attempt to incorporate event level correlations into the unbinned likelihood formalism, for example, by parametrising weight fluctuations or using augmented likelihoods that include correlation terms. Until such methods mature, the safest approach for unfolded data is to introduce correlations either at the binned level, which effectively relinquishes some of the benefits of unbinned methods in exchange for correct coverage of uncertainties, or by bootstrapping, which significantly increases the computational cost of the analysis.

VII.C.4 Global constraints.

Even when individual interactions are truly independent, certain *global* aspects of an analysis can couple what were otherwise independent events. A prime example is the presence of common normalisation factors or theoretical parameters that influence all events. For instance, the integrated luminosity of a dataset that is used to convert event counts to cross sections is usually known with some uncertainty. If treated as a nuisance parameter, that single parameter induces correlations between the rates of events in different signal regions. Effectively, an upward fluctuation in luminosity would scale up the expected counts for all events. In a frequentist formulation, events remain independent, conditional on a fixed luminosity, but when one considers the uncertainty on luminosity, the joint distribution of all events together has an additional covariance because they all scale together.

Similarly, theoretical model uncertainties, like parton distribution functions in a hadron collision, can correlate the kinematic distributions of all events. For example,

if a parton distribution function parameter shifts, it will concurrently affect the probability distributions of many events, creating a correlation in their fluctuations.

Usually these effects are handled by introducing correlated systematic shifts in the expectation rather than by accommodating the event correlations in the likelihood; nonetheless, such effects highlight that true data generating processes have an interdependence mediated through shared parameters. Reformulated in Bayesian terms, if one marginalises over an uncertain global parameter, the remaining distribution of events is no longer a simple product of independent probability density functions.

Beyond nuisance parameters, there are also physics scenarios that produce genuine correlations between events: for instance, certain new physics might produce clustered events. For example, the decay of a heavy state might yield two separate collision-like signals in the detector, or cosmic ray air shower might produce multiple spatially separated detector hits treated as separate events. These are exotic cases but illustrate the broader point that nature can produce data that do not align with the one-event-at-a-time paradigm.

The high level takeaway is that, in general, whenever a theoretical or methodological constraint ties the outcomes of different events, the independence assumption is violated. In the context of cross section measurements, one common manifestation is in combined or multidimensional fits. If one fits multiple distributions simultaneously with shared parameters, the events populating those distributions are effectively analysed in one joint likelihood. The independence assumption still might hold event-by-event, but events in different categories become statistically coupled through the shared parameter constraints. The customary solution is again to incorporate

those constraints at the likelihood level, via nuisance parameters, profile likelihood techniques, or covariance matrices across distributions. This ensures that while the calculation may still treat events as independent given the parameters, the final results correctly reflect the induced uncertainties.

In summary, the assumption that events are generated independently and identically is a powerful simplification that underlies many HEP analysis techniques, from classical maximum likelihood fits to cutting edge machine learning inference frameworks. These assumptions are approximately valid for raw collision data under well controlled conditions and have enabled analysts to exploit the full statistical power of unbinned data. Nonetheless, real world complexities such as detector imperfections, pileup, processing, and global effects produce correlations that demand a careful treatment beyond the idealised i.i.d. model.

In contemporary binned analyses, it is standard to incorporate such effects via covariance matrices and nuisance parameters rather than to abandon the independence assumption entirely. The challenge for the field moving forward is to extend extant statistical formalism and tools to unbinned data, so that one can continue to perform unbiased, efficient inference even as one performs complex global fits of unbinned unfolded data, where events are no longer perfectly independent.

Ultimately, recognising the limits of the independence assumptions and quantifying any resulting biases or miscoverage is essential for robust cross section measurements. By confronting these limitations, either by correcting them or explicitly incorporating them into the statistical model, HEP analyses can ensure that improvements in

methodology such as unbinned approaches remain scientifically reliable even in the face of complex data dependencies.

VII.D Consequences for inference.

For independent data points, the joint likelihood for N events factorises into a product of single event likelihoods, making the log-likelihood a sum over events, as shown in Equation (VII.4).

Equation (VII.4) underpins many inference techniques and allows the use of well known asymptotic statistical results like Wilks' theorem to efficiently compute confidence intervals and p -values in high energy physics analyses [570]. In the context of HEP experiments, this means that one can often streamline parameter extraction by summing likelihood contributions from each interaction event independently.

However, correlated data arise inevitably in the process of reconstructing physical (particle level or parton level) quantities from raw detector observations. In modern analyses of differential cross sections, a number of effects introduce non-negligible correlations among data points or among reconstructed events. The calibration performed to address the imperfections discussed in Section VII.C couples the statistical content of neighbouring bins. As a result, the final differential cross section values in adjacent kinematic bins are often statistically anti-correlated. Likewise, global efficiency or acceptance corrections (e.g. accounting for detector coverage) can introduce positive correlations across all bins by scaling yields coherently.

The residual effects of pileup tend to induce positive correlations between events. For instance, an increased number of simultaneous interactions tends to add extra low p_T particles across the event, causing all jet energies in that event to shift slightly. When such effects are corrected statistically by subtracting an average pileup contribution, the correction uncertainty is correlated across all events in a given

run. Moreover, pileup mitigation often relies on average profiles, meaning that any mismodeling of pileup leads to systematic shifts in distributions that are common to many events leading to positive correlations between event weights. Thus, while raw detector level data is physically independent, its inferred particle level properties carry a common uncertainty component from pileup treatments.

In binned unfolding, such as iterative Bayesian or regularised matrix inversion methods, adjacent bins often develop substantial anti-correlations due to the regularisation constraints, and all bins can share common normalisation systematics. In unbinned unfolding approaches where one obtains a reweighted set of events at particle level rather than fixed bins, these reweighted events carry event-by-event weight factors that are determined collectively to reproduce the measured detector data. Consequently, the unfolded events are statistically correlated with each other. If one unfolded event's weight fluctuates upward to fit a fluctuation in the data, another event's weight must adjust downward to compensate. The assumption of independent events is therefore violated for such unfolded particle level samples.

Many experimental systematic effects (energy scale and efficiency corrections, luminosity uncertainty, *et cetera*) induce fully or partially correlated uncertainties across multiple data points. For example, the uncertainty in the integrated luminosity, used to normalise cross sections is a single multiplicative factor affecting the normalisation of the cross section in every bin simultaneously. This manifests as a 100% bin-to-bin correlation for that component of uncertainty since all bins shift up or down together by the same relative factor. Similarly, a jet energy scale uncertainty might cause a shape distortion that correlates a rise in high- p_T bin yields with a fall in low- p_T

bin yields. These systematic effects mean that the total uncertainty covariance on the final results has significant off-diagonal terms. Even if the statistical fluctuations of events were independent, the presence of correlated systematics requires careful treatment in any inferential test or fit. The impact of various sources of uncertainty and their effects are summarised in Table VII.1.

Table VII.1: Sources of correlated uncertainties in cross section measurements and their impact on unfolded distributions. These correlations affect parameter extraction, uncertainty propagation, and confidence interval coverage in physics analyses.

Source	Impact on unfolded distributions
Detector resolution	Finite resolution causes true events to migrate between measurement bins, inducing statistical correlations in the unfolded spectrum. This migration induced covariance must be properly modelled to avoid biasing shape sensitive fits and to correctly propagate uncertainties through the unfolding procedure.
Pileup	Additional particles from simultaneous interactions contaminate all events within a luminosity block. Residual uncertainties after pileup subtraction create correlated systematic effects across events and bins, affecting both normalisation and shape measurements. The correlation strength typically scales with instantaneous luminosity and is therefore expected to be increasingly challenging to address in future high luminosity experiments.

Continued on next page

Table VII.1 – continued from previous page

Source	Impact on unfolded distributions
Unfolding	Smoothness constraints or penalty terms couple adjacent bin estimates or event weights, transforming independent Poisson counts into correlated random variables. The resulting off-diagonal covariance terms can be non-trivial, particularly for aggressive regularisation schemes. This coupling must be propagated to downstream analyses.
Systematics	Experimental systematics (luminosity calibration, trigger efficiency, energy scale) and theoretical uncertainties (PDF sets, scale variations) affect multiple bins coherently. These produce positive correlations that, if neglected, lead to underestimated uncertainties on integrated quantities and incorrect confidence intervals for extracted parameters.

VII.E Formalism.

Let $\{\nu_j\}_{j=1}^M$ be the unfolded, normalised particle level yield in bin j and $\mathbb{P}_j(\theta)$ be the theory prediction or fit template for that bin given parameters θ . The standard measure of agreement is the correlated χ^2 statistic described in Equation (VII.5), where Σ is the $M \times M$ covariance matrix of the measurement. The covariance matrix Σ encodes the variance of each bin along the diagonal, $\Sigma_{jj} = \text{Var}(\nu_j)$ and the covariance in the off-diagonal entries $\Sigma_{ij} = \text{Cov}(\nu_i, \nu_j)$.

In the limit of many events and Gaussian uncertainties, $\chi_{\text{full}}^2(\theta)$ is expected to follow a χ^2 distribution with $M - n_\theta$ degrees of freedom, where n_θ is the number of fitted parameters, for correctly estimated Σ . This forms the basis of parameter fits and theory tests using binned data. One can find the best fit parameters by minimising χ_{full}^2 , and one can evaluate goodness of fit or confidence intervals by referring to the χ^2 distribution [50, 59, 591, 593, 623]. Analyses that fit parton distribution functions or theoretical model parameters to data from multiple bins, or even multiple experiments, must propagate the full covariance; otherwise the fit would erroneously constrain parameters using spurious information.

To observe the impact of correlated versus uncorrelated treatments, consider a simple example. Let a measurement report two bins with identical central values and uncertainties, and let there be a 100% positive correlation between their uncertainties.⁵ If a theory predicts both bins' values higher than observed, say by $+1\sigma$ of the reported uncertainty in each bin, a proper correlated χ^2 calculation will find *no significant tension* because the two data points moved together could be explained by a single

⁵For example, a fully correlated normalisation error.

upward fluctuation or a slight systematic shift. In fact, the χ^2_{full} in this case might be of order 1, well within expectations.

In contrast, if one had naïvely ignored the correlation and treated the points as independent, the same $+1\sigma$ deviations in each bin would add in quadrature, giving $\chi^2_{\text{diag}} = (+1\sigma)^2 + (+1\sigma)^2 = 2\sigma^2$ for two degrees of freedom. The corresponding p -value would be significantly smaller, and one might incorrectly suspect a poor fit to the theory. In this way, neglecting positive correlations can inflate test statistics and lead to under-coverage of the confidence intervals [16]. Conversely neglecting anti-correlations in the data can deflate test statistics and lead to over-coverage of the confidence intervals. In realistic HEP data of course, there is no way *a priori* of predicting the sign of the misestimation without explicitly evaluating all the elements of the covariance.

This simple example underscores how a common systematic uncertainty should not be counted as independent evidence of mismatch in each bin. Only by using the full covariance does one obtain the correct statistical interpretation. More generally, the effect of data correlations on parameter extraction can be understood in terms of effective sample size and information content. If one has N independent data points, uncertainties on fitted parameters typically scale as $\sim 1/\sqrt{N}$, all else being equal. But if the N points have positive correlations, the effective amount of independent information is smaller than N . For instance, in the idealised case of N observations such that each pair of points has a uniform correlation coefficient ρ . The $N \times N$

covariance matrix is

$$\Sigma = \begin{bmatrix} 1 & \rho & \cdots & \rho \\ \rho & 1 & \cdots & \rho \\ \vdots & & \ddots & \vdots \\ \rho & \rho & \cdots & 1 \end{bmatrix} \quad (\text{VII.6})$$

The positive definiteness of covariance matrices requires that $\rho \in (-1/(N-1), 1]$. For such a system, the error of the mean, or any common strength parameter estimator is

$$\text{Err}(\bar{X}) = \sigma \sqrt{\frac{1}{N} + \left(1 - \frac{1}{N}\right)\rho} \quad \rho \in \left(-\frac{1}{N-1}, 1\right], \quad (\text{VII.7})$$

$$\text{Err}(\bar{X}) \Big|_{\rho=-\frac{1}{N-1}} = 0, \quad \text{Err}(\bar{X}) \Big|_{\rho=0} = \frac{\sigma}{\sqrt{N}}, \quad \text{Err}(\bar{X}) \Big|_{\rho=1} = \sigma. \quad (\text{VII.8})$$

where σ is the standard deviation of each individual point. The factor $1 + (N-1)\rho$ quantifies the change in variance due to correlation.

Theorem 13 (Equicorrelated sampling error and effective sample size). For every $N \in \mathbb{N}_{\geq 2}$, and for every set of real valued random variables $\{X_i\}_{i=1}^N$, if:

$$\begin{aligned} \exists(\mu, \sigma, \rho) \in \mathbb{R} \times (0, \infty) \times \left[-\frac{1}{N-1}, 1\right] \quad \forall i \in \mathbb{N}_{\leq N} \quad \mathbb{E}(X_i) = \mu \\ \wedge \forall i, j \in \mathbb{N}_{\leq N}^2 \quad \Sigma_{ij} = \sigma^2[(1-\rho)\delta_{ij} + \rho] \end{aligned} \quad (\text{VII.9})$$

i.e.

$$\Sigma_{ij} = \begin{cases} \sigma^2, & i = j, \\ \rho \sigma^2, & i \neq j. \end{cases} \quad (\text{VII.10})$$

Then:

1. $\Sigma \succeq 0$ iff $\Sigma \in \left[-\frac{1}{N-1}, 1\right]$ and $R(\rho) \succ 0$ iff $\rho \in \left(-\frac{1}{N-1}, 1\right)$.

2. For the sample mean $\bar{X} \equiv \frac{1}{N} \sum_{i=1}^N X_i$,

$$\text{Err}(\bar{X}) \equiv \sqrt{\text{Var}(\bar{X})} = \sigma \sqrt{\frac{1}{N} + \left(1 - \frac{1}{N}\right)\rho}. \quad (\text{VII.11})$$

Proof of Theorem 13. (1) The matrix Σ has eigenvalue $\sigma^2(1 + (N - 1)\rho)$ with eigenvector $\begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix}^T$ and eigenvalue $\sigma^2(1 - \rho)$ with multiplicity $N - 1$. Hence $\Sigma \succeq 0$ iff $\sigma^2(1 - \rho) \geq 0$ and $\sigma^2(1 + (N - 1)\rho) \geq 0$, i.e. $\rho \in [-1/(N - 1), 1]$, and $R(\rho) \succ 0$ iff both inequalities are strict.

(2) Using $\bar{X} = N^{-1} \sum_i X_i$,

$$\text{Var}(\bar{X}) = \frac{1}{N^2} \sum_{i,j=1}^N \Sigma_{ij} = \frac{N\sigma^2}{N^2} (1 + (N - 1)\rho) = \sigma^2 \left(\frac{1}{N} + \left(1 - \frac{1}{N}\right)\rho \right), \quad (\text{VII.12})$$

which yields the stated $\text{Err}(\bar{X})$. $\square$

Definition 47 (Effective sample size). The *effective size* of a sample is the effective number of independent points, $N_{\text{eff}} = N/[1 + (N - 1)\rho]$ [25, 260, 301, 583] in that sample accounting for correlations.

For $\rho > 0$, $N_{\text{eff}} < N$; in the extreme case of $\rho \rightarrow 1$ where all points completely correlated, $N_{\text{eff}} \rightarrow 1$ no matter how large N is. For $\rho = 0$, as one would expect, $N_{\text{eff}} = N$. For $\rho < 0$, $N_{\text{eff}} > N$, with the extreme case of $\lim_{\rho \rightarrow -\frac{1}{N-1}} N_{\text{eff}} = \infty$

This is precisely what would happen, for example, if a common normalisation uncertainty dominated. No matter how many bins are measured, if they all share the same normalisation shift, the overall normalisation is just one degree of freedom of uncertainty rather than N independent ones. Ignoring correlations amounts to assuming $N_{\text{eff}} = N$ which underestimates the true variance of combined measurements or fitted parameters. Consequently, the fit uncertainties can be grossly

underestimated if correlations are neglected. Confidence intervals derived under the false assumption of independent data will under-cover—the actual probability that the true value lies within the reported interval will be lower than the stated confidence level. Conversely, if anti-correlations are ignored, the fit uncertainties will be overestimated and confidence intervals will over-cover and degrade fit precision. For example, if unfolding induces anti-correlations between neighbouring bins, a fit that neglects those anti-correlations will effectively treat upward fluctuations in one bin and simultaneous downward fluctuations in the adjacent bin as if they were statistically independent deviations of opposite sign. The fit might then struggle to accommodate both, leading to a larger fit uncertainty. The concrete study of a Gaussian unfolding problem demonstrated in Section VII.G shows that using only the diagonal elements of the covariance matrix yields larger asymptotic errors on the fitted parameters compared to using the full covariance, particularly when the induced correlations are large. In this study, the “diagonal only” error bands are overly conservative because they double count fluctuations that in reality were constrained by correlation. In contrast, the full covariance analysis uncertainty estimate and the uncertainty estimate computed through pseudo-experiments, meant to represent bootstrapping, are in good agreement and show smaller uncertainties. These findings reinforce that incorporating the full covariance structure is critical for obtaining accurate uncertainty estimates, both to avoid underestimation in some scenarios and excessive conservatism in others.

From the perspective of bias in parameter extraction, the presence of correlated data does not inherently introduce bias in a fit provided the fitting procedure properly

accounts for those correlations. An unbiased estimator remains unbiased under linear transformations.⁶ That said, the processes that create correlated data can themselves introduce bias. For instance, unfolding methods that regularise strongly can produce a biased estimate of the true distribution, often biasing towards some smooth default model. If one then fits theory parameters to such unfolded data, the results might inherit this unfolding bias. Moreover, an inappropriate handling of correlations can indirectly lead to bias if the fitter effectively misweights parts of the data. In extreme cases, ignoring correlation could cause the fit to chase statistical fluctuations in one bin without realising that a correlated fluctuation in another bin is providing a counteracting influence. The binned study in Section VII.G.2 finds that the best fit values of parameters were very similar whether correlations were accounted for or not, and any small biases as a function of detector smearing were attributable to the overall analysis procedure rather than the correlation handling. This suggests that for reasonably well behaved problems, the central values of fits may not shift dramatically by ignoring correlations indicating that the maximum likelihood estimator for θ can remain approximately unbiased. The bigger impact is on the estimated uncertainties and the validity of statistical tests, not necessarily the point estimate itself. Still, in more complex or poorly modelled situations, neglecting correlations could conceivably pull the fit off the true value if the fitter is effectively solving the wrong optimisation problem.

One frontier where the correlated data problem has become especially salient is in the advent of unbinned unfolding and machine learning driven differential

⁶E.g. combining bins with weights given by Σ^{-1} .

measurements. New algorithms allow experiments to publish results in an unbinned form i.e. instead of histogramming the unfolded cross section in fixed bins, the result may be a reweighted set of events or a learned probability density function representing the particle level distribution. This is extremely powerful for downstream physics interpretation because it preserves maximum information without binning or averaging and, in principle, enables multidifferential or high dimensional comparisons that would be infeasible with coarse bins. However, such analyses face a serious challenge in suitably accounting for correlations in the unfolded data. All unfolding techniques, binned or unbinned, introduce correlations among the events or data points that make up the unfolded result. When one attempts to perform likelihood based inference directly on such an unbinned result, one is confronted by the fact that one does not have N independent draws from the true distribution, but rather a single, composite object that was sculpted by the full dataset. In absence of a known analytical likelihood for correlated events, one tempting but naïve approach is to ignore the correlations and plug the unfolded events, with their weights, into Equation (VII.4), treating them as if they were independent samples of $p(x | \theta)$. This is precisely the scenario in which the independence assumption is violated and can lead to pathological statistical conclusions.

The studies in Section VII.G.4 will quantitatively demonstrate the pitfalls of using unbinned unfolded data naïvely in likelihood fits. A Gaussian example is used to compare three approaches,

1. A traditional binned unfolding followed by a binned likelihood fit (using Equation (VII.5)),

2. An unbinned unfolding with the results binned *post hoc* for a covariance based fit, and
3. An unbinned unfolding with a fully unbinned fit ignoring event correlations i.e. treating the unfolded events as independent in Equation (VII.4).

An important observation from these investigations is that using numerical methods for inference is one possible way to work around this problem. If, instead of relying on asymptotic error formulas or nominal χ^2 distributions, one employs numerical methods such as bootstrapping or pseudo-experiments to directly compute the distribution of estimators or test statistics, then it is possible to obtain correct uncertainties when event level correlations are present. In the Gaussian toy study, when the unfolding + fit process is performed many many bootstrap replicas, the empirical spread of the fitted parameters (‘numerical uncertainty’) matches well with the full-covariance analytic result. In this way one can still use unbinned unfolded data for parameter extraction, but one must forgo simplistic formulas and instead rely on computationally intensive resampling to get trustworthy error bars.

Practically, this is a strong motivation either to re-bin the data, recovering a standard χ^2 approach with covariance, or to develop new statistical formalism that explicitly includes the correlation structure in unbinned likelihoods. At present, there is no known closed-form likelihood analogue to Equation (VII.4) that accommodates event-to-event correlations in an unbinned dataset. Developing such a formalism⁷ is left to future research. Until then, the safest course for analyses is either to stick

⁷For example, an approach using copulas to model the joint PDF of all events, or incorporating the covariance matrix into a generalised likelihood.

to the well-validated binned methods or to use numerical inference procedures for unbinned results.

It is worth highlighting that various HEP experiments are actively exploring these modern methods. For instance, the H1 collaboration and others have released unfolded results using machine learning that are effectively unbinned density estimates [209]. So far, no official parameter fits or new physics tests have been performed directly on such unbinned outputs.

The field is cognisant that while unbinned unfolded data contain rich information, ignoring correlations could yield misleading constraints or false signals. Hence correlations must be treated with the same rigour as any other aspect of the measurement. In practical terms for theory testing, this means that if one is given a published covariance matrix along with a set of data, one should incorporate it in the likelihood or χ^2 . If one only has an unbinned sample of weighted events, one should bin them with the provided weights, obtain an equivalent covariance estimator before fitting, or otherwise use provided ensemble replicas to assess uncertainty. One can always use techniques like the bootstrap to calibrate the confidence intervals and ensure correct coverage. This naturally comes at a computational cost, but it is necessary if the results are to be trustworthy.

In summary, the problem of correlated data in collider analyses is a critical consideration for precision measurements and ignoring these correlations can lead to bias in fitted values (if subtle), under- or over-estimated uncertainties, and confidence intervals or p -values that do not mean what we think they mean. As experiments push toward ever more differential and high-dimensional measurements (often aided

by machine learning), developing sound statistical tools to handle event correlations will be essential to fully realise the potential of these data for testing the Standard Model and searching for new physics. Correlated data have profound impacts on both parameter extraction and theory tests in HEP experiments. Any analysis that ignores the correlation structure in its data is at risk of drawing incorrect conclusions, whether that be an unrealistically precise measurement (uncertainties too small), a failure to detect a true deviation (uncertainties too large or biases), or a discrepancy due to mismodelled uncertainty. The challenges posed by correlated data become increasingly acute as we move toward more granular unbinned measurements. Ensuring statistical procedures remain robust in the face of these correlations is a key task for the interplay of machine learning techniques and rigorous uncertainty quantification in high energy physics.

VII.F Uncertainty quantification.

VII.F.1	Analytic approaches: asymptotic theory and its limitations.	427
VII.F.2	Godambe information: the sandwich estimator)	429
VII.F.3	Wilks' theorem violations and Bartlett corrections	432
VII.F.4	Numerical resampling approaches: bootstrapping and toy	
	Monte Carlo.	435
VII.F.4.i	Non-parametric bootstrap resampling.	435
VII.F.4.ii	Parametric bootstrap: Monte Carlo simulations.	440
VII.F.4.iii	Practical considerations and guidelines.	443
	VII.F.4.iii.a Number of replicas.	443
	VII.F.4.iii.b Computational cost and parallelisation.	444
	VII.F.4.iii.c Diagnostics for coverage and reliability.	444
	VII.F.4.iii.d Incorporating systematic uncertainties.	445
	VII.F.4.iii.e Reporting and using bootstrap results.	446
	VII.F.4.iii.f Common pitfalls.	447

In high energy physics analyses, quantifying uncertainties in measurements is as important as the central values themselves. As outlined in Sections VII.B to VII.D, real world measurements often violate these independence assumptions. This section will establish a rigorous framework for uncertainty quantification in the presence of such correlations, discussing both analytic asymptotic techniques and numerical resampling approaches, followed by a discussion of concrete guidance on implementing these methods, including the choice of replicas, computational cost trade-offs, and diagnostics to ensure correct coverage.

VII.F.1 Analytic approaches: asymptotic theory and its limitations.

Traditional uncertainty estimates in HEP rely heavily on analytic results from classical asymptotic theory. Under the independent, identically distributed (i.i.d.) data assumption (*vide* Section VII.B), the maximum likelihood estimator (MLE) of a parameter θ is consistent and asymptotically normal.

Theorem 14 (Asymptotic Normality of the MLE [39, 83, 84, 194, 277, 441, 570]). Consider a set of i.i.d. observations $\{X_j\}_{j=1}^N$ and a model parametrised by parameter

$\theta \in \Omega \subseteq \mathbb{R}^d$. The likelihood is

$$\mathcal{L}(\theta) = \prod_{i=1}^N p(X_i | \theta), \quad (\text{VII.13})$$

and the Fisher information matrix is

$$I(\theta) = \mathbb{E}_\theta [\nabla_\theta^2 (-\ln \mathcal{L}(\theta))] . \quad (\text{VII.14})$$

Under regularity conditions ensuring consistency and differentiability, the maximum likelihood estimator is distributed as

$$\lim_{N \rightarrow \infty} \hat{\theta} \sim \mathcal{N}(\theta_{\text{true}}, I(\theta_{\text{true}})^{-1}) \quad (\text{VII.15})$$

In practical terms, one often obtains the covariance matrix of estimators by inverting the Hessian of the log-likelihood at the optimum or via the observed information matrix.

Theorem 15 (Wilks [788]). Let $\{X_j\}_{j=1}^N$ be a set of i.i.d measurements with likelihood

$$L(\theta) = \prod_{i=1}^N p(X_i | \theta), \quad \theta \in \Omega \subseteq \mathbb{R}^d. \quad (\text{VII.16})$$

Let the null hypothesis be

$$H_0 : \theta \in \Theta_0, \quad \dim \Theta_0 = r, \quad (\text{VII.17})$$

and the alternative is $H_1 : \theta \in \Theta_1$, with $\Omega = \Theta_0 \cup \Theta_1$. The generalised likelihood ratio statistic is

$$\mathcal{L} = \frac{\sup_{\theta \in \Omega} L(\theta)}{\sup_{\theta \in \Theta_0} L(\theta)}. \quad (\text{VII.18})$$

Under standard regularity conditions and assuming H_0 is true,

$$\lim_{N \rightarrow \infty} -2 \ln \mathcal{L} \sim \chi_{d-r}^2 \quad (\text{VII.19})$$

This theorem underpins the common practice of deriving confidence intervals from profile likelihood scans [18, 166, 306, 430, 571]. For instance, the 1σ interval on a single parameter is given by the range $\Delta(-2 \ln \mathcal{L}) = 1$, assuming the asymptotic χ^2_1 distribution. These asymptotic formulae are analytically convenient and computationally cheap, requiring only the final fit result and local curvature information rather than additional dataset replicas.

However, the accuracy of asymptotic methods can deteriorate when their assumptions are violated, such as in scenarios with correlated or weighted events, model misspecification, or parameter boundaries. If events are correlated, or effectively so, due to shared systematic effects, the simple χ^2 approximations may no longer hold. For example, a composite likelihood or misspecified likelihood, where correlations among data points are ignored in the model, can yield a test statistic whose distribution deviates from χ^2 , often requiring a mixture of χ^2 distributions or an effective scale factor [18, 81, 118, 122, 227]. In such cases, naive use of $\Delta\chi^2 = 1$ for 1σ intervals can significantly misestimate the true coverage.

VII.F.2 Godambe information: the sandwich estimator)

When the model used for inference does not perfectly describe the data's correlation structure, one can apply the sandwich covariance estimator to obtain valid uncertainties.

Theorem 16 (Godambe Information [716]). Let $\{X_j\}_{j=1}^N$ be i.i.d. observations modelled by $p(x | \theta)$, $\theta \in \Omega \subseteq \mathbb{R}^d$. Consider an estimating function

$$U(\theta) = \sum_{j=1}^N u(X_j | \theta), \quad (\text{VII.20})$$

where

$$\forall \theta \in \Omega \quad \mathbb{E}_\theta [u(X_j | \theta)] = 0 \quad (\text{VII.21})$$

such that standard regularity conditions hold ensuring consistency and asymptotic normality of the root $\hat{\theta}$ solving $U(\theta) = 0$. Let

$$H(\theta) = -\mathbb{E}_\theta \left[\frac{\partial u(X_j | \theta)}{\partial \theta} \right] \quad (\text{the sensitivity matrix}), \text{ and} \quad (\text{VII.22})$$

$$J(\theta) = \text{Var}_\theta [u(X_j | \theta)] \quad (\text{the variability matrix}). \quad (\text{VII.23})$$

Then the *Godambe information matrix* can be computed as,

$$G(\theta) = H(\theta)^T J(\theta)^{-1} H(\theta). \quad (\text{VII.24})$$

and

$$\lim_{N \rightarrow \infty} \hat{\theta} \sim \mathcal{N} \left(\theta, \frac{1}{N} G(\theta)^{-1} \right). \quad (\text{VII.25})$$

That is to say, the asymptotic covariance of $\hat{\theta}$ is $N^{-1}G(\theta)^{-1}$.

Theorem 17 (Extension of Theorem 16 [381, 520, 582, 778]). If the estimating function is the score of a correctly specified likelihood,

$$u(X_j | \theta) = -\nabla_{\theta} \ln p(X_j | \theta), \quad (\text{VII.26})$$

then

$$H = -\mathbb{E}[\nabla^2 \ln \mathcal{L}], \quad (\text{VII.27})$$

$$J = \text{Var}(\nabla \ln \mathcal{L}), \quad (\text{VII.28})$$

and the sandwich (or Godambe) information matrix is given by

$$G(\hat{\theta}) = H(\hat{\theta}) J(\hat{\theta})^{-1} H(\hat{\theta}), \quad (\text{VII.29})$$

Intuitively, H^{-1} is the optimistic covariance estimate assuming the model is correct, while J captures the actual observed fluctuations of the score, which increase if data are correlated or the model is incorrect [331, 534]. The ‘sandwich’ $J^{-1} = H^{-1}GH^{-1}$ then provides a robust covariance for $\hat{\theta}$ that remains consistent even if the likelihood is misspecified [743, 763, 764]. In the limit of i.i.d. data with correct model, $H = J$ and the sandwich reduces to H^{-1} as expected. In the presence of event-to-event correlations, $J > H$ (in a matrix sense), meaning the sandwich variance $H^{-1}JH^{-1}$ is larger than the naive H^{-1} , reflecting the loss of information from the unmodelled dependencies. This approach rigorously generalises the notion of ‘effective degrees of freedom’ due to correlations.

While the Godambe information is rarely computed explicitly in everyday HEP analyses, it underlies the validity of techniques employed in composite likelihood fits [233, 603] and provides a formal path to include correlations analytically. J can

be computed by evaluating the empirical covariance of score contributions event-by-event, accounting for any known shared systematics, then used to adjust the reported errors. The challenge is that J is increasingly intractable to compute when the full joint distribution of data is complicated.

VII.F.3 Wilks' theorem violations and Bartlett corrections

Even when one can form a profile likelihood that accounts for some correlations via nuisance parameters, subtle deviations from asymptotic assumptions can persist. For instance, finite sample bias in the likelihood ratio statistic can lead to systematic under-coverage or over-coverage. A classic remedy in statistical theory is the Bartlett correction, which rescales the test statistic to better match the χ^2 distribution.

Theorem 18 (Bartlett correction [686, 775]). Let $\{X_j\}_{j=1}^N$ be i.i.d. according a model $p(x | \theta)$ and let the hypotheses be

$$H_0 : \theta \in \Theta_0 \quad H_1 : \theta \in \Theta_1 \quad \Omega = \Theta_0 \cup \Theta_1, \quad (\text{VII.30})$$

with $\dim \Omega = d$, $\dim \Theta_0 = r$, and the usual likelihood-ratio statistic

$$\mathcal{L} = \frac{\sup_{\theta \in \Theta_0} L(\theta)}{\sup_{\theta \in \Omega} L(\theta)}. \quad (\text{VII.31})$$

Under standard regularity conditions and assuming H_0 is true,

$$\exists c \in \mathbb{R} \quad c \frac{\cdot \mathbb{E}_{\Theta_0}[-\ln \mathcal{L}]}{d - r} = 1 + \frac{B}{N} + \mathcal{O}(N^{-2}), \quad (\text{VII.32})$$

with B determined by lower order cumulants of the log-likelihood.⁸

⁸B can be worked out via the third and fourth derivatives of the log-likelihood [774].

Theorem 19 (Asymptotics of the Bartlett correction). The Bartlett corrected statistic

$$\tilde{\mathcal{L}} = c \cdot \mathcal{L} \quad (\text{VII.33})$$

satisfies

$$\mathbb{E}_{\Theta_0}[\tilde{\mathcal{L}}] = d - r \quad \text{and} \quad \lim_{N \rightarrow \infty} \tilde{\mathcal{L}} \sim \chi^2_{d-r}. \quad (\text{VII.34})$$

In essence, one multiplies $-2\Delta \ln \mathcal{L}$ by a factor $c < 1$ determined from lower-order asymptotic expansions or auxiliary simulation [97, 320]⁹ such that $c \cdot (-2\Delta \ln \mathcal{L})$ has expectation equal to the nominal χ^2 degrees of freedom under H_0 .¹⁰

Bartlett corrections have been studied in various fields for improving likelihood ratio tests in small sample and correlated data situations in certain mixture models [232]. In practice, applying a Bartlett factor in HEP would require either analytical derivation for the specific model or an empirical calibration. One would have to fit a scale factor so that the distribution of $-2\Delta \ln \mathcal{L}$ from pseudo-experiments matches χ^2 . While not routinely done in most HEP measurements, this is one possible correction that can be applied when asymptotic results are suspect.

Another scenario of interest is the profile likelihood under model misspecification. If the true data generating process lies outside the assumed model family, the MLE will converge to the ‘closest’ point in parameter space, *the pseudo-true value* [51, 736] but the usual confidence intervals might be misleading [81, 215]. The profile likelihood curve may be too steep or too shallow relative to the actual sampling distribution of $\hat{\theta}$. In such cases, one can again use the sandwich variance to adjust uncertainties, or

⁹In practice, c is typically computed using bootstraps when analytic expressions are not available.

¹⁰This correction ensures nominal χ^2 coverage up to second order, mitigating the finite sample bias discussed.

else rely on a numerical calibration of the likelihood ratio [257]. For example, one could generate pseudo-datasets from a more complete model or from the measured data itself via resampling and directly measure the distribution of the profile $\Delta \ln \mathcal{L}$ statistic, using that to set confidence intervals rather than assuming a χ^2 [119, 554]. This is essentially a hybrid approach. One is still using the likelihood ratio as the test statistic, but determining its cutoff values by simulation. This technique often employed in searches for new physics when parameters are near physical boundaries or when the approximate independence of data does not hold.

In summary, analytic asymptotic methods provide a powerful toolkit for uncertainty quantification under idealised conditions of large samples, independent data, correctly specified models. Analytic handles like the Fisher information are used to propagate uncertainties; the profile likelihood method has been a workhorse for HEP measurements. Yet, in the presence of strong correlations or non-standard conditions, these methods can lead to errors in inference parameters.

As discussed in Sections VII.C and VII.D, ignoring an unfolding-induced covariance or a shared systematic can violate the regularity conditions of Wilks' theorem. The net effect can be either an underestimate or overestimate of uncertainties, depending on the nature of the correlation.

This is a clear indication that Wilks' theorem in its standard form is failing here. The take-home message is that one must either incorporate the correlation structure into the analytic calculation through approaches like the sandwich estimator or fall back on numerical methods to assess uncertainties.

VII.F.4 Numerical resampling approaches: bootstrapping and toy Monte Carlo.

Numerical uncertainty estimation techniques generate many ‘fake’ realisations of the experiment to directly empirically measure the spread of an estimator or unfolded distribution. These methods do not rely on an explicit formula for the variance; instead, they approximate the sampling distribution by Monte Carlo simulation or resampling. Two broad classes are widely used: non-parametric bootstrapping (resampling from the observed data itself) and parametric bootstrapping (also known as toy Monte Carlo pseudo-experiments, sampling from a parametrised model). Both have become indispensable in modern HEP analyses, especially when dealing with complex outputs where analytic propagation is intractable. This section will review these approaches in turn, including variations relevant to unbinned inference with correlations, and then discuss practical guidelines for their use.

VII.F.4.i Non-parametric bootstrap resampling.

The classical (non-parametric) bootstrap [85] is a generic tool to estimate the uncertainty of virtually any statistic by resampling the data. Given an original dataset of N events $\{x_1, \dots, x_N\}$, one generates a bootstrap replica by sampling N points with replacement from the original set. This procedure randomly selects events such that some original events may appear multiple times in the replica while others may be omitted, effectively drawing from the empirical distribution of the data. By repeating this B times, one can simulate the variability of any statistic: the

distribution of the statistic across the B replicas serves as an approximation for its true sampling distribution.

In context of unfolding or a cross-section measurement, the ‘estimator’ might be the unfolded spectrum itself or, more often, a physics parameter extracted in a subsequent fit. For example, Chapter V employed bootstrap resampling to propagate statistical uncertainty in the MOMENT UNFOLDING method. B bootstrap datasets were generated by resampling events and were unfolded. Then the standard deviation of the unfolded moments was computed across replicas to measure the uncertainty. This ensured that fluctuations in both the data and the simulation based correction were reflected in the final error bands.

One advantage of the non-parametric bootstrap is that it makes minimal assumptions about the underlying distribution. One does not need an analytic model of how data are distributed, only the empirically observed sample. This is particularly useful in high dimensional or highly non-Gaussian situations common in HEP. For instance, when the outcome of an analysis is a complicated function of many event properties and detector effects, the non-parametric bootstrap might be the only possible method to estimate uncertainties.

The bootstrap will faithfully capture the variance as long as the original dataset is representative of the true distribution. It also naturally incorporates correlations among observables. Each bootstrap replica is like a pseudo-experiment drawn under the null hypothesis that “the observed dataset is the true underlying distribution”. By analysing all replicas identically, one can extract not just variances for each observable but the full covariance matrix via the sample covariance of the replica

results [281]. As just two examples of many in [86, 300] the ATLAS Collaboration uses bootstrapping as a straightforward way to estimate statistical correlations between bins of an unfolded cross section. The non-parametric bootstrap can even be used for comparisons between two separate analyses that share data.

Because each event is resampled with a deterministic seed, the same random fluctuations can be propagated coherently into multiple analysis outcomes, allowing an estimate of their correlation [60]. This is a significant benefit over analytic approaches, where deriving the covariance between two complex observables might be extremely cumbersome.

Despite its generality, the non-parametric bootstrap must be applied with care in the presence of correlated events. The standard bootstrap assumes each observation is an independent draw from some distribution. If the true data have correlations,¹¹ a naïve bootstrap that samples individual data points independently will destroy the correlation structure. In such cases, one should resample at the level of the correlated unit. An example of this is when a collision event produces multiple jets whose properties are analysed (hence jets from the same event are correlated), the bootstrap should treat the entire event as the fundamental unit to resample. This is known as a cluster bootstrap or block bootstrap [243, 291, 609, 616, 724]. With sufficient computational resources available, one could resample whole event with replacement to build each replica dataset, thus preserving event correlations.

Similarly, if data have an inherent time or spatial correlation, one might sample temporal or spatial blocks of consecutive events. By doing so, the bootstrap replicas

¹¹E.g., an event consists of multiple particles' measurements, or there are clusters of events that fluctuate together due to a common systematic.

reflect the dependency structure present in the original data. Failing to do this can lead to misestimated uncertainties because the resampled replicas would appear overly variable or overly independent compared to reality, giving an impression of containing more information than actually available.

Another subtlety to be aware of is that the non-parametric bootstrap conditions on the observed sample size N . Each replica has exactly N events, though not N unique even. However, in many HEP measurements, N itself is a Poisson quantity. The number of events passing selection follows Poisson statistics with mean proportional to the true cross section and integrated luminosity. By fixing N , the bootstrap neglects the contribution of counting uncertainty. For very large N , this distinction is negligible because $\text{Poisson}(\lambda)$ fluctuations are proportional to $\sqrt{\lambda}$, which is small relative to λ . But for small datasets or when quoting total rate uncertainties, this matters.

A simple variant that addresses this is the so-called Poisson bootstrap [711, 738].

Definition 48 (Poisson bootstrap). The *Poisson bootstrap* is a variant of non-parametric bootstrap resampling in which instead of drawing exactly N samples each time, each original event, x_i , is assigned a $\text{Poisson}(1)$ random weight, w_{ij} in each replica, j [508, 578]. x_i is then included in replica j w_{ij} times.

The expected total count in any replica of a Poisson bootstrap is then also N (since $\sum_{i=1}^N \text{Poisson}_i(1) \sim \text{Poisson}(N)$, which has expected value N), but the total count of any individual replica is no longer fixed at N . It is now a Poisson random variable that fluctuates around N . This method has been used directly in particle

physics experiments [12, 48] because it mirrors the way a real experiment could have produced a more or fewer events than the mean.

The Poisson bootstrap has the additional advantage of simplifying code. One does not need to sample events one by one; instead, one can generate N Poisson random numbers and replicate event i that many times. Several HEP analyses have implemented exactly this approach, with a fixed random seed per event to ensure reproducibility, producing thousands of Poisson-weighted replicas to evaluate both per bin uncertainties and bin-to-bin correlations for unfolded spectra [88, 153, 154, 245, 282, 388, 494, 720]. Importantly, this approach implicitly treats the entire collection of events as coming from a Poisson process, which is appropriate for counts of collisions.

For unbinned unfolded data, where the result of an unfolding procedure is a continuous distribution or a set of weighted events at “truth level”, the bootstrap can be applied by resampling at the detector level before unfolding. In other words, one can generate bootstrap variations of the raw data (and even of the simulation used in the unfolding) and run the unfolding on each variation. This yields an ensemble of unfolded results, from which uncertainties can be derived. This is the strategy followed in this work for evaluating unfolding uncertainties. The same detector level data is fluctuated many times by resampling counts in each bin, run the unfolding algorithm each time and the spread in the unfolded solution is observed.

In this way applying the procedure to bootstrapped detector level samples yields many sets of weights, whose variance indicates the uncertainty on the underlying truth distribution. Because many modern methods lack simple closed-form error

propagation, such a bootstrap approach is often the only viable way to quantify uncertainties on the unfolded distribution. One must ensure, of course, that the computational cost is manageable. Unbinned unfolding with neural networks can be computationally expensive, so the number of bootstrap replicas B may be limited by available computational resources.

VII.F.4.ii Parametric bootstrap: Monte Carlo simulations.

The parametric bootstrap differs in that it assumes one has a generative model for the data. Instead of resampling the observed events, one simulates new datasets from a known or fitted probability distribution. In particle physics, this is synonymous with conducting many pseudo-experiments. The procedure is laid out in Algorithm 4. This approach is inherently model dependent—it relies on the assumption that our

Algorithm 4 Parametric bootstrap pseudo-experiments

Input: Observed data $X_1, \dots, X_N$, model $p(x \mid \theta)$, number of replicas B

Output: Bootstrap estimates $\{\hat{\theta}^{(b)}\}_{b=1}^B$ or distribution of a statistic

```

 $\hat{\theta} \leftarrow \arg \max_{\theta} \prod_{i=1}^n f(X_i; \theta)$ 
for  $b \leftarrow 1$  to  $B$  do
    Simulate pseudo-dataset  $X_1^{(b)}, \dots, X_n^{(b)} \stackrel{\text{i.i.d.}}{\sim} f(\cdot; \hat{\theta})$ 
    Compute  $\hat{\theta}^{(b)}$  by applying the same estimation procedure to  $X_{1:n}^{(b)}$ 
return  $\{\hat{\theta}^{(b)}\}_{b=1}^B$ 

```

fitted model is a good representation of the underlying truth. However, in many cases, especially when the goal is to test consistency with that model or to quote

uncertainties assuming the model, this is perfectly acceptable. In fact, it is often the only way to incorporate known physics processes and detector effects exactly as they enter the analysis.

One benefit of the parametric bootstrap is the ease of incorporating known systematic and correlation effects. Because one is generating data, one can build in whatever correlations we expect. For example, if two variables in each event have a known correlation, the simulation can sample them jointly rather than independently. If there is a nuisance parameter with some uncertainty, one can randomly vary that parameter for each pseudo-experiment, drawn from its uncertainty prior, so that the ensemble of pseudo-datasets reflects the lack of knowledge of the parameter. In this way, the resulting spread of outcomes includes the effect of that systematic.¹² Similarly, parametric generation naturally handles Poisson fluctuations in event counts. Each pseudo-experiment can draw $N_{\text{events}} \sim \text{Poisson}(\langle N \rangle)$ where $\langle N \rangle$ is the expected number given the cross section and luminosity. Thus, unlike the fixed-size non-parametric bootstrap, the parametric approach inherently includes the statistical uncertainty from the finite event count as well.

In HEP, the parametric bootstrap underlies many standard practices. For instance, the CL_s method for setting limits on new physics involves generating many pseudo-datasets under the background only and signal + background hypotheses to determine the distribution of a test statistic [567, 570, 576, 647, 667, 668, 726, 749, 752, 769]. Similarly, when experiments quote an unfolding covariance matrix from ‘toy MC,’

¹²This is analogous to the commonly used ‘toy Monte Carlo’ approach of varying nuisance parameters within their errors to see the impact on a fit result. Here one integrates it into the generation of each bootstrap.

they have typically taken the central unfolded result as truth, re-simulated many fake datasets through detector simulation, unfolded each, and computed the covariance of the ensemble. This can be interpreted as a parametric bootstrap of the unfolding result. Section VII.D discussed how ignoring certain correlations can undermine analytic formulae; the parametric bootstrap gives a robust alternative because it does not assume those formulae. Instead it measures the uncertainty by numerically. For example, if one is unsure about the validity of Wilks' theorem in a complex fit, one can generate toy experiments at the best fit parameters and validate empirically what fraction of the time $\Delta \ln \mathcal{L}$ exceeds a certain value.

The primary limitation of bootstrapping is the heavy computational load and potential model bias. Generating and analysing B pseudo-experiments can be extremely compute intensive. In practice, one often uses a fast approximation such as a parametrised smearing or a fast detector simulation to make this feasible. [281] suggests that on the order of $B = 1000$ replicas is a reasonable compromise in many cases, though even $B = 10\,000$ is used for certain jet cross section covariance evaluations.

The exact choice of B is particular to the specific analysis being conducted, and should be chosen by monitoring when the statistical error on the uncertainty estimate itself becomes negligible (*vide* Section VII.F.3). As for model bias, if the true data-generating process differs from the model used, the bootstrap will estimate the variance around the wrong central value. This is usually acceptable when one is interested in relative uncertainties or when the model has been tuned to data (e.g., using the unfolded result itself as the truth baseline). Nonetheless, it is prudent to

cross check. One can, for example, perform both a non-parametric and parametric bootstrap and ensure they yield compatible uncertainty estimates. Discrepancies might indicate sensitivity to modelling.

This section would be incomplete without noting a final variant: the hybrid bootstrap. In some cases, one might resample certain aspects from data and others from a model. For instance, one could bootstrap the residuals of a fit rather than raw data, which is common in regression contexts, to mimic new noise samples, though less common in HEP. Or one might resample events but also randomly fluctuate a global parameter for each replica. These approaches blur the line between non-parametric and parametric and can be tailored to specific correlation structures.

VII.F.4.iii Practical considerations and guidelines.

While bootstrapping and toy MC methods are conceptually straightforward, their successful deployment in a HEP analysis requires careful planning. Here are some high level practical guidelines and considerations for using these techniques to quantify uncertainties.

VII.F.4.iii.a Number of replicas.

The accuracy of bootstrap estimates improves with the number of replicas, but computational cost grows linearly. In practice, $B \sim \mathcal{O}(10^3)$ is common. A few hundred replicas may suffice for stable estimates of variances, but for reliable estimation of the full covariance matrix, especially in high dimensions, or tails of distributions, one might need an order of magnitude more. It is good practice to check convergence of

the uncertainty estimate: e.g. run $B = 200, 400$ and 800 and monitor if the reported uncertainty or key figure of merit changes appreciably. If it stabilises, B is sufficient.

The standard error on a bootstrap derived standard deviation is $\sigma/\sqrt{2(B-1)}$, so increasing B yields diminishing returns.

VII.F.4.iii.b Computational cost and parallelisation.

Each bootstrap replica involves an independent analysis of the dataset, and so parallelisation can offer overwhelming computational gains. One should exploit parallel computing¹³ to run many bootstraps concurrently. In a modern analysis framework, it is often feasible to distribute $B = 1000$ jobs to a cluster and retrieve results in a few hours, where running them serially would be prohibitive. If the analysis involves training a machine learning model for each replica, one might need to reduce B or use a pretrained model across replicas if appropriate.¹⁴

An alternative to reduce the computational cost is the bootstrapping of approximate models. For example, one could fit a faster surrogate model, like a parametric function, to each replica instead of rerunning a full simulation or complex inference. This sacrifices some fidelity but can massively speed up the procedure.

VII.F.4.iii.c Diagnostics for coverage and reliability.

After obtaining uncertainty estimates, whether analytic or bootstrap, it is crucial to validate that they have the correct coverage. Coverage means that a nominal 68%

¹³E.g. multithreading.

¹⁴Bearing in mind that if the model training itself is subject to statistical fluctuation, use of a pretrained model will fail to capture that source of variance.

confidence interval indeed contains the true value about 68% of the time in repeated experiments.

Often one can leverage simulation to perform a coverage study. This involves generating many mock datasets from a known truth and seeing how often the method's interval would capture that truth. Such studies can be done with the parametric bootstrap, treating the truth as known. If the intervals under-cover or over-cover significantly, one might adjust the method.¹⁵

Within a single dataset, one diagnostic test is to compare different methods. If bootstrap intervals are significantly different from asymptotic ones, it flags a potential issue with the latter. Discrepancies between bootstrap derived intervals and asymptotics require investigation. They could indicate anything from a bug in analysis to a breakdown of core assumptions. Conversely, agreement between them builds confidence that uncertainties are well understood. In the next few sections, we will see examples of this where the asymptotic formulae give systematically different errors from the bootstrap spread, indicating miscoverage. The recommendation then, echoing recent studies, is to trust the numerical approach or augmented analytic approach that accounts for correlations [629].

VII.F.4.iii.d Incorporating systematic uncertainties.

So far, the discussion has focussed on statistical uncertainty from limited data samples. In a full measurement, one must also account for systematic uncertainties.

¹⁶ Bootstrapping can be extended to some systematic effects. One can, in principle,

¹⁵E.g., One might use a different statistic, or apply a Bartlett-like scale factor.

¹⁶From detector calibration, theory model choices, *et cetera*.

treat a systematic variation as producing an alternate ‘truth’ and then bootstrap around that.

However, more common in HEP is to evaluate systematics separately by shifting calibrations and re-running the analysis a few times, and to then combine statistical and systematic errors. One must be careful not to mix these in the same replicas unless the systematic is itself statistical in nature. For example, the statistical error on a background estimation can be propagated by bootstrapping the background sample generation.

The overarching principle is to match the uncertainty evaluation method to the source of uncertainty, i.e. to use bootstraps or toys for random fluctuations, and to treat systematic shifts by exploring the parameter space of uncertainties, possibly with their own pseudo-experiments if they have uncertainty distributions.

VII.F.4.iii.e Reporting and using bootstrap results.

When using bootstrap ensembles, one often obtains a covariance matrix for multiple observables or parameters as the primary output. This covariance can be used in downstream fits like treating the unfolded spectrum and its bootstrap covariance as data for a theory fit. It is important to regularise or smooth the covariance if B is not extremely large, because being an empirical estimate, it can be noisy, and potentially not even positive definite if B is too low relative to number of bins.

Simple techniques include increasing B or applying shrinkage to the covariance matrix. Showing consistency between the diagonal errors from bootstrap and asymptotics, and quoting the correlation coefficients if relevant, can be very informative in

publications. Tables or matrices of bootstrap correlations might be provided in an Appendix or auxiliary material if they are of interest.

VII.F.4.iii.f Common pitfalls.

Bootstrap bias: The bootstrap can sometimes reveal bias in an estimator, e.g., the average of bootstrap estimates differs from the original estimate. If significant, one might use the bootstrap to correct the bias (bias-corrected estimates) or employ the BCa (bias-corrected and accelerated) percentile method for confidence intervals [499].

Small sample issues: In regions with extremely few events, bootstrap samples may not be representative because one keeps drawing the same few events in different orders. In such regimes, exact methods or simple conservative analytic bounds might be more reliable.

Randomness: One should ensure reproducibility by fixing the random number seeds for bootstrapping procedures. This allows others to regenerate the same replicas and verify results.

Interpreting bootstrap outcomes: One should remember that the bootstrap gives an estimate of uncertainty, not a guarantee. If the data are very sparse or the model very wrong, the bootstrap will dutifully reflect those. Results should always be interpreted in the physical context and cross verified with sanity checks.

In conclusion, bootstrapping versus asymptotics is not an either–or choice but a complementary set of tools. Asymptotic techniques are fast and illuminating; they are instructive for learning how the uncertainty scales with data size and often give insight into which parameters or features dominate the error. Numerical techniques, on the other hand, are robust and account for real–world complexities that elude analytic treatment. For cutting edge analyses a balanced approach might be most appropriate, where one uses analytic formulae as a first pass and consistency check, but relies on bootstrapping or pseudo–experiments for the final uncertainty quantification when data correlations or algorithmic complexities are significant to ensure that confidence intervals are trustworthy. Indeed, using numerical techniques can turn up hidden uncertainty contributions that asymptotic formulae would have glossed over. The next sections will apply this framework to concrete case studies, comparing binned vs. unbinned approaches and demonstrating how the choice of uncertainty quantification technique impacts the physics conclusions.

VII.G Case studies.

VII.G.1 Setup.	450
VII.G.2 Fully binned baseline.	450
VII.G.2.i Methodology.	451
VII.G.3 Correlation diagnostics after unbinned unfolding.	460
VII.G.3.i Pairwise weight–distance correlations	461
VII.G.3.ii Histogram covariance matrices.	464
VII.G.3.iii Implications.	465
VII.G.4 Unbinned unfolding: binned and unbinned inference.	466
VII.G.4.i Parameter bias.	468
VII.G.4.ii Uncertainty estimation and coverage.	470
VII.G.5 Extension to higher dimensions.	475
VII.G.5.i Evaluation.	479

VII.G.1 Setup.

A Truth distribution is constructed from a Gaussian with mean $\mu_T = 0.2$ and variance $\sigma_T^2 = 0.81$. 10^4 events are drawn from this distribution to serve as the dataset at the particle level. In addition, a larger Monte Carlo (MC) sample of 10^5 events is generated from a Gaussian with mean $\mu_G = 0.0$ and variance $\sigma_G^2 = 1.0$. This Generation sample plays two roles—it provides an initial guess of the truth distribution for unfolding, and it is used to derive the detector response. Notably, the Generation distribution differs from the true distribution, mimicking the realistic situation in which the nominal MC does not perfectly match reality.

Both the Truth and the Generation datasets are passed through a detector response model to simulate detector level measurements by smearing each particle level event with an independent Gaussian error of mean 0 and a certain resolution σ_{det} , representing finite detector resolution. Multiple smearing levels are examined, with σ_{det} ranging from 0, no smearing, i.e. a perfect detector, up to 0.75, substantial resolution degradation. The smeared Truth sample represents the detector level Data that an experiment would observe, while the smeared Generation sample represents the detector level Simulation.

VII.G.2 Fully binned baseline.

To establish a baseline, first consider this simple one dimensional unfolding scenario using a known Gaussian distribution and the traditional binned unfolding approach. In this fully binned context, generate Gaussian particle level and detector level data, apply Iterative Bayesian Unfolding (IBU), and perform a standard binned

parameter fit to the unfolded histogram. The goal is to characterise the bias, variance, and confidence interval coverage of the inferred Gaussian parameters under simple conditions, before proceeding to more advanced unbinned methods in subsequent sections.

This study intentionally uses a Gaussian example with known truth parameters so that any bias can be directly evaluated and so that many pseudo-experiments can be quickly generated to study statistical fluctuations. This demonstration will highlight the importance of properly accounting for the covariance matrix of the unfolded result when making inferences on physics parameters.


500 independent pseudo-datasets are generated from the Truth/Data model, each with a Poisson fluctuation in total event count with a expectation value of 10^4 events per dataset to simulate realistic statistical scatter.¹⁷

VII.G.2.i Methodology.

Bin all the datasets using a fixed histogram binning. For this example, 15 uniform bins are used spanning the range $[-3, 4]$.¹⁸ The choice of binning is kept constant for the unfolding procedure and subsequent analysis. An illustration of the distributions at truth level and detector level is shown in Figure VII.1. Figure VII.2a shows the Simulation distribution at smearnings $\sigma_{\text{det}} \in \{0.00, 0.25, 0.50, 0.75\}$ and the corresponding response matrices.

¹⁷The study as presented uses pseudo-experiments of a Poisson fluctuated size rather than a non-parametric Poisson bootstrap due to the ease of generating Gaussian examples, but essentially results were replicated using a non-parametric Poisson bootstrap proper, and the observations and analysis carry over essentially unchanged.

¹⁸The bin range chosen is asymmetric around 0 because $\mu_T = 0.2 > 0$.




(a) Particle level distributions comparing Truth (blue) with parameters $\mu_T = 0.2, \sigma_T^2 = 0.81$ with Generation (orange) with parameters $\mu_G = 0.0, \sigma_G^2 = 1.00$. The histograms represent normalised frequency distributions and the dashed curves show the analytic density.

(b) Detector level distributions after applying Gaussian smearing with resolution $\sigma_{\text{det}} = 1.0$. Data (blue) has parameters $\mu_D = 0.2$ and $\sigma_D^2 = 1.81$, while Simulation (orange) has $\mu_S = 0.0$ and $\sigma_S^2 = 2.00$.

Figure VII.1: Comparison of particle level and detector level distributions. **(a)** shows the particle level distributions where the Truth (blue histogram and dashed curve) differs from the Generation (orange histogram and dashed curve). **(b)** demonstrates the effect of finite detector resolution through Gaussian smearing with $\sigma_{\text{det}} = 1.0$, which convolves with the underlying distributions to produce broader observed distributions at the detector level.

Using the large MC sample, construct a detector response matrix mapping particle level bins to detector level bins. Each element R_{ij} of this matrix gives the probability for an event originating in truth bin i to be reconstructed in detector bin j . In practice, R_{ij} is obtained by binning the MC events by their true and smeared values. The response matrix, along with the smeared data histogram, serves as input to the



(a) Detector level distributions as a function of detector resolution σ_{det} . Step histograms show the normalised density of Simulation after applying Gaussian smearing with $\sigma_{\text{det}} \in \{0.00, 0.25, 0.50, 0.75\}$. Dashed curves represent the corresponding analytic Gaussian with variance $\sigma_G^2 + \sigma_{\text{det}}^2$.

(b) Discretised response matrices $R(\mu_i | \nu_j)$ for detector resolutions 0.00, 0.25, 0.50 and 0.75. Each row is normalised so that $\sum_i R(x_i | z_j) = 1$ for all j . Observe the transition from a diagonal matrix (perfect resolution) to increasingly off diagonal structures as detector smearing increases.

Figure VII.2: Characterisation of detector response effects through progressive smearing of Generation. **(a)** illustrates how Generation is broadened by convolution with Gaussians of width σ_{det} . The variance adds in quadrature: $\sigma_S^2 = \sigma_G^2 + \sigma_{\text{det}}^2$. **(b)** presents the corresponding response matrices, which encode the probability of bin migration due to detector effects. For perfect resolution ($\sigma_{\text{det}} = 0$), the response matrix is diagonal, indicating no bin migration. As detector resolution degrades, off diagonal elements become increasingly prominent, as events are reconstructed in neighbouring bins.

unfolding algorithm.¹⁹ Figure VII.2b shows the response matrices estimated using Generation–Simulation pairs for various levels of smearing.

Then apply the IBU algorithm to the binned Data in order to unfold the effects of smearing. Starting from the Generation distribution, IBU iteratively updates

¹⁹In Bayesian terms, the unsmeared MC histogram also provides a prior estimate for the truth distribution.

the estimate of the Truth histogram by comparing the Simulation and Data in the detector space and applying Bayes' theorem to reweight contributions back to particle space. In this study, IBU was iterated for five iterations.

This unfolded result consists of bin counts (or densities) for the truth histogram, along with an estimated covariance matrix for those bin counts. The bin-to-bin covariances arise from the finite statistics of the data and from the smearing correlations introduced by the unfolding procedure.

One can estimate the covariance matrix of the unfolded histogram by repeating the unfolding on many statistically independent pseudo-experiments, similar to bootstrapping. After unfolding, perform a binned fit to extract the parameters of the underlying distribution (in this case, the Gaussian's mean μ and variance σ^2). For each unfolded histogram, a χ^2 minimisation is used to fit a Gaussian model to the unfolded data. Importantly, the theoretical Gaussian model is integrated over each bin to yield the expected content in that bin for given values of μ and σ^2 , so that the comparison between the model and the unfolded histogram is exact with no interpolation or bin centre approximation.


The fit's χ^2 is computed using the full covariance matrix of the unfolded bins, thereby incorporating all statistical correlations between bins. The outcome of the fit is a pair of best-fit parameters $(\hat{\mu}, \hat{\sigma}^2)$ along with their asymptotic uncertainty estimates (1σ intervals) derived from the curvature of the χ^2 at the minimum. For comparison, this study also consider a *diagonal covariance* fit, repeating the same procedure but using only the diagonal elements of the covariance, functionally treating unfolded bin counts as if they were independent. This allows one to see the impact

of neglecting inter-bin correlations on the inferred uncertainties. To numerically evaluate the bias, variance, and confidence interval coverage of this procedure, the entire analysis is repeated for each pseudo-experiment.


This ensemble provides a distribution of fitted parameter values $\hat{\mu}$ and $\hat{\sigma}^2$ from which one can quantify the bias, which is the difference between the average fitted value and the true value, the variance or spread of the estimates (related to the expected statistical uncertainty), and the empirical coverage of confidence intervals. In practice, one could alternatively use bootstrapping on a single dataset to assess these metrics; indeed, this study has also checked that bootstrapped replicas yield consistent uncertainty estimates as the independent toys, confirming that the ensemble size is sufficient.

Use the above procedure to obtain an unfolded result and fit for each pseudo-experiment. One can now examine the bias and variance of the estimates.

The analysis is repeated for each value of the detector resolution $\sigma_{\text{det}} \in [0, 0.75]$ to see how worsening detector effects impact the inference. Figure VII.4a and Figure VII.4b summarise the accuracy or bias of the method, showing the mean fitted μ and σ^2 across 500 trials as a function of the detector smearing. Figure VII.3 focuses on the precision and uncertainty coverage, comparing the nominal 1σ errors from the fits to the actual spread of the results. These results demonstrate that the IBU unfolding followed by a proper binned fit recovers the true distribution parameters without significant bias. For all tested smearing levels, the average fitted mean $\mathbb{E}[\hat{\mu}]$ remains within statistical uncertainty of the true value, 0.2, and the average fitted variance $\mathbb{E}[\hat{\sigma}^2]$ within statistical uncertainty of 0.81. In Figure VII.4a and




(a) Asymptotic and bootstrap RMS uncertainty on the Gaussian mean μ as a function of σ_{det} . The lower panel displays the ratio of asymptotics to the bootstrap RMS, demonstrating systematic overestimation of uncertainties when correlations are ignored.




(b) Asymptotic and bootstrap RMS uncertainty on σ^2 as a function of σ_{det} . The diagonal approximation increasingly overestimates the uncertainty as resolution degrades, indicating that neglecting correlations leads to measurable over-coverage.

Figure VII.3: Comparison of uncertainty estimates for binned unfolded distribution parameters as a function of resolution. (a) and (b) show uncertainties for the mean μ and variance σ^2 , respectively. The analysis compares (i) asymptotics using the full covariance matrix (green circles), (ii) asymptotics using the diagonal covariance approximation that neglects off-diagonal terms (pink squares), and (iii) empirical validation through 500 pseudo-experiments (green stars). The excellent agreement between the full covariance asymptotic errors and bootstrap derived uncertainties validates the analytical error propagation. The diagonal approximation consistently overestimates uncertainties, with the discrepancy growing with detector smearing. Thus, proper treatment of the full covariance is essential for when extracting distribution parameters.



(a) Mean best-fit values of μ as a function of σ_{det} . The lower panel shows the ratio to truth, demonstrating unbiased parameter extraction within statistical uncertainties across all resolutions.



(b) Mean best-fit values of σ^2 as a function of σ_{det} . The slight systematic trend toward misestimation at large σ_{det} may arise from regularisation bias in the unfolding or binning effects.

Figure VII.4: Parameter extraction from unfolded distributions across varying detector resolutions. (a) and (b) display the mean best-fit values for the mean μ and variance σ^2 obtained by fitting unfolded distributions. The analysis demonstrates that both the full covariance (green circles) and diagonal approximation (pink squares) yield unbiased central values, recovering the true parameters (red dashed lines) within statistical uncertainties across the full range of detector resolutions. Error bars represent the standard error of the mean over the ensemble of pseudo-experiments. The ratio panels confirm parameter recovery within uncertainties, of truth values validating the unfolding procedure's ability to correct for detector effects without introducing systematic biases.

Figure VII.4b, the averaged best fit values lie on or around the horizontal red dashed lines marking the true parameters. The deviations are well within the statistical error bars²⁰, indicating no significant bias in the unfolding or fitting procedure. Notably, the choice of using the full covariance versus only diagonal uncertainties in the fit has no effect on the central values obtained. Both approaches yield correct $\hat{\mu}$ and $\hat{\sigma}^2$, which might be explained by the bias in this context being dominated by any unfolding imperfections and IBU, with sufficient iterations, being an asymptotically unbiased estimator of the truth distribution.

When using the full covariance matrix in the χ^2 fit, the asymptotic uncertainty estimates for $\hat{\mu}$ and $\hat{\sigma}^2$ are found to be in excellent agreement with the actual distribution of fit results across the ensemble. In Figure VII.3, the green circle markers show the average 1σ uncertainties from the fits (i.e. the fit errors from the covariance matrix of each fit) as a function of detector resolution. These are virtually indistinguishable from the green star markers, which indicate the empirical RMS standard deviation of the 500 fitted values at each resolution. In other words, the fit's error estimates accurately predict the trial-to-trial fluctuations of the outcomes. This agreement implies that the reported 68% confidence intervals have the correct coverage: approximately 68% of the pseudo-experiments' $\hat{\mu}$ (or $\hat{\sigma}^2$) results lie within $\pm 1\sigma$ of the true value, as expected for well calibrated uncertainties. As the detector resolution worsens, the uncertainty on the inferred parameters grows, reflecting the loss of information, but at each point, the full covariance χ^2 fit's uncertainty remains an accurate representation of the actual variance in the results.

²⁰The standard error of the mean over the 500 trials.

If one ignores the off-diagonal elements of the unfolded covariance matrix and assumes the bins are independent, the inferred uncertainties become misestimated. In this example, neglecting the negative bin-to-bin correlations from unfolding leads to an overestimation of the parameter errors. The pink square markers in Figure VII.3 show the average fit uncertainty on μ and σ^2 obtained when using only the diagonal uncertainties. These are significantly larger than both the true ensemble spread (green stars) and the full covariance asymptotic errors (green circles) once any non-zero smearing is present.

This overly conservative error estimate would lead to over-coverage of confidence intervals. The nominal 68% interval contains the true value substantially more than 68% of the time. This indicates an inefficient use of information. The fit is effectively double-counting fluctuations in each bin that in reality are anti-correlated with fluctuations in other bins. These findings reinforce that incorporating the full covariance matrix from the unfolding is essential to obtaining reliable and correctly sized confidence intervals in the fully binned approach, and suggests that this may be the case in unbinned approaches too.

In summary, this study in the fully binned regime demonstrates that a conventional unfolding approach (IBU) combined with rigorous statistical treatment produces unbiased and well calibrated inference of physics parameters. The unfolded results, when analysed with their full covariance matrix, give parameter estimates whose uncertainties accurately reflect the true spread across experiments, ensuring correct confidence interval coverage. The study also shows that simplifying assumptions like

treating unfolded bins as independent can skew the uncertainty estimates, in this case, making them overly conservative, even though the point estimates remain correct.

These findings provide a firm baseline for comparison with the unbinned methods discussed in the next sections. After a brief detour to discuss correlation diagnostics after unbinned unfolding, following section will investigate how unbinned unfolding techniques handle event correlations and whether they can achieve a similar level of statistical reliability as this fully binned approach. The lessons learned here, particularly the necessity of accounting for induced correlations, will carry forward as one transitions to unbinned inference on correlated data.


VII.G.3 Correlation diagnostics after unbinned unfolding.

Before confronting any inference procedure with an unbinned output, one must first quantify the statistical dependencies that the unfolding induces. This subsection introduces two complementary diagnostics that make those correlations both visible and quantifiable.

1. The **pairwise weight—distance correlation** $\rho(|x_i - x_j|)$ between all pairs unfolded events i and j , and
2. The **bin—bin covariance matrix**, Σ of a fine histogram binned after unfolding.

Both are evaluated for four detector resolutions 0.00, 0.25, 0.50 and 0.75 and for two distinct estimators, KDE and neural network (NN) conditional density estimators.

The results are displayed in Figure VII.5 and Figure VII.6, respectively, and form the empirical basis for the coverage study in Section VII.G.4.




(a) KDE based OMNIFOLD: Weight correlations decay exponentially with separation distance, maintaining values of ~ 0.1 at distances of 2–3 units for $\sigma_{\text{det}} = 0.75$. Figure created by Owen Long [44].

Figure VII.5: Average pairwise correlation between unfolding weights as a function of absolute distance $|x_i - x_j|$ between events in the observable space. Both panels show results $\sigma_{\text{det}} \in \{0.00, 0.25, 0.50, 0.75\}$. Error bars represent the RMS variation of correlation values within each distance bin. **(a)** shows the KDE approach which exhibits more localised correlations. **(b)** displays the neural network (NN) method, which produces longer-range correlations. (Continued on next page)

VII.G.3.i Pairwise weight–distance correlations

For every event, by considering the weight distributions across pseudo–experiments one can compute the correlation coefficient, ρ as defined in Definition 19. One can then average ρ_{ij} in bins of $\Delta x = |x_i - x_j|$, where x is the unfolded observable. A few different features of the resulting curves in Figure VII.5 (Figure VII.5a for KDE, Figure VII.5b for NN) are worth nothing.



(b) NN-based OMNIFOLD: Weight correlations exhibit rapid decay, becoming negligible beyond $|x_i - x_j| \approx 1$ for all detector resolutions. Figure created by Owen Long [44].

Figure VII.5: (continued from previous page) The characteristic correlation length scales approximately linearly with detector resolution. Figure created by Owen Long [44].

Perfect detector ($\sigma_{\text{det}} = 0$): One should expect the average pairwise correlation to be $\rho(|\Delta x|) \simeq 0$ for all separations $\Delta x > 0$, confirming that when no smearing is applied, OMNIFOLD reproduces the factorised likelihood limit in which event weights are statistically independent [64, 309]. At $\Delta x = 0$ we trivially should expect $\rho(0) = 1$ because every event is, of course, perfectly correlated with itself. In both Figures VII.5a and VII.5b, it is correctly the case that $\rho(0) = 1$ and $\rho(|\Delta x| \gg 1) \simeq 0$.

In the truly ideal i.i.d. case, one would expect the fall off from $\rho(0) = 1$ to $\rho(|\Delta x| > 0) = 0$ to resemble a Dirac δ at the origin. The KDE estimator indeed approaches this limit. Its kernel bandwidth determines a correlation

length $\ell \lesssim 0.15$. [620, 708]. By contrast, even with extensive experimentation with architecture and hyperparameter tuning one finds that NN priors leave a residual plateau $\rho \approx 0.75$ around $\ell \approx 1$, which then leads to the damped oscillatory structure of the correlation curve. One could speculate about reasons why unfolding with neural networks even in the perfect detector resolution case should lead to correlated weights. Two potential explanations are

1. *Network smoothness*: the shared weights in the neural network classifier impose a finite “receptive field,” [135, 280, 699] causing nearby events to receive similar gradient updates and hence correlated weights; and
2. *Global normalisation*: OMNIFOLD enforces $\sum_i w_i = N_{\text{MC}}$ at every iteration. This normalisation might be introducing a weak positive correlation of size $1/N$ even in the absence of detector effects [568, 789].

Although numerically tiny, this effect, whatever its underlying explanation sets the optimistic lower bound on variance reduction when working with neural networks. Thus even the idealised zero-smearing case does not yield perfectly independent weights in practice, a useful caution when quoting statistical uncertainties from inference using ML processed data.

Mild smearing ($\sigma_{\text{det}} = 0.25$): Short range correlations of order $\rho \gtrsim 0.5$ appear for $\Delta x \lesssim 0.4$. These are already sufficient to reduce the effective sample size

$$N_{\text{eff}} = \frac{N}{1 + (N - 1)\rho} \quad (\text{VII.35})$$

considerably. $\rho < 0$ for $\Delta \in [1, 2]$, as expected. For both the NN and the KDE, ρ exhibits the same damped oscillatory structure, as would be expected by the

normalisation imposed by the unfolding. ℓ_{KDE} is noticeably smaller than ℓ_{NN} . The contrast is less pronounced than in the perfect resolution case, yet still visible.

Moderate and large smearing ($\sigma_{\text{det}} \geq 0.5$): Correlations develop a broad plateau independent of method choice. In that regime N_{eff} collapses explaining why a naïve fit ignoring correlations would significantly misestimate errors.


The NN generally yields more long range, but slightly weaker correlations than the KDE at the same resolution, suggesting that a higher capacity estimator might be able to better absorb detector noise. However, the qualitative behaviour is identical; any non-zero smearing produces, long range weight correlations that violate the i.i.d. assumption built into standard unbinned likelihoods.

VII.G.3.ii Histogram covariance matrices.

Since most downstream analyses eventually bin the unfolded events, a histogram covariance matrix can be effective to visualise the correlations. The impact on such *binned* summaries can be visualised by filling a 40 bin histogram in the range $[-4, 4]$ for each pseudo-experiment and computing

$$\Sigma_{ab} = \mathbb{E} [(N_a - \bar{N}_a)(N_b - \bar{N}_b)]. \quad (\text{VII.36})$$

The matrices, shown in Figure VII.6, reinforce the observations from the correlation curves. Even at zero smearing, while KDE produces a nearly diagonal Σ , OMNIFOLD's weights exhibit small but noticable off diagonal elements. Growing the smearing leads to pronounced *anti-correlations* along the first off diagonal and coherent positive



(a) KDE-based OMNIFOLD: Covariance matrices exhibit concentrated correlation structures with significant values (> 0.1) confined primarily to nearest and next to nearest neighbour bins.


Figure VII.6: Covariance matrices of binned histograms constructed from unfolded distributions for $\sigma_{\text{det}} \in \{0.00, 0.25, 0.50, 0.75\}$. Each axis represents the covariance structure of a 40-bin histogram over $[-4, 4]$ in the observable space, with bin width $\Delta x = 0.2$. Matrix elements ρ_{ij} represent the correlation between bins i and j . **(a)** displays results from KDE and **(b)** shows NN results, both using OMNIFOLD. Both methods exhibit a clear progression from nearly diagonal matrices at perfect resolution to increasingly correlated structures as detector resolution degrades. The integrated correlation strength, sum of absolute off-diagonal elements, increases approximately quadratically with σ_{det} , reflecting the information loss that must be compensated through correlated adjustments. (Continued on next page)

correlations far from the diagonal. For $\sigma_{\text{det}} = 0.75$ the largest off diagonal coefficient reaches $|\rho_{ab}| \simeq 0.5$, echoing the plateau seen in $\rho(|\Delta x|)$. These patterns are almost identical for KDE and NN analyses.

Such strong off diagonal structure rigorously explains why a χ^2 fit that ignores covariances by using only the diagonal of Σ leads to significant miscoverage.

VII.G.3.iii Implications.

The correlation diagnostics presented here empirically validate the theoretical expectation that OMNIFOLD weights are correlated. Quantifying the strength and



(b) NN-based OMNIFOLD: Covariance matrices exhibit broad correlation bands extending ≥ 10 bins from the diagonal at $\sigma_{\text{det}} = 0.75$.

Figure VII.6: (Continued from previous page) These covariance matrices are essential inputs for rigorous statistical analysis of unfolded data, and must be propagated through any subsequent fits or hypothesis tests to obtain correct uncertainties. Figure produced by Owen Long [44].

range of those correlations enables a rough estimate of how badly an independence based error formula will fail. With these diagnostics in hand, one is now equipped to benchmark the statistical performance of binned and unbinned inference workflows on unbinned unfolded data.

VII.G.4 Unbinned unfolding: binned and unbinned inference.

Having established a fully binned baseline in Section VII.G.2, this section examines two inference workflows that utilise unbinned unfolding on the same Gaussian data. Both approaches start by unfolding the detector level data without binning, producing a set of weighted events at truth level. Because the deconvolution process induces non-negligible correlations between event weights as evidenced by the covariance patterns in Figures VII.5 and VII.6, these workflows differ in how they handle those correlations.

When performing binned inference on unbinned unfolded data, the unfolded weighted events are aggregated into a histogram, and a binned template fit is performed using a χ^2 statistic that includes the full binned covariance matrix. In this workflow the unfolded distribution is fit to the parametric Gaussian model with mean μ and variance σ^2 by integrating the model prediction over the histogram bins. The covariance matrix of the bins is estimated from an ensemble of repeated pseudo-experiments.

One then obtains best fit parameters by χ^2 minimisation and determines their uncertainties via the usual $\Delta\chi^2 = 1$ criterion. This approach is analogous to a standard binned analysis except that the data have been unfolded using an unbinned method. In this workflow offers a statistically rigorous method to propagate the full unfolding-induced covariance into the fit.

Alternatively, one could attempt to fit the parametric model directly to the weighted events using an unbinned maximum likelihood procedure that ignores inter-event correlations. One constructs the weighted negative log-likelihood,

$$\text{NLL}(\theta) = - \sum_{i=1}^N w_i \ln p(x_i | \theta), \quad (\text{VII.37})$$

where x_i and w_i are the kinematic value and weight of event i , and $p(x_i|\theta)$ is the model density for parameters $\theta = (\mu, \sigma^2)$.

This unbinned likelihood sum is maximised to obtain the best fit $\hat{\theta}$. The Hessian (curvature) of the NLL at the optimum provides an asymptotic error estimate for θ . Equivalently, one finds the $\Delta \ln \mathcal{L} = 0.5$ offset for a 1σ interval.

This workflow assumes statistical independence of events [530, 565, 666]. Thus, while method the latter method yields a fit and a formal error bar, these must be interpreted with caution since the underlying likelihood model is misspecified.

Both of the above approaches are applied to the same Gaussian datasets described earlier in Section VII.G.1. Unfolding is performed with the OMNIFOLD algorithm. In the one dimensional setting a KDE based implementation is also tested. For each resolution setting, the 500 pseudo-data samples are unfolded and then subjected to both the aforementioned inference procedures. This allows a comparison of the bias, uncertainty estimation, and coverage of the two workflows against the fully binned baseline.

VII.G.4.i Parameter bias.

Both unbinned workflows yield fitted parameter values that are consistent with those from the binned baseline. There is no appreciable bias introduced by the unbinned unfolding step or the choice of inference method. Figure VII.7 (bottom panels) shows the mean fitted μ and σ^2 as a function of detector smearing for each method. These results demonstrate the strength of unbinned methods when feasible, while validating that properly implemented binned unfolding with full covariance treatment can achieve comparable performance.

The bias panels (bottom) reveal that all methods maintain bias levels below 1% of the parameter values for $\sigma_{\text{det}} < 0.5$, with UBU exhibiting the most stable performance. At all detector resolutions, IBU shows comparable or larger bias in variance estimation than UBU, but the inferred parameters are still well within statistical uncertainties.




Figure VII.7: Comparison of inference uncertainty and bias for distribution parameters extracted from unfolded data as a function of σ_{det} . (Top left) uncertainty on μ , (top right) uncertainty on σ^2 , (bottom left) bias in μ estimation, and (bottom right) bias in σ^2 estimation. Three unfolding and inference approaches are compared: (i) Unbinned Unfolding (UBU) with unbinned maximum likelihood (ML) fitting shown as blue circles, (ii) binned Iterative Bayesian Unfolding (IBU) with χ^2 fitting using the full covariance matrix (orange triangles), and (iii) IBU with χ^2 fitting using only diagonal covariance elements (green squares). The blue shaded band in each panel represents $\pm$ RMS inference uncertainty from the UBU+ML approach. Figure created by Owen Long [44].

Indeed, approaches produce estimates of μ and σ^2 across the range of smearings within the statistical uncertainties (blue shaded).

Notably, the bias observed in the naive unbinned ML fit is the same as that in the proper χ^2 fit that accounts for the covariance of a given dataset. This can be explained as both fits ultimately maximising a likelihood, or minimising a χ^2 , to match the unfolded distribution to the model; if the model is correctly specified, the MLEs should coincide. Thus, unbinned unfolding does not itself induce a bias in the extracted physics parameters verifying the claim that the OMNIFOLD procedure correctly reproduces the shape of the true distribution on average. This is demonstrated by the agreement of fitted values with the true parameter shown by the horizontal lines in Figure VII.7 and the baseline results in Figure VII.4a.

VII.G.4.ii Uncertainty estimation and coverage.

In sharp contrast to the agreement in central values, the two workflows differ markedly in their reported uncertainties and the statistical coverage of those uncertainties. The binned χ^2 approach with full covariance proves to be consistent with the baseline in its uncertainty estimates. For each smearing level, the asymptotic 1σ errors obtained from the $\Delta\chi^2 = 1$ criterion agree well with the empirical spread (RMS) of the fitted parameters over the 500 pseudo-experiments. This is illustrated in the top panels of Figure VII.7: the curve corresponding to the full covariance χ^2 fit lies on top of the true uncertainty obtained from the pseudo data ensemble (star markers), for both μ (left) and σ^2 (right). In other words, binned inference on

unbinned data using a χ^2 fit that incorporates the full covariance matrix produces accurate confidence intervals that maintain the nominal coverage.

These uncertainty panels (top) demonstrate that an unbinned asymptotic maximum likelihood (ML) fit on unbinned unfolded (UBU) data consistently underestimates statistical uncertainties across all detector resolutions, and is most similar to the results of performing an asymptotic binned χ^2 fit on the same UBU results while ignoring off diagonal covariance element. The uncertainty degradation with increasing σ_{det} follows the expected $\sqrt{1 + \frac{\sigma_{\text{det}}^2}{\sigma_{\text{true}}^2}}$ scaling for both parameters.

This behaviour is as expected and desired. By incorporating the complete covariance matrix of the unfolded histogram, the binned asymptotic χ^2 fit on UBU results correctly accounts for the event correlations in the statistical inference, and precisely matches the RMS uncertainty computed numerically on UBU results.

In one sense, these results agree with the fully binned study. Recall that in the baseline binned analysis, using the full covariance matrix yielded uncertainty estimates consistent with the bootstrap pseudo-data spread, whereas using only diagonal uncertainties did not, as shown in Figure VII.3. Figure VII.7 explicitly confirms that this is true in the unbinned case too. In another sense though, the unbinned results invert the binned picture. Notice that if one performs a binned fit to IBU results but incorrectly ignores off diagonal bin correlations, the uncertainties are overestimated and the χ^2 fit yields unnecessarily large error bars, analogous to the diagonal-covariance points in Figure VII.3. This over-conservative result similarly arises from double-counting the anti-correlated fluctuations in each bin and leads to over-coverage.

By contrast, the naïve unbinned ML inference on the unbinned data fails undercovers. As shown in Figure VII.7 (top panels, orange points), the asymptotic errors reported by the unbinned likelihood fit are dramatically smaller than the true spread of the fit results, except in the trivial case of zero smearing. This suggests more strongly that it is difficult *a priori* to determine the sign of miscoverage due to ignored correlations.

Intriguingly, the naive unbinned error bars hardly change with detector resolution—the orange curve in Figure VII.7 is nearly flat—indicating that the fit is seemingly just as ‘precise’ for a very smeared dataset as for a perfect detector. This unphysical result could be a direct consequence of ignoring the event correlations. The maximum likelihood fit treats each weighted event as independent information, thereby overestimating the effective sample size. Intuitively, when events are strongly correlated, the true number of independent degrees of freedom is smaller than N ; but the naive NLL sum scales like N , yielding an misestimate of the error on $\hat{\theta} = (\mu, \sigma^2)$.

In this Gaussian experiment, the effect is severe even at moderate smearing. For instance, at $\sigma_{\text{det}} = 0.5$ the unbinned ML formula underestimates the uncertainty on μ by roughly a factor of two compared to the bootstrap truth, and at $\sigma_{\text{det}} = 0.75$ the discrepancy is even larger. This breakdown of the asymptotic approximation in the presence of inter-event correlations is the essential failure mode of the naive unbinned approach. It should be emphasised that when the detector resolution is perfect, the unfolding induces no correlations and indeed the unbinned ML errors do coincide with the correct uncertainties, and all methods become equivalent in this

limit. But for any non-zero smearing, the independence assumption is violated and the standard likelihood formulae no longer hold.

The practical implication of these findings is that naively applying unbinned inference to unfolded data can produce misleadingly constraints, even though the fit may appear to converge normally. In this study, the fully unbinned workflow would have erroneously suggested measurement insensitive to detector smearing, whereas in truth the uncertainties should grow with smearing, as correctly reflected by methods that account for correlations. This highlights the necessity of handling the event correlations in some way, potentially through the hybrid approach. By introducing a binning at the final inference stage and using a covariance matrix, one essentially restores statistical consistency, at limits the impact of binning artifacts. Such an approach largely retains the benefits of unbinned unfolding, with no information loss up to the point of inference, while yielding parameter uncertainties and coverage properties in line with a rigorous frequentist construction. Hence this ‘unbinned unfolding + binned fit’ should be more precise than a fully binned analysis.

For example, in Figure VII.7 the full covariance UBU results achieve a similar or smaller uncertainty than the traditional IBU based results at each smearing value. This suggests an advantage to delaying any binning as late in the analysis chain as possible, consistent with the intuition that using unbinned distributions throughout the unfolding can preserve more information for the final fit. One should note, however, that in the 1D case this advantage is modest. UBU and IBU approaches do not differ vastly in precision. In higher dimensions the unbinned approach is expected

to drastically outperform a binned analysis (since binning in many dimensions is impractical or introduces large discretisation errors) [130, 165, 250].

While the hybrid method provides a sound stopgap, it somewhat undermines the original motivation for unbinned methods, which aims to avoid histogramming altogether. The fully unbinned method, on the other hand, fully actualises the principle underlying the unbinned paradigm but lacks a statistical asymptotic formalism that accounts for correlations. The miscoverage observed for the naive unbinned fit underlines the need for correlation aware unbinned inference techniques. In other words, if one wishes to perform event level likelihood fits on unfolded data, one must either develop a formalism to incorporate the event-to-event covariance information into the likelihood or rely on numerical methods to calibrate the uncertainties. In this study, one does find that if one estimates uncertainties numerically through pseudo-experiments or bootstraps, the precision of the unbinned fit can be evaluated correctly. This is evidenced by the fact that the RMS spread of the naive ML fit outcomes (the blue line in Figure VII.7) does increase with smearing and matches the covariance-fit results.

Despite this, using brute force pseudo-experiments or bootstraps to determine errors is computationally expensive and may not always be feasible. Besides, errors computed through bootstrapping offer no analytical understanding of the uncertainty. Nonetheless, until a theoretical framework is developed to handle correlations at the likelihood level, any unbinned inference on unfolded events should be performed through bootstrapping.

In summary, the comparison of binned versus unbinned inference after unbinned unfolding reveals that both workflows yield unbiased parameter estimates, but only the approaches that account for correlations produce reliable uncertainties and coverage. A naïve approach of treating weighted unfolded events as independent leads to misestimated uncertainties and hence significant miscoverage. These results vividly demonstrate the statistical pitfalls of ignoring induced correlations, and they motivate the correlation-aware inference strategies discussed in the next section. As unbinned unfolding techniques become increasingly prevalent in extracting cross sections [13, 89, 129], developing a robust framework to correctly propagate uncertainties through the unbinned pipeline is essential. This study provides a first quantitative glimpse of the issue, showing that while unbinned unfolding can preserve accuracy and potentially improve precision, one must incorporate the full correlation structure to achieve valid statistical inference. This will be crucial for ensuring proper coverage and trustworthy uncertainty estimates in future high dimensional unfolding analyses.

VII.G.5 Extension to higher dimensions.

The diagnostics of Sections VII.G.3 and VII.G.4 established that even in one dimension the naïve unbinned likelihood approach seriously underestimates uncertainties once unfolding procedure induces correlated event weights. A natural next question is whether that pathology grows, diminishes, or saturates as one moves to multidifferential measurements. To answer this question, one can therefore repeat the Gaussian study in two, four and six dimensions, keeping the experimental setup identical, except for adjusting the numerical parameters that define the Gaussians.

Table VII.2: Gaussian model parameters for 2 dimensional unfolding validation studies. Generation parameters define the Generation distribution, Detector parameters specify the smearing, and Truth parameters define the Truth distribution. $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ list the component means and standard deviations, and $\boldsymbol{\rho}$ is the correlation matrix. The mismatch between Generation and Truth parameters tests the unfolding method.

	Generation	Detector	Truth
$\boldsymbol{\mu}$	$\begin{bmatrix} 0.0 \\ 1.0 \end{bmatrix}$		$\begin{bmatrix} 0.2 \\ 0.8 \end{bmatrix}$
$\boldsymbol{\sigma}$	$\begin{bmatrix} 1.0 \\ 1.5 \end{bmatrix}$	$\begin{bmatrix} 0.5 \\ 0.8 \end{bmatrix}$	$\begin{bmatrix} 0.9 \\ 1.3 \end{bmatrix}$
$\boldsymbol{\rho}$	$\begin{bmatrix} 1.0 & -0.6 \\ -0.6 & 1.0 \end{bmatrix}$		$\begin{bmatrix} 1.0 & -0.6 \\ -0.6 & 1.0 \end{bmatrix}$

The Gaussian parameters used in the 2, 4 and 6 dimensional studies are listed in Tables VII.2, VII.3 and VII.4 respectively. Detector resolutions are scaled coordinate wise so that the signal to noise ratio in every dimension matches the $d = 1$ baseline.

Table VII.3: Gaussian model parameters for 4 dimensional unfolding validation studies. See Table VII.2 for parameter definitions.

	Generation	Detector	Truth
μ	$\begin{bmatrix} 1.0 \\ 0.0 \\ -0.5 \\ 0.5 \end{bmatrix}$		$\begin{bmatrix} 0.8 \\ 0.1 \\ -0.6 \\ 0.7 \end{bmatrix}$
σ	$\begin{bmatrix} 1.0 \\ 0.7 \\ 1.1 \\ 0.8 \end{bmatrix}$	$\begin{bmatrix} 0.4 \\ 0.5 \\ 0.6 \\ 0.3 \end{bmatrix}$	$\begin{bmatrix} 0.8 \\ 0.6 \\ 1.0 \\ 0.6 \end{bmatrix}$
ρ	$\begin{bmatrix} 1.0 & 0.1 & -0.2 & 0.3 \\ 0.1 & 1.0 & 0.0 & 0.1 \\ -0.2 & 0.0 & 1.0 & 0.7 \\ 0.3 & 0.1 & 0.7 & 1.0 \end{bmatrix}$		$\begin{bmatrix} 1.0 & 0.0 & -0.3 & 0.4 \\ 0.0 & 1.0 & 0.2 & 0.0 \\ -0.3 & 0.2 & 1.0 & 0.5 \\ 0.4 & 0.0 & 0.5 & 1.0 \end{bmatrix}$

Table VII.4: Gaussian model parameters for 6 dimensional unfolding validation studies. See Table VII.2 for parameter definitions.

	Generation	Detector	Truth
μ	$\begin{bmatrix} 1.0 \\ 0.0 \\ -0.5 \\ 0.5 \\ -1.0 \\ 0.3 \end{bmatrix}$		$\begin{bmatrix} 0.8 \\ 0.1 \\ -0.6 \\ 0.7 \\ -0.8 \\ 0.1 \end{bmatrix}$
σ	$\begin{bmatrix} 1.0 \\ 0.7 \\ 1.1 \\ 0.8 \\ 1.2 \\ 1.4 \end{bmatrix}$	$\begin{bmatrix} 0.4 \\ 0.5 \\ 0.6 \\ 0.3 \\ 0.4 \\ 0.4 \end{bmatrix}$	$\begin{bmatrix} 0.8 \\ 0.6 \\ 1.0 \\ 0.6 \\ 1.0 \\ 1.1 \end{bmatrix}$
ρ	$\begin{bmatrix} 1.0 & 0.1 & 0.2 & -0.3 & 0.0 & 0.0 \\ 0.1 & 1.0 & 0.0 & -0.2 & 0.3 & 0.1 \\ 0.2 & 0.0 & 1.0 & 0.1 & -0.2 & 0.3 \\ -0.3 & -0.2 & 0.1 & 1.0 & 0.1 & 0.0 \\ 0.0 & 0.3 & -0.2 & 0.1 & 1.0 & 0.7 \\ 0.0 & 0.1 & 0.3 & 0.0 & 0.7 & 1.0 \end{bmatrix}$		$\begin{bmatrix} 1.0 & 0.0 & 0.2 & -0.2 & 0.1 & 0.0 \\ 0.0 & 1.0 & 0.0 & -0.1 & 0.2 & 0.0 \\ 0.2 & 0.0 & 1.0 & 0.0 & -0.3 & 0.4 \\ -0.2 & -0.1 & 0.0 & 1.0 & 0.2 & 0.0 \\ 0.1 & 0.2 & -0.3 & 0.2 & 1.0 & 0.5 \\ 0.0 & 0.0 & 0.4 & 0.0 & 0.5 & 1.0 \end{bmatrix}$




Figure VII.8: Comparison of empirical RMS parameter uncertainties with mean asymptotic uncertainty estimates for all model parameters across $d = 1, 2, 4$ and 6 . Each point represents a single parameter (mean or covariance matrix element) from the multivariate Gaussian. The blue circles denote components of the mean, and the orange diamonds denote components of the covariance. The grey dashed line indicates perfect agreement, while the green solid line shows the least squares fit for each dimension. The systematic deviation above the unit line demonstrates that asymptotic uncertainties consistently underestimate the true parameter uncertainties. The increasing scatter with dimension indicates heterogeneous behaviour across different parameter types, with off diagonal correlation parameters typically showing larger deviations than mean parameters. Figure created by Owen Long [44].

VII.G.5.i Evaluation.

For every fit parameter, that is, all μ_k and all unique σ_{kl}^2 , compute the empirical RMS is computed over 500 pseudo-experiments, and the average asymptotic error reported by the naïve unbinned likelihood Hessian. If the likelihood were well calibrated, these two quantities should coincide. Figure VII.8 plots RMS versus asymptotic error for all parameters in $d \in \{1, 2, 4, 6\}$. The dashed diagonal marks perfect agreement and the solid green line has a slope fixed to the mean of the RMS/asymptotic error ratio for that dimension.

In each case, the RMS uncertainty is higher than the asymptotic uncertainty. The ratio of the RMS uncertainty to the asymptotic uncertainty is roughly the same for

all of the model parameters with the ratio ranging from 1.18 to 1.28. The analytic error bars understate the true variance by a factor of about 3 to 4, consistent with the effective sample size argument [286, 574, 583], *vide* Section VII.E.

Figure VII.9 repeats the $d = 6$ study with detector resolution scaled by 0, 1, 2 and 3 relative to the σ_{det} in Table VII.4. The slope grows with the smearing factor, confirming that weight correlations, not intrinsic variance, drive the failure [238].

In high dimensions, the mapped weight vector $w(\mathbf{x})$ is a smooth function on a space where almost all pairs of points are distant. Hence the classifier must assign similar gradients to a larger neighbourhood, inflating correlations. [73] explains this phenomenon for KDEs, where excessively wide kernels force longer range weight correlations. The same effect applies to NNs as well because global normalisation $\sum_i w_i = N_{\text{MC}}$ adds an $\mathcal{O}(1/N)$ positive correlation to every pair of events [505].

In six dimensions the naïve asymptotic formula would underestimate statistical errors for several covariance elements, while the RMS spread shows true uncertainties. Either a binned-covariance approach or a correlation-aware unbinned likelihood [418] is necessary. Recent proposals such as Schrödinger bridge unfolding or low-rank covariance compression [415] offer interesting possibilities for exploration.




Figure VII.9: Comparison of empirical RMS uncertainties with asymptotic estimates for the six dimensional unfolding study for 1, 2, 3 and $4 \times \sigma_{\text{det}}$ listed in Table VII.4. The grey dashed line indicates perfect agreement, while green line shows least squares fits for each smearing level. The RMS/asymptotic ratio increases monotonically with detector smearing, from 1.15 at perfect resolution to 1.45 at $3 \times$ smearing, demonstrating that asymptotic uncertainties become increasingly optimistic as detector effects worsen. This degradation arises from increased correlations induced by the unfolding process. The vertical spread of points also increases with smearing, indicating that the reliability of asymptotic estimates becomes more parameter dependent. Correlation parameters (off-diagonal elements of ρ) show particularly poor estimation at high smearing, with some parameters exhibiting RMS/asymptotic ratios exceeding 2.0. Careful uncertainty validation is essential when applying unbinned unfolding to experiments with significant detector effects. Figure created by Owen Long [44].

VII.H Conclusions and outlook.

This chapter presented a detailed study of parameter inference performed on an unbinned unfolded dataset, using a controlled Gaussian simulation. Employing a simplified scenario where the true distribution and detector response are known analytically facilitates isolating and rigorously examining the statistical subtleties introduced by unbinned unfolding.

The Gaussian studies demonstrated that preserving event level information until the final fit can indeed confer a tangible advantage: an unbinned unfolding followed by binned parameter estimation was observed to outperform a fully binned analysis in terms of fit precision. This validates the intuitive principle that one should delay information reduction as much as possible in the analysis chain, thereby maximising the use of available information. Importantly, however, the findings also exposed critical caveats that must temper this optimism, especially when interpreting unbinned results with standard inference techniques.

A central observation is that the output of an unfolding procedure violates the usual assumption of statistical independence among events. In these studies, the iterative OMNIFOLD procedure produces a weighted set of ‘unfolded’ events that are correlated with one another. A naïve unbinned likelihood based inference, which treats the unfolded events as if they were independent observations, fails to yield reliable uncertainty estimates. In particular, when event-to-event correlations are ignored, the standard asymptotic formulae for parameter uncertainties²³ become invalid.

²³Derived from a Fisher information or Hessian approximation.

Concretely, the naïve application of an unbinned maximum likelihood fit significantly underestimates the true uncertainty on the fitted parameters in Gaussian examples. This underestimation was evidenced by comparing the analytic errors to the empirical spread of fit results across many pseudo-experiments. The analytic errors, assuming independence, were significantly smaller than the actual RMS of the fitted parameters.

Such a discrepancy is a direct manifestation of ignoring the induced correlations and a clear warning that conventional unbinned inference tools cannot be directly applied to unbinned unfolded data. In short, the lack of statistical independence introduced by unbinned unfolding can invalidate classical error estimates. Thus the entire inferential procedure must be handled with care.

On the positive side, the investigation also highlighted that correlation aware approaches can effectively restore valid inference, albeit with practical limitations. When one incorporates the inter-event correlations into the analysis, one finds that the derived uncertainties align with the true performance of the fit. In the Gaussian studies, the numerically estimated uncertainties, obtained from the spread of outcomes over many trials, agree well with a covariance informed analytical approach, confirming that the primary cause of the asymptotic formula breakdown was the neglect of correlations. This serves as an important proof of principle. If one properly accounts for the covariance between events²⁴ in the unbinned dataset, inference methods can yield correct confidence intervals and hypothesis tests even in the unbinned paradigm. However, it should also be emphasised that currently available unbinned covariance aware treatments are computationally challenging to scale.

²⁴Or as the case may be, between event weights.

In a binned analysis the covariance matrix is of manageable size, but in an unbinned analysis even if a suitable covariance kernel were developed, one in principle faces an $N \times N$ covariance kernel among N events or event weights, which is intractable for the large event samples typical of HEP data. The use of bootstrap ensembles to numerically evaluate uncertainties, while effective for a toy study, is expensive for high dimensional or high multiplicity data. Thus, the practical implementation of correlation aware unbinned inference encounters serious scaling limitations, underscoring the need for new strategies to handle or approximate these large covariance structures.

Crucially, the results presented here should be viewed as a cautionary case study rather than a definitive statement on all unbinned analyses. They were obtained in an idealised Gaussian context, with complete knowledge of the generative process. This controlled setup allows one to rigorously identify potential pitfalls and verify proposed solutions, but it also means that the qualitative behaviours observed might not universally translate to real data. In particular, the observation that ignoring correlations had little impact on the central values and empirical precision of the fit in the Gaussian example may be a fortunate consequence of the symmetry and simplicity of that example. One should not assume this will hold in general.

Therefore, while this study provides a rigorous proof of principle and important warnings, further work is required to generalise of these findings. The takeaway is the set of statistical issues that could plausibly afflict unbinned cross section measurements, and the verification of their existence in a controlled setting. It now remains to investigate how these issues manifest in realistic analyses.

Looking ahead, this work motivates several important future directions for both methodology and application. An immediate next step is to extend these diagnostic studies to real HEP observables and data. It will be invaluable to apply similar techniques (e.g. pseudodata experiments, bootstrap uncertainty evaluations, and covariance measurements) to a realistic experimental unfolding scenario, for example, a differential cross section measurement published with unbinned results, to evaluate the size of event level correlations and to quantify their impact on parameter fits. Such studies on real or high fidelity simulated data would confirm whether the cautionary lessons from the Gaussian example are broadly applicable, and could reveal any additional complications arising from more complex data structure or detector effects.

Another critical line of research is to develop new statistical frameworks for unbinned inference that explicitly account for event correlations. This could involve formulating modified likelihood functions or test statistics that include correlation terms, or designing hybrid approaches that retain the advantages of unbinned data while imposing an effective covariance model. The ultimate goal would be to have an unbinned inference methodology that is correlation aware by construction, obviating the need for *ad hoc* binning or numerical uncertainty estimates.

In tandem with this, there is a clear need for scalable covariance approximation techniques. Instead of attempting to handle a full covariance kernel of enormous dimension, one might seek low rank representations, clustering of events into groups with approximate independence, or other dimensionality reduction methods to capture the dominant correlation effects without the full cost. Research into such approximations, potentially informed by the machine learning algorithms used in unbinned

unfolding, will be essential to make correlation aware inference feasible for large datasets.

Finally, further explorations of the interplay between machine learning regularisation and inference accuracy could be informative. A deeper understanding of how the ML aspects affect downstream inference, for instance, whether a stronger regularisation might reduce variance at the cost of introducing bias or correlations, would be extremely valuable. By addressing these open questions, the community can build on the foundation laid in this work.

In summary, this chapter has established a foundational understanding of unbinned inference on correlated data, highlighting both the potential benefits of fully unbinned analyses and the new statistical challenges they pose. The Gaussian studies act as a proof of principle, rigorously demonstrating that unbinned unfolding methods can be combined with parameter fits, but also warning that ignoring induced correlations can invalidate conventional uncertainty estimates. These conclusions, drawn in a simplified setting, strongly motivate the development of improved tools and methods before applying unbinned inference to precision measurements. As the field moves toward ever more complex and high dimensional data analyses, the insights gained here will guide the creation of a robust statistical framework for unbinned cross section measurements, one that maximises information usage while properly accounting for the complex correlations introduced by the unfolding process. The lessons of this chapter therefore serve as both a caution and a call to action, laying the groundwork for more reliable unbinned inference techniques in high energy physics.

Chapter VIII

Symmetries in data: connections to unfolding challenges.

VIII.A Symmetries and unfolding.	490
VIII.A.1 The complementary nature of symmetry discovery and unfolding.	491
VIII.A.2 Symmetry aware cross sections.	493
VIII.B Formalism and importance.	500
VIII.B.1 Fundamental symmetries in HEP.	503
VIII.B.2 Symmetries in detector response functions	511
VIII.B.3 How symmetries manifest in measured cross sections.	521
VIII.B.4 Challenges in identifying symmetries from noisy data.	525
VIII.C Statistical definition of dataset symmetries.	529
VIII.C.1 Distinction between point and dataset symmetries.	530

VIII.C.2 Inertial reference densities and their theoretical role.	531
VIII.C.3 Inertial densities and relative symmetry.	533
VIII.C.4 Applications to particle physics.	536
VIII.D SYMMETRYGAN: Discovering symmetries with adversarial learning.	538
VIII.D.1 The SYMMETRYGAN architecture.	539
VIII.D.2 Machine learning with inertial restrictions.	541
VIII.D.3 Deep learning implementation details.	545
VIII.D.4 Verification	547
VIII.D.5 Other symmetry discovery methods.	547
VIII.E Empirical experiments.	550
VIII.E.1 Gaussian experiments.	551
VIII.E.2 Particle physics experiments.	558
VIII.E.3 Interpreting discovered symmetries.	566
VIII.F Towards symmetry inference.	569
VIII.G Symmetry informed unfolding	571
VIII.G.1 Symmetry preserving neural network architectures.	573
VIII.G.2 Reducing dimensionality through symmetry identification. .	574
VIII.G.3 Hidden symmetries and emergent simplicity.	576
VIII.G.4 Applications for unfolding jet substructure variables.	578
VIII.H Improved precision using symmetries	580
VIII.H.1 Data augmentation using discovered symmetries.	581
VIII.H.2 Symmetry constrained unfolding.	583
VIII.I Conclusion.	585

<i>Chapter VIII. Symmetries in data: connections to unfolding challenges.</i>	489

VIII.I.1 Beyond linear symmetries.	586

VIII.I.2 Approximate symmetries and symmetry breaking.	587

VIII.I.3 A unified framework.	589

VIII.A Symmetries and unfolding.

VIII.A.1 The complementary nature of symmetry discovery and un-	
folding.	491
VIII.A.2 Symmetry aware cross sections.	493

VIII.A.1 The complementary nature of symmetry discovery and unfolding.

Definition 49 (Symmetry discovery). *Symmetry discovery* is the process of identifying invariant transformations of a dataset, that is to say, operations under which the probability distribution of the dataset remains unchanged in a statistical sense [76, 138, 229].

Unfolding, as discussed, refers to the inverse problem of inferring underlying truth level distributions from observed detector level data, accounting for distortions due to limited resolution and acceptance. At first glance, these two tasks appear distinct; one concerns recovering numerical distributions, while the other uncovers structural invariances. However, symmetry discovery and unfolding are in fact complementary facets of data driven inference, and integrating the two can yield deeper insights and improved measurements.

From a conceptual standpoint, both tasks share the common goal of revealing hidden truth from observed data. Unfolding endeavours to undo the convolution with the detector response and recover the true differential cross section or underlying distribution that generated the measurements. Symmetry discovery seeks to reveal underlying structures or invariances in the data—patterns that persist under transformations, reflecting fundamental symmetries of the physical process or the measurement apparatus.

In practice, these goals can be intertwined. If one discovers a symmetry in the dataset, that knowledge can constrain the unfolding procedure by reducing the

effective degrees of freedom in the solution space. Conversely, a properly unfolded distribution is expected to manifest latent symmetries that may have been obscured by detector effects in the raw data. Thus, identifying a symmetry and unfolding a distribution reinforce one another. The former provides a guiding principle or constraint for the latter, while the latter provides a cleaner canvas on which the former can be observed.

Imposing a symmetry as a prior constraint in unfolding can be seen as a form of physically motivated regularisation. For example, architectures that conserve the 4-momentum or enforce Lorentz invariance by design, as discussed in Chapter III, effectively impose such constraints, narrowing the set of viable solutions. By restricting solutions to those that respect a discovered or expected invariance, one reduces the space of admissible unfolded distributions to those that are physically plausible, which leads to improved stability and fidelity of the results [99].

This principle has been implicitly utilised in classical unfolding and simulation based calibration. For example, one could assume certain symmetries such as isotropy or detector uniformity when designing the response matrix or when combining symmetric regions of phase space to reduce uncertainties. With a data driven symmetry discovery tool in hand, one need not rely solely on presumed symmetries. Instead, one can verify them empirically or even discover unexpected invariances. In turn, these empirically verified symmetries can be fed back into the inference pipeline to sharpen measurements, such that a discovered symmetry can inform the unfolding algorithm so that the final measured distribution upholds the invariance.

It is instructive to compare symmetry discovery and unfolding side by side to appreciate their complementary roles. Table VIII.1 summarises the differences and points of contact between the two.

While unfolding typically requires an explicit model of the measurement process¹ and often relies on supervised learning or iterative inversion techniques, symmetry discovery can be pursued in an unsupervised manner, requiring only the dataset and a class of transformations to probe. The outcome of unfolding is a corrected distribution intended for direct physical interpretation. The outcome of symmetry discovery is a characterisation of invariances, a set of transformations T such that the dataset's distribution is invariant under T within statistical uncertainties. These outcomes are different in nature, but they are mutually beneficial.

Knowledge of invariant structure can guide numerical inference, and conversely, obtaining a more accurate numerical distribution makes it easier to discern subtle invariant patterns. In essence, symmetry discovery and unfolding can form a feedback loop in the broader endeavour of measurement and inference, where each can improve the other.

VIII.A.2 Symmetry aware cross sections.

Since differential cross section measurements lie at the core of particle physics, achieving high precision in these measurements is essential, as even subtle deviations between the measured spectra and theory predictions can signal new physics or the

¹E.g., a response matrix or a parametrised detector simulation.

need for refined models. In this context, symmetries play a pivotal role in both the design and interpretation of cross section measurements.

Many physical processes come with known symmetry expectations. For instance, in proton–proton collisions producing particle jets, one expects azimuthal symmetry about the beam axis, i.e. the physics is invariant under rotations in the plane perpendicular to the beam. Consequently, the differential cross section should, after correcting for detector non-uniformities, be independent of the absolute azimuthal angle ϕ of a jet or dijet system [37, 80, 197, 597].

Likewise, for processes initiated by identical colliding particles, one often anticipates a symmetry between forward and backward directions. In a symmetric proton–proton collider, this implies that the rapidity distribution of a centrally produced system ² should be symmetric about zero rapidity [323, 425, 579]. Such symmetry means that the cross section for producing a system at rapidity $+y$ is the same as at $-y$, all else being equal. If a measured differential cross section exhibits a significant asymmetry in these variables after unfolding and acceptance corrections, it would either indicate a previously unaccounted detector bias or hint at a physical effect, both of which are of interest to investigate.

²E.g., dijet pair.

Table VIII.1: Relationship between unfolding and symmetry discovery approaches in HEP. While unfolding aims to correct detector effects to recover true physical distributions, symmetry discovery seeks to identify invariance properties directly from data. These complementary techniques can be combined synergistically: unfolding provides cleaner distributions where symmetries become more apparent, while discovered symmetries can constrain and regularise the unfolding process. This bidirectional relationship enables more robust extraction of physical information from experimental data, particularly in scenarios where detector effects might obscure underlying symmetries or where symmetry constraints can help guide unfolding.

Aspect	Unfolding	Symmetry discovery
Goal	Reconstruct Truth from Data by inverting detector response	Identify transformation groups under which distributions remain invariant
Inputs	• Measured data • Detector response matrix • Regularisation scheme	• Measured or unfolded data • Parametrised groups • Invariance test statistics
Outputs	• Differential cross sections • Probability density • Uncertainties/correlations	• Symmetry generators • Invariance confidence levels • Conserved quantities

Table VIII.1 *continued from previous page*

Aspect	Unfolding	Symmetry discovery
Basis	Incorporates domain knowledge through priors/constraints to regularise ill posed inverse problem	Tests generic transformation classes allowing data driven discovery of invariances
Approach	Statistical inference problem. Invert $p(x) = \int r(x z)p(z)dz$	Hypothesis testing to evaluate $p(x) \stackrel{?}{=} p(g(x)) \cdot g'(x) $
Benefit	Produces distributions where true physical symmetries manifest more clearly, enabling validation of theoretical predictions and discovery of emergent invariances	Discovered symmetries constrain unfolding solution space, reducing regularisation dependence and improving stability of deconvolution

Being symmetry aware in a measurement can mean two things. First, verifying that expected symmetries are indeed present, within uncertainties, in the data, and second, leveraging those symmetries to improve the measurement. On the verification side, symmetry considerations provide valuable consistency checks. Experiments can test whether their unfolded distributions respect fundamental symmetries, and a failure to observe an expected symmetry is a red flag, prompting scrutiny of systematic effects or potential new physics contributions.

On the other hand, when a symmetry is confirmed, one can exploit it to gain statistical and systematic advantages. For example, if a distribution is believed

to be symmetric in a certain variable, one can “augment” or combine data from symmetric regions, effectively doubling the effective statistics for that distribution. A common practice is to report differential cross sections as a function of $|y|$ or other symmetry reduced variables, which assumes $y \leftrightarrow -y$ symmetry and thereby reduces statistical fluctuations [458, 511, 532]. By incorporating symmetry in this manner, uncertainties can be reduced and the measurement becomes more robust against localised fluctuations.

Symmetry aware analysis also serves to impose physically motivated constraints that guard against overfitting noise in the unfolding process. If one has discovered that a distribution must be invariant under a transformation,³ imposing this invariance in the unfolding procedure will tie neural network parameters or bin values together that would otherwise float independently. This effectively decreases the number of free parameters describing the unfolded result, acting as a regulariser that prefers solutions consistent with the symmetry. The net effect is an improvement in the precision of the measured cross section and a reduction in spurious oscillatory features that might arise from statistical fluctuations. Moreover, by reducing the dependence on bins with low occupancy (because they are combined with their symmetric counterparts), binned symmetry aware unfolding can also mitigate the impact of detector acceptance edges or inefficiencies in specific regions of phase space.

It is important to note, however, that any symmetry based constraint should be applied with careful consideration. One must ensure that the symmetry is either theoretically well founded or empirically validated, lest one impose a false invariance

³E.g., rotating the entire event by some angle, or exchanging two identical particles in the final state.

and obscure a genuine asymmetry. This caution further motivates the need for data driven symmetry discovery and validation tools.

One can use methods like SYMMETRYGAN [225] to verify whether the data uphold the symmetry to a high degree of confidence. Only then would one proceed to incorporate that symmetry into the unfolding process or in the presentation of results. Thus symmetry aware differential cross section measurements can harness known, or discovered invariances to enhance precision and reliability, while simultaneously providing a framework to detect symmetry violations that could point to new phenomena.

The research presented in this chapter builds upon the foundations laid in earlier chapters of this dissertation, extending the paradigm of symmetry utilisation in the context of measurement and unfolding. Chapter II, in its survey of existing techniques, provides a statistical foundation for incorporating known symmetries into data analyses. It also highlights the challenges such approaches face, such as the need for careful validation of assumptions. Chapter III includes references to how *a priori* known symmetries can be hard coded into machine learning models, introducing Lorentz group equivariant networks that guarantee Lorentz invariance in the unfolding of particle physics data [314]. The inclusion of symmetry constraints in unfolding models described in Chapters IV to VI through preserving physical invariants like momentum or charge conservation in the generative model could help reduce the solution space and lead to more physically plausible unfolded results.

All of these strategies rely on prior knowledge of the symmetry. This chapter shifts to an data driven approach that provides a novel, flexible, and fully differentiable deep

learning based method to discover symmetries directly from data using adversarial learning, which then might allow one to leverage those discovered symmetries to inform the unfolding process. This perspective emphasises the overarching theme of the dissertation, the interplay of measurement and inference, by using data driven insights in the form of symmetry discovery to enhance the core measurement task of differential cross section unfolding.

The remainder of this chapter is organised as follows. Sections VIII.B and VIII.C introduce the formal statistical definition of a dataset symmetry, addressing subtleties like Jacobian volume effects via the concept of an inertial reference density. They also provides a brief overview of the symmetries most relevant to HEP. Section VIII.D presents the SYMMETRYGAN framework, which employs a generative adversarial network to automatically learn symmetry transformations from data. Section VIII.E validates this approach on illustrative examples and then applies SYMMETRYGAN to simulated dijet events, demonstrating how it can uncover physically meaningful symmetries in collider data. Section VIII.G discusses how the discovered symmetry information can be used to constrain unfolding problems: we outline methods to incorporate symmetry constraints into the unfolding procedure to reduce uncertainties and bias. In Section VIII.H is a discussion of potential symmetry aware unfolding methodologies and how enforcing the symmetries identified by SYMMETRYGAN can improve the precision of differential cross section measurements. Finally, the chapter concludes by highlighting how the insights gained here connect back to the broader narrative of the dissertation, reinforcing the benefits of combining machine learning driven discovery with principled measurement techniques.

VIII.B Formalism and importance.

VIII.B.1 Fundamental symmetries in HEP.	503
VIII.B.2 Symmetries in detector response functions	511
VIII.B.2.i Spatial uniformity and rotational symmetry.	512
VIII.B.2.ii Polar coverage and boost invariance.	513
VIII.B.2.iii Resolution effects and approximate invariance.	514
VIII.B.2.iv Mirror and charge symmetry.	516
VIII.B.2.v Permutation symmetry and identical particles.	517
VIII.B.3 How symmetries manifest in measured cross sections.	521
VIII.B.3.i Exact symmetries and flat distributions.	521
VIII.B.3.ii Symmetries in kinematic shapes.	523
VIII.B.3.iii Interplay of physical and detector symmetries.	524
VIII.B.4 Challenges in identifying symmetries from noisy data.	525
VIII.B.4.i Dimensionality challenges.	527

In physics and statistics, a symmetry refers to an invariance of a system or dataset under a well-defined transformation.

Definition 50 (Symmetry under a group.). Let G be a group of transformations (continuous or discrete) acting on a space of states or observations X . A physical system or probability distribution is *symmetric under G* if applying any transformation $g \in G$ leaves the relevant observables unchanged.

Definition 51 (Symmetry of a function). In group theoretic terms, a *function* $f : X \rightarrow Y$ is *symmetric* under a group G there exists a group action of G on X

$$\begin{aligned} \Phi : G \times X &\rightarrow X \\ (g, x) &\mapsto g(x) \end{aligned} \tag{VIII.1}$$

such that for every $x \in X$ and every $g \in G$, the value of the function $f(x) = f(g(x))$.

Definition 52 (Symmetry of a probability density function). If $p(x)$ denotes a probability density on X , a group G is a *symmetry of p* if

$$\forall g \in G \int p(x) \, dx = \int p(g(x)) \, d(g(x)). \tag{VIII.2}$$

In measure theoretic language, a symmetry corresponds to the existence of an invariant measure.

Definition 53 (Invariant measure). For every measure space (X, Σ, μ) and for every group G acting on X through a Σ -measurable action Φ , μ is called an *invariant measure under G* if $\forall (g, A) \in G \times \Sigma$

1. $\mu(g(A)) = \mu(A)$.

$$2. \mu(g(A)) = \mu(A)$$

Equivalently, using the pushforward $g_{\#}(\mu(A)) := \mu(g^{-1}(A))$, μ is defined to be invariant under G if

$$\forall (A, g) \in \Sigma \times G \quad g_{\#}(\mu(A)) = \mu(A). \quad (\text{VIII.3})$$

These definitions encompass both continuous symmetries⁴ and discrete symmetries.⁵

Symmetry principles lie at the heart of modern particle physics and strongly influence every aspect of the field. At the theoretical level, fundamental symmetries constrain the form of physical laws and often correspond to conserved quantities or selection rules. At the data level, symmetries, and their breaking, shape the distributions of observed events and can be leveraged for more efficient data analysis. In the context of HEP experiments, many observables are governed by symmetries of the underlying theory as well as symmetries introduced by the detector and measurement process. It is therefore crucial to articulate how these symmetries operate both in ideal physics scenarios and in real observations.

This section provides a rigorous overview of symmetries relevant to particle physics and measurements. It begins with the fundamental symmetries in particle physics that underlie observable phenomena in Section VIII.B.1. Section VIII.B.2 discusses symmetries in detector response functions and how the measurement apparatus can preserve or violate underlying invariances. Next, Section VIII.B.3 examines how symmetries manifest in measured cross sections and data distributions, clarifying the translation from physical symmetry to statistical patterns in experimental histograms.

⁴Lie groups, such as rotations depending on a continuous angle parameter.

⁵Groups constructed as Jordan–Hölder extensions [jordan_traite_nodate., 673, 799] e.g., a mirror reflection or a permutation of identical objects.

Finally, Section VIII.B.4 highlights the challenges in identifying symmetries from noisy data, setting the stage for data-driven symmetry discovery techniques. This foundation will be essential for later sections that introduce methods like SYMMETRYGAN for learning symmetries from data.

VIII.B.1 Fundamental symmetries in HEP.

Particle physics is built upon a framework of symmetries that determine the allowed forms of interactions and the conservation laws observed in experiments. Spacetime symmetries, in particular subgroups of the Poincaré group, are foundational. The Poincaré group includes continuous Lorentz invariance (rotations and boosts) and translations (in spacetime). Lorentz invariance implies that the laws of physics take the same form in any inertial reference frame. Equivalently, physical observables can be expressed in terms of Lorentz invariant quantities⁶ that remain unchanged under boosts or rotations.

For example, the Mandelstam variables s , t , u in a scattering process or the decay angle distribution in a particle's rest frame are formulated to respect Lorentz symmetry. In practice, exact Lorentz invariance means there is no preferred direction or absolute velocity in the underlying theory. A given interaction process should yield identical outcomes whether the laboratory frame is, say, Earth bound or boosted to a constant velocity with respect to the Earth. As a consequence of Lorentz symmetry and spatial isotropy, angular momentum and linear momentum are conserved in

⁶E.g., invariant masses, angles, and dimensionless ratios.

isolated systems⁷. Time translation symmetry leads to energy conservation, ensuring that system's total energy and the collision centre of mass energy are fixed constants of motion. These spacetime symmetries are exact symmetries of all known fundamental interactions and provide the basis for defining covariant formalisms in quantum field theory.

Beyond spacetime, the internal symmetries of the Standard Model dictate the spectrum of particles and their interactions. Chief among these is the gauge symmetry group $SU(3)_C \times SU(2)_L \times U(1)_Y$, which defines quantum chromodynamics and electroweak theory. Gauge symmetries are local symmetries that entail the introduction of gauge bosons. Although these are internal symmetries rather than symmetries of observable spacetime, they have observable consequences such as electric charge conservation, associated with $U(1)_Y$ hypercharge symmetry, and the existence of multiple particle generations. The gauge symmetries of the Standard Model are spontaneously broken in certain cases.⁸ However, even broken symmetries leave remnant effects, such as the custodial symmetry in the Higgs sector or approximate conservation of isospin in QCD. These internal symmetries set selection rules. Processes that violate gauge charge conservation are forbidden and decays proceed only via symmetric channels.

Alongside continuous symmetries, several discrete symmetries play a crucial role in particle physics. The most prominent are C (charge conjugation, exchanging particles with their antiparticles), P (parity, spatial inversion or mirror reflection), and T (time

⁷Noether's theorem associates these conservations with rotational and translational symmetry, respectively [2].

⁸One of the most notable examples is the electroweak $SU(2)_L \times U(1)_Y$ breaking to $U(1)_{\text{EM}}$ via the Higgs mechanism, which introduces masses for the $W^\pm$ and Z bosons and differentiates the electromagnetic and weak interactions.

reversal). Each of these can be considered a transformation that might leave the fundamental laws invariant.

In the Standard Model, CP symmetry is approximately a symmetry of electromagnetic and strong interactions, but notably broken in weak interactions. This manifests as differences in the behaviour of matter and antimatter. The most notable instance of this is the well known CP violation in neutral kaon and B -meson decays. Violation of CP means those processes occur at different rates or with different phase relationships than their CP-mirrored counterparts [694]. If CP were an exact symmetry of the dynamics, one would expect, for example, the angular distribution of decay products in a mirror reflected process, swapping particles for antiparticles, to be identical to the original. The observed deviations are vital clues about the underlying physics.

Parity by itself is also violated maximally in the weak interaction.⁹ On the other hand, the strong and electromagnetic interactions conserve parity, so for many processes, especially at high energies where electroweak effects are subdominant, it is a good symmetry. A process governed by QCD, like multijet production, should occur equally in a configuration and its mirror reflected image, unless the experimental setup selects a handedness. Charge conjugation is likewise not a symmetry of the full Standard Model, since, for example, there are no right handed neutrinos to pair with left handed ones under C, but for purely electromagnetic processes C symmetry implies, that producing a negatively charged particle is as likely as producing the corresponding positively charged antiparticle under equivalent conditions.

⁹Classic examples are the left handed nature of neutrinos and the parity asymmetric angular distribution of electrons in polarised ^{60}Co beta decay [773].

Importantly, the combination CPT is believed to be an exact symmetry of local quantum field theory. CPT symmetry implies, for instance, that particle and antiparticle masses and lifetimes are exactly equal [680]. While CPT is typically not directly tested by single distribution symmetries in HEP experiments, it provides a fundamental consistency check on any observed CP or T violation. Table VIII.2 summarises these fundamental symmetries, their group theoretic character, and their status in the Standard Model.

Table VIII.2: Summary of fundamental symmetries relevant to particle physics, their group theoretic structure, and experimental signatures. The table includes spacetime symmetries, gauge symmetries, discrete symmetries, and quantum statistical symmetries. “Charge” refers broadly to conserved quantities arising from continuous symmetries via Noether’s theorem or to quantum numbers constrained by discrete symmetries. The ‘Status’ column indicates whether each symmetry is exact, approximate, or explicitly broken within the Standard Model. Observable signatures provide experimental handles for testing these symmetries. The hierarchy of symmetry breaking guides the search for BSM physics through precision tests of invariance.

Symmetry	Group	Charge	Signatures	Status
Lorentz	$SO(3, 1)$	$x^{[\mu, p^\nu]}$	Invariant mass; angular distributions	Exact [577]
Translation	$\mathbb{R}^{1,3}$	p^μ	Missing momentum; vertex momentum conservation	Exact [639]
Gauge	$SU(3)_C$ $\times SU(2)_L$ $\times U(1)_Y$	Colour isospin	Jet colour flow patterns; W boson charge asymmetry;	Exact [700]

Continued on next page

Table VIII.2 continued from previous page

Symmetry	Group	Charge	Signatures	Status
Electroweak	$\langle H \rangle \neq 0$	M_W, M_Z	$\rho = M_W^2 / (M_Z^2 \cos^2 \theta_W)$	Broken [765]
Charge (C)	$\mathbb{Z}_2 : \psi \rightarrow \mathcal{C}\bar{\psi}^T$	$\alpha_{q\pm} \leftrightarrow \alpha_{\bar{q}\mp}$	e^+/e^- production ratios; $\pi^0 \rightarrow \gamma\gamma, \pi^0 \not\rightarrow 3\gamma$	QED/QCD, not weak [773]
Parity (P)	$\mathbb{Z}_2 : \vec{x} \rightarrow -\vec{x}$	$\eta_P = \pm 1$	Neutrino helicity; asymmetric β decay	QED/QCD, not weak [776]
CP	$\mathbb{Z}_2 \times \mathbb{Z}_2$	δ_{CP}	$B^0\text{-}\bar{B}^0$ mixing asymmetry; kaon decay: $\epsilon_K \sim 10^{-3}$	Violated [633] $\delta_{CP} \approx 68 \text{ deg}$
Time (T)	$\mathbb{Z}_2 : t \rightarrow -t$	$\mathcal{M}$	Electric dipole moments; $K_L \rightarrow \pi^+\pi^-e^+e^-$ decay	Violated [772] (CPT thm)

Continued on next page

Table VIII.2 continued from previous page

Symmetry	Group	Charge	Signatures	Status
CPT	$\mathbb{Z}_2$ (anti-unitary)	$m, \tau_{q\pm} = m, \tau_{\bar{q}\mp}$	$ m_p - m_{\bar{p}} /m_p < 10^{-10}$ $ \tau_\mu - \tau_{\bar{\mu}} /\tau_\mu < 10^{-5}$	Exact [725]
Permutation	S_n/A_n	Bose/Fermi stats	HBT correlations in $\pi^+\pi^+$; Pauli exclusion in spectra	Exact [787]

Beyond the Standard Model's built in symmetries, there are approximate global symmetries that often prove useful in particle physics. Examples include isospin symmetry, an $SU(2)$ symmetry treating up and down quarks as identical in the limit of equal masses, and flavour symmetries, like the $SU(3)$ of the light uds quarks, both of which are not exact, but underlie patterns in hadron production and decay. For instance, isospin symmetry implies that processes differing only by swapping an up quark with a down quark, such as producing a proton versus a neutron, have nearly equal cross sections, up to corrections from the up–down mass difference or electromagnetic effects.

Similarly, the universality of physical laws under interchange of identical particles leads to permutation symmetry. If two particles of the same type appear in a final state, the probability distribution is invariant under exchanging them. In quantum theories this is enforced by (anti)symmetrisation of identical particle states. In HEP observables, permutation symmetry means that one cannot physically distinguish, say, which of two identical jets in an event is ‘jet 1’ or ‘jet 2’—any labelling is arbitrary and the underlying physics treats the two jets on equal footing. When calculating cross sections, this symmetry is accounted for by dividing by the a symmetry factor to avoid over counting identical configurations.

In practical analyses analysis one often has to impose an ordering, such as ‘leading’ and ‘subleading’ jet by momentum, for convenience, but the fundamental permutation invariance implies that any physical conclusion should not depend on this arbitrary ordering.

In summary, fundamental symmetries, Poincaré (Lorentz and translations), gauge invariances, discrete symmetries like CPT, and permutation invariance for identical particles, provide a set of invariance principles for particle interactions. These symmetries constrain the form of theoretical cross sections and transition rates.

Many measurable quantities in experiments such as cross sections, angular distributions, *et cetera*, either reflect these symmetries, when they hold, or provide avenues to detect symmetry breaking when deviations are observed. However, the symmetries of nature at the fundamental level are not always manifest in what detectors actually record. We next turn to how the detector response and measurement process can modify or obscure these symmetries.

VIII.B.2 Symmetries in detector response functions

The detector response function $r(x \mid z)$ describes the probability of observing a measurement outcome x given a true state z . An ideal detector with perfect resolution would preserve all physical symmetries present at the particle level. In reality, detectors often break or reduce symmetries that the underlying physics possesses. Understanding which symmetries are preserved, approximated, or lost convolution with $r(x \mid z)$ is crucial for interpreting measured data. Here we discuss several common invariances and how they are affected by realistic detectors in their response.

VIII.B.2.i Spatial uniformity and rotational symmetry.

Many detectors are designed with a roughly cylindrical geometry around a natural axis, such as the beam axis in collider experiments, aiming for azimuthal symmetry. Hence they are designed to provide close to uniform coverage in the plane normal to the axis.

Ideally, if the physical process yields a uniform distribution in the azimuthal angle ϕ (i.e. no preferred direction around the axial line), a perfectly symmetric detector would register an equal number of events in each azimuthal segment. In practice, certain asymmetries are present in any measurement.

For example, a detector may have support structures or cabling at certain angles, or irregular segmentation, leading to variation in efficiency with ϕ . The electromagnetic calorimeter (ECAL), as an illustration, might be segmented into modules that cover specific ϕ slices. Given this, events falling into the gap between modules could be recorded with lower efficiency or energy resolution, creating a ϕ -dependence in the observed data even if the true distribution was uniform. Detectors often have periodic segmentation, meaning continuous rotational invariance is broken down to a discrete rotation symmetry, invariant only under rotations corresponding to full module spacings.

As a concrete example, imagine a detector with 360 identical modules each covering $\Delta\phi = 1^\circ$. This detector is invariant under rotation by multiples of 1 deg increments, but a rotation by an arbitrary angle (say 0.5 deg) would lead to a different alignment of a particle's trajectory with respect to module boundaries, yielding a measurably different response [66]. Thus, the continuous symmetry is reduced to a discrete one,

and even that discrete symmetry may be imperfect if modules are not exactly identical or have time varying efficiency.

Thus azimuthal symmetry at the physics level is usually preserved approximately by detector design, but slight non-uniformities in ϕ response are common and must be accounted for either via calibration or acceptance corrections.

VIII.B.2.ii Polar coverage and boost invariance.

Often detectors in HEP experiments, especially collider experiments, also have limited coverage in the polar direction, along the beam axis. No real detector covers the full 4π solid angle; there is always a cut off at some polar angle (or pseudorange η) beyond which particles escape detection. In particular, the forward regions close to a beam are notoriously difficult to instrument.

This breaks the full spherical symmetry of space. A process that is symmetric under arbitrary rotations, such as a perfectly isotropic decay in its rest frame, will not appear isotropic in the laboratory measurement if a significant portion of the solid angle is unobserved. Detectors are typically ideally symmetric under rotations about the beam axis but not under arbitrary rotations that tilt the beam axis, because the beam direction is a fixed axis of symmetry.

In other words, the presence of beams singles out a preferred direction, the beam axis $\hat{z}$, and detectors are built around this axis. Consequently, the data may reflect cylindrical symmetry, invariant under $SO(2)$ rotations around $\hat{z}$, but not full $SO(3)$ rotational symmetry.

This also has consequences for Lorentz boost invariance. While the underlying physics is Lorentz invariant, the detector is a fixed apparatus in one frame. A boost along the beam direction, i.e. a change to the reference frame moving with respect to the collision, will generally change how events are distributed relative to the detector acceptance.

As an example, consider a boost that causes particles to have higher longitudinal momentum; in the lab frame, more particles will end up at small polar angles, closer to the beam line, where detection efficiency is lower, thus the observed distribution of, say, pseudorapidity η will shift. The detector has a finite acceptance in η , so a Lorentz boost that moves events into the far forward region will result in a fraction of events being lost. Therefore, the measured distributions are not invariant under Lorentz boosts, even though the underlying parton level kinematics can be expressed in Lorentz invariant terms.

Thus the physical construction of detectors break global translational and boost symmetry by virtue of being static and having edges. A high energy interaction viewed in different inertial frames is physically identical, but a detector at rest in the lab frame will record it differently unless one corrects for acceptance and inefficiencies.

VIII.B.2.iii Resolution effects and approximate invariance.

Even if a symmetry could hold in principle, the resolution and threshold effects of detectors often spoil exact invariance. A salient example arises with Lorentz invariance and invariant mass reconstruction.

As a thought experiment, imagine a two body decay producing a pair of muons, such as $Z^0 \rightarrow \mu^+ \mu^-$. The true invariant mass of the muon pair is fixed, irrespective of the Z boson's momentum, because it is simply a Lorentz scalar. However, a detector measures muon momenta with finite precision, and that precision typically degrades at high momentum¹⁰. If the Z boson is produced nearly at rest in the lab, its decay muons have moderate momenta and the detector might reconstruct the invariant mass with a narrow resolution. If instead the Z is produced with a large boost, the muons each have higher lab frame momenta, and the detector's momentum resolution broadens the reconstructed mass distribution. The result is that the distribution of reconstructed $m_{\mu\mu}$ for boosted Z events is broader (and potentially biased) compared to that for non-boosted events.

Thus, a Lorentz boost, which should not matter to an ideal measurement, actually changes the statistical distribution of an observable due to detector response. This is an example of an approximately respected symmetry. At low boost the symmetry holds well, but at high boost the symmetry is effectively broken by detector effects.

Similarly, thresholds in detector sensitivity (e.g. a calorimeter that only records energy above some minimum) can break symmetry under transformations that redistribute energy. A detector that is equally efficient for electrons and positrons should demonstrate C symmetry. However, if the process produces a broad energy spectrum, imposing a cut on low energy particles can introduce a bias. Low energy e^+ are more likely to be lost than e^- due to their different interaction rates with detector

¹⁰Tracking detectors determine momentum from curvature in a magnetic field, which becomes very small for high momentum muons, leading to larger relative uncertainty in the measurement.

material. In such cases, even an underlying physical symmetry may not yield equal measured counts.

VIII.B.2.iv Mirror and charge symmetry.

Detectors are not usually built to be fully symmetric under parity inversion or charge conjugation, even though one often assume these symmetries for the relevant physics should reflect in data. A parity inversion would swap the ‘forward’ and ‘backward’ directions in the detector. If the detector has identical coverage in the forward ($+z$) and backward ($-z$) hemispheres, one could say that it is parity symmetric with respect to the interaction point.

Many detectors strive for this by having symmetric endcaps on both sides of the interaction region. However, even then, subtle asymmetries can exist because it is not feasible to prevent one side from having a slightly different material distribution or a different calibration from the other. As a result, a process that is forward–backward symmetric in its physics¹¹ might show a forward–backward asymmetry in the raw data. Experiments typically correct for such differences by equalizing calibrations, but the point stands that the intrinsic detector response can break the symmetry.

Likewise, charge conjugation symmetry in detection would mean the detector is equally sensitive to positive and negative charges. While the detector electronics and geometry generally don’t prefer one charge sign, magnetic fields introduce a notable asymmetry, because charged particles bend in opposite directions, and this can lead to charge dependent acceptance. In a magnetic spectrometer, positively

¹¹In $p - p$ collisions for example, the two beam directions are equivalent so the distribution of particles as a function of rapidity y should be symmetric about $y = 0$.

charged particles bend in one direction and negatively charged particle in the opposite direction. If the acceptance boundaries, like the edge of the detector volume, cut off tracks in one curvature direction more than the other, one will observe a difference in detection rates even if production is symmetric.

Another example is that the different interaction of e^+ and e^- with matter could lead to different detection efficiencies. These are lower order effects, but they illustrate that a detector is a physical object that need not respect the abstract symmetries of the theory. Careful simulations and calibrations are performed to quantify and mitigate these asymmetries in HEP experiments.

VIII.B.2.v Permutation symmetry and identical particles.

Physically, as noted, swapping two identical particles should change nothing in an ideal measurement. Detectors, however, could introduce differences. Two identical particles (say two photons) that go into different regions of the detector can have their energies might be measured with different resolutions or one might pass quality cuts and the other fail due to region specific noise. As a result, the joint distribution of the two particle system in the measured data might not be symmetric under exchange, even though it was at truth level.

As a another simple example, consider two jets in an event for which, at the particle level, the probability $\mathbb{P}(E_1, \eta_1; E_2, \eta_2)$ is symmetric under $(1 \leftrightarrow 2)$. After detection, suppose that jet 1 falls in the central barrel, with excellent energy resolution and jet 2 falls in the forward region, with poorer resolution and lower efficiency. The measured energies $E_1^{(\text{meas})}$ and $E_2^{(\text{meas})}$ will have different response smearing. If one

then orders jets by measured energy and calls the highest energy jet the leading jet, the distribution of leading and subleading jets will not mirror one another exactly. Effectively, the detector induced asymmetry has assigned labels to the jets where none existed. Analysts must be wary of these effects; one option used is to ‘symmetrize’ the analysis if possible to recover the permutation symmetry that the physics assures.

Table VIII.3 summarises a few key examples of how an ideal symmetry at the particle level can be broken or reduced by detector effects. These examples illustrate why fully accounting for detector response is essential when testing physical symmetry hypotheses with data.

Despite these challenges, experimentalists strive to design detectors with as much symmetry as feasible and to correct for known asymmetries. For instance, collider detectors often have nearly full 2π azimuthal coverage and layered symmetries, segmenting in ϕ and η uniformly, specifically to preserve rotational invariance and facilitate combining data over symmetric regions. Detector simulation and calibration are used to quantify symmetry breaking. If a ϕ –dependence is observed in calibration data, it can be corrected so that the final analysis treats those variations as a systematic uncertainty or removes them. Nonetheless, the reality remains that physical symmetries can fail to translate into measured symmetries.

Table VIII.3: Ideal SM symmetries and detector induced symmetry breaking in HEP experiments. While fundamental interactions may respect certain symmetries exactly, detector geometries, material distributions and reconstruction algorithms violate these symmetries at measurement level. The table illustrates how common detector limitations transform SM symmetries. These effects must be carefully modelled in simulations and corrected through calibration to extract the underlying physics.

Symmetry	Ideal outcome		Detector effect
Azimuthal rotation $\phi \rightarrow \phi + \alpha$	Rotational invariance:	uniform distribution	Discrete n -fold symmetry: detector segmentation creates ϕ -dependent acceptance. Dead material between modules introduces periodic inefficiencies. Triggers based on detector regions further break symmetry
	$\phi \in [0, 2\pi]$.	No preferred transverse direction	
Lorentz boost $\beta_z \rightarrow \beta_z + \Delta\beta$	Physics invariant under boosts;	invariant longitudinal cross sections expressible in terms of Lorentz scalars.	Lab frame dependence: fixed detector geometry defines preferred frame. Forward boosts push particles beyond acceptance. Resolution degrades for highly boosted objects. Trigger thresholds defined in lab frame are not invariant

Continued on next page

Table VIII.3 continued from previous page

Symmetry	Ideal outcome	Detector effect
Parity $\mathbf{x} \rightarrow -\mathbf{x}$	For P conserving processes: equal rates and distributions for original and parity transformed events	Geometric asymmetry: detectors rarely possess reflection symmetry about $z = 0$. Forward and backward regions have different instrumentation, acceptance, and resolution. Magnetic field direction picks out handedness
Charge conjugation $q \rightarrow -\bar{q}$	C symmetric interactions produce particles and antiparticles with identical rates and kinematic distributions	Charge dependent efficiency: opposite charges bend oppositely in magnetic field, sampling different detector regions. Material interactions differ (e.g., K^+ vs K^- nuclear cross sections). Trigger and particle ID algorithms may have charge bias
Permutation $(i \leftrightarrow j)$	For identical particles: joint distribution has exchange symmetry	Position dependent response: particles in different regions experience different resolutions, efficiencies, and systematics. Ordering by p_T obscures underlying symmetry. Combinatorial background differs for same region vs different region pairs

In the language of probability distributions, if $p(z)$ is invariant under transformation T , but the response $r(x | z)$ is not invariant in the corresponding way, then the folded

distribution $p(x) = \int r(x | z) p(z), dz$ will not be invariant under T applied to x . Only if both p_{truth} and r share the symmetry T will p_{data} exhibit it. This conceptual understanding is vital when one interprets measured cross sections and tries to infer or discover symmetries from data.

VIII.B.3 How symmetries manifest in measured cross sections.

Given the above considerations, one can now examine how symmetries and symmetry violations are reflected in the measured distributions that experiments record and report. A measured cross section differential in some observable is effectively a statistical aggregate of many collision events, after selection cuts and corrections. If the underlying physics possesses a symmetry, one might expect the differential cross section to reflect that, provided the measurement process does not hide or distort it. In practice, one observes in histograms a mixture of genuine physical symmetry patterns and effects of detector acceptance or sample selection.

VIII.B.3.i Exact symmetries and flat distributions.

A hallmark of a symmetry in a distribution is a repeated pattern indicating invariance. For example, consider azimuthal invariance in a proton–proton collision. Since the colliding protons provide a cylindrically symmetric initial state of two identical beams head on, no physics process at the parton level prefers a particular ϕ direction. Consequently, the true differential cross section $d\sigma/d\phi$ for an inclusive process is invariant under the transformation $\phi \mapsto \phi + \delta\phi$ (aside from small QED

effects or residual detector magnetization influences). If the detector has uniform ϕ coverage and the analysis has no ϕ -dependent cuts, the measured distribution of events as a function of ϕ should be approximately flat. Any significant deviation from flatness might indicate an instrumental problem or a selection bias.

Once the symmetry has been established, one might combine data from all ϕ slices (since they are equivalent) to improve statistical precision, effectively using the symmetry to gather more data. However, as noted, small modulations can appear if certain detector modules deviate from the rest; these are corrected or quoted as systematic uncertainties.

Another example is rapidities in symmetric collisions: in a $p - p$ collider, for example, at equal beam energies, the centre of mass frame coincides with the lab frame, and the process is symmetric under exchanging the two beam directions. This implies that the distribution of particles in rapidity y is symmetric about $y = 0$ for processes that do not involve a bias.¹² This is why measurements of inclusive jet or hadron yields often present results as a function of $|y|$ or $|\eta|$, the absolute value of rapidity or pseudorapidity, invoking the symmetry $y \leftrightarrow -y$ to double the statistics and simplify presentation. The physical symmetry (identical proton beams) justifies this, and one checks that, within uncertainties, the $+y$ and $-y$ distributions are consistent before merging. Thus, a symmetry in initial conditions and dynamics (here, invariance under $y \rightarrow -y$) leads to a clear symmetry in the measured cross section (equal yields for $\pm y$).

¹²For instance, pure QCD dijet production should yield a symmetric $d\sigma/dy$ for jets, with equal activity in the forward ($+y$) and backward ($-y$) hemispheres.

VIII.B.3.ii Symmetries in kinematic shapes.

Symmetries often impose recognizable shapes or constraints on distributions. For example, energy and momentum conservation require that for each event, the vector sum of momenta of final state particles equals that of initial state. As a result, distributions of total transverse momentum in events would be expected to peak at zero, and any significant imbalance indicates e.g. neutrinos or detector holes. This is not a symmetry in the sense of a group acting on one event's space, but rather a deterministic constraint on the ensemble. The distribution of missing momentum should be centered at zero and isotropic in azimuth. Experiments can thus verify that the missing transverse momentum vector has no preferred direction to validate rotational symmetry and momentum conservation in aggregate.

If one measures the transverse momentum spectrum of the first jet against that of the second jet in dijet events (with jets ordered by p_T), there is no fundamental reason for these spectra to differ except for the ordering bias. The leading jet p_T distribution will be harder by construction, and the subleading softer, but any jet is equally likely to be at a given p_T as its partner, aside from that ordering. This symmetry can therefore be verified through the similarity between the distribution of subleading jet p_T and the leading jet p_T distribution of a lower energy subset, or by symmetrising the dataset by swapping jets event by event and seeing no change in overall two jet correlation distributions.

In summary, wherever a symmetry exists, one finds redundancies or equalities in the measured spectra: sections of phase space that should mirror other sections. Experimental analyses often exploit this to measure detector backgrounds, by assuming

that an uninstrumented region should have similar counts as a well instrumented region after normalisation.

VIII.B.3.iii Interplay of physical and detector symmetries.

It is important to disentangle which symmetries in a measured cross section come from physics and which from measurement procedure. An analysis might impose a cut that itself introduces a symmetry or asymmetry. When presenting a measured cross section, unfolding detector effects to reconstruct particle level distributions to the extent possible, to report a cross section as it would appear with an ideal detector, can restore the symmetries that belong to the physics by removing the distortions of measurement [702]. For example, if the raw data show a ϕ -dependence due to detector inefficiency, the unfolded cross section as a function of ϕ should be flat, with larger uncertainties reflecting the correction.

In this sense, symmetries provide a consistency check. If after unfolding one still sees a symmetry violation in a quantity that should be symmetric by physics, the unfolding procedure might be flawed. Conversely, if a symmetry is expected to be broken by physics, one must be careful to ensure the detector is not distorting the size of the asymmetry. For instance, measuring a forward–backward asymmetry in top quark production requires excellent control of any detector differences between the forward and backward directions so that the observed asymmetry can be trusted as a physical sign of potential weak interaction interference (or potentially new physics).

A physical symmetry is a property of the underlying probability law. It requires equal probabilities for events and their transformed versions. A statistical symmetry

of a dataset entails that the finite sample of observed data appears invariant under some transformation, within the limits of noise, so that with large data, one expects the symmetry to become apparent as the fluctuations average out. If deviations persist significantly beyond expected fluctuations, that flags either a real symmetry violation or unaccounted systematics.

VIII.B.4 Challenges in identifying symmetries from noisy data.

Identifying symmetries in experimental data is not always straightforward. Noisy data, stemming from finite statistics, background processes, and detector imperfections, can obscure or mimic symmetry signals. This section outlines the main challenges one faces in discerning true invariances or symmetry violations within collider datasets, and the need for methods like the SYMMETRYGAN approach developed later in this work to address these challenges.

A fundamental challenge is that any empirical distribution has random fluctuations. If an underlying distribution is perfectly symmetric (say truly uniform in ϕ), a finite sample will still exhibit some variation across ϕ bins. Hence any symmetry discovery method must have a mechanism to distinguish a real asymmetry from a mere fluctuation.

Conversely, an underlying asymmetry can be washed out by limited statistics. This is especially pertinent in searches for new symmetries or violations. The signals are often at the level of small deviations and can be difficult to detect over statistical

fluctuation. Moreover, multiple comparisons increase the probability that one finds an apparent “symmetric pattern” in some projection of the data purely by chance.

As discussed, detector effects can induce or conceal asymmetries. Often the largest uncertainties in measuring symmetry come from how well we understand the detector. For example, in measuring a forward–backward asymmetry, uncertainties in the relative efficiency of the forward and backward region directly translate to uncertainty in the asymmetry observable. If those efficiencies are poorly known, one might not be able to distinguish a symmetry violation from detector effects.

Similarly, backgrounds and other processes that mimic the signal might not share the symmetry of the signal. Suppose one is looking for a symmetry in a certain particle decay distribution; if there is a significant background from a different process that does not respect that symmetry, the combined data will appear to break the symmetry even if the signal alone is symmetric. Careful background subtraction or isolation is required. In practice, identifying a symmetry often involves comparing two distributions (e.g. $P(x)$ vs $P(Tx)$ for some transformation T) and seeing if they differ. If they do, one must estimate if the difference is due to known systematic effects. This typically demands high-precision calibration. For instance, to confirm CP symmetry in production of particle vs antiparticle to the 10^{-3} level, one needs detector efficiencies known to better than 10^{-3} between positively and negatively charged particle detection [535].

VIII.B.4.i Dimensionality challenges.

HEP interaction events are often high dimensional, consisting of many particles with various kinematic attributes. A symmetry might not be evident in any single one dimensional projection, but rather in a complicated combination of variables. For example, Lorentz invariance is best seen when considering all 4-momenta together or invariants like masses; a naive look at just one momentum component would not show it. Permutation symmetry in a multijet event is a property of the joint distribution of all jet momenta, not necessarily obvious if one only looks at single jet spectra.

This is where machine learning methods become attractive, because they can, in principle, detect subtle patterns in high dimensional data. However, even ML models need guidance. The space of possible transformations is huge, and hence searching it naively for invariances is intractable. Hence traditional methods often restrict attention to physically motivated symmetry transformations.¹³

Since scanning for symmetries by comparing all possible pairs of transformed distributions is computationally impossible, as data volumes grow and analysis spaces become more complex, more automated symmetry discovery mechanisms are needed. The SYMMETRYGAN approach discussed in this chapter is one attempt to automate the discovery of symmetries by leveraging generative adversarial networks. Conceptually, a SYMMETRYGAN trains a generator (applying candidate transformations) against a discriminator to test if the transformed data is statistically indistinguishable from the original data. If the generator generates a transformation under which the discriminator is maximally confounded, that transformation corresponds to a

¹³Rotations, reflections, boosts, particle exchanges, *et cetera*.

symmetry of the data distribution. Implementing this is challenging: the model must search a continuous space of transformations, handle approximate symmetries, and avoid trivial solutions. A careful choice of network architecture using known mathematical structure is needed to make such learning robust.

In summary, identifying symmetries from noisy HEP data requires

1. Sufficient statistics and rigorous statistical tests to differentiate real invariances from fluctuations,
2. Precise control of detector systematics,
3. Methods to probe high dimensional and subtle symmetry patterns that might elude simple projection analyses, and
4. Methodological consideration to handling approximate symmetries in a principled way.

These challenges motivate the development of tools like SYMMETRYGAN, which will be introduced in Section VIII.D. Such tools aim to combine physical insights with machine learning’s ability to detect patterns, thereby providing a statistical discovery framework for symmetries. SYMMETRYGAN and similar approaches offer a promising path to unveil symmetries that are latent in complex data. The rigorous understanding of symmetry and symmetry-breaking provided in this section will form the foundation on which those computational methods build, ensuring that any discovered ‘symmetry’ is physically meaningful and relevant to the challenges of unfolding and analysing HEP data.

VIII.C Statistical definition of dataset symmetries.

VIII.C.1 Distinction between point and dataset symmetries.	530
VIII.C.2 Inertial reference densities and their theoretical role.	531
VIII.C.3 Inertial densities and relative symmetry.	533
VIII.C.3.i Failure of the naïve PDF preserving criterion.	534
VIII.C.3.ii Relative symmetry.	534
VIII.C.3.iii Choice of the inertial density.	535
VIII.C.4 Applications to particle physics.	536

The concept of symmetry in physics typically evokes images of rotational invariance in crystals, parity conservation in weak interactions, or gauge transformations in field theory. Yet when one turns one’s attention to experimental data, especially the high dimensional datasets emerging from modern HEP experiments, the notion of symmetry becomes surprisingly subtle. What does it mean for a collection of measured events to possess a symmetry? This question, deceptively simple in appearance, reveals profound connections between statistical inference, group theory, probability theory and physics.

VIII.C.1 Distinction between point and dataset symmetries.

Statistical symmetries in datasets represent a fundamental departure from traditional geometric symmetries, requiring careful consideration of probability measures, transformation Jacobians and reference densities. This section explores the mathematical foundations, practical applications and machine learning approaches to understanding and leveraging dataset symmetries. The distinction between symmetries of individual data elements and entire datasets, formalised in Definitions 50 to 53 lies at the heart of statistical theory of symmetries.

Unlike point symmetries, where transformations either preserve or don’t preserve specific locations, dataset symmetries must account for how probability densities transform under coordinate changes. This fundamental difference necessitates the introduction of inertial reference densities to properly define statistical symmetries.

VIII.C.2 Inertial reference densities and their theoretical role.

The concept of inertial reference densities emerges as a necessary theoretical construct, analogous to inertial frames in classical mechanics [741, 795, 803]. These reference densities provide a baseline against which statistical symmetries can be meaningfully defined, resolving the ambiguity inherent in coordinate transformations of probability measures.

Definition 54 (Reference density). A *reference density* $\rho(x)$ is a probability density function that establishes a canonical measure for comparing probability distributions by enabling the definition of relative entropy

$$H_\rho[X] = -\mathbb{E}_\mu \left[\ln \frac{d\mu}{d\rho} \right]. \quad (\text{VIII.4})$$

Additionally, it provides a coordinate independent way to specify symmetry transformations, ensuring that symmetry definitions remain consistent across different parametrisations of the same statistical manifold. This can be analogised to the way in which phase space shifting operations that leave the Gibbs integration measure invariant can be understood as gauge transformations, with the reference density playing the role of a gauge fixing condition [126].

The mathematical machinery for statistical symmetries centres on how probability densities transform under coordinate changes. Under a transformation $g : Y \rightarrow X$, probability densities transform as

$$p_Y(y) = p_X(g(y)) |\det g'(y)|. \quad (\text{VIII.5})$$

This transformation law, involving the absolute value of the Jacobian determinant $J = \det g'(y)$, [801] ensures probability conservation. The Jacobian factor measures how volumes scale under transformation. $|\det(J)| = 1$ characterises volume preserving transformations like rotations and reflections.

Definition 55 (Diffeomorphism invariant probability densities.). Let $X \subseteq \mathbb{R}^n$ be an open set and Φ_G be a C^1 action of a group G on X . Let dx denote the Lebesgue measure on X . A probability density $p : X \rightarrow [0, \infty)$ is diffeomorphism invariant under the action of Φ_G if

$$\forall \text{ measurable } A \subseteq X, \forall g \in G, \mu_p(g \cdot A) = \mu_p(A), \quad (\text{VIII.6})$$

where

$$\mu_p(A) = \int_A p(x) dx. \quad (\text{VIII.7})$$

This definition implies that, for almost every $x \in X$ and for every $g \in G$,

$$p(g(x)) |\det Dg(x)| = p(x), \quad (\text{VIII.8})$$

where $Dg(x)$ is the Jacobian of the diffeomorphism $x \mapsto g(x)$.

This invariance is strictly stronger than mere ‘point symmetry,’ since it requires that the *entire* probability measure be preserved under *all* elements of G , compensating exactly for any local volume distortion. Furthermore,

Definition 56 (Symmetric action of a group [627].). Let $(\Omega, \mathcal{F}, \mu)$ be a probability space and G a group. G is defined to have a *symmetric action* on the measure space

$(\Omega, \mathcal{F})$ if there exists a function

$$\phi : G \times \Omega \rightarrow \Omega, \quad (g, \omega) \mapsto g \cdot \omega \quad (\text{VIII.9})$$

such that

1.

$$\forall \omega \in \Omega \quad \phi(1, \omega) = \omega \quad (\text{VIII.10})$$

2.

$$\forall \omega \in \Omega \quad \forall g_1, g_2 \in G \quad \phi(g_1 g_2, \omega) = \phi(g_1(\phi(g_2, \omega))) \quad (\text{VIII.11})$$

3.

$$\forall g \in G \quad \forall A \in \mathcal{F} \quad \mu(g \cdot A) = \mu(A), \quad (\text{VIII.12})$$

Using the orbit–stabiliser theorem, one can decompose the action of such a group into orbits and stabilisers. Representation theory then enables the systematic construction of invariant functions and the decomposition of function spaces into irreducible components.

VIII.C.3 Inertial densities and relative symmetry.

In the preceding discussion, we demonstrated that any candidate symmetry transformation $g : \mathbb{R}^n \rightarrow \mathbb{R}^n$ preserve a target probability density p . However, preservation of p alone is too weak to capture the notion of symmetry familiar from physics. In particular, in one dimension every continuous distribution admits a ‘quantile flip’ map that preserves its density, yet such maps do not correspond to genuine physical symmetries.

VIII.C.3.i Failure of the naïve PDF preserving criterion.

Let X be a real random variable with cumulative distribution function (CDF) F . By the probability integral transform,

$$U = F(X) \sim \mathcal{U}[0, 1], \quad (\text{VIII.13})$$

and conversely, for $U \sim \mathcal{U}[0, 1]$, the inverse transform

$$X = F^{-1}(U) \quad (\text{VIII.14})$$

recovers the original distribution of X [100]. The uniform distribution on $[0, 1]$ admits a non-trivial involution

$$\tilde{g}(u) = 1 - u, \quad \tilde{g} \circ \tilde{g} = \mathbb{1}, \quad (\text{VIII.15})$$

which clearly preserves the uniform density. Transporting $\tilde{g}$ back to the domain of X defines the quantile involution map

$$g(x) = F^{-1}(\tilde{g}(F(x))), \quad (\text{VIII.16})$$

which satisfies

$$p(g(x)) |g'(x)| = p(x). \quad (\text{VIII.17})$$

Thus every univariate random variable admits a $\mathbb{Z}_2$ ‘symmetry’ of the form Equation (VIII.16). Yet these transformations depend explicitly on p and do not reflect any model independent invariance.

VIII.C.3.ii Relative symmetry.

To eliminate such spurious symmetries, one must introduce a *reference* or *inertial* density $p_I : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$.

Definition 57 (Relative symmetry.). A diffeomorphism g is defined to be a symmetry of a density p *relative* to p_I if both p and p_I are preserved.

$$p(g(x)) |\det Dg(x)| = p(x), \quad p_I(g(x)) |\det Dg(x)| = p_I(x). \quad (\text{VIII.18})$$

In this formulation, g must belong to the intersection of the symmetry groups of p and p_I . Quantile based maps like the one described in Equation (VIII.16) will generally fail to preserve any fixed p_I across different target densities.

VIII.C.3.iii Choice of the inertial density.

The analogy to inertial frames in mechanics suggests choosing p_I to reflect the ambient geometry of the data. For many physical applications in $\mathbb{R}^n$, or $\mathbb{R}^{p,q}$ in a space of signature (p, q) , the most natural inertial density is the *constant* (Lebesgue) density:

$$p_I(x) \propto 1, \quad x \in \mathbb{R}^{p,q}. \quad (\text{VIII.19})$$

Under this choice, the relative symmetry condition Equation (VIII.18) reduces to

$$p(g(x)) |\det Dg(x)| = p(x), \quad |\det Dg(x)| = 1, \quad (\text{VIII.20})$$

so that genuine symmetries must be volume preserving transformations under which p is invariant. The unimodular subgroup of the affine group $\text{Aff}(p, q)$ that acts by linear maps with unit Jacobian determinant is the special affine group $\text{ASL}^\pm(p, q)$. In fact, there is a canonical isomorphism between all such groups of the same dimension $d = p + q$, and so these groups can be referred to simply as $\text{ASL}^\pm(d)$ independent of signature. This choice of p_I recovers the standard symmetry groups employed in particle physics and continuum mechanics.

Requiring joint preservation of p and a fixed inertial density p_I provides a rigorous definition of symmetry that excludes trivial, density dependent involutions and aligns with the model independent invariances central to physics.

VIII.C.4 Applications to particle physics.

Particle physics provides a rich domain for applying statistical symmetries, particularly in detector calibration, data unfolding, and inference tasks at collider experiments [355]. In detector response modelling, statistical symmetries constrain how identical particles must be treated [804]. Permutation symmetry requires that detector analysis respect the indistinguishability of identical particles, affecting event reconstruction algorithms, particle identification methods, and background estimation techniques [31]. Response matrices used in unfolding procedures must respect particle exchange symmetries, with regularisation procedures preserving these symmetry properties.

Unfolding represents an interesting application where statistical symmetries could guide the correction of detector effects. The process must preserve the statistical symmetries of the underlying physics, maintaining correlations between identical particles and ensuring conservation laws implied by symmetries remain valid. Full phase space unfolding could be facilitated by statistical symmetries that constrain unfolding procedures in high dimensional phase spaces. Background estimation methods also exploit symmetry properties through carefully designed control regions, while systematic uncertainties account for potential symmetry violations.

In this way statistical symmetries fundamentally shape measurement and inference tasks through multiple pathways. Phase space shifting operations that preserve physically meaningful quantities have even been used outside of HEP to develop new computational approaches for molecular simulations [284]. These symmetries lead to exact correlation relations between interactions and observable properties, offering systematic ways to derive new statistical relationships.

In machine learning too, encoding known symmetries dramatically improves performance. Symmetry encoding reduces sample complexity exponentially, with multidimensional symmetries providing disproportionately large returns [202]. Equivariant neural networks, designed to respect known symmetries through architectural constraints, achieve superior data efficiency and generalisation [96].

The theoretical framework connects reference measures, symmetry principles, and statistical inference in profound ways. Reference measure selection can be guided by symmetry considerations, leading to principled approaches for prior selection in Bayesian inference [771]. Symmetry principles suggest natural classes of invariant estimators, while statistical analogues of physical conservation laws emerge from symmetry considerations, constraining the behaviour of statistical systems [692]. This unified framework offers concrete advantages for computational statistics and machine learning while opening new directions for both theoretical development and practical applications.

VIII.D SYMMETRYGAN: Discovering symmetries with adversarial learning.

VIII.D.1 The SYMMETRYGAN architecture.	539
VIII.D.2 Machine learning with inertial restrictions.	541
VIII.D.3 Deep learning implementation details.	545
VIII.D.4 Verification	547
VIII.D.5 Other symmetry discovery methods.	547

The challenge of automatically discovering symmetries in high dimensional data represents a fundamental problem at the intersection of physics and machine learning. While physicists have long relied on theoretical insight to identify symmetries, the explosion of complex data from modern experiments demands automated approaches. SYMMETRYGAN emerges as an elegant solution, leveraging the power of adversarial learning to discover hidden invariances without prior knowledge of their existence.

The core insight driving a SYMMETRYGAN is deceptively simple. If a transformation truly represents a symmetry of a dataset, then a well trained discriminator should not be able to distinguish between the original data and its transformed counterpart. This principle, when combined with a framework of inertial reference densities, yields a powerful methodology for symmetry discovery that is both theoretically sound and practically effective.

VIII.D.1 The SYMMETRYGAN architecture.

The SYMMETRYGAN framework modifies the traditional generative adversarial network architecture in a fundamental way. Rather than mapping from random noise to structured data, SYMMETRYGAN maps from data to data, learning transformations that preserve the underlying probability distribution. Like a traditional GAN, however, SYMMETRYGAN consists of two neural networks engaged in adversarial training.

1. A generator $g : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that learns symmetry transformations, and
2. A discriminator $d : \mathbb{R}^n \rightarrow [0, 1]$ that attempts to distinguish original from transformed data.

The generator network g parametrises potential symmetry transformations of the input space. Unlike traditional GANs where the generator creates new samples, here it generates transformations that are applied to existing data points. The discriminator d plays its familiar adversarial role, attempting to classify whether a given sample comes from the original dataset or has been transformed by g .

The training objective is related but different from standard GANs. The loss function is

$$L[g, d] = -\frac{1}{N} \sum_{i=1}^N \ln d(x_i) + \ln(1 - d(g(x_i))) . \quad (\text{VIII.21})$$

This formulation differs from the usual binary cross entropy in that the same data samples appear in both terms, creating a self-adversarial structure. The generator is trained to produce transformations that maximally confound the discriminator, while the discriminator learns to detect deviations from the true data distribution.

By varying the loss with respect to the neural networks, one can show that at the theoretical optimum, the optimal discriminator is

$$d_*(x) = \frac{p(x)}{p(x) + p(g_*(x))|g'_*(x)|} = \frac{1}{2} , \quad (\text{VIII.22})$$

and the optimal generator obeys

$$p(x) = p(g_*(x))|g'_*(x)| . \quad (\text{VIII.23})$$

Hence in equilibrium, the discriminator is maximally confounded with equal probability of classifying any sample as original or transformed, and the generator obeys precisely Definition 52.

The architecture naturally handles both discrete and continuous symmetries. For discrete symmetries like reflections or cyclic groups, the generator learns specific

transformation matrices. For continuous symmetries like rotations, it learns individual Lie group elements, parametrised by continuous parameters. This flexibility makes SYMMETRYGANs applicable across diverse physical systems. Figure VIII.1 shows the modified GAN architecture of a SYMMETRYGAN.

VIII.D.2 Machine learning with inertial restrictions.

The incorporation of inertial reference densities, as established in Section VIII.C.3, presents both theoretical necessity and practical challenges. These restrictions can be implemented through three distinct approaches.

Simultaneous discrimination: In this approach, discriminators evaluate transformations on both the target dataset and samples from the inertial density. The loss function is extended to include terms from the inertial distribution p_I

$$\tilde{L}[g, d, d_I] = L[g, d] + \lambda L_I[g, d_I] \quad (\text{VIII.24})$$

where d_I is a separate discriminator network specialized for the inertial distribution. This method offers maximal flexibility, allowing discovery of symmetries even when the inertial distribution's symmetries are not known analytically. However, it requires the ability to sample from p_I which becomes problematic for improper priors like uniform distributions on $\mathbb{R}^n$.

Two stage selection: This approach involves first identifying all PDF preserving maps, then filtering for those preserving the inertial density. The two stage approach decouples the symmetry discovery problem into training multiple generators to find PDF preserving transformations of the target data and *post*




Figure VIII.1: Schematic comparison of training architectures for (top) a standard Generative Adversarial Network (GAN) and (bottom) the proposed SymmetryGAN framework for automated symmetry discovery in particle physics data. In a standard GAN, the generator $g : \mathcal{Z} \rightarrow \mathcal{X}$ maps random noise $z \sim p_z$ to synthetic data samples, while the discriminator $d : \mathcal{X} \rightarrow [0, 1]$ attempts to distinguish real data $x \sim p_{\text{data}}$ from generated samples $g(z)$. The adversarial training objective drives g to match the data distribution without explicit physics constraints. The SymmetryGAN architecture modifies this framework with a symmetry transformation module $g_\theta : \mathcal{X} \rightarrow \mathcal{X}$, parametrized by learnable parameters θ , which acts on both real samples. The discriminator receives transformed pairs $(x, T_\theta(x))$, learning to identify distribution preserving transformations. When the discriminator cannot distinguish between a sample and its transformed version, the transformation T_θ represents an symmetry of the distribution. The inertial reference dataset, while not explicitly shown in this diagram, is incorporated directly into the generator architecture in Section VIII.E

hoc verifying which transformations also preserve p_I . While conceptually clean, this method proves computationally wasteful, as the space of PDF preserving maps vastly exceeds the space of true symmetries.

Upfront restriction: Here one constrains the generator architecture to only produce transformations that preserve the inertial density by construction in the first place, using prior knowledge about p_I .

It is this third approach that is adopted in the SYMMETRYGAN implementation. When the inertial distribution is uniform on $\mathbb{R}^n$ the generator should be restricted to equiareal maps—transformations with unit Jacobian determinant, as shown in Equation (VIII.20). The intersection of the equiareal group and affine group $\text{Aff}(n)$ is the affine special linear group $\text{ASL}^\pm(n)$ and its extensions.

The upfront restriction method offers computational efficiency gains by searching only among valid symmetries in addition to theoretical guarantees that discovered transformations are true symmetries. It also simplifies training without multiple discriminators or *post hoc* verification.

The restriction to affine equiareal maps, while limiting, captures a rich class of symmetries including rotations, reflections, shears, and all their compositions.

Theorem 20 (Iwasawa Decomposition of the Affine Group [781]). Let

$$\text{Aff}(n) = GL(n, \mathbb{R}) \ltimes \mathbb{R}^n \quad (\text{VIII.25})$$

be the real affine group. Define the subgroups

$$\begin{aligned} Z &= \{\pm I_n\} \cong \mathbb{Z}_2, \\ K &= SO(n), \\ A &= \{D(\mathbf{r}) = \text{diag}(r_1, \dots, r_n) \mid r_i \in (0, \infty)\}, \\ N &= \{S(u) \mid S_{ii} = 1, S_{i < j} \in \mathbb{R}, S_{i > j} = 0\}, \\ V &= \{\mathbf{v} \mid v_i \in \mathbb{R}\} \cong \mathbb{R}^n. \end{aligned} \quad (\text{VIII.26})$$

Then every element $M \in \text{Aff}_n(\mathbb{R})$ admits a unique factorisation

$$M = \sqrt{|\det M|} (\pm I_n)^{\frac{1 - \text{sgn}(\det M)}{2}} R(\theta) D(\mathbf{r}) S(u) + \mathbf{v}, \quad (\text{VIII.27})$$

where $\pm I_n \in Z$, $R(\theta) \in K$, $D(\mathbf{r}) \in A$, $S(u) \in N$, $\mathbf{v} \in V$.

In two dimensions these matrices take the form

$$I_2 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad (\text{VIII.28})$$

$$R(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}, \quad (\text{VIII.29})$$

$$D(r) = \begin{bmatrix} r & 0 \\ 0 & r^{-1} \end{bmatrix}, \quad (\text{VIII.30})$$

$$S(u) = \begin{bmatrix} 1 & u \\ 0 & 1 \end{bmatrix}, \quad (\text{VIII.31})$$

$$v = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}. \quad (\text{VIII.32})$$

Hence every $g \in \text{Aff}(2)$ admits the six parameter representation

$$g(x) = \sqrt{|d|} \left(I_2 \right)^{\frac{1-\text{sgn}(d)}{2}} R(\theta) D(r) S(u) x + v, \quad d = \det g, \quad (\text{VIII.33})$$

The Iwasawa parametrisation very well suited to gradient based optimisation over the space of affine symmetries.

VIII.D.3 Deep learning implementation details.

The practical implementation of SYMMETRYGAN requires careful attention to architectural choices, training dynamics, and numerical stability. The framework's success depends on balancing the adversarial training process while maintaining the theoretical properties that enable symmetry discovery.

For linear symmetries, the generator directly parametrises transformation matrices. Rather than using fully connected layers, the generator outputs parameters of the Iwasawa decomposition, ensuring all generated transformations lie within the constrained search space. This architectural choice provides interpretability, where each parameter has clear geometric meaning, stability, by avoiding ill conditioned matrices through structured parametrisation, and efficiency, by reducing the parameter count compared to unconstrained matrices from n^2 to $n(n-1)/2$.

To investigate specific symmetry subgroups, the loss function can further incorporate additional constraints. For example, to find cyclic symmetries $\mathbb{Z}_q$ the loss can be augmented to encourage $g^q = \mathbb{1}$.

$$L_{\text{cyclic}}[g, d] = L_{\text{BCE}}[g, d] + \alpha \sum_i \|g^q(x_i) - x_i\|^2. \quad (\text{VIII.34})$$

The discriminator architecture should follow standard practices for the data domain. For low dimensional examples, simple feedforward networks with 2 – 3 hidden layers suffice.

Analysing the loss landscape reveals its topological structure. For simple distributions like Gaussians, the loss landscape contains distinct maxima corresponding to each symmetry, separated by deep valleys. This topology explains why random initialisation reliably discovers different symmetries. The generator converges to the nearest local maximum, which corresponds to a true symmetry, and there is no reason to expect that the nearest local maximum for any given random initialisation should be the identity map.

The framework demonstrates remarkable robustness to hyperparameter choices. Across experiments with Gaussian mixtures and particle physics data, a range of hyperparameter settings achieve consistent symmetry discovery.

VIII.D.4 Verification

Several complementary checks can be employed to determine whether the discovered transformation constitutes a true symmetry.

Visual inspection: Examination of low dimensional slices of the original and transformed data provides a preliminary, coarse grained, qualitative assessment of whether the transformation plausibly preserves the structure of the data.

Loss value: For true symmetries, the binary cross entropy loss function converges to $2 \ln 2$. Measuring the loss value for a discovered transformation therefore offers a quantitative criterion for symmetry verification.

Distribution matching: Statistical divergences between X and $g(X)$, such as the Kullback–Leibler divergence, quantify whether the two datasets share the same probability density.

Invariant verification: Quantities known to be invariant under the symmetry, such as statistical moments, should be identical for both distributions.

VIII.D.5 Other symmetry discovery methods.

The landscape of symmetry discovery in machine learning has evolved considerably, with approaches ranging from classical statistical methods to modern deep learning

architectures. Understanding SYMMETRYGAN’s position within this ecosystem illuminates both its unique contributions and its connections to broader methodological trends.

Classical statistical approaches traditionally rely on hypothesis testing and moment matching. These methods typically test specific, predefined symmetry hypotheses by relying on low dimensional projections or summary statistics. They struggle with high dimensional data or complex symmetry groups and require extensive domain knowledge to formulate appropriate tests.

The Cramer–von Mises [791, 792] and Anderson–Darling tests [777] exemplify this approach to examine whether transformed data follows the same distribution as the original. While rigorous, these methods scale poorly and cannot discover unexpected symmetries, they can only validate known symmetries.

Machine learning methods introduced automation to the symmetry discovery process. Methods like Group-Invariant Autoencoders [182] learn representations invariant to known symmetry groups but cannot discover unknown symmetries. Spatial Transformer Networks [492] learn task specific transformations but don’t explicitly identify symmetries *per se*. Maximum Mean Discrepancy approaches [472] and other kernel methods can test for the existence of symmetry but struggle with continuous groups.

Contemporary deep learning methods have introduced a range of innovations to the field. Lie Group Learning Networks directly parametrise Lie algebra generators [178]. Such methods have excellent theoretical grounding, and physically interpretable parameters. They are, however, restricted to continuous symmetries, and require

complex optimisation for convergence. Latent LieGAN (LaLiGAN) enables the discovery nonlinear symmetries via latent space linearisation [146]. It made strides in handling non-linear group actions, but the additional complexity leads to potential latent space artifacts

A SYMMETRYGAN both draws on and innovates on several important trends in modern machine learning for HEP. Methods have broadly benefited from incorporating domain knowledge through architectural constraints and loss terms. In general the rise of equivariant learning has driven the move beyond invariance to discover the underlying symmetry structure. SYMMETRYGANs also exemplify the trend towards interpretability by producing human understandable symmetry parameters rather than opaque features. The method is however limited in its restriction to parametrisable symmetry classes. It also exclusively enables symmetry discovery, i.e. the task of identifying elements of the symmetry group of a dataset. Although the work does provide some steps towards symmetry inference, the task of identifying and classifying the symmetry group as a whole, it remains challenging to infer complete symmetry group structure from discovered elements.

The field continues to evolve rapidly, with SYMMETRYGANs representing a significant step toward automated, theoretically grounded symmetry discovery. Its success in particle physics applications, detailed in the following section, demonstrates the practical value of this principled approach to uncovering hidden invariances in complex data.

VIII.E Empirical experiments.

VIII.E.1 Gaussian experiments.	551
VIII.E.1.i One dimensional Gaussian.	551
VIII.E.1.ii Two dimensional Gaussians.	553
VIII.E.1.iii Gaussian Mixture Models	556
VIII.E.2 Particle physics experiments.	558
VIII.E.3 Interpreting discovered symmetries.	566

The theoretical machinery of a SYMMETRYGAN, with its rigorous statistical framework and inertial reference densities, proves its worth through empirical validation. Moving from abstract formulations to concrete applications reveals both the method’s effectiveness and challenges that might drive future research. This section details the application of this method on carefully controlled Gaussian experiments to complex collider data, demonstrating how SYMMETRYGANs bridge the gap between mathematical elegance and practical discovery.

VIII.E.1 Gaussian experiments.

The choice to begin with Gaussian distributions reflects more than mere mathematical convenience; it represents a strategic approach to validating a novel methodology. Gaussians offer analytically tractable loss landscapes that allow direct comparison between theoretical predictions and empirical results, serving as a crucial proof of concept before tackling high dimensional HEP data.

VIII.E.1.i One dimensional Gaussian.

The simplest non-trivial test case consists of a Gaussian distribution $\mathcal{N}(0.5, 1.0)$ with a reflection symmetry about $x = 0.5$. This apparently simple example harbours surprising richness. The disconnected nature of the symmetry space in even this simple example demonstrates the training dynamics in a SYMMETRYGAN.

The distribution admits precisely two linear symmetries: the identity transformation $g(x) = x$ and the reflection $g(x) = 1 - x$. When the generator is parametrised as an arbitrary one dimensional affine transformation, $g(x) = b + cx$, the analytic loss




Figure VIII.2: Analytic loss landscape for the one dimensional normal distribution $\mathcal{N}((0.5, 1))$ under affine transformations $g(x) = b + cx$. The landscape exhibits two distinct maxima (indicated by stars) corresponding to the identity transformation $(b, c) = (0, 1)$ and reflection $(b, c) = (1, -1)$, separated by a deep valley at $c = 0$. This topological barrier creates a disconnected solution space that deterministically routes optimisation trajectories based on initial parameter values.

landscape has the topology shown in Figure VIII.2. The two maxima corresponding to these symmetries are separated by a deep valley at $c = 0$, creating a disconnected solution space. The empirical results validate the theoretical predictions with remarkable precision. Random initialisation of parameters $(b_i, c_i) \sim \mathcal{U}([-5, 5]^2)$ leads to convergence at one of two distinct clusters: $(b_f, c_f) = (0, 1)$ for the identity or $(1, -1)$ for the reflection as can be seen in Figure VIII.3. The loss barrier at $c = 0$ acts as a watershed, deterministically routing the optimisation based on the sign of




Figure VIII.3: Empirical validation of SYMMETRYGAN on 1 dimensional Gaussian data showing convergence to discrete symmetries based on initial conditions. Random initialisation with $(b_i, c_i) \sim \mathcal{U}([-5, 5]^2)$ leads to convergence at two distinct fixed points. **(a)** Distribution of final parameters (b_f, c_f) reveals two clusters at $(0, 1)$ and $(1, -1)$, corresponding to the identity and reflection symmetries respectively. **(b)** Deterministic routing based on initial slope: $c_i > 0$ converges to identity ($c_f = 1$) while $c_i < 0$ converges to reflection ($c_f = -1$), with a sharp phase transition at $c_i = 0$. **(c)** Final intercept b_f shows no correlation with initial value b_i , demonstrating that the loss landscape has no barrier in intercept space.

the initial slope. This behaviour illuminates that the loss landscape topology directly determines the discoverable symmetry structure.

VIII.E.1.ii Two dimensional Gaussians.

Two dimensional Gaussian examples increase the complexity while maintaining analytical tractability. Consider first the standard bivariate normal $N_{1,1} = \mathcal{N}(\mathbf{0}, \mathbb{I}_2)$, which possesses the full orthogonal group $O(2)$ as its symmetry group. When the

generator is restricted to the form

$$g(x) = \begin{bmatrix} c & s \\ -s & c \end{bmatrix} x \quad (\text{VIII.35})$$

SYMMETRYGAN would be expected to discover that only transformations satisfying $c^2 + s^2 = 1$ represent true symmetries. As can be seen in Figure VIII.4a, the empirical results demonstrate the method's ability to learn this constraint without explicit enforcement. Starting from random initialisations within $[-1, 1]^2$, the learned parameters consistently converge to the unit circle, validating that SYMMETRYGAN can discover not just discrete symmetries but continuous symmetry manifolds. This represents a significant achievement: the neural network learns the algebraic constraint defining $SO(2)$ purely from data.

The two dimensional Gaussian example also allows us to test the approach described in Equation (VIII.34). Figure VIII.5 shows the discovered symmetries for the $N_{1,1}$ distribution when the loss function is augmented with the constraint Equation (VIII.34), for $q = 2, 3$ and 7 , with $\alpha = 0.1$. Both the analytic loss landscape and SYMMETRYGAN's output confirm that the continuous $SO(2)$ symmetry is split into discrete symmetries at the q^{th} roots of unity upon the inclusion of the mean squared error term.

One could also consider a two dimensional Gaussian with a non-trivial covariance matrix, such as $N_{1,2} = \mathcal{N}(\mathbf{0}, \text{diag}(1, 2))$. The full symmetry group of this distribution is highly non-trivial and is described below, but among other features, it contains as a subgroup the Klein 4-group $V_4 = \{\mathbb{1}, -\mathbb{1}, \sigma_3, -\sigma_3\}$ for the Pauli matrix σ_3 . When SYMMETRYGAN is limited to transformations of the form Equation (VIII.35),




Figure VIII.4: Analytic loss landscapes for two dimensional Gaussian distributions under rotation like transformations $g(\mathbf{x}) = \begin{pmatrix} c & -s \\ s & c \end{pmatrix} \mathbf{x}$, with empirically discovered symmetries overlaid as scatter points. (a) Isotropic Gaussian $\mathcal{N}(\mathbf{0}, \mathbb{I})$ exhibits continuous $SO(2)$ symmetry, with loss maxima forming a circle at $c^2 + s^2 = 1$. The discovered symmetries densely populate this manifold, confirming the theoretical prediction. (b) Anisotropic Gaussian $\mathcal{N}(\mathbf{0}, \text{diag}(1, 2))$ possesses only discrete Klein four-group symmetries V_4 , with isolated maxima at $(c, s) \in \{(\pm 1, 0), (0, \pm 1)\}$. The symmetry breaking from continuous to discrete group structure is induced by the non-uniform covariance. Darker regions indicate higher loss values; stars mark theoretical optima.

it discovers the Klein 4-group V_4 as demonstrated in Figure VIII.4b. When the generator is allowed to search for general linear transformations inside $\text{Aff}(2)$, it discovers the full symmetry group that comprises the aforementioned Klein 4-group as well as transformations involving the dilatation matrix $D(r)$ described in Theorem 20.

Figure VIII.6 shows the discovered symmetries for the $N_{1,1}$ distribution and Figure VIII.7 shows the discovered symmetries for $N_{1,2}$ when the generator is parametrised to discover general affine transformations in $\text{Aff}(2)$.




Figure VIII.5: Controlled symmetry breaking in the isotropic Gaussian $\mathcal{N}(\mathbf{0}, \mathbb{I})$ through cyclic constraint enforcement. The loss function is augmented with a mean squared error term $\alpha \|g^q(\mathbf{x}) - \mathbf{x}\|^2$ that enforces q -fold cyclic symmetry. The continuous $SO(2)$ symmetry observed in Figure VIII.4a is discretised into cyclic subgroups: **(a)** $\mathbb{Z}_2$ symmetry with $q = 2$, showing two antipodal maxima at $(c, s) = (\pm 1, 0)$; **(b)** $\mathbb{Z}_3$ symmetry with $q = 3$, displaying three maxima at the cube roots of unity; **(c)** $\mathbb{Z}_7$ symmetry with $q = 7$, exhibiting seven equally spaced maxima on the unit circle. The discovered symmetries (scatter points) precisely align with the q^{th} roots of unity, validating the theoretical prediction that the cyclic constraint induces symmetry discovery of discrete components. Darker regions indicate higher loss values, stars indicate analytic maxima.

VIII.E.1.iii Gaussian Mixture Models

The power of a SYMMETRYGAN becomes more apparent when applied to Gaussian mixture models with complex symmetry groups. Three such examples are considered, inspired by the the examples in [398]. The first is a one dimensional bimodal distribution,

$$B(x) = \frac{1}{2}\mathcal{N}(-1, 1) + \frac{1}{2}\mathcal{N}(1, 1) \quad (\text{VIII.36})$$

which possesses the cyclic group $\mathbb{Z}_2 = \{x \mapsto \pm x\}$ as its symmetry group. The discovered symmetries are shown in Figure VIII.8. The second is a two dimensional




Figure VIII.6: Symmetry discovery in the full affine group $\text{Aff}(2)$ for the isotropic Gaussian $\mathcal{N}(\mathbf{0}, \mathbb{I})$, decomposed via Iwasawa parametrization $g = \sqrt{|d|} I^{\frac{1-\text{sgn}(d)}{2}} \cdot R(\theta) \cdot D(r) \cdot S(u) + \mathbf{v}$. Three orthogonal slices through the 6 dimensional loss landscape reveal the symmetry structure: **(a)** Determinant-rotation space (d, θ) shows continuous maxima along $d = \pm 1$ (vertical red lines) for all rotation angles. **(b)** Dilation-shear space (r, u) exhibits a unique maximum at $(r, u) = (1, 0)$ (red star), indicating that only uniform scaling preserves the isotropic structure. **(c)** Translation space (v_1, v_2) peaks at the origin (red star), verifying that the centred Gaussian admits no translational symmetries. Empirically discovered transformations (yellow points) cluster precisely at the theoretical optima, validating the analytic predictions.

octagonal distribution,

$$O(x) = \frac{1}{8} \sum_{i=1}^8 \mathcal{N}(\cos \frac{2\pi i}{8}, 0.1) \times \mathcal{N}(\sin \frac{2\pi i}{8}, 0.1) \quad (\text{VIII.37})$$

which has the dihedral group of order 16, D_8 , as its symmetry group, a non-commutative discrete group containing both rotations and reflections. The third is a 5×5 square distribution,

$$S(x) = \frac{1}{25} \sum_{i=0}^4 \sum_{j=0}^4 \mathcal{N}(i-2, 0.1) \times \mathcal{N}(j-2, 0.1) \quad (\text{VIII.38})$$

which has the dihedral group of order 8, D_4 , as its symmetry group, also a non-commutative discrete group containing both rotations and reflections. The left column




Figure VIII.7: Symmetry discovery in $\text{Aff}(2)$ for the anisotropic Gaussian $\mathcal{N}(\mathbf{0}, \text{diag}(1, 2))$, revealing discrete symmetry breaking compared to the isotropic case (Figure VIII.6). The anisotropic covariance structure constrains the symmetry group to a supergroup of the Klein 4-group V_4 : **(a)** Determinant-shear space (d, u) exhibits two isolated maxima (red stars) at $(d, u) \in \{(1, 0), (-1, 0)\}$, corresponding to orientation preserving and orientation reversing transformations with no shear. **(b)** Dilation-rotation space (r, θ) shows four discrete maxima (red stars) at $(r, \theta) \in \{(1, 0), (1, \pi), (\sqrt{2}, \pi/2), (\sqrt{2}, 3\pi/2)\}$, representing identity, π -rotation, and the two coordinate involutions scaled by $\sqrt{2}$ to account for the eigenvalue ratio. **(c)** Translation space (v_1, v_2) remains peaked at the origin (red star), confirming no translational symmetries. The discovered transformations (yellow points) cluster at these discrete locations, demonstrating how anisotropy breaks continuous symmetries into a discrete group.

of Figure VIII.9 shows data from each of these distributions, the middle column shows the discovered rotations, and the right column shows the discovered reflections. In each case SYMMETRYGAN is able to correctly discover the symmetries of the distribution.

VIII.E.2 Particle physics experiments.

The transition from Gaussian toy models to HEP data represents a significant increase in complexity.




Figure VIII.8: Symmetry discovery for a one dimensional bimodal Gaussian mixture $p(x) = \frac{1}{2}\mathcal{N}(-1, 1) + \frac{1}{2}\mathcal{N}(1, 1)$ possessing $\mathbb{Z}_2$ symmetry. **(a)** Probability density function showing two symmetric modes at $x = \pm 1$ with equal weights and variances, creating a distribution invariant under reflection about the origin. **(b)** Analytic loss landscape in the (b, c) parameter space for linear transformations $g(x) = b + cx$, with empirically discovered symmetries overlaid (blue points). The landscape exhibits two distinct maxima at $(b, c) = (0, 1)$ and $(0, -1)$, corresponding to the identity and reflection transformations respectively. The discovered symmetries cluster precisely at these theoretical optima, confirming that the SYMMETRYGAN correctly identifies the cyclic group $\mathbb{Z}_2 = \{x \mapsto \pm x\}$ as the complete symmetry group of the distribution.

This study focuses on dijet data. Jets are the collimated streams of hadrons that emerge when high energy quarks or gluons undergo fragmentation, and two jet final states constitute one of the dominant signatures recorded at the LHC. Given an infrared- and collinear-safe clustering algorithm, every jet is assigned a unique 4-momentum. The ensuing kinematic ensemble permits a systematic search for symmetries in the momentum distributions.




Figure VIII.9: Symmetry discovery for two dimensional Gaussian mixture models with non-commutative discrete symmetry groups. **Top row:** Octagonal distribution $p(\mathbf{x}) = \frac{1}{8} \sum_{k=0}^7 \mathcal{N}(\mu_k, 0.1\mathbb{I})$ where $\mu_k = (\cos \frac{2\pi k}{8}, \sin \frac{2\pi k}{8})$, possessing dihedral symmetry D_8 . **(a)** Probability density showing eight equally-spaced modes on the unit circle. **(b)** Discovered rotations cluster at multiples of $\frac{\pi}{4}$, corresponding to the eight rotational symmetries. **(c)** Discovered reflections appear at eight angles, with antipodal points representing the same reflection axis. **Bottom row:** Square lattice distribution $p(\mathbf{x}) = \frac{1}{25} \sum_{i,j=-2}^2 \mathcal{N}((i,j), 0.1\mathbb{I})$ with dihedral symmetry D_4 . **(d)** Probability density showing a 5×5 grid of Gaussian modes. **(e)** Discovered rotations concentrate at $0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}$, identifying the four rotational symmetries. **(f)** Discovered reflections cluster at four angles corresponding to horizontal, vertical, and diagonal reflection axes. The precise alignment of discovered symmetries with theoretical predictions demonstrates SYMMETRYGAN's ability to identify complex non-commutative groups.

This analysis employs the background dijet sample released for the LHC Olympics anomaly detection challenge [287, 374]. Parton showering and hadronisation are simulated with PYTHIA 8.219 [612], while detector effects are modelled using DELPHES 3.4.1 [496, 513]. Final state particles reconstructed by DELPHES are clustered into anti- k_T jets with radius parameter $R = 1$ using FASTJET 3.3.0 [547, 607, 608]. Events must pass a single jet trigger requiring $p_T > 1.2$ TeV; the subsequent analysis is restricted to the two highest- p_T jets in each event, where transverse momentum is defined by $p_T = \sqrt{p_x^2 + p_y^2}$. This experiment was also repeated on the transverse momenta of dijet events from CMS Open Data [378]. SYMMETRYGANs produced similar results on both datasets.

Each event is characterised by the transverse momentum components

$$\mathbf{x} = \begin{bmatrix} p_{1x} \\ p_{1y} \\ p_{2x} \\ p_{2y} \end{bmatrix} \quad (\text{VIII.39})$$

The initial exploration considers transformations from $SO(2) \times SO(2)$, independently rotating each jet in the transverse plane, i.e., the generator is restricted to the block diagonal form

$$g_{\theta_1, \theta_2}(\mathbf{x}) = \begin{bmatrix} R(\theta_1) & 0 \\ 0 & R(\theta_2) \end{bmatrix} \mathbf{x} \quad (\text{VIII.40})$$

Momentum conservation predicts that only simultaneous rotations $\theta_1 = \theta_2$ should preserve the distribution. The empirical results confirm this beautifully: starting from random initializations in $[0, 2\pi)^2$, the learned parameters cluster along the diagonal




Figure VIII.10: Symmetry discovery on transverse momenta of dijet events using a SYMMETRYGAN with a $SO(2) \times SO(2)$ search space, where each jet undergoes independent azimuthal rotation $g_{\theta_1, \theta_2}(\mathbf{p}) = (R(\theta_1)\mathbf{p}_1, R(\theta_2)\mathbf{p}_2)$. **(a)** Distribution of converged parameters (θ_1^f, θ_2^f) from independent training runs initialised uniformly in $[0, 2\pi)^2$. The discovered symmetries cluster along the diagonal $\theta_1 = \theta_2$ (blue line), confirming that only simultaneous rotations preserve the dijet distribution, a direct consequence of transverse momentum conservation. **(b)** Symmetry discovery map $\Omega : (\theta_1^i, \theta_2^i) \mapsto \theta^f$ revealing the loss landscape dynamics. The final rotation angle follows $\theta^f = \frac{\theta_1^i + \theta_2^i}{2}$ when $|\theta_1^i - \theta_2^i| < \pi$ (shown in green), and $\theta^f = \frac{\theta_1^i + \theta_2^i}{2} - \pi$ when $|\theta_1^i - \theta_2^i| > \pi$ (shown in blue). This bisection rule represents the path of steepest ascent in the loss landscape, demonstrating how gradient dynamics naturally discover the physical constraint without prior knowledge.

$\theta_1 = \theta_2$. Figure VIII.10a shows the discovered symmetries for the $SO(2) \times SO(2)$ case. The symmetry discovery map for this system reveals unexpected elegance. The




Figure VIII.11: Visual validation of discovered symmetry in dijet events through transverse momentum distributions. **(a)** Original LHC Olympics showing the two leading jets' momenta (p_x, p_y) in the transverse plane. The circular structure reflects the detector's azimuthal symmetry. **(b)** Same events transformed by a discovered generator from SYMMETRYGAN, showing identical statistical properties. The preservation of the ring structure and angular distribution confirms that the discovered transformation represents a genuine symmetry of the dijet system. Each point represents one event.

mapping from initial to final parameters follows:

$$\Omega(\theta_1, \theta_2) = \begin{cases} \frac{\theta_1 + \theta_2}{2} & |\theta_1 - \theta_2| < \pi \\ \frac{\theta_1 + \theta_2}{2} - \pi & |\theta_1 - \theta_2| > \pi \end{cases} \quad (\text{VIII.41})$$

This bisection of the angular difference represents the path of steepest ascent in the loss landscape, demonstrating how gradient dynamics naturally discover the constraint imposed by momentum conservation, illustrated in Figure VIII.10b.

Expanding to the full $SO(4)$ search space introduces greater complexity. An arbitrary element of $SO(4)$ is the composition of six independent rotations, in




Figure VIII.12: **(a)** and **(b)** Azimuthal angle distributions after applying a random $SO(4)$ rotation to dijet events, demonstrating violation of symmetry. Transformed angles $\tilde{\phi}_1, \tilde{\phi}_2$ exhibit clear non-uniformity. $D_{KL}(p||\mathcal{U}) \approx 0.15$ quantifies this deviation from uniformity, confirming that arbitrary $SO(4)$ rotations do not preserve the dijet distribution. **(c)** and **(d)** Azimuthal angle distributions after applying a discovered symmetry. Transformed azimuthal angles $\tilde{\phi}_1, \tilde{\phi}_2$ exhibit uniformity within statistical fluctuations. $D_{KL}(p||\mathcal{U}) < 10^{-2}$ confirms agreement with the uniform distribution.

the six possible orthogonal 2-planes. The six parameter group does not admit a simple visualisation, and the discovered symmetries no longer lie in any two or three dimensional subspace. Validation requires more sophisticated statistical tests.

Visual inspection can, however, provide a preliminary qualitative plausibility check. Projections of the original and transformed distributions onto physically meaningful observables (jet p_T , azimuthal angles) show excellent agreement Figure VIII.11. Having passed the visual test, one can now apply more discerning tests to ensure the symmetries are not spurious. The transformed azimuthal angles

$$\begin{bmatrix} \tilde{\phi}_1 \\ \tilde{\phi}_2 \end{bmatrix} = \begin{bmatrix} \arctan \frac{g(\theta_1, \theta_2)(p_{1y})}{g(\theta_1, \theta_2)(p_{1x})} \\ \arctan \frac{g(\theta_1, \theta_2)(p_{2y})}{g(\theta_1, \theta_2)(p_{2x})} \end{bmatrix} \quad (\text{VIII.42})$$

should remain uniformly distributed only for true symmetries, and not for any other transformation. This is indeed the case as is shown in Figure VIII.12.

Computing the KL divergence between the transformed and original distributions is another way to test for symmetry. The KL divergence should be zero for true symmetries, and non-zero for any other transformation. While this is never exactly true, Figure VIII.13 shows the KL divergence for the transformations identified by the SYMMETRYGAN approach the irreducible statistical noise floor, while arbitrary transformations from $SO(4)$ clearly do not.

Finally, *post hoc* discriminators trained on discovered transformations achieve losses within 1% of the theoretical optimum $2 \ln 2$ confirming these represent genuine symmetries, illustrated in Figure VIII.14. The discovered subgroup of $SO(4)$ includes the expected $SO(2)$ rotations reflecting cylindrical detector geometry, but also unexpected combinations that would be difficult to identify through physical intuition alone.




Figure VIII.13: Quantitative validation of discovered symmetries via Kullback–Leibler divergence analysis. Histograms show $D_{KL}(\phi||\tilde{\phi})$ between original and transformed azimuthal angle distributions for **(a)** the leading jet and **(b)** the subleading jet. Two distinct populations emerge: discovered symmetries (red) with means of 0.0058 and 0.0090 respectively, approaching the irreducible statistical noise floor (green) at $D_{KL} \approx 10^{-3}$ arising from finite sample size. Random $SO(4)$ rotations (blue) exhibit significantly higher divergences with means of 0.37 and 0.34, confirming that they violate the symmetry. Note the logarithmic scale emphasises the low divergence region, and so the red and green clusters are closer than they may appear.

VIII.E.3 Interpreting discovered symmetries.

The interpretation of discovered symmetries requires bridging the gap between abstract mathematical transformations and physical understanding. This translation process reveals both the power and limitations of automated symmetry discovery. Each discovered symmetry transformation corresponds to a conservation law or invariance principle in the underlying physics. For the dijet system, the discovered azimuthal




Figure VIII.14: Discriminator loss distributions validating discovered symmetries against theoretical optimum. The histogram shows *post hoc* discriminator losses when trained to distinguish original from transformed dijet events. Discovered symmetries (blue) achieve losses tightly clustered around $2 \ln 2 \approx 1.386$ (dashed line), the theoretical minimum for perfect symmetries where transformed and original distributions are indistinguishable. Random $SO(4)$ rotations (black) yield significantly lower losses, indicating the discriminator easily identifies these as non-symmetries. The discovered transformations achieve losses within 1% of the theoretical optimum, providing independent confirmation that the SYMMETRYGAN identifies genuine invariances of the dijet system without prior physics knowledge.

rotations connect directly to angular momentum conservation about the beam axis. The continuous $SO(2)$ symmetry reflects the absence of any preferred direction in the transverse plane—a fundamental consequence of the cylindrical geometry of both the collision process and the detector.

More subtle are the discovered discrete symmetries. The non- $SO(2)$ transformations found combine specific overall rotations with jet exchanges. One such example is

$$g_{\mathbb{Z}_2} \begin{bmatrix} p_{1x} \\ p_{1y} \\ p_{2x} \\ p_{2y} \end{bmatrix} = \begin{bmatrix} -p_{2x} \\ -p_{2y} \\ -p_{1x} \\ -p_{1y} \end{bmatrix}, \quad (\text{VIII.43})$$

a transformation that simultaneously rotates by π and exchanges the two jets—a symmetry that follows from the identical nature of QCD jets but may not be obvious without a systematic search. The relationship between discovered symmetries and known physics provides crucial validation. Each symmetry found by SYMMETRYGAN corresponds to either:

- A fundamental symmetry of physics (Lorentz invariance, gauge symmetry)
- A symmetry of the experimental setup (cylindrical detector geometry)
- An emergent symmetry from kinematic constraints (momentum conservation)

No spurious symmetries appear, demonstrating the method’s reliability.

The next challenge in this space is how to move from discovering individual symmetry transformations to understanding the complete symmetry group structure?

VIII.F Towards symmetry inference.

The distinction between symmetry discovery and symmetry inference represents a fundamental challenge in the SYMMETRYGAN framework. While the method excels at finding individual elements of a symmetry group, reconstructing the complete group structure requires additional theoretical machinery.

The task of symmetry inference can be formulated as “Given a set of discovered symmetry transformations $\{g_1, g_2, \dots, g_n\}$, determine the minimal group G containing all these elements.” Three complementary approaches emerge for tackling this challenge.

Discrete subgroup analysis: By modifying the loss function to enforce specific algebraic relations, a SYMMETRYGAN can systematically probe for discrete subgroups. A SYMMETRYGAN trained with the cyclic group constraint, Equation (VIII.34), successfully identifies Z_q subgroups within larger symmetry groups, as seen in Figure VIII.5. This approach extends naturally to other discrete groups through appropriate constraint terms, though non-Abelian groups require more sophisticated loss modifications.

Group composition methods: The closure property of groups suggests an iterative approach. Given discovered symmetries $\{g_i\}$, one can attempt to compute all possible compositions $g_i \circ g_j$ to expand the known group elements. For continuous groups like $SO(2)$, a single irrational rotation angle generates a dense subgroup, effectively recovering the entire group through composition [488].

However, numerical precision limits this approach. Each composition compounds errors, leading to degradation after multiple iterations. The practical limit appears to be 10 to 20 compositions before accumulated errors overwhelm the signal.

The Symmetry Discovery Map: Perhaps the most promising direction involves learning the complete mapping from initialisation space to symmetry space:

$$\Omega : \mathbb{R}^k \rightarrow \mathcal{G} \tag{VIII.44}$$

where $\mathcal{G}$ represents the symmetry group manifold. This map encodes the complete symmetry structure implicitly, as its image is precisely the symmetry group.

Preliminary attempts to learn Ω have had limited success. The main challenges are that the map must begin near the identity to preserve the connection between initial and final parameters, and the min–max nature of GAN training becomes more complex when optimising over maps on the space of transformations rather than transformations themselves. Despite these challenges, successful learning of even approximate symmetry discovery maps would provide a powerful tool to characterise symmetry groups.

VIII.G Symmetry informed unfolding

VIII.G.1 Symmetry preserving neural network architectures.	573
VIII.G.2 Reducing dimensionality through symmetry identification. .	574
VIII.G.3 Hidden symmetries and emergent simplicity.	576
VIII.G.3.i Formalism	576
VIII.G.4 Applications for unfolding jet substructure variables. . . .	578

The fundamental challenge of unfolding, recovering truth level distributions from detector smeared observations, stems from its ill posed nature. Multiple truth distributions can yield identical observations after detector effects, creating an under-determined system primed for additional constraints. Symmetry provides a natural regularisation scheme for unfolding problems. When a physical process respects a symmetry, that invariance constrains the space of valid solutions, effectively regularising the inverse problem without introducing artificial biases. Consider the standard unfolding problem formulated as an optimisation,

$$\min_{p(z)} \left[\chi^2 \left\{ p(x), \int r(x | z) p(z) dz \right\} + \lambda \cdot \mathcal{R}(p(z)) \right] , \quad (\text{VIII.45})$$

where r represents is detector response and $\mathcal{R}$ is a regularisation term.

Traditional approaches use generic smoothness penalties, but symmetries provide physically motivated constraints. For a symmetry group G discovered through SYMMETRYGAN, one could construct the symmetry preserving regularisation,

$$\mathcal{R}_{\text{sym}}(p) = \sum_{g \in G} \|p - g \cdot p\|^2 . \quad (\text{VIII.46})$$

This penalty vanishes for distributions respecting the symmetry while penalising asymmetric solutions. Such a regularisation term would enforce invariances that the underlying physics demands rather than imposing arbitrary smoothness constraints.

However such an implementation would have to address subtle challenges. For continuous symmetries like $SO(2)$, the sum over group elements becomes an integral requiring careful handling. A possible solution is to select group elements that uniformly cover the symmetry manifold. For $SO(2)$, for instance, one could sample angles $\theta_i = 2\pi i/N$ for $i = 0, \dots, N - 1$. However, any such choice would be *ad hoc* and

would introduce artifacts into the unfolding process, because one would be treating the data as though it had the symmetry $\mathbb{Z}_N$ rather than $SO(2)$, in this example.

For non-compact groups, the challenge is further compounded by the fact that there no longer exists a equivariant sampling strategy in the first place. Methods such as importance sampling based on data distribution do exist [780], but they remain to be tested in the physics context.

The theoretical foundation of this approach connects to information theory. Each symmetry represents a constraint reducing the effective dimensionality of the solution space. For a distribution over N bins with a $\mathbb{Z}_k$ cyclic symmetry, the degrees of freedom reduce from N to N/k , improving the effective statistics by a factor of k . This dimensionality reduction should translate directly to improved unfolding stability.

VIII.G.1 Symmetry preserving neural network architectures.

The architectural encoding of symmetries represents a paradigm shift from *post hoc* constraints to intrinsic guarantees.

Definition 58 (Equivariant neural network.). A neural network f_θ is said to be *equivariant* under the action group G on its input space X if

$$\forall g \in G \quad \forall x \in X \quad f_\theta(g \cdot x) = g \cdot f_\theta(x) . \quad (\text{VIII.47})$$

For unfolding applications, one could seek architectures where the mapping from observed to truth level distributions preserves known symmetries. The Group Convolutional Neural Network (G-CNN) framework provides one example of such an approach [464]. For a symmetry group G acting on the input space, G-CNN layers

take the form

$$f * \psi = \int_H f(h)\psi(g^{-1}h)dh, \quad (\text{VIII.48})$$

where $*$ denotes group convolution and $h \in H$ is the stabiliser subgroup.

This formulation guarantees equivariance by construction, not training.

The LGN (Lorentz Group Network) architecture [314] constructs features from Lorentz scalars and 4-vectors in order to preserve Lorentz invariance. For identical particle handling, Deep Sets [416] and Set2Graph [346] provide principled approaches in which

$$f(\{x_1, \dots, x_n\}) = \rho \left(\sum_i \phi(x_i) \right), \quad (\text{VIII.49})$$

where ϕ and ρ are learned functions, guaranteeing invariance to particle ordering.

For processes involving gauge bosons, architectures must respect gauge transformations. The Gauge Equivariant Mesh CNN [283] provides a framework for $SU(3) \times SU(2) \times U(1)$ invariance.

A hybrid architecture for symmetry aware unfolding could combine these elements. The architecture would then necessarily ensure that certain symmetries would be present in the unfolded distribution. However, the practical implementation of such an architecture remains an open question.

VIII.G.2 Reducing dimensionality through symmetry identification.

The curse of dimensionality poses a significant challenge for unfolding problems, especially for binned approaches. Each additional degree of freedom exponentially

increases the dimension of the solution space, degrading statistical power and amplifying instabilities. Symmetry identification offers a principled method for revealing redundancies that can be eliminated without information loss.

For a distribution with symmetry group G , the effective dimensionality of the distribution can be reduced from the ambient space dimension to the dimension of the quotient space X/G .

Definition 59 (Group invariant function). A function $f : X \rightarrow \mathbb{R}$ is G -invariant if there exists a function $\tilde{f} : X/G \rightarrow \mathbb{R}$ with $f = \tilde{f} \circ \pi$ where $\pi : X \rightarrow X/G$ is the canonical projection.

This factorisation is the key to dimensionality reduction. One needs only unfold $\tilde{f}$ on the lower-dimensional quotient space. The statistical gain from this procedure scales dramatically with dimensionality. For a d -dimensional space with a k -parameter continuous symmetry group, the effective dimensionality reduces to $d - k$, yielding an effective statistics increase of

$$N_{\text{eff}} \sim \frac{N}{C_d}, \quad (\text{VIII.50})$$

where C_d is a complexity measure (e.g. covering number [523], VC dimension [760], local intrinsic dimensionality [255], or metric entropy [457]).

Dimensionality reduction reduces C_d , which increases N_{eff} . Any such process of symmetry based projection, however, requires careful treatment of statistical uncertainties.

VIII.G.3 Hidden symmetries and emergent simplicity.

Often, the most powerful dimensionality reductions come from unexpected symmetries. The discovery of approximate custodial symmetry in electroweak physics simplified phenomenology dramatically. Similarly, emergent symmetries in complex physical ensembles can enable radical dimensionality reduction [611].

The quantitative impact of symmetry based dimensionality reduction admits a rigorous mathematical treatment.

VIII.G.3.i Formalism

Definition 60 (*G*-Equivariant Fisher Information). For a likelihood $\mathcal{L}(x \mid \theta)$ whose parameter space carries an action of a finite or Lie group G , the *G-equivariant Fisher information matrix* is

$$[\mathcal{I}_G(\theta)]_{ij} = \mathbb{E}\left[\frac{\partial \ln \mathcal{L}}{\partial \theta_i} \frac{\partial \ln \mathcal{L}}{\partial \theta_j}\right] - \sum_{g \in G} w_g \mathbb{E}\left[\frac{\partial \ln \mathcal{L}}{\partial \theta_i} \frac{\partial \ln \mathcal{L}(g \cdot \theta)}{\partial \theta_j}\right], \quad (\text{VIII.51})$$

where $\{w_g\}_{g \in G}$ are averaging weights satisfying $\sum_{g \in G} w_g = 1$. Equivalently, $\mathcal{I}_G(\theta) = P_G \mathcal{I}(\theta) P_G^T$ where P_G is the orthogonal projection operator onto the G -invariant subspace [342].

Theorem 21 (Symmetry-Adapted Cramér–Rao Bound [705, 761, 782–785]). Let $\hat{\theta}$ be any unbiased estimator of $\theta \in \mathbb{R}^d$ based on N i.i.d. samples, and assume the statistical model $p(x \mid \theta)$ is G -equivariant. Then the covariance matrix obeys the

lower bound

$$\Sigma(\hat{\theta}) \succeq \frac{1}{N} \mathcal{I}_G^{-1}(\theta), \quad (\text{VIII.52})$$

where $\succeq$ denotes the Loewner partial order. The lower bound is saturated if and only if the score vectors projected onto the invariant subspace are linear in $(\hat{\theta} - \theta)$.

Proof of Theorem 21. (Sketch) The classical information inequality states that $\Sigma(\hat{\theta}) \succeq N^{-1} \mathcal{I}^{-1}(\theta)$. Because the model and the estimator respect the group action, both sides commute with P_G . Projecting the score function onto the invariant subspace and using $P_G^2 = P_G$ yields $\text{Cov}(\hat{\theta}) \succeq N^{-1} (P_G \mathcal{I} P_G)^{-1} = N^{-1} \mathcal{I}_G^{-1}$. $\square$

Theorem 22 (Variance Reduction for Finite Symmetries). For a discrete group G of order $|G|$ if

$$\rho = \frac{\mathbb{E} \left[\frac{\partial \ln \mathcal{L}}{\partial \theta} \frac{\partial \ln \mathcal{L} \circ g}{\partial \theta} \right]}{\mathbb{E} \left[\left(\frac{\partial \ln \mathcal{L}}{\partial \theta} \right)^2 \right]}; \quad (\text{VIII.53})$$

for the average pairwise correlation of score functions across group orbits, then

$$\frac{\text{Var}(\hat{\theta}_{\text{sym}})}{\text{Var}(\hat{\theta}_{\text{unconstrained}})}, = \frac{1}{|G|} + \frac{|G| - 1}{|G|} \rho. \quad (\text{VIII.54})$$

This means that symmetry is maximally beneficial ($\rho \rightarrow 0$) when orbit scores are uncorrelated, recovering the $1/|G|$ ideal reduction factor.

When G is continuous, replacing the group sum with Haar integration yields an integral projector and analogous results. The practical implication is that enforcing equivariance in unfolding algorithms tightens the attainable RMS by precisely the factor dictated by group averaged score correlations.

VIII.G.4 Applications for unfolding jet substructure variables.

Consider unfolding jet substructure observables like N-subjettiness ratios τ_{21} . The IRC safety of these observables implies approximate scale invariance. For a scaling symmetry with parameter λ :

$$\tau_{21}(\lambda p_T, \lambda m_{\text{jet}}) = \tau_{21}(p_T, m_{\text{jet}}) \quad (\text{VIII.55})$$

This symmetry reduces the 2D unfolding problem in (p_T, m_{jet}) to a 1D problem in the ratio m_{jet}/p_T . The maximal possible theoretical precision gain is

$$\sigma_{\tau_{21}}^{\text{sym}} = \frac{\sigma_{\tau_{21}}^{2\text{D}}}{\sqrt{\ln(p_T^{\text{max}}/p_T^{\text{min}})}}. \quad (\text{VIII.56})$$

For typical jet p_T ranges spanning a few orders of magnitude, this offers a tangible improvement in precision.

Such analyses also provides additional insights. Each symmetry can theoretically reduce the Kullback–Leibler divergence between truth and reconstructed distributions by as much as

$$\Delta D_{KL} = \frac{1}{2} \ln |G| - \frac{1}{2N} \text{Tr}(\mathcal{I}_G^{-1} \mathcal{I}_{\text{unconstrained}}) \quad (\text{VIII.57})$$

This reduction could directly translate to improved unfolding fidelity, with larger symmetry groups yielding greater gains.

Such thought experiments point toward a broader principle. The most powerful machine learning approaches for physics don’t just learn from data but incorporate fundamental physical principles into their architecture. Discovering and enforcing symmetries can transforms unfolding from a purely statistical exercise into a physics

informed inference problem, achieving results that respect both the data and the underlying theoretical framework.

VIII.H Improved precision using symmetries

VIII.H.1 Data augmentation using discovered symmetries. 581

VIII.H.2 Symmetry constrained unfolding. 583

The transformation of symmetry from abstract mathematical concept to concrete measurement tool represents one of the most profound opportunities in modern particle physics analysis. This section speculates about how integrating symmetry awareness and unfolding might affect the ability to extract truth from data. These ideas, while not yet implemented, are intended to reflect on potential future directions that translate discovered invariances into measurement precision.

VIII.H.1 Data augmentation using discovered symmetries.

The concept of data augmentation through symmetry sits at the intersection of theoretical physics and practical statistics. The flip side of the dimensionality reduction discussed in Section VIII.G is that when a dataset respects certain transformations, the effective statistics of that dataset can be multiplied without collecting a single additional event. This is not merely a computational trick, it reflects a deep truth about the information structure of symmetric systems.

Symmetry based data augmentation is the process of generating statistically equivalent synthetic data by applying discovered symmetry transformations to existing events, effectively increasing sample size while preserving all physical properties. Traditional data augmentation in machine learning often relies on heuristic transformations. These modifications hope to capture invariances without formal guarantees. By contrast, symmetry based augmentation leverages proven invariances discovered through methods like SYMMETRYGAN. Every augmented event is, by construction, as physically valid as the original.

As with the dimension reduction discussed in Section VIII.G, the the sampling strategy can introduce challenges. For finite groups, one might include all group elements, and for compact continuous groups so long as one samples carefully from the uniform measure on the group manifold to avoid bias, the artifacts introduced will, in the very least, be unbiased. For non-compact groups, the uniform measure is not well defined, and there is no way to guarantee that the artifacts introduced are unbiased.

The power of symmetry based augmentation extends beyond simple counting statistics. Symmetry augmentation can heal detector defects. Imagine a calorimeter module at $\phi = \pi/2$ suffering from reduced efficiency. With symmetry augmentation, one can rotate events from healthy detector regions to fill the gap:

The corrected distribution would be given by

$$n_{\text{corrected}}(\phi = \pi/2) = \frac{1}{|G|} \sum_{g \in G} \epsilon(g^{-1} \cdot \phi) \cdot n_{\text{observed}}(g^{-1} \cdot \phi) , \quad (\text{VIII.58})$$

where $\epsilon(\phi)$ represents the ϕ -dependent efficiency.

This is more than mere interpolation. It is the use of physical symmetry to transport information across phase space.

The interplay with machine learning architectures reveals another dimension. Neural networks trained on symmetry augmented data should exhibit

Faster convergence: The effective sample size increase should translate directly to fewer training epochs.

Better generalisation: The network ‘sees’ the full symmetry orbit during training, preventing overfitting to specific configurations.

Automatic equivariance: Even non-equivariant architectures should learn equivariant information from augmented data.

For all the benefits that symmetry aware augmentation might provide, not all symmetries admit straightforward augmentation. Even with simple discrete symmetries like a $\mathbb{Z}_2$ while one can invert event coordinates, care is needed with derived quantities. Or for instance, a jet clustering algorithm might produce different multiplicities when applied to parity reflected events, breaking the supposed equivalence. In this sense the augmentation must respect the full analysis chain, not just the raw kinematics, and symmetries that admit all the steps of the analysis chain are, naturally, more difficult to find.

VIII.H.2 Symmetry constrained unfolding.

A variant of MOMENT UNFOLDING where the functional basis is restricted to symmetry respecting functions could guide choices of basis functions grounded in physical principles. Chapter V discussed the p_T dependent Boltzmann-like basis functions, $\beta_a(p_T)$, which were approximated as linear functions of jet p_T ,

$$\beta_a(p_T) = \beta_a^{(0)} + \beta_a^{(1)} p_T, \quad (\text{VIII.59})$$

to constrain the flexibility of the generator network. A symmetry aware parametrisation of the $\beta_a(p_T)$ could provide a more principled way to constrain the generator network, while still achieving the necessary regularisation.

Implementing a similar approach in RAN architectures would require careful attention to the increased complexity of the architecture, but a symmetry aware

regularisation scheme could provide improved conditioning, by making the response matrix block-diagonal in symmetry sectors, and imposing physical consistency constraints where unfolded distributions automatically respect, for example, IRC safety, providing smoothness within symmetry sectors without violating invariance.

The quantitative impact of symmetry awareness on measurement precision admits rigorous analysis through the lens of information theory and statistical estimation. The information content of data about parameters θ under symmetry group G is captured by the equivariant Fisher information matrix $\mathcal{I}_G(\theta)$, which bounds the achievable precision of any unbiased estimator. The fundamental result stems from the Equivariant Cramér-Rao Bound discussed in Section VIII.G.

In addition, symmetry awareness provides more than variance reduction. It fundamentally alters the bias-variance tradeoff. Traditional unfolding must balance bias from regularisation and variance from statistical fluctuations.

Symmetry constraints modify the bias-variance tradeoff,

$$\text{Bias}_{\text{sym}} = \frac{\text{Var}_0}{n \cdot S(\mathcal{G})} , \quad (\text{VIII.60})$$

where $S(\mathcal{G})$ is the ‘symmetry factor’ quantifying variance reduction, so that the total error becomes

$$\text{MSE}_{\text{total}} = \text{Bias}_{\text{sym}}^2 + \frac{\text{Var}_0}{n \cdot S(\mathcal{G})} . \quad (\text{VIII.61})$$

In this sense, symmetry aware unfolding is not only about improved statistics, but also about extracting the maximum physical information from finite data.

VIII.I Conclusion.

VIII.I.1 Beyond linear symmetries.	586
VIII.I.1.i HEP applications.	587
VIII.I.2 Approximate symmetries and symmetry breaking.	587
VIII.I.3 A unified framework.	589

This chapter began with the seemingly simple question of what it means for a dataset to possess symmetry, only to wade through layers of mathematical formalism involving Jacobian factors, inertial reference densities, and the interplay between statistical and physical invariances. This final section looks forward to new territory ready for exploration.

VIII.I.1 Beyond linear symmetries.

The restriction to affine equiareal maps, while yielding impressive results, also represents a clear limitation of current analyses. Nature's symmetries extend far beyond the affine group, encompassing transformations that twist, fold, and reshape phase space in ways that linear maps cannot capture. Consider the Hénon map [753], a deceptively simple area preserving transformation,

$$g(x, y) = (x, y - x^2) . \quad (\text{VIII.62})$$

The Hénon map is but one of a vast and relatively uncharted space of the full (nonlinear) *equiareal group*. The structure of the equiareal group is complex enough to allow for a rich theory of symmetries and their associated invariances with the right computational tools.

In particle physics, such transformations emerge naturally, and are ubiquitous, from the rapidity transformation in relativistic kinematics, to the symplectic structure imposed by Poisson brackets on classical phase space and the corresponding commutation relations of quantum mechanics. Normalising flows, on account of their bijective nature, could be a natural fit for exploring symplectic geometries

and their associated symmetries. That said, the greater challenge lies not in the architecture but in the search space. Linear transformations form a finite dimensional Lie group, amenable to systematic exploration. Non-linear symplectic maps form an infinite dimensional space that lacks a known parametrisation or classification. Recent work on Hamiltonian neural networks [367] suggests one possible path to describe transformations through generating functions that automatically preserve symplectic structure.

VIII.I.1.i HEP applications.

A few examples of particle physics applications of non-linear symmetries are

Jet substructure: The Lund plane reveals approximate scale invariant symmetries best expressed through logarithmic transformations [396].

Heavy flavor physics: Dalitz plot analyses exhibit non-linear symmetries from resonance dynamics [657].

Multiparticle correlations: Collective flow in heavy ion collisions follows non-linear hydrodynamic symmetries [474].

The path forward requires synergy between differential geometry, machine learning, and physics intuition, a synthesis that will push the boundaries of all three fields.

VIII.I.2 Approximate symmetries and symmetry breaking.

Perfect symmetries, while mathematically elegant, rarely survive contact with experimental reality. The universe abounds in approximate symmetries—patterns

that almost hold, invariances with small violations, conservation laws with small deviations. Understanding and exploiting these near-symmetries represents perhaps the most important frontier for practical applications.

An approximate symmetry is a transformation under which probability distributions remain invariant up to small, controlled deviations, characterised by a symmetry-breaking parameter $\epsilon \ll 1$. The mathematical framework for approximate symmetries builds on perturbation theory. A transformation g_ϵ depending on breaking parameter ϵ is an approximate symmetry if

$$p(g_\epsilon(x))|g'_\epsilon(x)| = p(x) + \epsilon \Delta p(x) + O(\epsilon^2) . \quad (\text{VIII.63})$$

The firstorder breaking term $\Delta p(x)$ encodes how symmetry violation depends on phase space location, crucial information for precision measurements.

Consider isospin symmetry in nuclear physics—preserved at low energies and progressively violated as electroweak effects become important. A symmetry aware unfolding framework must gracefully handle this energy dependent breaking:

$$\mathcal{L}_{\text{approx}} = \mathcal{L}_{\text{data}} + \lambda(E) \mathcal{R}_{\text{isospin}} \quad (\text{VIII.64})$$

where $\lambda(E) \propto \exp(-E/\Lambda_{\text{QCD}})$ weakens the constraint at high energies [168].

A soft symmetry breaking framework could be promising. Instead of binary classification (symmetric or not symmetric), one could attempt to quantify the degree of symmetry through continuous measures, such as

$$S(p, g) = 1 - \frac{D_{KL}(p \| g \cdot p)}{\ln 2} . \quad (\text{VIII.65})$$

Such a symmetry score $S \in [0, 1]$ could enable gradient based optimisation while naturally handling approximate invariances.

VIII.I.3 A unified framework.

Measurement, inference, and physical understanding all form parts of a coherent self-consistent cycle of inference. Symmetries discovered from data improve unfolding precision, and better unfolding reveals cleaner symmetry patterns. Every symmetry provides a consistency check, where violations of expected symmetries signal systematic uncertainties, methodological gaps, or perhaps, new physics. The discovered symmetries are also themselves physics outputs. An understanding of which symmetries hold and which ones are broken in any given situation are not only of applied interest, but also of fundamental theoretical interest in the study of the laws of physics.

Looking ahead, several concrete developments seem within reach including automated analysis pipelines that start with raw detector data, automatically discover relevant symmetries, perform symmetry aware unfolding, quantify uncertainties, all in a self-consistent unbinned chain of inference. The marriage of machine learning's ability to detect and learn patterns with physics' principled reasoning promises new paradigms of analysis, where every symmetry discovered is both a practical tool and a theoretical insight.

Chapter IX

Conclusion.

IX.A	A synthesis of unfolding methods.	592
IX.B	A unified framework for modern unfolding.	595
IX.B.1	Overview.	596
IX.B.2	The interplay of physics and learning.	597
IX.B.3	Common themes across methods.	598
IX.C	Advances in cross section measurement.	603
IX.C.1	Summary of contributions.	604
IX.C.2	Impact and advances.	606
IX.D	Applications to current and future experiments.	610
IX.E	Open challenges and future directions.	616
IX.E.1	Challenges in unfolding methodology.	617
IX.E.2	Handling correlations in unbinned inference.	618
IX.E.3	Background treatment in complex environments.	618

IX.E.4	Systematic uncertainty propagation in ML methods.	619
IX.E.5	Future directions.	620
IX.E.6	Use on real experimental data.	622
IX.F	Closing: The broader vision	624

IX.A A synthesis of unfolding methods.

The measurement of differential cross sections lies at the heart of experimental particle physics, providing our most direct window into the fundamental interactions governing nature. Yet between the elegant simplicity of theoretical predictions and the messy reality of experimental observation stands the formidable challenge of *unfolding*, the process of extracting particle-level truth from detector-level measurements. This dissertation has explored a progressive sequence of machine learning approaches to this inverse problem, each method expanding our capacity to extract physical insight from increasingly complex experimental data. Every measurement in particle physics is filtered through imperfect detectors that smear, misdirect, and incompletely capture the particles we seek to study. Fundamentally, unfolding seeks to invert this detector response function: given observed data $p(x)$, recover the true distribution $p(z)$ that nature produced.

Traditional approaches to unfolding have relied on *discretisation* i.e. dividing continuous distributions into bins and solving a finite dimensional linear inverse problem. While mathematically tractable, this strategy imposes severe limitations. Information is lost through binning, correlations between observables become difficult to preserve, and the curse of dimensionality restricts most analyses to one or two variables. As we push toward precision measurements of increasingly subtle phenomena, such as jet substructure, effective field theory parameters, or signatures of new physics, these limitations become untenable. This work presents a progressive evolution of unfolding methodologies, each building upon the insights of its predecessors while addressing specific limitations.

NEURAL POSTERIOR UNFOLDING (NPU) enhanced traditional binned approaches by replacing point estimates with full posterior distributions, leveraging normalising flows to capture complex uncertainty structures while maintaining the computational advantages of amortised inference. NPU as a method embodies the recognition that unfolding is fundamentally about *uncertainty quantification*, not just central value estimation, and highlights the necessity of machine learning based approaches as the future of unfolding.

MOMENT UNFOLDING took a different path, abandoning bins entirely in favor of directly deconvolving distribution moments. By reformulating unfolding as a moment-matching problem solved through adversarial training, this approach demonstrated that meaningful physical quantities could be extracted without discretization, through a method that was simultaneously light, fast, and precise, opening the door to truly multidimensional analyses. **Reweighting Adversarial Networks** completed the transition to fully unbinned methods, learning to reweight simulated events to match observed data distributions. This approach preserves the complete information content of individual events, enabling extraction of arbitrary observables post-unfolding—a capability impossible with any binned method. The investigation of **unbinned correlations** revealed both the promise and peril of these advanced methods. While unbinned techniques maximise information extraction, they introduce subtle statistical correlations between events that invalidate standard uncertainty estimation procedures. Understanding and addressing these correlations proves essential for rigorous statistical inference.

Each methodological advance expands what can be meaningfully extracted from experimental data. Where binned methods see discrete histograms, unbinned approaches reveal continuous probability densities. Where traditional techniques handle single observables, modern methods unfold entire phase spaces.

This progression is not merely technical innovation. It represents a fundamental shift in the conceptualisation measurement in particle physics. Rather than forcing nature's continuous phenomena into predetermined bins, the methods presented in this dissertation let the data speak in its native language of probability densities and event-level information. The implications extend beyond improved precision; entirely new classes of measurements are opened up by methods that are more naturally suited to the data.

The methods developed here address the twin challenges facing modern particle physics, *higher precision* for constraining fundamental parameters and *dimensional scalability* for analysing complex final states. As the era of high luminosity physics dawns, with its unprecedented data volumes and discovery potential, these capabilities transform from useful tools to essential infrastructure for the field's continued progress.

IX.B A unified framework for modern unfolding.

IX.B.1	Overview.	596
IX.B.2	The interplay of physics and learning.	597
IX.B.3	Common themes across methods.	598
IX.B.3.i	Regularisation.	598
IX.B.3.ii	Uncertainty quantification.	599
IX.B.3.iii	Computational trade-offs.	600

IX.B.1 Overview.

Despite their diversity, the machine learning approaches developed in this dissertation emerge from a common theoretical foundation that reconceptualises unfolding not simply as the inversion of point bin count estimates, but as a problem of learning probability distributions under constraints. This unified perspective illuminates both the connections between methods and the fundamental principles governing successful unfolding strategies.

All unfolding methods can be understood as solving the same optimisation problem of finding the particle level distribution $p(z)$ that, when forward-folded through the detector response, best matches the observed data while respecting domain knowledge and maintaining statistical regularity. What distinguishes each approach is not the problem they solve, but rather how they parametrise the solution space and what constraints they impose.

NPU parametrises the posterior distribution directly using normalising flows. MOMENT UNFOLDING constrains only certain integral properties of the distribution while leaving the detailed shape free. RANs learn reweighting functions at particle level that transform detector level simulation to match data through more subtle constraints governed by the spectral normalisation and the Wasserstein metric. Each choice reflects a different balance between flexibility and regularisation, between data driven discovery and physics informed constraints.

The methods detailed in Chapter VII detail the challenges of the downstream analysis of unbinned unfolded data, and propose methods to address them. If the field is to adopt unbinned methods, as seems inevitable, it must also establish a

statistically rigorous framework for downstream inference and uncertainty estimation on such data.

IX.B.2 The interplay of physics and learning.

Traditional unfolding methods embed physics knowledge explicitly through response matrices computed from detailed simulation, regularisation terms encoding smoothness assumptions, or iterative procedures with early stopping before convergence. Machine learning approaches have sometimes been criticised for being opaque and for abandoning field specific insights in favour of pure data-driven optimisation. The methods presented in this dissertation seek to provide an interpretable suite of methods developed by engaging in a more nuanced balance between domain knowledge and statistical learning.

Each method presented involves a careful attempt to connect novel statistical and computational insights to the physics represented in the data. In NPU, the structure of the normalising flow itself encodes invariances. The transformation must be bijective, preserving the total probability. The choice of base distribution and flow architecture implicitly constrains which solutions are easily reachable, embedding a form of physics motivated regularisation. The method learns from data, but within a framework shaped by physical considerations. MOMENT UNFOLDING's design is even more explicitly motivated by physics. The decision to unfold moments itself is physics motivated, in that the ability of lattice QCD calculations to predict the scaling law behaviour of various jet observables only at the level of moments sets the quantities of interest. The architecture too, inspired by and structured like the Boltzmann

distribution, is manifestly physics inspired. The solution to a statistical mechanics problem in the 19th century inspired a solution to a cross section measurement problem in the 21st century. RANs too exemplify this synthesis in their network architecture. Further, both MOMENT UNFOLDING and RAN learn to reweight events from physics simulation, encoding the field’s best understanding of the underlying processes.

For each method presented, physics provides the scaffolding for the machine learning approach.

IX.B.3 Common themes across methods.

A few overarching themes emerge as valuable lessons from the studies in this dissertation.

IX.B.3.i Regularisation.

Any proposed unfolding method must think deeply and carefully about *regularisation*. Although the specific strategies for regularisation may manifest differently in each method, regularisation is an unavoidable universal feature of unfolding, and unacknowledged subtle forms of regularisation in analysis methods can have perilous effects on the statistical assumptions that undergird measurements.

Where traditional methods add explicit regularisation terms as smoothness penalties, ML approaches tend to encode regularisation through architectural choices such as network depth, activation functions, training procedures. This form of regularisation

is implicit but no less real. That said, machine learning approaches can also be regularised more explicitly, through the use of regularisation terms in the loss function.

NPU regularises through the constraint on normalising flows that they only learn bijective maps, and through early stopping. MOMENT UNFOLDING has a highly constrained generator that is only able to generate distributions that match the moments of the data, with a small fraction of the capacity of a fully flexible generator. The RAN method uses the Wasserstein distance and gradient penalties to ensure stable training. In addition to this, the unconstrained generator in a RAN in comparison to the constrained generator of MOMENT UNFOLDING requires further regularisation through methods such as spectral normalisation or entropy regularisation in order to prevent training instability. Despite their diversity, regularisation in all these forms serves the same purpose—preventing the ill posed inverse problem from amplifying statistical fluctuations into unphysical features.

IX.B.3.ii Uncertainty quantification.

Uncertainty quantification is a central and essential component of any unfolding method for cross section measurements. Where traditional methods often provide only point estimates with error bars on bin counts, modern approaches like NPU naturally and directly model the posterior distribution as an intrinsic feature of the method. MOMENT UNFOLDING and RAN preserve event level information, enabling nuanced uncertainty propagation.

In general, the shift from binned to unbinned unfolded cross section results poses a fundamental challenge to uncertainty quantification for parameter estimation,

and demands a new statistical framework that can appropriately account for the correlations introduced by the unfolding process.

IX.B.3.iii Computational trade-offs.

This leads to the third theme of *computational trade-offs* as an unavoidable reality of any unfolding method. Every method navigates the tension between expressiveness and tractability, between bias and variance, between speed and precision.

More flexible models can capture complex distributions but require more data and computation, as well as more complex regularisation schemes. Constrained models train efficiently but are limited in the information they can capture. Traditional unfolding methods are orders of magnitude faster than their unbinned counterparts and have well known asymptotics for parameter estimation and uncertainty quantification, but they only provide point estimates of bin counts, introduce artifacts into the unfolded data results and are limited to low dimensional analysis. Modern ML based approaches are significantly more computationally expensive, and pose challenges for downstream inference, but are able to jointly unfold multiple observables, thereby capturing inter-observables correlations, and do not introduce binning artifacts into the unfolded results.

Within the methodological framework of any given method too, the same trade-offs manifest more locally. Whether in the choice of computational resources devoted to training a neural network, or specific binning scheme chosen for a traditional method, every analysis involves almost a fractal of increasingly low level and local trade-offs that guide and shape the analysis.

These common themes express a few important lessons for the development and evaluation of any unfolding method. In both the development and evaluation of any unfolding method *prior dependence* remains important in some form. Even methods that make theoretical guarantees of converging to unbiased estimators, make those guarantees based on assumptions such as infinite, independent, and identically distributed data, training until theoretical convergence, *et cetera*. In practice, the output of every method depends on the initial simulation or assumption in some way. The key, therefore, is to make this dependence explicit and testable rather than buried in algorithmic choices. Stress tests with deliberately misspecified initial conditions help quantify robustness.

The architecture of the method matters as much as the algorithm. The choice of neural network structure, loss functions and training procedures embeds implicit assumptions about the problem. The most successful methods align these architectural properties with known physics constraints.

Finally, the testing of modern unfolding methods requires the development of suitable validation procedures that go beyond traditional closure tests. The methods require validation approaches that test robustness to misspecification, stability under resampling, and consistency across observables. Unbinned unfolded data requires particularly careful treatment of statistical dependencies during parameter inference to avoid misestimation of coverage. Successful unfolding requires carefully balancing flexibility with constraint, expressiveness with regularisation, and physics insight with data driven discovery. The methods in this dissertation explore different points in

this space, but all navigate the same fundamental tensions inherent in extracting truth from imperfect measurement.

IX.C Advances in cross section measurement.

IX.C.1	Summary of contributions.	604
IX.C.1.i	NEURAL POSTERIOR UNFOLDING (NPU)	604
IX.C.1.ii	MOMENT UNFOLDING	604
IX.C.1.iii	Reweighting Adversarial Networks	605
IX.C.1.iv	Unbinned inference on correlated data.	605
IX.C.1.v	SYMMETRYGAN	606
IX.C.2	Impact and advances.	606

This dissertation presents major methodological advances that collectively advance the ability to extract precise physical information for cross section measurements. Each contribution addresses specific limitations of existing approaches while opening new avenues for physics analysis.

IX.C.1 Summary of contributions.

The novel contributions developed in this work span the full spectrum of unfolding challenges, from enhancing traditional binned approaches to enabling entirely new measurement paradigms.

IX.C.1.i NEURAL POSTERIOR UNFOLDING (NPU)

This method addresses a significant limitation in extant binned unfolding methods by replacing traditional point estimators of bin counts with conditional normalising flows that directly model the posterior distribution. A key innovation lies in using amortised inference. Once trained, NPU processes new datasets in seconds rather than minutes, achieving a $13\times$ speed up over Fully Bayesian Unfolding while maintaining identical statistical coverage. NPU’s normalising flow architecture naturally handles degenerate response matrices where traditional methods fail catastrophically, providing proper uncertainty quantification in unconstrained regions of phase space.

IX.C.1.ii MOMENT UNFOLDING

Rather than reconstructing entire distributions, this approach directly unfolds the first n statistical moments using a Boltzmann distribution inspired reweighting

function. The generator architecture,

$$g(z, \boldsymbol{\beta}) = \frac{1}{P(\boldsymbol{\beta})} \exp \left(\sum_{a=1}^n \beta_a z^a \right), \quad (\text{IX.1})$$

embeds the maximum entropy principle, ensuring that the unfolded distribution is the most conservative distribution consistent with measured moments. This targeted approach achieves comparable or better precision than OMNIFOLD for moment extraction while requiring about 1% of the computational resources.

IX.C.1.iii Reweighting Adversarial Networks

Extending MOMENT UNFOLDING to the continuum limit, a RAN learns a flexible reweighting function that morphs simulation to match data across the entire phase space.

It's novel contribution is its ability to unfold entire (multidimensional) spectra while performing reweighting at particle level using gradients from the Wasserstein distance between measured data and simulation at detector level. In this way RANs enable robust unfolding even when detector level distributions have minimal overlapping support. In jet substructure measurements, RANs achieve comparable or better accuracy than OMNIFOLD for challenging observables with a variety of different features and scaling behaviours, while being an order of magnitude more computationally efficient.

IX.C.1.iv Unbinned inference on correlated data.

This work provides the first comprehensive framework for understanding how unfolding induced correlations between events fundamentally challenges the indepen-

dence assumption underlying standard likelihood construction for unbinned data. The analysis reveals that naive unbinned inference can significantly misestimate uncertainties with the misestimation magnified with worsening detector resolution. The proposed correlation aware inference methods restore proper statistical coverage while preserving the precision advantages of unbinned approaches.

IX.C.1.v SYMMETRYGAN

While not implemented for unfolding in this dissertation, SYMMETRYGAN introduces a differentiable framework for discovering continuous symmetries directly from data. The method employs adversarial training to learn transformations that leave the data distribution invariant, providing a novel, flexible and fully differentiable framework for discovering symmetries directly from data. This opens the door to physics-informed unfolding that automatically incorporates discovered invariances.

IX.C.2 Impact and advances.

These advances collectively enhance the ability to measure differential cross sections across multiple dimensions.

NPU solves the catastrophic failure of traditional unfolding methods to provide proper uncertainty quantification in unconstrained regions of phase space by enabling the extraction of the full posterior distribution of the unfolded data. While methods such as Fully Bayesian Unfolded are also able to provide proper uncertainty quantification in unconstrained regions of phase space, they are significantly more computationally expensive. By using amortised inference, NPU is able to process

new datasets in seconds rather than minutes, achieving a $13\times$ speedup over Fully Bayesian Unfolding while maintaining identical statistical coverage. The amortised inference capability transforms how experimentalists can handle systematic uncertainties. Where traditional methods require complete re-training for each systematic variation, NPU processes hundreds of variations in the time previously needed to process one. This property is very valuable for precision measurements where systematic uncertainties dominate and require extensive variation studies.

MOMENT UNFOLDING targets the statistical moments of cross section measurements, which are the quantities that are most often used in physics analyses, enabling direct comparison between experimental data and first-principles calculations. By unfolding the moments, MOMENT UNFOLDING is able to provide a prediction that is significantly more precise than IBU, while requiring only 1% of the computational resources OMNIFOLD requires to achieve the same precision. It also abandons the arbitrary iteration cut-offs of iterative unfolding methods, and instead uses an adversarial training procedure to unfold the data by training a single pair of networks to convergence. As a fast, light, and precise unbinned unfolding method, it is a natural candidate for unfolding substructure observables.

RANs extend MOMENT UNFOLDING to the continuum limit, enabling the unfolding of the full distribution of multidimensional, multi-observable distributions. By opening up the multidimensional phase space, eliminating the curse of dimensionality inherent to binned methods, RANs enable the study of correlations in high dimensional joint densities that are lost in projection during binned unfolding. RANs are particularly useful for observables where the full distribution of the cross section is necessary to

extract the physics of interest. Like MOMENT UNFOLDING, a RAN is non-iterative, and learns perturbations to the MCMC, effectively initiating the networks with a huge swathe of physics knowledge rather than having them learn to generate the entire distribution from scratch. RANs also offers a fundamental improvement over existing methods by being able to unfold distributions with poor detector level overlap as long as they have good particle level overlap, enabling the exploration of previously inaccessible regions of phase space.

The theoretical understanding of event correlations developed in Chapter VII is essential for extracting reliable physics parameters from unbinned data and ensuring valid inference. Without an suitable rigorous framework for parameter estimation and inference, the predictions from unbinned methods would be plagued by incorrect downstream estimates. SYMMETRYGAN demonstrates a pathway towards physics informed ML for physics. The ability to discover symmetries from data suggests a future where machine learning methods automatically incorporate physical constraints. For cross section measurements, this could reduce effective dimensionality and improve precision.

These methods together create a comprehensive toolkit that addresses major challenges in modern cross section measurements: precision, dimensionality, efficiency, and theoretical interpretability. The methods are not merely incremental improvements but represent fundamental advances in how we extract physics from data. Looking forward, these contributions establish the foundation for the next generation of cross section measurements. As experiments generate data at unprecedented rates and the need for ever more differential measurements grows, the ability to efficiently unfold

high dimensional distributions while maintaining rigorous uncertainty quantification will transition from advantageous to essential. The methods developed here are a step towards ensuring that data analysis capabilities keep pace with both experimental advances and the expanding scope of theoretical questions they seek to answer.

IX.D Applications to current and future experiments.

The methodological advances developed in this dissertation provide novel and effective tools for analyses at ongoing and upcoming high energy physics experiments. While the studies in this dissertation were conducted on simulated proton–proton collisions, the approaches are general and directly applicable to real experimental analyses at modern experiments. As experimental collaborations grapple with the challenges of unfolding in complex environments, the tools presented here offer novel approaches to emerging challenges.

In looking forward to the High–Luminosity LHC (HL–LHC) and future experiments, this dissertation provides a suite of methods well suited to the challenges of ever more voluminous data and increasing precise and subtle analysis.

Although these methods have not yet interfaced with specific experimental analyses but there is no fundamental barrier to doing so. Each method addresses particular needs faced by experiments, as summarised in Table IX.1, and the methods are primed for validation and adoption by the experimental community.

Table IX.1: Potential applications of the methods developed in this dissertation to current and future experimental analyses. Each approach addresses specific challenges: NPU for rapid, rigorous uncertainty quantification; MOMENT UNFOLDING for fast comparison to theory moments; RAN for full spectrum unfolding in high dimensions; and SYMMETRYGAN for uncovering and leveraging symmetries in data.

Method	Potential use	Key features
NPU	Precision measurements with extensive systematic variations (e.g. differential cross sections at LHC/HL-LHC).	Provides full posterior for uncertainties; amortized inference enables rapid re-unfolding for many systematic scenarios.
MOMENT UNFOLDING	Theoretically-targeted observables (e.g. jet substructure moments, EFT parameter fits).	Unfolds only the moments of interest with maximal efficiency; avoids binning and captures essential physics features.
RAN	High-dimensional distributions and complex final states (e.g. multi-jet kinematics, full jet substructure shapes).	Preserves complete event information, enabling truly unbinned, multi-variable unfolding; opens previously intractable phase-space regions to measurement
Unbinned Correlation Analysis	All unbinned parameter estimation tasks (fits to unfolded distributions, differential cross section measurements, <i>et cetera</i>)	Provides a framework for understanding and quantifying the correlations introduced by the unfolding process, and for developing methods to account for them.
SYMMETRYGAN	Data-driven symmetry discovery in collider data	Identifies hidden invariances directly from data; can inform and constrain unfolding or other analyses by incorporating newly found symmetries

One of the most immediate impacts of these methods will be in enabling high-dimensional and high-precision measurements that have traditionally been beyond reach. Modern HEP analyses increasingly consider multiple correlated observables, but unfolding such multi-differential distributions has been infeasible with conventional techniques. By eliminating binning and harnessing the full event information, an approach like the RAN approach allows experiments to unfold an entire phase space at once. This means that an analysis need not project onto one variable at a time. Instead, collaborations could directly measure joint distributions, preserving correlations that are essential for precise theory comparisons. With RANs, such richly differential results become attainable, enabling more incisive tests of QCD models and parton-shower Monte Carlo tuning.

Likewise, previously hidden features of the data—a slight excess of events in a tail here, a subtle correlation between angles and energies there—can be faithfully captured. In the search for new physics, these capabilities translate into a higher discovery potential. If a Beyond Standard Model (BSM) signal manifests as a distortion of a kinematic shape or a novel correlation between observables, an unbinned unfolding will not wash it out. In short, by addressing the limitations of low dimensional histograms, RANs, and unbinned methods generally, empower experiments to measure what was once unmeasurable.

In parallel, these methods address the pressing demand for *improved precision and uncertainty management* in current and future experiments. The HL-LHC will deliver an enormous dataset,¹ driving statistical uncertainties to unprecedented lows. The

¹On the order of 3 ab^{-1} of proton-proton collisions.

challenge is that systematic uncertainties, detector calibrations, modelling of physics processes, *et cetera*, will become the dominant limiting factor for many measurements. NPU directly speaks to this challenge. By outputting the full posterior distribution over unfolded results, NPU provides a complete uncertainty quantification that experiments can propagate into any downstream inference. This is a departure from the status quo of error bars and covariance matrices—it is a richer characterisation of uncertainty, crucial when subtle effects are being probed. Moreover, NPU’s amortised inference brings a transformative practical advantage. Once trained, the model can re-unfold new data in seconds. In an experimental workflow, this means that hundreds of systematic variations² can be processed almost instantly, where with traditional unfolding each would require a full re-run of an iterative algorithm or regularised matrix inversion.

The advantages of these approaches are not confined to proton–proton collisions or high–luminosity environments. Because they make minimal assumptions about the underlying physics (aside from requiring accurate simulation for training), they can be applied across a wide range of experimental frontiers. For instance, heavy ion collision experiments at RHIC and the LHC provide an entirely different testing ground where one is interested in the modifications of jets and particle distributions by the quark–gluon plasma. Unfolding in heavy ion is notoriously difficult due to the immense underlying event background and the fact that the ‘truth’ itself (the hard scatter before quenching) is not directly observable. While these methods have not yet been tested on heavy ion data, real or simulated, one can envision using

²E.g., varying tracking efficiency, jet energy scale, background shapes, *et cetera*

a technique like RAN to reweight a simulated dataset to match the distributions observed in heavy ion collisions, and potentially, through allowing negative weights, simultaneously enable background subtraction as well.

In doing so, the algorithm would effectively be learning the distribution of medium-induced energy loss and distortion that needs to be applied to events to reproduce the heavy ion measurements. The outcome could be an unfolded characterisation of quenching e.g., a probability distribution of energy loss for jets, or the full distribution of some jet shape after correcting for detector and medium effects. Such an analysis would offer far more detail than comparing binned pp vs. heavy ion spectra. It could reveal, for example, how the tail of the jet mass distribution is softened by the medium across the entire shape, information crucial for understanding the underlying QCD dynamics. Similarly, in the realm of lepton colliders,³ these methods can be invaluable. Although lepton collisions have cleaner environments, the emphasis there will be on ultra-precise measurements of Standard Model processes.⁴ Unbinned unfolding can maximise the information extracted from each event, expanding the precision frontier. For example, an e^+e^- collider experiment could unfold the angular distributions of decay products in a multi-body Higgs decay with full correlations, achieving a level of differential precision that would sharpen sensitivity to subtle effects.⁵ In short, from the hottest quark-gluon plasmas to the cleanest lepton collisions, the techniques presented here can broaden the reach of analyses to newer frontiers.

³Such as a future e^+e^- Higgs factory or a high energy muon collider.

⁴Like exquisite determinations of Higgs boson properties or multi-jet event shapes with minimal background.

⁵Such as higher order electroweak corrections or contact interactions.

Finally, the incorporation of machine learning into experimental workflows opens the door to qualitatively new ways of doing physics. The SYMMETRYGAN concept discussed in this dissertation exemplifies how ML can go beyond numerical inversion tasks and actually contribute to knowledge discovery. In an experimental setting, one could imagine using a SYMMETRYGAN (or future variants of it) on collected data to discover symmetries or invariants that were not incorporated into the simulation *a priori*.

In summary, grounding these unfolding methods in experimental reality shows their potential to revolutionise how measurements are done at HEP experiments. From the ongoing experimental analyses to future studies, the ability to unfold high dimensional distributions with quantified uncertainties will enable the study of deeper questions to obtain sharper answers.

The work presented here embraces modern machine learning solutions in order to fully exploit the forthcoming data deluge. By doing so, one can transform the wall of experimental challenges—high dimensionality, huge data volume, complex detector effects—into a new vista of opportunities for discovery. The next generation of experiments will not only collect more data, but, if armed with methods like NPU, MOMENT UNFOLDING, RAN, and SYMMETRYGAN, will *learn more* from the data they collect. The exploration of fundamental physics requires not only cutting edge instruments, but also cutting edge algorithms.

IX.E Open challenges and future directions.

IX.E.1	Challenges in unfolding methodology.	617
IX.E.2	Handling correlations in unbinned inference.	618
IX.E.3	Background treatment in complex environments.	618
IX.E.4	Systematic uncertainty propagation in ML methods.	619
IX.E.5	Future directions.	620
IX.E.5.i	Joint unbinned unfolding and nuisance parameter profiling	620
IX.E.5.ii	Theoretical guarantees for ML based unfolding. . .	621
IX.E.6	Use on real experimental data.	622

Alongside the advances presented in this dissertation, several open challenges have been posed in order to fully realise the potential of modern unfolding methods for cross section measurements. These challenges highlight areas where further research is needed and naturally suggest promising future directions. This section outlines key outstanding issues in unfolding methodology and then discusses possible research avenues to address them.

IX.E.1 Challenges in unfolding methodology.

As the data to be unfolded increases in complexity, it requires even more sophisticated ML methods to confront fundamental issues of stability and bias. Regularisation choices and model architectures embed implicit assumptions that can influence results, yet there is a lack of strong guarantees on estimator optimality or unbiasedness beyond theoretical idealisations for most ML based approaches. Establishing confidence that an unbinned unfolded result is ‘correct’⁶ is challenging without closed form estimators or asymptotic theorems. Moreover, robust validation of ML unfolding is non-trivial. Traditional closure tests and simple goodness-of-fit checks may not capture subtle failures.⁷ The community currently lacks standardised benchmarks and statistical tests tailored to unbinned, high-dimensional unfolding outputs. In short, while new methods greatly expand the field’s capabilities, providing robust guarantees for their reliability and understanding their failure modes remain open problems.

⁶I.e. consistent with the truth within uncertainties.

⁷E.g. mode collapse or oscillatory artifacts.

IX.E.2 Handling correlations in unbinned inference.

Since unbinned unfolding procedures induce non-trivial correlations between events, when unfolded events⁸ are subsequently used in a likelihood fit or other inference, the usual assumption of independent data points is violated. Ignoring these correlations can bias parameter estimates and misestimate confidence intervals, as demonstrated in Chapter VII. Although Section VII.G.4 proposed initial solutions to restore proper coverage, implementing these in practice poses challenges, and does not compare favourably to the asymptotic results of Section VII.F available for binned inference.

The statistical framework for downstream inference on unbinned unfolded data is still underdeveloped. Analysts must choose between *ad hoc* fixes, like binning the unfolded output with a covariance matrix or performing computationally expensive bootstraps, or risk statistical inconsistency. Developing a rigorous yet practical approach to incorporate inter-event correlations into standard inference procedures is an open problem that must be solved for unbinned methods to be fully used.

IX.E.3 Background treatment in complex environments.

Unfolding analyses in realistic experimental environments must contend with significant background processes and pile-up contributions. In simple terms, one needs to disentangle the “truth” signal distribution from various sources of contamination in the observed data. Traditional approaches handle this by subtracting estimated backgrounds or by unfolding only after tight event selection, but these steps can introduce their own uncertainties and biases. In the context of ML-based unfolding,

⁸Often represented as weighted events.

a more integrated approach is desirable, one that could, for instance, simultaneously learn to subtract background and unfold the signal. However, performing background estimation jointly with unfolding is non-trivial. In analyses where backgrounds are large as a fraction of the signal⁹, the unfolding algorithm might be able to allocate negative weights to simultaneously subtract background and unfold the signal, an ability beyond the scope of most current methods. The interplay between detector smearing and background subtraction is delicate, and naive approaches can amplify fluctuations or distort the unfolded spectrum. Developing methods that robustly handle background-contaminated data remains an open challenge.

IX.E.4 Systematic uncertainty propagation in ML methods.

Precision measurements require careful accounting of systematic uncertainties¹⁰ in the final results. In traditional unfolding, systematics are often propagated by repeating the unfolding with varied inputs¹¹ and then updating the covariance matrix of the unfolded bins. ML-based unfolding complicates this procedure. Retraining a complex model for every systematic variation can be computationally prohibitive and may suffer from training instabilities or statistical fluctuations. While methods like NPU ameliorate this by amortizing inference and enabling fast re-evaluation of different datasets once a model is trained, a truly seamless incorporation of systematics remains challenging. One open issue is how to profile nuisance parameters¹². Ideally, one would want to estimate nuisance parameters from unfolded data, but also one

⁹E.g. underlying event in heavy ion collisions or large QCD backgrounds in certain searches.

¹⁰Detector calibration, modeling inaccuracies, theory uncertainties, *et cetera*.

¹¹E.g., shifting the detector response or kinematic calibrations.

¹²Which encode systematic effects

would want to unfold data with the nuisance parameters profiled. [151] offers a promising approach to profile and unfold simultaneously in a fully unbinned manner, but further work is needed in this direction.

IX.E.5 Future directions.

To address the challenges outlined above and to fully leverage the methods developed in this dissertation, several concrete directions for future research are described below. These are largely concrete near-term goals and build directly on the foundation laid by our results, aiming to close the gap between methodological innovation and practical application in experimental analyses.

IX.E.5.i Joint unbinned unfolding and nuisance parameter profiling

A promising avenue is to integrate unfolding with the simultaneous estimation of nuisance parameters in a unified inference framework. Instead of the traditional two-step approach, where one wants to unfold before one profiles and one wants to profile before one unfolds, one could perform a *joint fit* that includes both the unfolded spectrum and the nuisance parameters. For example, [151] trains a reweighting network that not only matches detector-level distributions but is explicitly conditioned on nuisance parameters,¹³ updating those parameters in tandem with the unfolding until a combined optimum is reached. By profiling systematics during the unfolding, the resulting particle-level distribution would come with systematics already ‘folded in’ to its uncertainty bands, obviating the need for *post hoc* variation tests. Further work

¹³Such as energy scale shifts or resolution uncertainties.

on such a joint unfolding+profiling algorithm would allow experiments to extract truth-level distributions and constrain nuisance parameters in one go, while operating with unbinned data throughout. This approach can greatly streamline analyses. Instead of rerunning an unfolding for each systematic scenario, a single integrated procedure would provide the best-fit truth distribution with profiled uncertainties. A fully realized implementation of this approach would mark a significant advancement in how experimental collaborations handle systematics.

IX.E.5.ii Theoretical guarantees for ML based unfolding.

As ML methods become central to unfolding, it is important to develop a rigorous theoretical understanding of their behaviour. Future work should aim to establish conditions under which ML-based unfolding approaches are statistically consistent and unbiased, and to derive performance guarantees¹⁴ analogous to classical estimators. Mathematical proofs of convergence provide experimentalists with the confidence to use novel methods.

Similarly, formalizing uncertainty quantification in ML unfolding is a critical goal. An example of this is demonstrating that the network’s output, when combined with a certain resampling or ensembling technique, yields confidence intervals with correct coverage. Developing analytic approximations for error propagation¹⁵ could provide rigorous estimates of uncertainty without relying purely on brute-force methods. The end result of this line of research would be a set of theoretical tools and theorems that lend credence to ML unfolding in the eyes of the broader community. By underpinning

¹⁴E.g. convergence rates or coverage properties.

¹⁵Perhaps via the Fisher information or the delta method applied to network parameters.

algorithms with mathematical guarantees ¹⁶, ML-based unfolding methods can be elevated from powerful proofs of concept to well-founded statistical procedures.

IX.E.6 Use on real experimental data.

Ultimately, the true test and value of these unfolding methods will be realized by applying them to real collider data. A crucial next step is to deploy the developed techniques in actual analyses within experimental collaborations, both present and future. Using real data will stress test the methods and demonstrate their effectiveness in ways that go beyond simulated case studies. By collaborating with experimental teams, one can identify and overcome possible limitations in these methods, for instance, integrating these methods into the experiment’s software framework, ensuring that they meet data privacy and reproducibility standards, and ensuring results are interpretable to a broad audience of physicists. Early applications could involve ‘pilot’ measurements of well-known Standard Model distributions, allowing direct comparison of ML unfolding results with those from established methods as a cross-check. Such studies would build trust and demonstrate the advantages in practice. One possible example is showing that an unbinned ML method can reproduce a known cross section result with smaller uncertainties or reveal subtle shape features that a binned analysis missed. Moreover, applying these methods to real data paves the way for discovery potential. By handling high-dimensional data, one might uncover features hinting at new physics that would have been diluted in a traditional analysis. In summary, transferring these unfolding innovations from simulated benchmarks to

¹⁶Or at least clearly defined reliability criteria.

frontline experimental measurements is a high-impact future direction. Success in this endeavour will not only enrich the physics output of ongoing experiments but also validate the maturity of ML based unfolding as a standard tool in the particle physics arsenal.

IX.F Closing: The broader vision

In closing, it is fitting to step back and consider the broader trajectory of the field and how the advancements in this dissertation contribute to the broader endeavour of particle physics. The story of particle physics is a story of the continually improving ability to observe and measure the universe at more and more extreme scales. History has shown that each leap in experimental precision and capability—from cloud chambers to multi-tiered hermetic detectors, from cathode ray tubes to the LHC—has enabled profound discoveries and deeper tests of theory.

The field is now entering an era where precision measurement capabilities are reaching unprecedented levels. The experiments of the future promise enormous data volumes and finer experimental granularity than ever before. However, realising the discovery potential latent in these data requires commensurate advances in how the data is analysed and interpreted. Here lies the broader significance of the modern unfolding methods for cross section measurements developed in this work.

By improving the faithfulness and resolution of our measurements, we empower the next wave of physics insights. In a very real sense, enhanced measurement techniques drive physics discovery. They sharpen our eyes to see subtle effects and bolster our confidence to claim new phenomena or stringent limits. The pursuit of precision is not just an academic exercise; it is the pursuit of the ability to detect deviations from our current understanding of the laws of physics.

Unfolding, in particular, stands as a critical bridge between experiment and theory. Experimental detectors do not measure the ‘pure’ physics quantities we care about. They measure smeared, biased impressions of those quantities. Unfolding is the

procedure by which we translate these impressions back into a representation of nature’s truth, allowing meaningful comparisons with theoretical predictions. Simple bin-by-bin corrections or rudimentary matrix inversions are fundamentally limited in the amount of information they can encapsulate, often with uncontrolled uncertainties in degenerate regions of phase space. The innovations in this dissertation contribute broader and more stable methods to the field. By leveraging machine learning and statistical techniques, we can now transmit essentially the full differential information of events from the experimental side to the theory side. Instead of comparisons limited to a handful of histogram bins, theorists can be handed continuous distributions or high-dimensional correlations that reflect the true particle-level reality with unprecedented fidelity. This tightens the feedback loop between phenomenology and experiment. Models can be adjusted to fit detailed data, and data can be examined for patterns suggested by cutting-edge models. In essence, a robust unfolding framework ensures that experimentalists and theorists speak the same language, enabling the field to move toward a more synchronized and holistic understanding of results. Any tension between observed data and theoretical prediction can be pinpointed with greater accuracy, and potential signs of new physics can be interpreted without ambiguity stemming from detector effects.

Perhaps most importantly, the work presented here is part of a broader vision to extract maximal physics information from increasingly complex data. Modern high energy physics experiments confront us with vast datasets and intricate final states, from jets with complex substructure, to events with multiple interacting systems,¹⁷ to

¹⁷As in heavy ion collisions.

subtle kinematic corners where new physics might hide. Traditional analysis methods, while reliable, often force simplifications such as selecting a subset of “interesting” events, projecting multivariate problems onto single variables, or aggregating data into coarse bins. These compromises, born of necessity, mean that a wealth of information can remain untapped or obscured. The approaches developed in this dissertation seek to turn this limitation into opportunity. By embracing unbinned and high-dimensional analysis, we treat the data in its most natural and information-rich form. Ideas like adversarial reweighting, normalising flow posteriors, and Boltzmann inspired model architectures allow us to preserve structural information that would otherwise be lost in formats natural to the data. As a result, analyses can be performed without discarding complexity. Correlations between observables, full spectrum shapes, and subtle distribution tails are all retained and utilised. The broader vision is that, as a community, we move toward data analysis paradigms where nothing essential is neglected. Every bit of pattern in the data is either explained by our theories or highlighted as a clue to something new. This maximally informative approach will be indispensable in an age where discoveries may manifest as small deviations spread across multidimensional observables rather than obvious bumps in a single histogram. Extracting everything the data has to offer ensures that comparisons to theory are limited only by nature’s laws and energy scales, not by analysis tools.

Looking ahead, the trajectory of particle physics will increasingly rely on such sophisticated analysis frameworks. The next generation of experiments and observations¹⁸ will produce data of staggering complexity and volume. The work of building

¹⁸Whether at energy frontier colliders, precision intensity frontier experiments, or astroparticle detectors.

‘tools for discovery’ is never truly finished—it evolves in tandem with experimental progress. The unfolding methodologies advanced here exemplify the kind of toolset that will become necessary in this evolving landscape. They point to a future in which machine learning and physics insights coalesce to handle the challenges of big data without sacrificing scientific rigour, in which an analyst armed with such methods can ask more of the data. The broader vision, therefore, is one of a field that is limited less by the tools available to analyse data and more by the questions to asked of it. Turning formidable analysis challenges into solvable problems enhances the field’s capacity to recognise the faint whispers of new physics or to map the detailed contours of known phenomena.

To conclude, this dissertation has aimed to push forward the boundary of methods to extract truth from experimental measurements, but it also underscores a philosophy for the field’s future. Precision measurements and advanced inference techniques are synergistic drivers of discovery. Neither alone is sufficient, but together they allow the study of regimes that were previously inaccessible. Unfolding serves as a microcosm of this principle, in that it takes the raw, distorted output from detectors and refines it into a form where the underlying physics can be discerned and interpreted. Refining the unfolding process itself effectively sharpens one of the fundamental tools that enable scientific insight. As the community builds on and adopts these methods, we can be optimistic that the gap between experimental data and theoretical understanding will continue to narrow through deeper measurements, smarter analyses, and hopefully, extraordinary discoveries.

References

- [1] F. Muheim. *Nuclear and Particle Physics: Measurements and Theory*. Lecture notes, School of Physics and Astronomy, University of Edinburgh. n.d. URL: http://www2.ph.ed.ac.uk/~muheim/lectures/particle_phys/ (visited on 08/14/2025).
- [2] *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* (n.d.). Annual scientific reports and papers; various volumes available online.
- [3] A. Andreassen, P. Komiske, E. Metodiev, B. Nachman, and J. Thaler. “Pythia/Herwig + Delphes Jet Datasets for OmniFold Unfolding”. In: (). DOI: 10.5281/zenodo.3548091. URL: <https://zenodo.org/records/3548091>.
- [4] P. Avery and A. Korytov. “Cross section, Flux, Luminosity, Scattering Rates”. In: ().
- [5] J. Chung and K. Palmer. “A Hybrid Lsmr Algorithm For Large-scale Tikhonov Regularization”. In: (). DOI: 10.1137/140975024. URL: <http://www.siam.org/journals/sisc/37-5/97502.html>.

- [6] *Deep Neural Networks for particle reconstruction in high-granularity calorimeters*. en. URL: <https://ep-news.web.cern.ch/content/deep-neural-networks-particle-reconstruction-high-granularity-calorimeters> (visited on 07/15/2025).
- [7] *Environmental awareness: the challenges of CERN's IT infrastructure* CERN. URL: <https://home.cern/news/news/cern/environmental-awareness-challenges-cerns-it-infrastructure>.
- [8] *Implications of First LHC Data*. en. URL: <https://indico.cern.ch/event/94815/contributions/1282695/> (visited on 07/13/2025).
- [9] M. L. Mangano. "Introduction To Qcd". In: ().
- [10] *Measure Theory 1*. en. URL: <https://www.lulu.com/shop/david-fremlin/measure-theory-1/hardcover/product-14nmenep.html> (visited on 08/11/2025).
- [11] *Omnifold: OmniFold, a Library to Perform Unbinned and High-Dimensional Unfolding for HEP*. Version 0.1.36. URL: <https://github.com/ViniciusMikuni/omnifold> (visited on 07/24/2025).
- [12] G. Aad et al. "Evidence for the collective nature of radial flow in Pb+Pb collisions with the ATLAS detector". In: (Mar. 2025). arXiv: 2503.24125 [nucl-ex].
- [13] G. Aad et al. "Measurement of jet track functions in pp collisions at s=13 TeV with the ATLAS detector". In: *Phys. Lett. B* 868 (2025), p. 139680. DOI: 10.1016/j.physletb.2025.139680. arXiv: 2502.02062 [hep-ex].

- [14] R. Aaij et al. “Search for the lepton-flavour-violating decays $B^0 \rightarrow K^{*0\pm} e^\mp$ ”. In: (June 2025). arXiv: 2506.15347 [hep-ex].
- [15] A. Albert, S. Alves, M. André, M. Ardid, S. Ardid, J. .-. J. Aubert, J. Aublin, B. Baret, S. Basa, Y. Becherini, et al. “The ANTARES detector: Two decades of neutrino searches in the Mediterranean Sea”. In: *Physics Reports* 1121-1124 (June 2025), pp. 1–46. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2025.04.001. URL: <https://www.sciencedirect.com/science/article/pii/S0370157325001450> (visited on 07/13/2025).
- [16] C.-A. Alexe, J. Bendavid, L. Bianchini, and D. Bruschini. “Under-coverage in high-statistics counting experiments with finite MC samples”. In: (June 2025). DOI: 10.48550/arXiv.2401.10542. URL: <http://arxiv.org/abs/2401.10542> (visited on 07/22/2025).
- [17] M. Algren, T. Golling, F. A. Di Bello, and C. Pollard. “Mind the Gap: Navigating Inference with Optimal Transport Maps”. In: (July 2025). arXiv: 2507.08867 [physics.data-an].
- [18] S. Ansarifard. “A Likelihood Ratio Framework for Highly Motivated Subdominant Signals”. In: (May 2025). DOI: 10.48550/arXiv.2505.07698. URL: <http://arxiv.org/abs/2505.07698> (visited on 07/22/2025).
- [19] Anthropic. *Claude 4 Opus*. Large language model. June 2025. URL: <https://www.anthropic.com/claude>.

- [20] J. Y. Araz, D. Athanasakos, M. Ploskon, and F. Ringer. “Graph theory inspired anomaly detection at the LHC”. In: (June 2025). URL: <https://arxiv.org/pdf/2506.19920>.
- [21] M. Arratia, R. Milton, V. Mikuni, B. Nachman, J. Pan, F. Torales Acosta, and T. Wamorkar. “Towards Unfolding All Particles in High Q² DIS Events”. In: *H1 Prelim* 25.031 (Mar. 2025). URL: <https://www-h1.desy.de/h1/www/publications/htmlsplit/H1prelim-25-031.long.html>.
- [22] ATLAS Collaboration. “Electroweak, QCD and flavour physics studies with ATLAS data from Run 2 of the LHC”. In: *Physics Reports* 1116 (Apr. 2025), pp. 57–126. DOI: 10.1016/j.physrep.2024.12.003. URL: <http://arxiv.org/abs/2404.06829><http://dx.doi.org/10.1016/j.physrep.2024.12.003>.
- [23] D. Babusci et al. “Measurement of $\eta \rightarrow \pi^0 \gamma \gamma$ branching fraction with the KLOE detector”. In: (May 2025). arXiv: 2505.09285 [hep-ex].
- [24] C. Baldenegro, A. Soto-Ontoso, and G. Soyez. “Secondary Lund jet plane as a gluon-enriched sample”. In: *Jhep* 07 (2025), p. 088. DOI: 10.1007/jhep07(2025)088. arXiv: 2412.14247 [hep-ph].
- [25] M. R. Baumann, M. Taljaard, P. J. Heagerty, M. O. Harhay, G. Tong, R. Wang, and F. Li. *A tutorial on conducting sample size and power calculations for detecting treatment effect heterogeneity in cluster randomized trials*. arXiv:2501.18383 [stat]. Jan. 2025. DOI: 10.48550/arXiv.2501.18383. URL: <http://arxiv.org/abs/2501.18383> (visited on 07/22/2025).

- [26] L. Benato, C. Giordano, C. Krause, A. Li, R. Schöffbeck, D. Schwarz, M. Shooshtari, and D. Wang. “Unbinned Inclusive Cross-Section Measurements with Machine-Learned Systematic Uncertainties”. In: (May 2025). DOI: 10.48550/arXiv.2505.05544. arXiv: 2505.05544 [hep-ph]. (Visited on 07/19/2025).
- [27] G. Bohm and G. Zech. “Introduction to Statistics and Data Analysis for Physicists”. In: *Introduction to Statistics and Data Analysis for Physicists* (Aug. 2025). DOI: 10.1142/14343.
- [28] A. Buckley, J. Butterworth, J. Egan, C. Gutsche, S. Jeon, M. Habedank, T. Procter, P. Wang, Y. Yeh, and L. Yue. “Constraints On New Theories Using Rivet : CONTUR version 3 release note”. en. In: *arXiv e-prints* (May 2025), arXiv:2505.09272. DOI: 10.48550/arXiv.2505.09272. URL: <https://ui.adsabs.harvard.edu/abs/2025arXiv250509272B/abstract> (visited on 07/13/2025).
- [29] L. Builtjes, S. Caron, P. Moskvitina, C. Nellist, R. R. d. Austri, R. Verheyen, and Z. Zhang. *Attention to the strengths of physical interactions: Transformer and graph-based event classification for particle physics experiments*. arXiv:2211.05143 [hep-ph]. Jan. 2025. DOI: 10.48550/arXiv.2211.05143. URL: <http://arxiv.org/abs/2211.05143> (visited on 07/17/2025).
- [30] A. Butter, S. Diefenbacher, N. Huetsch, V. Mikuni, B. Nachman, S. P. Schweitzer, and T. Plehn. “Generative unfolding with distribution mapping”. In: *SciPost Phys.* 18.6 (June 2025), p. 200. ISSN: 25424653. DOI: 10.21468/scipostphys.18.6.200.

- [31] P. Calvo-Barlés, S. G. Rodrigo, and L. Martín-Moreno. “Machine learning for detection of equivariant finite symmetry groups in dynamical systems”. en. In: *Machine Learning: Science and Technology* 6.2 (June 2025). Publisher: IOP Publishing, p. 025058. ISSN: 2632-2153. DOI: 10.1088/2632-2153/ade04d. URL: <https://dx.doi.org/10.1088/2632-2153/ade04d> (visited on 07/24/2025).
- [32] F. Canelli et al. “A Practical Guide to Unbinned Unfolding”. In: (July 2025). arXiv: 2507.09582 [hep-ph].
- [33] E. Canonero and G. Cowan. “Correlated Systematic Uncertainties and Errors-on-Errors in Measurement Combinations: Methodology and Application to the 7-8 TeV ATLAS-CMS Top Quark Mass Combination”. In: 85.2 (Feb. 2025), p. 156. DOI: 10.1140/epjc/s10052-025-13884-w. URL: <http://arxiv.org/abs/2407.05322> (visited on 07/21/2025).
- [34] S. Caron, N. Dobрева, A. F. Sánchez, J. D. Martín-Guerrero, U. Odyurt, R. R. d. A. Bazan, Z. Wolffs, and Y. Zhao. “TrackFormers: In Search of Transformer-Based Particle Tracking for the High-Luminosity LHC Era”. In: *The European Physical Journal C* 85.4 (Apr. 2025). arXiv:2407.07179 [hep-ex], p. 460. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-025-14156-3. URL: <http://arxiv.org/abs/2407.07179> (visited on 07/17/2025).
- [35] V. Chekhovsky et al. “Measurement of event shapes in minimum-bias events from proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: (May 2025). arXiv: 2505.17850 [hep-ex].

- [36] V. Chekhovsky et al. “Search for medium effects using jet axis decorrelation in inclusive jets from PbPb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV”. In: *Jhep* 06 (2025), p. 120. DOI: 10.1007/jhep06(2025)120. arXiv: 2502.13020 [nucl-ex].
- [37] H. Chen, H. Ruan, and H. X. Zhu. “Energy-Energy Correlator at Hadron Colliders: Celestial Blocks and Singularities”. In: (May 2025). arXiv:2505.16753 [hep-ph]. DOI: 10.48550/arXiv.2505.16753. URL: <http://arxiv.org/abs/2505.16753> (visited on 07/23/2025).
- [38] A. Chiefa, M. N. Costantini, J. Cruz-Martinez, E. R. Nocera, T. R. Rabemananjara, J. Rojo, T. Sharma, R. Stegeman, and M. Ubiali. “Parton distributions confront LHC Run II data: a quantitative appraisal”. en. In: *Journal of High Energy Physics* 2025.7 (July 2025), p. 67. ISSN: 1029-8479. DOI: 10.1007/jhep07(2025)067. URL: [https://doi.org/10.1007/JHEP07\(2025\)067](https://doi.org/10.1007/JHEP07(2025)067) (visited on 07/13/2025).
- [39] C. H. Chong and F. Mies. *Likelihood asymptotics of stationary Gaussian arrays*. Feb. 2025. DOI: 10.48550/arXiv.2502.09229. URL: <http://arxiv.org/abs/2502.09229> (visited on 07/22/2025).
- [40] A. T. L. A. S. Collaboration. “Measurement of Jet Track Functions in $\sqrt{s}=13$ TeV Collisions at $\sqrt{s}=13$ TeV with the ATLAS Detector”. In: *Physics Letters B* 868 (June 30, 2025), p. 139680. ISSN: 03702693. DOI: 10.1016/j.physletb.2025.139680. arXiv: 2502.02062 [hep-ex]. URL: <http://arxiv.org/abs/2502.02062> (visited on 07/24/2025).

- [41] A. Collaboration. “Measurement of jet track functions in $\sqrt{s}=13$ TeV collisions at $\sqrt{s}=13$ TeV with the ATLAS detector”. In: *Physics Letters B* (June 2025). arXiv:2502.02062 [hep-ex], p. 139680. ISSN: 03702693. DOI: 10.1016/j.physletb.2025.139680. URL: <http://arxiv.org/abs/2502.02062> (visited on 07/15/2025).
- [42] W. Commons. *File:Standard Model of Elementary Particles.svg* — *Wikimedia Commons, the free media repository*. [Online; accessed 12-July-2025]. 2025. URL: <https://commons.wikimedia.org/w/index.php?title=File:Standard%5C%5FModel%5C%5Fof%5C%5FElementary%5C%5FParticles.svg&oldid=1013094830%7D>.
- [43] A. Correia, F. I. Giasemis, N. Garroum, V. V. Gligorov, and B. Granado. “Graph Neural Network-Based Pipeline for Track Finding in the Velo at LHCb”. In: *Connecting The Dots 2023*. June 2025. arXiv: 2406.12869 [physics.ins-det].
- [44] K. Desai, O. Long, and B. Nachman. “Unbinned Inference with Correlated Events”. In: *E-print* (Apr. 2025). arXiv: 2504.14072 [physics.data-an].
- [45] Desai, Krish and Nachman, Benjamin and Thaler Jesse. *MomentUnfolding: Deconvoluting the Statistical Moments of Collider Data Using a GAN-like Method*. Version v1.0. GitHub repository. Aug. 14, 2025. URL: <https://github.com/HEP-GAN/MomentUnfolding>.
- [46] *dmlc/xgboost*. original-date: 2014-02-06T17:28:03Z. Aug. 2025. URL: <https://github.com/dmlc/xgboost> (visited on 08/11/2025).

- [47] A. Falcão and A. Takacs. “High-Dimensional Unfolding in Large Backgrounds”. In: (July 2025). arXiv: 2507.06291 [hep-ph].
- [48] L. Favaro, R. Kogler, A. Paasch, S. P. Schweitzer, T. Plehn, and D. Schwarz. “How to Unfold Top Decays”. In: (Feb. 2025). DOI: 10.48550/arXiv.2501.12363. URL: <http://arxiv.org/abs/2501.12363> (visited on 07/21/2025).
- [49] R. Finotello, V. Lahoeche, and D. O. Samary. *Functional Renormalization for Signal Detection: Dimensional Analysis and Dimensional Phase Transition for Nearly Continuous Spectra Effective Field Theory*. arXiv:2507.01064 [physics] version: 1. June 2025. DOI: 10.48550/arXiv.2507.01064. URL: <http://arxiv.org/abs/2507.01064> (visited on 07/14/2025).
- [50] G. K. Goswami and A. Pradhan. “Constraining a $f(R, L_m)$ Gravity Cosmological Model with Observational Data”. In: (July 2025). DOI: 10.48550/arXiv.2505.18226. URL: <http://arxiv.org/abs/2505.18226> (visited on 07/22/2025).
- [51] M. Guardiani, V. Eberle, M. Westerkamp, J. Rüstig, P. Frank, and T. Enßlin. “Latent-space Field Tension for Astrophysical Component Detection An application to X-ray imaging”. In: (June 2025). URL: <https://arxiv.org/abs/2506.20758>.
- [52] HEP ML Community. *A Living Review of Machine Learning for Particle Physics*. Living document, continuously updated. 2025. URL: <https://iml-wg.github.io/HEPML-LivingReview/> (visited on 08/14/2025).

- [53] D. B. Horoshko, A. B. Mikhalychev, F. Jelezko, and P. P. Kuzhir. “Weak convexity of Fisher information matrix and superresolved localization of blinking sources of light”. In: 7.3 (July 2025), p. 033065. DOI: 10.1103/t48h-6hcf. URL: <http://arxiv.org/abs/2501.14674> (visited on 07/21/2025).
- [54] R. G. Huang, A. Cudd, M. Kawaue, T. Kikawa, B. Nachman, V. Mikuni, and C. Wilkinson. “Machine Learning-Assisted Unfolding for Neutrino Cross-section Measurements with the OmniFold Technique”. In: (June 2025). DOI: 10.1103/sp1f-n9k2. URL: <http://arxiv.org/abs/2504.06857><http://dx.doi.org/10.1103/sp1f-n9k2>.
- [55] N. Huetsch et al. “The landscape of unfolding with machine learning”. In: *SciPost Phys.* 18.2 (2025), p. 070. DOI: 10.21468/SciPostPhys.18.2.070. arXiv: 2404.18807 [hep-ph].
- [56] G. Jung et al. “Roadmap on machine learning glassy dynamics”. In: *Nature Rev. Phys.* 7.2 (2025), pp. 91–104. DOI: 10.1038/s42254-024-00791-4. arXiv: 2311.14752 [cond-mat.soft].
- [57] H. Kheddar, Y. Himeur, A. Amira, and R. Soualah. “Image and Point-cloud Classification for Jet Analysis in High-Energy Physics: A survey”. In: *Frontiers of Physics* 20.3 (Feb. 2025). DOI: 10.15302/frontphys.2025.035301. URL: <http://arxiv.org/abs/2403.11934><http://dx.doi.org/10.15302/frontphys.2025.035301>.
- [58] V. Kuleshov and S. Deshpande. *Calibrated and Sharp Uncertainties in Deep Learning via Density Estimation*. arXiv:2112.07184 [cs] version: 3. May 2025.

- DOI: 10.48550/arXiv.2112.07184. URL: <http://arxiv.org/abs/2112.07184> (visited on 07/15/2025).
- [59] K. S. Kuzmin, N. M. Levashko, and M. I. Krivoruchenko. “Electromagnetic nucleon form factors in the extended vector meson dominance model”. In: *Phys. Rev. D* 111.1 (2025), p. 013004. DOI: 10.1103/PhysRevD.111.013004. arXiv: 2412.13150 [hep-ph].
- [60] P. L’Ecuyer. *Stochastic Simulation and Monte Carlo Methods*. Lecture notes for IFT 6561. Montréal, QC: Université de Montréal, 2025. URL: <https://www.iro.umontreal.ca/~lecuyer/ift6561/book.pdf> (visited on 08/14/2025).
- [61] C. Langenbruch. “MoreFit: A More Optimised, Rapid and Efficient Fit”. In: (May 2025). DOI: 10.48550/arXiv.2505.12414. URL: <http://arxiv.org/abs/2505.12414> (visited on 07/21/2025).
- [62] V. Mikuni and B. Nachman. “Method to simultaneously facilitate all jet physics tasks”. In: *Phys.Rev.D* 111.5 (Mar. 2025). ISSN: 24700029. DOI: 10.1103/physrevd.111.054015.
- [63] V. Mikuni and B. Nachman. “Solving key challenges in collider physics with foundation models”. In: *Physical Review D* 111.5 (Mar. 2025), p. L051504. ISSN: 24700029. DOI: 10.1103/physrevd.111.1051504/figures/2/medium. URL: <https://journals.aps.org/prd/abstract/10.1103/PhysRevD.111.L051504>.

- [64] R. Milton, V. Mikuni, T. Lee, M. Arratia, T. Wamorkar, and B. Nachman. “Tools for unbinned unfolding”. In: *Jinst* 20.05 (2025), P05034. DOI: 10.1088/1748-0221/20/05/p05034. arXiv: 2503.09720 [hep-ph].
- [65] S. Mishra-Sharma. *smsharma/awesome-neural-sbi*. original-date: 2023-01-20T19:48:13Z. July 2025. URL: <https://github.com/smsharma/awesome-neural-sbi> (visited on 07/21/2025).
- [66] S. Nabat, A. Ghosh, E. Witkowski, G. Kasieczka, and D. Whiteson. “Learning broken symmetries with approximate invariance”. In: *Physical Review D* 111.7 (Apr. 2025), p. 072002. DOI: 10.1103/PhysRevD.111.072002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.111.072002> (visited on 07/23/2025).
- [67] D. Nair. “Development of a Deep Neural Network Double b-tagger for Boosted Topologies using CMS Open Data”. Master’s thesis. National Institute of Technology, Jamshedpur, May 15, 2025. URL: https://chep.iisc.ac.in/Personnel/pages/jyothsna/thesis-repo/DevNair_MSc_Thesis.pdf.
- [68] B. Palumbo, P. Massa, and F. Benvenuto. *Convergence rates for Tikhonov regularization on compact sets: application to neural networks*. arXiv:2505.19936 [math]. May 2025. DOI: 10.48550/arXiv.2505.19936. URL: <http://arxiv.org/abs/2505.19936> (visited on 07/13/2025).
- [69] Y. Pan, C. Wang, Z. Wu, Q. Wang, M. Zhang, and Z. Xu. *IDInit: A Universal and Stable Initialization Method for Neural Network Training*. Mar. 2025. DOI:

- 10.48550/arXiv.2503.04626. URL: <http://arxiv.org/abs/2503.04626> (visited on 07/21/2025).
- [70] H. Peng and J. Zhang. “Efficient, scalable emulation of stochastic simulators: A mixture density network based surrogate modeling framework”. In: *Reliability Engineering & System Safety* 257 (May 2025), p. 110806. ISSN: 0951-8320. DOI: 10.1016/j.ress.2025.110806. URL: <https://www.sciencedirect.com/science/article/pii/S0951832025000092> (visited on 07/15/2025).
- [71] V. Prassa. “Bayesian mixture density networks for evaluating neutron-induced fission charge yields”. In: *Physical Review C* 112.1 (July 2025). Publisher: American Physical Society, p. 014607. DOI: 10.1103/gd81-zzmy. URL: <https://link.aps.org/doi/10.1103/gd81-zzmy> (visited on 07/15/2025).
- [72] A. C. Rodriguez et al. “A Link Between White Dwarf Pulsars and Polars: Multiwavelength Observations of the 9.36-minute Period Variable Gaia22ayj”. In: *Publ. Astron. Soc. Pac.* 137.2 (2025), p. 024202. DOI: 10.1088/1538-3873/adb0f1. arXiv: 2501.01490 [astro-ph.SR].
- [73] RoX818. *Kernel Density Estimation: Non-Parametric Probability Explained*. en-US. Section: AI Fundamentals & Education. Jan. 2025. URL: <https://aicompetence.org/kernel-density-estimation-non-parametric-probability/> (visited on 07/23/2025).
- [74] R. Schöfbeck. “Refinable modeling for unbinned SMEFT analyses”. In: *Machine Learning: Science and Technology* 6.1 (Jan. 2025), p. 015007. ISSN: 2632-2153.

- DOI: 10.1088/2632-2153/ad9fd1. URL: <http://arxiv.org/abs/2406.19076> (visited on 07/21/2025).
- [75] Scikit-learn Developers. *sklearn.decomposition.TruncatedSVD* — *Scikit-learn API Documentation*. API reference for the Truncated SVD transformer in version 1.7.1. Scikit-learn. 2025. URL: <https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.TruncatedSVD.html> (visited on 08/14/2025).
- [76] B. Shaw, S. Kunapuli, A. Magner, and K. R. Moon. *Lie Group Symmetry Discovery and Enforcement Using Vector Fields*. arXiv:2505.08219 [stat]. May 2025. DOI: 10.48550/arXiv.2505.08219. URL: <http://arxiv.org/abs/2505.08219> (visited on 07/23/2025).
- [77] A. Shmakov, K. Greif, M. J. Fenton, A. Ghosh, P. Baldi, and D. Whiteson. “Full Event Particle-Level Unfolding with Variable-Length Latent Variational Diffusion”. In: *SciPost Physics* 18.4 (Jan. 2025). DOI: 10.21468/SciPostPhys.18.4.117. URL: <http://arxiv.org/abs/2404.14332><http://dx.doi.org/10.21468/SciPostPhys.18.4.117>.
- [78] A. Shmakov, K. T. Greif, M. J. Fenton, A. Ghosh, P. Baldi, and D. Whiteson. “Full event particle-level unfolding with variable-length latent variational diffusion”. In: *SciPost Phys.* 18.4 (2025), p. 117. DOI: 10.21468/SciPostPhys.18.4.117. arXiv: 2404.14332 [hep-ex].
- [79] *simulation-based-inference/simulation-based-inference.github.io*. original-date: 2022-02-21T15:36:55Z. July 2025. URL: <https://github.com/simulation-based-inference/simulation-based-inference.github.io>.

- based-inference/simulation-based-inference.github.io (visited on 07/21/2025).
- [80] J. Spinner, L. Favaro, P. Lippmann, S. Pitz, G. Gerhartz, T. Plehn, and F. A. Hamprecht. “Lorentz Local Canonicalization: How to Make Any Network Lorentz-Equivariant”. In: (May 2025). arXiv:2505.20280 [stat.ML]. DOI: 10.48550/arXiv.2505.20280. URL: <http://arxiv.org/abs/2505.20280> (visited on 07/23/2025).
 - [81] S. Tang, Y. Wang, and C. Jin. *Is Elo Rating Reliable? A Study Under Model Misspecification*. arXiv:2502.10985 [cs]. Feb. 2025. DOI: 10.48550/arXiv.2502.10985. URL: <http://arxiv.org/abs/2502.10985> (visited on 07/22/2025).
 - [82] Y.-S. Ting. “Statistical Machine Learning for Astronomy – A Textbook”. In: (June 2025). arXiv: 2506.12230 [astro-ph.IM].
 - [83] Y.-h. Tseng. *Inverse Gaussian Distribution, Introduction and Applications: Comprehensive Analysis of Power Plant Performance: A Study of Combined Cycle and Nuclear Power Plant*. Jan. 2025. DOI: 10.48550/arXiv.2501.14820. URL: <http://arxiv.org/abs/2501.14820> (visited on 07/22/2025).
 - [84] D. Xie. *Local False Sign Rate and the Role of Prior Covariance Rank in Multivariate Empirical Bayes Multiple Testing*. arXiv:2502.16118 [stat]. Apr. 2025. DOI: 10.48550/arXiv.2502.16118. URL: <http://arxiv.org/abs/2502.16118> (visited on 07/22/2025).
 - [85] E. Zivot. *8.6 The Nonparametric Bootstrap / Introduction to Computational Finance and Financial Econometrics with R*. June 2025. URL: <https://bookdown>.

- org / compfinezbook / introcompfinr / The - Nonparametric - Bootstrap .
html (visited on 07/22/2025).
- [86] G. Aad et al. “Underlying-event studies with strange hadrons in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *The European Physical Journal C* 2024 84:12 84.12 (Dec. 2024), pp. 1–29. ISSN: 1434-6052. DOI: 10.1140/EPJC/S10052-024-13243-1. URL: <https://link.springer.com/article/10.1140/epjc/s10052-024-13243-1>.
- [87] G. Aad et al. “Measurements of inclusive and differential cross-sections of $t\bar{t}\gamma$ production in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Jhep* 10 (2024), p. 191. DOI: 10.1007/jhep10(2024)191. arXiv: 2403.09452 [hep-ex].
- [88] G. Aad et al. “Search for pair production of higgsinos in events with two Higgs bosons and missing transverse momentum in s=13 TeV pp collisions at the ATLAS experiment”. In: *Phys. Rev. D* 109.11 (2024), p. 112011. DOI: 10.1103/PhysRevD.109.112011. arXiv: 2401.14922 [hep-ex].
- [89] G. Aad et al. “Simultaneous Unbinned Differential Cross-Section Measurement of Twenty-Four Z+jets Kinematic Observables with the ATLAS Detector”. In: *Phys. Rev. Lett.* 133.26 (Dec. 2024), p. 261803. DOI: 10.1103/PhysRevLett.133.261803. URL: <https://arxiv.org/abs/2405.20041>.
- [90] R. Aaij et al. “Measurement of the Z boson production cross-section in pp collisions at $\sqrt{s} = 5.02$ TeV”. In: *Jhep* 02 (2024), p. 070. DOI: 10.1007/jhep02(2024)070. arXiv: 2308.12940 [hep-ex].

- [91] F. T. Acosta, K. Desai, V. Mikuni, B. Nachman, and J. Pan. “Multidimensional Deconvolution with Profiling”. In: *NeurIPS. ML4PS*. 177. 2024.
- [92] Anthropic. *Claude 3.5 Sonnet*. Version 3.5. June 2024. URL: <https://www.anthropic.com/claude>.
- [93] ATLAS Collaboration, G. Aad, B. Abbott, K. Abeling, N. Abicht, S. Abidi, A. Aboulhorma, H. Abramowicz, H. Abreu, Y. Abulaiti, et al. “Measurement of jet substructure in boosted t events with the ATLAS detector using 140 fb^{-1} of 13 TeV pp collisions”. In: *Physical Review D* 109.11 (June 2024). Publisher: American Physical Society, p. 112016. DOI: 10.1103/PhysRevD.109.112016. URL: <https://link.aps.org/doi/10.1103/PhysRevD.109.112016> (visited on 07/13/2025).
- [94] M. T. Augustine. *A Survey on Universal Approximation Theorems*. arXiv:2407.12895 [cs] version: 1. July 2024. DOI: 10.48550/arXiv.2407.12895. URL: <http://arxiv.org/abs/2407.12895> (visited on 07/15/2025).
- [95] M. Backes, A. Butter, M. Dunford, and B. Malaescu. “An unfolding method based on conditional invertible neural networks (cINN) using iterative training”. In: *SciPost Phys. Core* 7.1 (2024), p. 007. DOI: 10.21468/scipostphyscore.7.1.007. arXiv: 2212.08674 [hep-ph].
- [96] R. Basheer and D. Mishra. *Current Symmetry Group Equivariant Convolution Frameworks for Representation Learning*. arXiv:2409.07327 [cs]. Sept. 2024.

- DOI: 10.48550/arXiv.2409.07327. URL: <http://arxiv.org/abs/2409.07327> (visited on 07/24/2025).
- [97] H. Battey, P. McCullagh, and D. Xiang. *Non-standard boundary behaviour in binary mixture models*. Aug. 2024. DOI: 10.48550/arXiv.2407.20162. URL: <http://arxiv.org/abs/2407.20162> (visited on 07/22/2025).
- [98] V. Belis, P. Odagiu, and T. K. Aarrestad. “Machine learning for anomaly detection in particle physics”. In: *Reviews in Physics* 12 (Dec. 2024), p. 100091. ISSN: 2405-4283. DOI: 10.1016/j.revip.2024.100091. URL: <https://www.sciencedirect.com/science/article/pii/S2405428324000017>.
- [99] J. Brehmer, V. Bresó, P. d. Haan, T. Plehn, H. Qu, J. Spinner, and J. Thaler. “A Lorentz-Equivariant Transformer for All of the LHC”. In: (Dec. 2024). DOI: 10.48550/arXiv.2411.00446. URL: <http://arxiv.org/abs/2411.00446> (visited on 07/23/2025).
- [100] G. Casella and R. L. Berger. “Statistical Inference, Second Edition”. In: *Statistical Inference, Second Edition* (Jan. 2024), pp. 1–535. DOI: 10.1201/9781003456285/STATISTICAL-INFERENCE-ROGER-BERGER-GEORGE-CASELLA/RIGHTS-AND-PERMISSIONS. URL: <https://www.taylorfrancis.com/books/mono/10.1201/9781003456285/statistical-inference-george-casella-roger-berger>.
- [101] C. Chatterjee, A. Petulante, K. Jani, J. Spencer-Smith, Y. Hu, R. Lau, H. Fu, T. Hoang, S. C. Zhao, and S. Deshmukh. “Pre-trained Audio Transformer

- as a Foundational AI Tool for Gravitational Waves”. In: (Dec. 2024). arXiv: 2412.20789 [gr-qc].
- [102] CMS Collaboration. “Luminosity determination using Z boson production at the CMS experiment”. In: *European Physical Journal C* 84.1 (Jan. 2024). DOI: 10.1140/epjc/s10052-023-12268-2. URL: <http://arxiv.org/abs/2309.01008><http://dx.doi.org/10.1140/epjc/s10052-023-12268-2>.
- [103] A. Collaboration. “A simultaneous unbinned differential cross section measurement of twenty-four $Z\gamma$ jets kinematic observables with the ATLAS detector”. In: *Physical Review Letters* 133.26 (Dec. 2024). arXiv:2405.20041 [hep-ex], p. 261803. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.133.261803. URL: <http://arxiv.org/abs/2405.20041> (visited on 07/15/2025).
- [104] C. M. S. Collaboration. “The CMS statistical analysis and combination tool: COMBINE”. In: *Computing and Software for Big Science* 8.1 (Nov. 2024), p. 19. ISSN: 2510-2036, 2510-2044. DOI: 10.1007/s41781-024-00121-4. URL: <http://arxiv.org/abs/2404.06614> (visited on 07/21/2025).
- [105] T. H. collaboration, V. Andreev, M. Arratia, A. Baghdasaryan, A. Baty, K. Begzsuren, A. Bolz, V. Boudry, G. Brandt, D. Britzger, et al. *Machine Learning-Assisted Measurement of Lepton-Jet Azimuthal Angular Asymmetries in Deep-Inelastic Scattering at HERA*. arXiv:2412.14092 [hep-ex]. Dec. 2024. DOI: 10.48550/arXiv.2412.14092. URL: <http://arxiv.org/abs/2412.14092> (visited on 07/15/2025).

- [106] M. Crispim Romão, J. Guilherme Milhano, and M. van Leeuwen. “Jet sub-structure observables for jet quenching in quark gluon plasma: A machine learning driven analysis”. In: *SciPost Phys* 16 (2024), p. 15. DOI: 10.21468/SciPostPhys.16.1.015.
- [107] F. D’Angelo, M. Andriushchenko, A. Varre, and N. Flammarion. *Why Do We Need Weight Decay in Modern Deep Learning?* arXiv:2310.04415 [cs]. Nov. 2024. DOI: 10.48550/arXiv.2310.04415. URL: <http://arxiv.org/abs/2310.04415> (visited on 07/20/2025).
- [108] B. Delaney, N. Schulte, G. Ciezarek, N. Nolte, M. Williams, and J. Albrecht. “Applications of Lipschitz neural networks to the Run 3 LHCb trigger system”. In: *EPJ Web Conf.* 295 (2024), p. 09005. DOI: 10.1051/epjconf/202429509005. arXiv: 2312.14265 [hep-ex].
- [109] H. Deng, Y. Yang, J. Li, C. Chen, W. Jiang, and S. Pu. *Fast Updating Truncated SVD for Representation Learning with Sparse Matrices.* arXiv:2401.09703 [math]. Jan. 2024. DOI: 10.48550/arXiv.2401.09703. URL: <http://arxiv.org/abs/2401.09703> (visited on 07/13/2025).
- [110] K. Desai, B. Nachman, and J. Thaler. “Moment extraction using an unfolding protocol without binning”. In: *Physical Review D* 110.11 (2024), p. 116013.
- [111] S. Diefenbacher, G.-H. Liu, V. Mikuni, B. Nachman, and W. Nie. “Improving generative model-based unfolding with Schrödinger bridges”. In: *Phys. Rev. D* 109.7 (2024), p. 076011. DOI: 10.1103/PhysRevD.109.076011. arXiv: 2308.12351 [hep-ph].

- [112] H. Du, C. Krause, V. Mikuni, B. Nachman, I. Pang, and D. Shih. *Unifying Simulation and Inference with Normalizing Flows*. arXiv:2404.18992 [hep-ph] version: 1. Apr. 2024. DOI: 10.48550/arXiv.2404.18992. URL: <http://arxiv.org/abs/2404.18992> (visited on 07/15/2025).
- [113] N. D. Gagunashvili. “Data Unfolding with Mean Integrated Square Error Optimization”. In: (2024).
- [114] M. F. Giannelli and R. Zhang. “CaloShowerGAN, a Generative Adversarial Networks model for fast calorimeter shower simulation”. In: *The European Physical Journal Plus* 139.7 (July 2024). arXiv:2309.06515 [physics], p. 597. ISSN: 2190-5444. DOI: 10.1140/epjp/s13360-024-05397-4. URL: <http://arxiv.org/abs/2309.06515> (visited on 07/17/2025).
- [115] G. Grosso. “Anomaly-aware summary statistic from data batches”. In: *Jhep* 12 (2024), p. 093. DOI: 10.1007/jhep12(2024)093. arXiv: 2407.01249 [hep-ex].
- [116] B. Hashemi and C. Krause. “Deep Generative Models for Detector Signature Simulation: A Taxonomic Review”. In: *Reviews in Physics* 12 (Dec. 2024). arXiv:2312.09597 [physics], p. 100092. ISSN: 24054283. DOI: 10.1016/j.revip.2024.100092. URL: <http://arxiv.org/abs/2312.09597> (visited on 07/16/2025).
- [117] N. Huetsch, J. M. Villadamigo, A. Shmakov, S. Diefenbacher, V. Mikuni, T. Heime, M. Fenton, K. Greif, B. Nachman, D. Whiteson, et al. “The Landscape of Unfolding with Machine Learning”. In: *SciPost Physics* 18.2 (May 2024).

- DOI: 10.21468/SciPostPhys.18.2.070. URL: <http://arxiv.org/abs/2404.18807><http://dx.doi.org/10.21468/SciPostPhys.18.2.070>.
- [118] H. Jamil, I. Moustaki, and C. Skinner. *Pairwise likelihood estimation and limited information goodness-of-fit test statistics for binary factor analysis models under complex survey sampling*. arXiv:2311.02543 [stat]. July 2024. DOI: 10.48550/arXiv.2311.02543. URL: <http://arxiv.org/abs/2311.02543> (visited on 07/22/2025).
- [119] Y. Jiang, C. Liu, and H. Zhang. *Finite Sample Valid Inference via Calibrated Bootstrap*. Aug. 2024. DOI: 10.48550/arXiv.2408.16763. URL: <http://arxiv.org/abs/2408.16763> (visited on 07/22/2025).
- [120] T. Kutz, J. R. Pybus, D. W. Upton, C. Cotton, A. Deshpande, A. Deur, W. B. Li, D. Nguyen, M. Nycz, and X. Zheng. “High precision measurements of α_s at the future EIC”. In: *Phys. Rev. D* 110.7 (2024), p. 074004. DOI: 10.1103/PhysRevD.110.074004. arXiv: 2406.05591 [nucl-th].
- [121] A. J. Larkoski. “QCD Masterclass Lectures on Jet Physics and Machine Learning”. In: *Eur. Phys. J. C* 84.10 (Oct. 28, 2024), p. 1117. DOI: 10.1140/epjc/s10052-024-13341-0. arXiv: 2407.04897 [hep-ph]. URL: <http://arxiv.org/abs/2407.04897> (visited on 07/24/2025).
- [122] W. Li, R. White, and D. Prangle. *Optimal combination of composite likelihoods using approximate Bayesian computation with application to state-space models*. Sept. 2024. DOI: 10.48550/arXiv.2404.02313. URL: <http://arxiv.org/abs/2404.02313> (visited on 07/22/2025).

- [123] T. Mengel, P. Steffanic, C. Hughes, A. C. O. Da Silva, and C. Nattrass. “Multiplicity Based Background Subtraction for Jets in Heavy Ion Collisions”. In: (Feb. 2024). arXiv: 2402.10945 [hep-ex].
- [124] E. M. Metodiev, J. Thaler, and R. Wynne. “Anomaly detection in collider physics via factorized observables”. In: *Physical Review D* 110.5 (Sept. 2024), p. 055012. ISSN: 24700029. DOI: 10.1103/physrevd.110.055012/supplemental.pdf. URL: <https://journals.aps.org/prd/abstract/10.1103/PhysRevD.110.055012>.
- [125] S. Miao, Z. Lu, M. Liu, J. Duarte, and P. Li. *Locality-Sensitive Hashing-Based Efficient Point Transformer with Applications in High-Energy Physics*. arXiv:2402.12535 [cs]. June 2024. DOI: 10.48550/arXiv.2402.12535. URL: <http://arxiv.org/abs/2402.12535> (visited on 07/17/2025).
- [126] J. Müller, S. Hermann, F. Sammüller, and M. Schmidt. “Gauge Invariance of Equilibrium Statistical Mechanics”. In: *Physical Review Letters* 133.21 (Nov. 2024). arXiv:2406.19235 [cond-mat.stat-mech], p. 217101. DOI: 10.1103/PhysRevLett.133.217101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.133.217101> (visited on 07/24/2025).
- [127] S. Navas, C. Amsler, T. Gutsche, W. Vogelsang, C. Hanhart, U. G. Meißner, J. J. Hernández-Rey, A. Pich, C. Lourenço, A. Ceccucci, et al. “Review of particle physics”. In: *Physical Review D* 110.3 (Aug. 2024), p. 51. ISSN: 24700029. DOI: 10.1103/physrevd.110.030001. URL: <https://pdg.lbl.gov/2025/reviews/kinematics%5C%5Fand%5C%5Fcross%5C%5Fsections.html>.

- [128] R. Newbury, J. Collins, K. He, J. Pan, I. Posner, D. Howard, and A. Cosgun. *A Review of Differentiable Simulators*. July 8, 2024. DOI: 10.48550/arXiv.2407.05560. arXiv: 2407.05560 [cs]. URL: <http://arxiv.org/abs/2407.05560> (visited on 07/24/2025). Pre-published.
- [129] “Observation of $\gamma\gamma \rightarrow \tau\tau$ in proton-proton collisions and limits on the anomalous electromagnetic moments of the tau lepton”. In: (2024).
- [130] C.-C. Pan, X. Dong, Y.-C. Sun, A.-Y. Cheng, A.-B. Wang, Y.-X. Hu, and H. Cai. “SwdFold: A Reweighting and Unfolding method based on Optimal Transport Theory”. In: (June 2024). URL: <https://arxiv.org/abs/2406.01635>.
- [131] J. Pan, M. LeBlanc, B. Nachman, M. Pettee, J. Pan, and D. Gillberg. “Jet Substructure Measurements in Multijet Production with the ATLAS Experiment”. In: (2024). URL: <https://cds.cern.ch/record/2909063>.
- [132] T. Pani. *Generalized angularities measurements from STAR at $\sqrt{s_{NN}} = 200$ GeV*. arXiv:2403.13921 [nucl-ex]. Mar. 2024. DOI: 10.48550/arXiv.2403.13921. URL: <http://arxiv.org/abs/2403.13921> (visited on 07/15/2025).
- [133] C. Pazos, S. Aeron, P.-H. Beauchemin, V. Croft, Z. Huan, M. Klassen, and T. Wongjirad. *Towards Universal Unfolding of Detector Effects in High-Energy Physics using Denoising Diffusion Probabilistic Models*. Nov. 2024. DOI: 10.48550/arXiv.2406.01507. URL: <http://arxiv.org/abs/2406.01507> (visited on 07/20/2025).

- [134] M. Ressegotti. “Search for heavy, long-lived, charged particles with large ionisation energy loss and time-of-flight with the ATLAS experiment”. In: *PoS Eps-hep2023* (2024), p. 484. DOI: 10.22323/1.449.0484.
- [135] H. Saratchandran, T. X. Wang, and S. Lucey. *Weight Conditioning for Smooth Optimization of Neural Networks*. arXiv:2409.03424 [cs]. Sept. 2024. DOI: 10.48550/arXiv.2409.03424. URL: <http://arxiv.org/abs/2409.03424> (visited on 07/22/2025).
- [136] A. Segura and A. C. Parra. “A Practical Guide to Statistical Techniques in Particle Physics”. In: (Nov. 2024). URL: <http://arxiv.org/abs/2411.00706>.
- [137] A. Segura and A. C. Parra. “A Practical Guide to Statistical Techniques in Particle Physics”. In: (Nov. 2024). arXiv: 2411.00706 [hep-ph].
- [138] B. Shaw, A. Magner, and K. R. Moon. *Symmetry Discovery Beyond Affine Transformations*. arXiv:2406.03619 [cs]. Nov. 2024. DOI: 10.48550/arXiv.2406.03619. URL: <http://arxiv.org/abs/2406.03619> (visited on 07/23/2025).
- [139] A. Simonelli. “Analytic Solutions of the DGLAP Evolution and Theoretical Uncertainties”. In: 84.8 (Aug. 2024), p. 867. DOI: 10.1140/epjc/s10052-024-13169-8. arXiv: 2401.13663 [hep-ph]. (Visited on 07/18/2025).
- [140] G. Singh, J. Rembser, L. Moneta, D. Lange, and V. Vassilev. “Making Likelihood Calculations Fast: Automatic Differentiation Applied to RooFit”. In: *EPJ Web of Conferences* 295 (May 2024). Ed. by R. De Vita, X. Espinal, P. Laycock, and O. Shadura, p. 06014. ISSN: 2100-014X. DOI: 10.1051/

- epjconf/202429506014. URL: <https://www.epj-conferences.org/10.1051/epjconf/202429506014>.
- [141] J. Spinner, V. Bresó, P. de Haan, T. Plehn, J. Thaler, and J. Brehmer. “Lorentz-Equivariant Geometric Algebra Transformers for High-Energy Physics”. In: *Advances in Neural Information Processing Systems* 37 (May 2024). ISSN: 10495258. URL: <https://arxiv.org/pdf/2405.14806>.
- [142] S. V. Stroud, P. Duckett, M. Hart, N. Pond, S. Rettie, G. Facini, and T. Scanlon. *Transformers for Charged Particle Track Reconstruction in High Energy Physics*. arXiv:2411.07149 [hep-ex]. Nov. 2024. DOI: 10.48550/arXiv.2411.07149. URL: <http://arxiv.org/abs/2411.07149> (visited on 07/17/2025).
- [143] S. Van Stroud, P. Duckett, M. Hart, N. Pond, S. Rettie, G. Facini, and T. Scanlon. “Transformers for Charged Particle Track Reconstruction in High Energy Physics”. In: (Nov. 2024). URL: <https://arxiv.org/pdf/2411.07149>.
- [144] A. Wang, A. Gandrakota, J. Ngadiuba, V. Sahu, P. Bhatnagar, E. E. Khoda, and J. Duarte. “Interpreting Transformers for Jet Tagging”. In: (Dec. 2024). URL: <https://arxiv.org/pdf/2412.03673>.
- [145] Z. Xu and Y. Li. “Learning Physical Simulation with Message Passing Transformer”. In: (June 2024). URL: <https://arxiv.org/pdf/2406.06060>.
- [146] J. Yang, N. Dehmamy, R. Walters, and R. Yu. *Latent Space Symmetry Discovery*. Aug. 2024. DOI: 10.48550/arXiv.2310.00105. URL: <http://arxiv.org/abs/2310.00105> (visited on 07/24/2025).

- [147] X. G. Yang, Z. Quan, Y. W. Dong, M. Xu, C. Zhang, J. J. Wang, C. L. Liao, Q. Wu, J. Y. Sun, X. Liu, et al. “Application of a deep learning method for shower axis reconstruction in a 3D imaging calorimeter”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1066 (Sept. 2024), p. 169571. ISSN: 0168-9002. DOI: 10.1016/j.nima.2024.169571. URL: <https://www.sciencedirect.com/science/article/pii/S0168900224004972> (visited on 07/15/2025).
- [148] A. Yue, H. Jia, and J. Gonski. *Autoencoders for At-Source Data Reduction and Anomaly Detection in High Energy Particle Detectors*. arXiv:2411.01118 [hep-ex] version: 1. Nov. 2024. DOI: 10.48550/arXiv.2411.01118. URL: <http://arxiv.org/abs/2411.01118> (visited on 07/16/2025).
- [149] J. Zeng, Y. Wang, A. M. Tartakovsky, and D. Barajas-Solano. *Solving High-dimensional Inverse Problems Using Amortized Likelihood-free Inference with Noisy and Incomplete Data*. arXiv:2412.04565 [cs]. Dec. 2024. DOI: 10.48550/arXiv.2412.04565. URL: <http://arxiv.org/abs/2412.04565> (visited on 07/16/2025).
- [150] Z. Zhao, F. Mokhtar, R. Kansal, H. Li, and J. Duarte. *Large-Scale Pre-training and Finetuning for Efficient Jet Classification in Particle Physics*. arXiv:2408.09343 [hep-ex]. Aug. 2024. DOI: 10.48550/arXiv.2408.09343. URL: <http://arxiv.org/abs/2408.09343> (visited on 07/17/2025).
- [151] H. Zhu, K. Desai, M. Kuusela, V. Mikuni, B. Nachman, and L. Wasserman. “Multidimensional Deconvolution with Profiling”. In: *NeurIPS. ML4PS*. 150.

2024. DOI: <https://doi.org/10.48550/arXiv.2409.10421>. arXiv: 0902.0885 [hep-ph].
- [152] G. Aad et al. “Measurements of the suppression and correlations of dijets in Pb+Pb collisions at sNN=5.02 TeV”. In: *Phys. Rev. C* 107.5 (2023). [Erratum: *Phys.Rev.C* 109, 029901 (2024)], p. 054908. DOI: 10.1103/PhysRevC.107.054908. arXiv: 2205.00682 [nucl-ex].
- [153] G. Aad et al. “Search for nonresonant pair production of Higgs bosons in the bb^-bb^- final state in pp collisions at s=13 TeV with the ATLAS detector”. In: *Phys. Rev. D* 108.5 (2023), p. 052003. DOI: 10.1103/PhysRevD.108.052003. arXiv: 2301.03212 [hep-ex].
- [154] M. Abdallah et al. “Higher-order cumulants and correlation functions of proton multiplicity distributions in sNN=3 GeV Au+Au collisions at the RHIC STAR experiment”. In: *Phys. Rev. C* 107.2 (2023), p. 024908. DOI: 10.1103/PhysRevC.107.024908. arXiv: 2209.11940 [nucl-ex].
- [155] S. Acharya et al. “Multiplicity dependence of charged-particle production in pp, p-Pb, Xe-Xe and Pb-Pb collisions at the LHC”. In: *Phys. Lett. B* 845 (2023). [Erratum: *Phys.Lett.B* 853, 138700 (2024)], p. 138110. DOI: 10.1016/j.physletb.2024.138700. arXiv: 2211.15326 [nucl-ex].
- [156] J. Ackerschott, R. K. Barman, D. Gonçalves, T. Heimel, and T. Plehn. “Returning CP-Observables to The Frames They Belong”. In: *SciPost Physics* 17.1 (July 2023). DOI: 10.21468/SciPostPhys.17.1.001. URL: <http://arxiv.org/abs/2307.12345>.

- org/abs/2308.00027%20http://dx.doi.org/10.21468/SciPostPhys.17.1.001.
- [157] T. Alghamdi et al. “Toward a generative modeling analysis of CLAS exclusive 2π photoproduction”. In: *Phys. Rev. D* 108.9 (2023), p. 094030. DOI: 10.1103/PhysRevD.108.094030. arXiv: 2307.04450 [hep-ph].
- [158] M. Algren, T. Golling, M. Guth, C. Pollard, and J. A. Raine. “Flow Away your Differences: Conditional Normalizing Flows as an Improvement to Reweighting”. In: (Apr. 2023). URL: <https://arxiv.org/pdf/2304.14963>.
- [159] R. G. Ambrosio, J. ter Hoeve, M. Madigan, J. Rojo, and V. Sanz. “Unbinned Multivariate Observables for Global SMEFT Analyses from Machine Learning”. In: *Journal of High Energy Physics* 03.3 (Mar. 2023), p. 033. ISSN: 1029-8479. DOI: 10.1007/jhep03(2023)033. arXiv: 2211.02058 [hep-ph]. (Visited on 07/19/2025).
- [160] M. Bandieramonte, M. Faucci Giannelli, F. A. Corchia, R. Zhang, J. F. Beirer, M. Duehrssen-Debling, C. Young, H. Day-Hall, D. Salamani, J. D. Sorenson, et al. *FastCaloGAN: A FAST SIMULATION OF THE ATLAS CALORIMETER WITH GANs*. 2023. URL: <https://cds.cern.ch/record/2857787> (visited on 07/17/2025).
- [161] D. Bank, N. Koenigstein, and R. Giryes. “Autoencoders”. In: *Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook, Third Edition* (Jan. 2023), pp. 353–374. DOI: 10.1007/978-3-031-24628-

- 9{_}16. URL: <https://link.springer.com/chapter/10.1007/978-3-031-24628-9%5C%5F16>.
- [162] N. Berger. “Simplified likelihoods using linearized systematic uncertainties”. en. In: *Journal of High Energy Physics* 2023.4 (Apr. 2023), p. 84. ISSN: 1029-8479. DOI: 10.1007/jhep04(2023)084. URL: [https://doi.org/10.1007/JHEP04\(2023\)084](https://doi.org/10.1007/JHEP04(2023)084) (visited on 07/14/2025).
- [163] P. Buterus and H. Sambale. “Some notes on moment inequalities for heavy-tailed distributions”. In: (Aug. 2023). URL: <https://arxiv.org/pdf/2308.14410>.
- [164] A. Butter, B. M. Dillon, T. Plehn, and L. Vogel. “Performance versus resilience in modern quark-gluon tagging”. In: *SciPost Phys. Core* 6 (2023), p. 85. DOI: 10.21468/SciPostPhysCore.6.4.085.
- [165] A. Butter, T. Plehn, S. Schumann, S. Badger, S. Caron, K. Cranmer, F. A. Di Bello, E. Dreyer, S. Forte, S. Ganguly, D. Gonçalves, E. Gross, T. Heimel, G. Heinrich, L. Heinrich, A. Held, S. Höche, J. N. Howard, P. Ilten, J. Isaacson, T. Janßen, S. Jones, M. Kado, M. Kagan, G. Kasieczka, F. Kling, S. Kraml, C. Krause, F. Krauss, K. Kröninger, R. K. Barman, M. Luchmann, V. Magerya, D. Maitre, B. Malaescu, F. Maltoni, T. Martini, O. Mattelaer, B. Nachman, S. Pitz, J. Rojo, M. Schwartz, D. Shih, F. Siegert, R. Stegeman, B. Stienen, J. Thaler, R. Verheyen, D. Whiteson, R. Winterhalder, and J. Zupan. “Machine learning and LHC event generation”. In: *SciPost Phys.* 14.4 (Apr. 2023), p. 079. ISSN: 25424653. DOI: 10.21468/SCIPOSTPHYS.14.4.079.

- [166] E. Canonero, A. R. Brazzale, and G. Cowan. “Higher-order asymptotic corrections and their application to the Gamma Variance Model”. en. In: *The European Physical Journal C* 83.12 (Dec. 2023), p. 1100. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-023-12263-7. URL: <https://doi.org/10.1140/epjc/s10052-023-12263-7> (visited on 07/21/2025).
- [167] J. Chan and B. Nachman. “Unbinned profiled unfolding”. In: *Physical Review D* 108.1 (July 2023). Publisher: American Physical Society, p. 016002. DOI: 10.1103/PhysRevD.108.016002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.108.016002> (visited on 07/13/2025).
- [168] N. Chreim, C. Hoelbling, N. Pielmeier, and L. Varnhorst. “Isospin Breaking Effects in the 2-Flavor Schwinger Model”. In: *PoS*. Vol. LATTICE2022. Feb. 2023, p. 293. DOI: 10.22323/1.430.0293. URL: <https://doi.org/10.22323/1.430.0293> (visited on 07/24/2025).
- [169] L. Collaboration. “Multidifferential study of identified charged hadron distributions in Z -tagged jets in proton-proton collisions at $\sqrt{s}=13$ TeV”. In: *Physical Review D* 108.3 (Aug. 2023). arXiv:2208.11691 [hep-ex], p. L031103. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.108.L031103. URL: <http://arxiv.org/abs/2208.11691> (visited on 07/15/2025).
- [170] T. H. collaboration, V. Andreev, M. Arratia, A. Baghdasaryan, A. Baty, K. Begzsuren, A. Bolz, V. Boudry, G. Brandt, D. Britzger, et al. “Unbinned Deep Learning Jet Substructure Measurement in High Q^2 ep collisions at HERA”. In: *Physics Letters B* 844 (Sept. 2023). arXiv:2303.13620 [hep-ex],

- p. 138101. ISSN: 03702693. DOI: 10.1016/j.physletb.2023.138101. URL: <http://arxiv.org/abs/2303.13620> (visited on 07/15/2025).
- [171] A. Corso, D. Karamadian, R. Valentin, M. Cooper, and M. J. Kochenderfer. “A holistic assessment of the reliability of machine learning systems”. (In Submission). 2023. DOI: 10.48550/arxiv.2307.10586.
- [172] M. Dasgupta, B. K. El-Menoufi, and J. Helliwell. “QCD resummation for groomed jet observables at NNLL+NLO”. In: *Journal of High Energy Physics* 01.1 (Jan. 2023), p. 045. ISSN: 1029-8479. DOI: 10.1007/jhep01(2023)045. URL: <http://arxiv.org/abs/2211.03820> (visited on 07/19/2025).
- [173] B. Delattre, Q. Barthélemy, A. Araujo, and A. Allauzen. *Efficient Bound of Lipschitz Constant for Convolutional Layers by Gram Iteration*. June 2023. DOI: 10.48550/arXiv.2305.16173. URL: <http://arxiv.org/abs/2305.16173> (visited on 07/20/2025).
- [174] H. Do, Y. Kim, and G. G. Lee. “Prompt- and Trait Relation-aware Cross-prompt Essay Trait Scoring”. In: *Findings of the Association for Computational Linguistics: ACL 2023*. Association for Computational Linguistics, 2023, pp. 1538–1551. DOI: 10.18653/v1/2023.findings-acl.98. URL: <http://dx.doi.org/10.18653/v1/2023.findings-acl.98>.
- [175] A. Döpp, C. Eberle, S. Howard, F. Irshad, J. Lin, and M. Streeter. “Data-driven science and machine learning methods in laser–plasma physics”. en. In: *High Power Laser Science and Engineering* 11 (Jan. 2023), e55. ISSN: 2095-4719, 2052-3289. DOI: 10.1017/hpl.2023.47. URL: <https://www.cambridge.org/>

- core/journals/high-power-laser-science-and-engineering/article/datadriven-science-and-machine-learning-methods-in-laserplasma-physics/B50C69868941B26062ECF6AFCF2BF3B9# (visited on 07/20/2025).
- [176] T. Dorigo et al. “Toward the end-to-end optimization of particle physics instruments with differentiable programming”. In: *Rev. Phys.* 10 (2023), p. 100085. DOI: 10.1016/j.revip.2023.100085. arXiv: 2203.13818 [physics.ins-det].
- [177] M. Erdmann, K. Hafner, J. Schulte, and M. Straub. “Autoencoder-extended Conditional Invertible Neural Networks for Unfolding Signal Traces”. en. In: *Journal of Physics: Conference Series* 2438.1 (Feb. 2023). Publisher: IOP Publishing, p. 012072. ISSN: 1742-6596. DOI: 10.1088/1742-6596/2438/1/012072. URL: <https://dx.doi.org/10.1088/1742-6596/2438/1/012072> (visited on 07/16/2025).
- [178] A. Gabel, V. Klein, R. Valperga, J. S. W. Lamb, K. Webster, R. Quax, and E. Gavves. *Learning Lie Group Symmetry Transformations with Neural Networks*. arXiv:2307.01583 [cs]. July 2023. DOI: 10.48550/arXiv.2307.01583. URL: <http://arxiv.org/abs/2307.01583> (visited on 07/24/2025).
- [179] J. García Pardinas, M. Calvi, J. Eschle, A. Mauri, S. Meloni, M. Mozzanica, and N. Serra. “GNN for Deep Full Event Interpretation and Hierarchical Reconstruction of Heavy-Hadron Decays in Proton–Proton Collisions”. In: *Comput. Softw. Big Sci.* 7.1 (2023), p. 12. DOI: 10.1007/s41781-023-00107-8. arXiv: 2304.08610 [hep-ex].

- [180] A. Goujon, A. Etemadi, and M. Unser. *On the Number of Regions of Piecewise Linear Neural Networks*. arXiv:2206.08615 [cs]. Dec. 2023. DOI: 10.48550/arXiv.2206.08615. URL: <http://arxiv.org/abs/2206.08615> (visited on 07/15/2025).
- [181] S. S. Griбанov and A. S. Popov. “Kinematic and vertex fitting package for the CMD-3 experiment”. In: *Journal of Instrumentation* 18.05 (May 2023), P05030. ISSN: 1748-0221. DOI: 10.1088/1748-0221/18/05/P05030. URL: <http://arxiv.org/abs/2208.11569> (visited on 07/21/2025).
- [182] Z. Hao, R. Kansal, J. Duarte, and N. Chernyavskaya. “Lorentz group equivariant autoencoders”. In: *The European Physical Journal C* 2023 83:6 83.6 (June 2023), pp. 1–15. ISSN: 1434-6052. DOI: 10.1140/EPJC/S10052-023-11633-5. URL: <https://link.springer.com/article/10.1140/epjc/s10052-023-11633-5>.
- [183] D. Hendrycks and K. Gimpel. *Gaussian Error Linear Units (GELUs)*. arXiv:1606.08415 [cs]. June 2023. DOI: 10.48550/arXiv.1606.08415. URL: <http://arxiv.org/abs/1606.08415> (visited on 07/15/2025).
- [184] D. Holmberg, D. Golubovic, and H. Kirschenmann. “Jet energy calibration with deep learning as a Kubeflow pipeline”. In: *Computing and Software for Big Science* 7.1 (Sept. 2023). DOI: 10.1007/s41781-023-00103-y. URL: <http://arxiv.org/abs/2308.12724><http://dx.doi.org/10.1007/s41781-023-00103-y>.

- [185] T. Iwata and A. Kumagai. *Meta-learning to Calibrate Gaussian Processes with Deep Kernels for Regression Uncertainty Estimation*. arXiv:2312.07952 [stat]. Dec. 2023. DOI: 10.48550/arXiv.2312.07952. URL: <http://arxiv.org/abs/2312.07952> (visited on 07/15/2025).
- [186] Z.-B. Kang, A. J. Larkoski, and J. Yang. “Towards a Nonperturbative Formulation of the Jet Charge”. en. In: *Physical Review Letters* 130.15 (Apr. 2023), p. 151901. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.130.151901. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.151901> (visited on 07/19/2025).
- [187] R. Kansal, A. Li, J. Duarte, N. Chernyavskaya, M. Pierini, B. Orzari, and T. Tomei. “Evaluating generative models in high energy physics”. In: *Physical Review D* 107.7 (Apr. 2023), p. 076017. ISSN: 24700029. DOI: 10.1103/physrevd.107.076017/figures/10/medium. URL: <https://journals.aps.org/prd/abstract/10.1103/PhysRevD.107.076017>.
- [188] Y. Ke, I. Luise, Q. Buat, and G. Piacquadio. “Recent Developments on the Statistical Treatment of Flavour Tagging Uncertainties in ATLAS”. In: *PoS Lhcb2022* (2023), p. 322. DOI: 10.22323/1.422.0322.
- [189] C. Krause and D. Shih. “Fast and accurate simulations of calorimeter showers with normalizing flows”. In: *Phys. Rev. D* 107.11 (2023). _eprint: 2106.05285, p. 113003. DOI: 10.1103/PhysRevD.107.113003.

- [190] Y. Lee and F. C. Park. *On Explicit Curvature Regularization in Deep Generative Models*. arXiv:2309.10237 [cs]. Sept. 2023. DOI: 10.48550/arXiv.2309.10237. URL: <http://arxiv.org/abs/2309.10237> (visited on 07/14/2025).
- [191] R. Liu, A. Gandrakota, J. Ngadiuba, M. Spiropulu, and J.-R. Vlimant. *Fast Particle-based Anomaly Detection Algorithm with Variational Autoencoder*. arXiv:2311.17162 [hep-ex]. Nov. 2023. DOI: 10.48550/arXiv.2311.17162. URL: <http://arxiv.org/abs/2311.17162> (visited on 07/16/2025).
- [192] M. Lourens, I. Sinayskiy, D. K. Park, C. Blank, and F. Petruccione. “Hierarchical quantum circuit representations for neural architecture search”. In: *npj Quantum Inf.* 9.1 (2023), p. 79. DOI: 10.1038/s41534-023-00747-z. arXiv: 2210.15073 [quant-ph].
- [193] A. Maevskiy, F. Ratnikov, A. Zinchenko, V. Riabov, A. Sukhorosov, and D. Evdokimov. “Generative Adversarial Networks for the fast simulation of the Time Projection Chamber responses at the MPD detector”. In: *Journal of Physics: Conference Series* 2438.1 (Feb. 2023). arXiv:2203.16355 [physics], p. 012087. ISSN: 1742-6588, 1742-6596. DOI: 10.1088/1742-6596/2438/1/012087. URL: <http://arxiv.org/abs/2203.16355> (visited on 07/17/2025).
- [194] A. McCormack and P. Hoff. *Information Geometry and Asymptotics for Kronecker Covariances*. arXiv:2308.02260 [math]. Aug. 2023. DOI: 10.48550/arXiv.2308.02260. URL: <http://arxiv.org/abs/2308.02260> (visited on 07/22/2025).

- [195] K. Mehta, A. Mahajan, and P. Kumar. “Effects of Spectral Normalization in Multi-Agent Reinforcement Learning”. In: *2023 International Joint Conference on Neural Networks (IJCNN)*. June 2023, pp. 1–8. DOI: 10.1109/ijcnn54540.2023.10191226. URL: <https://ieeexplore.ieee.org/document/10191226> (visited on 07/20/2025).
- [196] S. Munikoti, D. Agarwal, L. Das, and B. Natarajan. “A General Framework for Quantifying Aleatoric and Epistemic Uncertainty in Graph Neural Networks”. In: *Neurocomputing* 521 (Feb. 2023), pp. 1–10. ISSN: 0925-2312. DOI: 10.1016/j.neucom.2022.11.049. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0925231222014424> (visited on 07/24/2025).
- [197] K. Pedersen, M. Mangedarage, and Z. Sullivan. “Harnessing the global correlations of the QCD power spectrum”. In: (Dec. 2023). DOI: 10.48550/arXiv.2312.08627. URL: <http://arxiv.org/abs/2312.08627> (visited on 07/23/2025).
- [198] A. Rogachev and F. Ratnikov. “GAN with an Auxiliary Regressor for the Fast Simulation of the Electromagnetic Calorimeter Response”. In: *Journal of Physics: Conference Series* 2438.1 (Feb. 2023). arXiv:2207.06329 [physics], p. 012086. ISSN: 1742-6588, 1742-6596. DOI: 10.1088/1742-6596/2438/1/012086. URL: <http://arxiv.org/abs/2207.06329> (visited on 07/17/2025).
- [199] C. Sauer. “Measurement of the Triple-Differential Cross-Section for the Production of Multijet Events using 139 fb⁻¹ of Proton-Proton Collision Data at $\sqrt{s} = 13$ TeV with the ATLAS Detector to Disentangle Quarks and Gluons at the Large Hadron Collider”. In: (2023). DOI: 10.11588/heidok.00033941.

- [200] Y. Song. *Measurement of CollinearDrop jet mass and its correlation with SoftDrop groomed jet substructure observables in $\sqrt{s}=200$ GeV pp collisions by STAR*. arXiv:2307.07718 [nucl-ex]. July 2023. DOI: 10.48550/arXiv.2307.07718. URL: <http://arxiv.org/abs/2307.07718> (visited on 07/15/2025).
- [201] A. Stéphanovitch, U. Tanielian, B. Cadre, N. Klutchnikoff, and G. Biau. *Optimal 1-Wasserstein Distance for WGANs*. arXiv:2201.02824 [stat]. Oct. 2023. DOI: 10.48550/arXiv.2201.02824. URL: <http://arxiv.org/abs/2201.02824> (visited on 07/17/2025).
- [202] B. Tahmasebi and S. Jegelka. *The Exact Sample Complexity Gain from Invariances for Kernel Regression*. arXiv:2303.14269 [cs]. Nov. 2023. DOI: 10.48550/arXiv.2303.14269. URL: <http://arxiv.org/abs/2303.14269> (visited on 07/24/2025).
- [203] “Transformer Neural Networks for Identifying Boosted Higgs Bosons decaying into $b\bar{b}$ and $c\bar{c}$ in ATLAS”. In: (Aug. 2023). URL: <https://cds.cern.ch/record/2866601>.
- [204] L. Vaslin, V. Barra, and J. Donini. “GAN-AE: an anomaly detection algorithm for New Physics search in LHC data”. en. In: *The European Physical Journal C* 83.11 (Nov. 2023), p. 1008. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-023-12169-4. URL: <https://doi.org/10.1140/epjc/s10052-023-12169-4> (visited on 07/17/2025).

- [205] T. Xia. *Penalty Gradient Normalization for Generative Adversarial Networks*. arXiv:2306.13576 [cs]. June 2023. DOI: 10.48550/arXiv.2306.13576. URL: <http://arxiv.org/abs/2306.13576> (visited on 07/20/2025).
- [206] A. Xu, S. Han, X. Ju, and H. Wang. “Generative Machine Learning for Detector Response Modeling with a Conditional Normalizing Flow”. In: (Mar. 2023). URL: <https://arxiv.org/pdf/2303.10148>.
- [207] R. Abdul Khalek et al. “Science Requirements and Detector Concepts for the Electron-Ion Collider: EIC Yellow Report”. In: *Nucl. Phys. A* 1026 (2022), p. 122447. DOI: 10.1016/j.nuclphysa.2022.122447. arXiv: 2103.05419 [physics.ins-det].
- [208] A. Accardi et al. “Opportunities for precision QCD physics in hadronization at Belle II – a snowmass whitepaper”. In: *Snowmass 2021*. Apr. 2022. arXiv: 2204.02280 [hep-ex].
- [209] V. Andreev et al. “Measurement of Lepton-Jet Correlation in Deep-Inelastic Scattering with the H1 Detector Using Machine Learning for Unfolding”. In: *Phys. Rev. Lett.* 128.13 (2022), p. 132002. DOI: 10.1103/PhysRevLett.128.132002. arXiv: 2108.12376 [hep-ex].
- [210] M. Arratia et al. “Publishing unbinned differential cross section results”. In: *Jinst* 17.01 (2022), P01024. DOI: 10.1088/1748-0221/17/01/p01024. arXiv: 2109.13243 [hep-ph].
- [211] M. Arratia, D. Britzger, O. Long, and B. Nachman. “Optimizing observables with machine learning for better unfolding”. en. In: *Journal of Instrumentation*

- 17.07 (July 2022). Publisher: IOP Publishing, P07009. ISSN: 1748-0221. DOI: 10.1088/1748-0221/17/07/p07009. URL: <https://dx.doi.org/10.1088/1748-0221/17/07/P07009> (visited on 07/13/2025).
- [212] M. Bellagente. “Go with the Flow: Normalising Flows applications for High Energy Physics”. In: (2022). Publisher: Heidelberg University Library. DOI: 10.11588/heidok.00031386. URL: <https://archiv.ub.uni-heidelberg.de/volltextserver/id/eprint/31386> (visited on 07/15/2025).
- [213] C. Bierlich, S. Chakraborty, N. Desai, L. Gellersen, I. Helenius, P. Ilten, L. Lönnblad, S. Mrenna, S. Prestel, C. T. Preuss, et al. “A comprehensive guide to the physics and usage of PYTHIA 8.3”. In: *SciPost Physics Codebases* (Nov. 2022), p. 8. DOI: 10.21468/SciPostPhysCodeb.8. URL: <https://scipost.org/10.21468/SciPostPhysCodeb.8> (visited on 07/13/2025).
- [214] C. Bonneville and C. Earls. “Bayesian Deep Learning for Partial Differential Equation Parameter Discovery with Sparse and Noisy Data”. In: *Journal of Computational Physics: X* 16 (Sept. 2022), p. 100115. DOI: 10.1016/j.jcp.2022.100115. URL: <https://ui.adsabs.harvard.edu/abs/2022JCPX...1600115B/abstract> (visited on 07/24/2025).
- [215] E. Bortolato and L. Ventura. *On approximate robust confidence distributions*. Dec. 2022. DOI: 10.48550/arXiv.2212.09552. URL: <http://arxiv.org/abs/2212.09552> (visited on 07/22/2025).
- [216] D. Britzger. “The Linear Template Fit”. en. In: *The European Physical Journal C* 82.8 (Aug. 2022), p. 731. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-

- 022-10581-w. URL: <https://doi.org/10.1140/epjc/s10052-022-10581-w> (visited on 07/14/2025).
- [217] M. Caselle, E. Cellini, A. Nada, and M. Panero. “Stochastic normalizing flows for lattice field theory”. In: (Oct. 2022), pp. 8–13. URL: <https://arxiv.org/abs/2210.03139v1>.
- [218] J. Chen, Y. Yu, and Y. Liu. “Physics-guided mixture density networks for uncertainty quantification”. In: *Reliability Engineering & System Safety* 228 (Dec. 2022), p. 108823. ISSN: 0951-8320. DOI: 10.1016/j.ress.2022.108823. URL: <https://www.sciencedirect.com/science/article/pii/S0951832022004422> (visited on 07/15/2025).
- [219] A. Collaboration. “Measurement of the groomed jet radius and momentum splitting fraction in pp and Pb-Pb collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV”. In: *Physical Review Letters* 128.10 (Mar. 2022), p. 102001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.128.102001. URL: <http://arxiv.org/abs/2107.12984> (visited on 07/19/2025).
- [220] A. Collaboration. “Measurements of the groomed and ungroomed jet angularities in pp collisions at $\sqrt{s} = 5.02$ TeV”. In: *Journal of High Energy Physics* 05.5 (May 2022), p. 061. ISSN: 1029-8479. DOI: 10.1007/jhep05(2022)061. URL: <http://arxiv.org/abs/2107.11303> (visited on 07/19/2025).
- [221] H. Collaboration, V. Andreev, M. Arratia, A. Baghdasaryan, A. Baty, K. Begzsuren, A. Belousov, A. Bolz, V. Boudry, G. Brandt, et al. *Measurement of lepton-jet correlation in deep-inelastic scattering with the H1 detector using*

- machine learning for unfolding*. arXiv:2108.12376 [hep-ex]. Apr. 2022. DOI: 10.48550/arXiv.2108.12376. URL: <http://arxiv.org/abs/2108.12376> (visited on 07/15/2025).
- [222] Z. Cui, T. Gao, K. Talamadupula, and Q. Ji. *Knowledge-augmented Deep Learning and Its Applications: A Survey*. Nov. 2022. DOI: 10.48550/arXiv.2212.00017. URL: <http://arxiv.org/abs/2212.00017> (visited on 07/20/2025).
- [223] M. Dai, H. Hang, and A. Srivastava. *Rethinking Multidimensional Discriminator Output for Generative Adversarial Networks*. July 2022. DOI: 10.48550/arXiv.2109.03378. URL: <http://arxiv.org/abs/2109.03378> (visited on 07/21/2025).
- [224] D. Darulis, R. Tyson, D. G. Ireland, D. I. Glazier, B. McKinnon, and P. Pauli. “Machine Learned Particle Detector Simulations”. In: (July 2022). URL: <https://arxiv.org/pdf/2207.11254>.
- [225] K. Desai, B. Nachman, and J. Thaler. “SymmetryGAN: Symmetry discovery with deep learning”. In: *Phys. Rev. D* 105 (9 May 2022), p. 096031. DOI: 10.1103/PhysRevD.105.096031. URL: <https://link.aps.org/doi/10.1103/PhysRevD.105.096031>.
- [226] Y. Fan, Y. Li, N. Xue, and D. Ding. “Analysis of Regularized Poisson GLM Spike-Train Modeling”. en. In: *Journal of Physics: Conference Series* 2173.1 (Jan. 2022). Publisher: IOP Publishing, p. 012019. ISSN: 1742-6596. DOI: 10.1088/1742-6596/2173/1/012019. URL: <https://dx.doi.org/10.1088/1742-6596/2173/1/012019> (visited on 07/14/2025).

- [227] A. Fowlie. “Comment on "Accumulating Evidence for the Associate Production of a Neutral Scalar with Mass around 151 GeV"”. In: *Physics Letters B* 827 (Apr. 2022), p. 136936. ISSN: 03702693. DOI: 10.1016/j.physletb.2022.136936. URL: <http://arxiv.org/abs/2109.13426> (visited on 07/22/2025).
- [228] S. Gong, Q. Meng, J. Zhang, H. Qu, C. Li, S. Qian, W. Du, Z.-M. Ma, and T.-Y. Liu. “An efficient Lorentz equivariant graph neural network for jet tagging”. In: *JHEP* 07 (2022), p. 030. DOI: 10.1007/JHEP07(2022)030. arXiv: 2201.08187 [hep-ph].
- [229] C. Hagemeyer. *Learning Linear Symmetries in Data Using Moment Matching*. arXiv:2204.01213 [cs]. Apr. 2022. DOI: 10.48550/arXiv.2204.01213. URL: <http://arxiv.org/abs/2204.01213> (visited on 07/23/2025).
- [230] J. N. Howard. “Advancing Particle Physics with Sophisticated Computational Frameworks”. en. PhD thesis. UC Irvine, 2022. URL: <https://escholarship.org/uc/item/4cs3f5fp> (visited on 07/20/2025).
- [231] J. Hron, R. Novak, J. Pennington, and J. Sohl-Dickstein. *Wide Bayesian Neural Networks Have a Simple Weight Posterior: Theory and Accelerated Sampling*. June 15, 2022. DOI: 10.48550/arXiv.2206.07673. arXiv: 2206.07673 [stat]. URL: <http://arxiv.org/abs/2206.07673> (visited on 07/24/2025). Pre-published.
- [232] Z. Hu, P. Li, D. Follmann, and J. Qin. “A mixture distribution approach for assessing genetic impact from twin study”. eng. In: *Statistics in Medicine* 41.14 (June 2022), pp. 2513–2522. ISSN: 1097-0258. DOI: 10.1002/sim.9367.

- [233] Z. Huang and D. Ferrari. “Fast Construction of Optimal Composite Likelihoods”. In: (2022).
- [234] T. S. Humble, A. Delgado, R. Pooser, C. Seck, R. Bennink, V. Leyton-Ortega, C.-C. J. Wang, E. Dumitrescu, T. Morris, K. Hamilton, et al. “Snowmass White Paper: Quantum Computing Systems and Software for High-energy Physics Research”. In: (Mar. 2022). DOI: 10.48550/arXiv.2203.07091. arXiv: 2203.07091 [quant-ph]. URL: <http://arxiv.org/abs/2203.07091> (visited on 07/20/2025).
- [235] L. Jiang, Y. Chen, and L. Ding. *Algorithmic Regularization in Model-free Overparametrized Asymmetric Matrix Factorization*. Sept. 2022. DOI: 10.48550/arXiv.2203.02839. URL: <http://arxiv.org/abs/2203.02839> (visited on 07/21/2025).
- [236] G. R. Khattak, S. Vallecorsa, F. Carminati, and G. M. Khan. “Fast simulation of a high granularity calorimeter by generative adversarial networks”. en. In: *The European Physical Journal C* 82.4 (Apr. 2022), p. 386. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-022-10258-4. URL: <https://doi.org/10.1140/epjc/s10052-022-10258-4> (visited on 07/17/2025).
- [237] P.-Y. Lagrave and E. Tron. “Equivariant Neural Networks and Differential Invariants Theory for Solving Partial Differential Equations”. en. In: *Physical Sciences Forum* 5.1 (2022), p. 13. ISSN: 2673-9984. DOI: 10.3390/psf2022005013. URL: <https://www.mdpi.com/2673-9984/5/1/13> (visited on 07/20/2025).

- [238] C. Langenbruch. “Parameter uncertainties in weighted unbinned maximum likelihood fits”. en. In: *The European Physical Journal C* 82.5 (May 2022), p. 393. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-022-10254-8. URL: <https://doi.org/10.1140/epjc/s10052-022-10254-8> (visited on 07/23/2025).
- [239] S. G. Lee, W. Ping, B. Ginsburg, B. Catanzaro, and S. Yoon. “BigVGAN: A Universal Neural Vocoder with Large-Scale Training”. In: *11th International Conference on Learning Representations, ICLR 2023* (June 2022). URL: <https://arxiv.org/pdf/2206.04658>.
- [240] M. Letizia, G. Losapio, M. Rando, G. Grosso, A. Wulzer, M. Pierini, M. Zanetti, and L. Rosasco. “Learning new physics efficiently with nonparametric methods”. en. In: *The European Physical Journal C* 82.10 (Oct. 2022), p. 879. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-022-10830-y. URL: <http://arxiv.org/abs/2204.02317> (visited on 07/21/2025).
- [241] X. Li, K. Mimasu, K. Yamashita, C. Yang, C. Zhang, and S.-Y. Zhou. “Moments for positivity: using Drell-Yan data to test positivity bounds and reverse-engineer new physics”. In: *Journal of High Energy Physics* 2022.10 (Oct. 2022). DOI: 10.1007/jhep10(2022)107. URL: <http://arxiv.org/abs/2204.13121> [http://dx.doi.org/10.1007/JHEP10\(2022\)107](http://dx.doi.org/10.1007/JHEP10(2022)107).
- [242] Z. Li, V. Heyraud, K. Donatella, Z. Denis, and C. Ciuti. “Machine learning via relativity-inspired quantum dynamics”. In: *Phys. Rev. A* 106.3 (2022), p. 032413. DOI: 10.1103/PhysRevA.106.032413. arXiv: 2205.07925 [quant-ph].

- [243] J. G. MacKinnon, M. Ø. Nielsen, and M. D. Webb. *Cluster-Robust Inference: A Guide to Empirical Practice*. May 2022. DOI: 10.48550/arXiv.2205.03285. URL: <http://arxiv.org/abs/2205.03285> (visited on 07/22/2025).
- [244] B. Maier, S. M. Narayanan, G. de Castro, M. Goncharov, C. Paus, and M. Schott. “Pile-up mitigation using attention”. In: *Mach. Learn. Sci. Tech.* 3.2 (June 2022), p. 025012. DOI: 10.1088/2632-2153/ac7198. URL: <https://arxiv.org/abs/2107.02779>.
- [245] T. Manole, P. Bryant, J. Alison, M. Kuusela, and L. Wasserman. “Background Modeling for Double Higgs Boson Production: Density Ratios and Optimal Transport”. In: (Aug. 2022). arXiv: 2208.02807 [stat.AP].
- [246] L. Meunier, B. Delattre, A. Araujo, and A. Allauzen. *A Dynamical System Perspective for Lipschitz Neural Networks*. Feb. 2022. DOI: 10.48550/arXiv.2110.12690. URL: <http://arxiv.org/abs/2110.12690> (visited on 07/20/2025).
- [247] J. G. Milhano and K. Zapp. “Improved background subtraction and a fresh look at jet sub-structure in JEWEL”. In: *Eur. Phys. J. C* 82.11 (2022), p. 1010. DOI: 10.1140/epjc/s10052-022-10954-1. arXiv: 2207.14814 [hep-ph].
- [248] R. Moodie. “Optimising hadronic collider simulations using amplitude neural networks”. In: *Journal of Physics: Conference Series* 2438.1 (Aug. 2022). DOI: 10.1088/1742-6596/2438/1/012149. URL: <http://arxiv.org/abs/2202.04506>
<http://dx.doi.org/10.1088/1742-6596/2438/1/012149>.
- [249] Particle Data Group, R. L. Workman, V. D. Burkert, V. Crede, E. Klempt, U. Thoma, L. Tiator, K. Agashe, G. Aielli, B. C. Allanach, et al. “Review of

- Particle Physics”. In: *Progress of Theoretical and Experimental Physics* 2022.8 (Aug. 2022), p. 083c01. ISSN: 2050-3911. DOI: 10.1093/ptep/ptac097. URL: <https://doi.org/10.1093/ptep/ptac097> (visited on 07/13/2025).
- [250] T. Plehn, A. Butter, B. Dillon, T. Heimel, C. Krause, and R. Winterhalder. “Modern Machine Learning for LHC Physicists”. In: (Nov. 2022). URL: <https://arxiv.org/abs/2211.01421>.
- [251] H. Qu, C. Li, and S. Qian. “Particle Transformer for Jet Tagging”. In: *Proceedings of Machine Learning Research* 162 (Feb. 2022), pp. 18281–18292. ISSN: 26403498. URL: <https://arxiv.org/pdf/2202.03772>.
- [252] A. Rubbia. *Phenomenology of Particle Physics*. Cambridge, UK: Cambridge University Press, June 2022. ISBN: 978-1-316-51934-9.
- [253] R. Saleh Moghaddam and M. E. Zomorrodian. “Comparison of quark jets with gluon jets and measurement of the coupling constant”. In: *Indian Journal of Physics* 96.11 (Sept. 2022), pp. 3311–3319. ISSN: 09749845. DOI: 10.1007/s12648-022-02290-7.
- [254] C. F. G. d. Santos and J. P. Papa. “Avoiding Overfitting: A Survey on Regularization Methods for Convolutional Neural Networks”. In: *ACM Computing Surveys* 54.10 s (Jan. 2022). DOI: 10.1145/3510413. URL: <http://arxiv.org/abs/2201.03299><http://dx.doi.org/10.1145/3510413>.
- [255] P. Tempczyk, R. Michaluk, Ł. Garncarek, P. Spurek, J. Tabor, and A. Goliński. *LIDL: Local Intrinsic Dimension Estimation Using Approximate Likelihood*.

- arXiv:2206.14882 [stat]. July 2022. DOI: 10.48550/arXiv.2206.14882. URL: <http://arxiv.org/abs/2206.14882> (visited on 07/24/2025).
- [256] S. Thais, P. Calafiura, G. Chachamis, G. Dezoort, J. Duarte, S. Ganguly, M. Kagan, D. Murnane, M. S. Neubauer, and K. Terao. “Graph Neural Networks in Particle Physics: Implementations, Innovations, and Challenges”. In: (2022).
- [257] Y. Tian, B. E. Shepherd, C. Li, D. Zeng, and J. J. Schildcrout. *Analyzing Clustered Continuous Response Variables with Ordinal Regression Models*. July 2022. DOI: 10.48550/arXiv.2207.08325. URL: <http://arxiv.org/abs/2207.08325> (visited on 07/22/2025).
- [258] M. O. Turkoglu, A. Becker, H. A. Gündüz, M. Rezaei, B. Bischl, R. C. Daudt, S. D’Aronco, J. D. Wegner, and K. Schindler. *FiLM-Ensemble: Probabilistic Deep Learning via Feature-wise Linear Modulation*. 2022. arXiv: 2206.00050 [cs.LG]. URL: <https://arxiv.org/abs/2206.00050>.
- [259] X. Wang and M. Zhang. “How Powerful Are Spectral Graph Neural Networks”. In: *Proceedings of the 39th International Conference on Machine Learning*. International Conference on Machine Learning. PMLR, June 28, 2022, pp. 23341–23362. URL: <https://proceedings.mlr.press/v162/wang22am.html> (visited on 07/24/2025).
- [260] X. Wang, E. L. Turner, J. S. Preisser, and F. Li. “Power considerations for generalized estimating equations analyses of four-level cluster randomized trials”. In: *Biometrical Journal* 64.4 (Apr. 2022), pp. 663–680. ISSN: 0323-3847,

- 1521-4036. DOI: 10.1002/bimj.202100081. URL: <http://arxiv.org/abs/2108.11466> (visited on 07/22/2025).
- [261] J. Willard, X. Jia, S. Xu, M. Steinbach, and V. Kumar. *Integrating Scientific Knowledge with Machine Learning for Engineering and Environmental Systems*. arXiv:2003.04919 [physics]. Mar. 2022. DOI: 10.48550/arXiv.2003.04919. URL: <http://arxiv.org/abs/2003.04919> (visited on 07/15/2025).
- [262] R. Winterhalder, V. Magerya, E. Villa, S. P. Jones, M. Kerner, A. Butter, G. Heinrich, and T. Plehn. “Targeting multi-loop integrals with neural networks”. In: *SciPost Phys.* 12.4 (Apr. 2022), p. 129. ISSN: 25424653. DOI: 10.21468/scipostphys.12.4.129.
- [263] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097.
- [264] Y. Xia and N. Zabaras. “Bayesian multiscale deep generative model for the solution of high-dimensional inverse problems”. In: *Journal of Computational Physics* 455 (Apr. 2022), p. 111008. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2022.111008. URL: <https://www.sciencedirect.com/science/article/pii/S0021999122000705> (visited on 07/14/2025).
- [265] J. Zhao, F. Schäfer, and A. Anandkumar. *ZerO Initialization: Initializing Neural Networks with only Zeros and Ones*. Nov. 2022. DOI: 10.48550/arXiv.2110.12661. URL: <http://arxiv.org/abs/2110.12661> (visited on 07/21/2025).

- [266] G. Aad et al. “Measurements of top-quark pair single- and double-differential cross-sections in the all-hadronic channel in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector”. In: *Jhep* 01 (2021), p. 033. DOI: 10.1007/jhep01(2021)033. arXiv: 2006.09274 [hep-ex].
- [267] G. Anantha Padmanabha and N. Zabarar. “Solving inverse problems using conditional invertible neural networks”. In: *Journal of Computational Physics* 433 (May 2021), p. 110194. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2021.110194.
- [268] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, A. Suresh, and J. Thaler. “Scaffolding Simulations with Deep Learning for High-dimensional Deconvolution”. In: (May 2021). DOI: 10.48550/arXiv.2105.04448. URL: <http://arxiv.org/abs/2105.04448> (visited on 07/21/2025).
- [269] P. Baldi. “Deep Learning in Science”. In: *Deep Learning in Science* (Apr. 2021). DOI: 10.1017/9781108955652. URL: <https://www.cambridge.org/core/books/deep-learning-in-science/5A620F1E5DB54A7FB1D3E3C1A80E0860>.
- [270] P. Baron. “Comparison of Machine Learning Approach to Other Commonly Used Unfolding Methods”. In: *Acta Phys. Polon. B* 52.8 (2021), p. 863. DOI: 10.5506/APhysPolB.52.863. arXiv: 2104.03036 [hep-ex].
- [271] E. Buhmann, S. Diefenbacher, E. Eren, F. Gaede, D. Hundhausen, G. Kasieczka, W. Korcar, K. Krüger, P. McKeown, and L. Rustige. *Hadrons, Better, Faster, Stronger*. arXiv:2112.09709 [physics]. Dec. 2021. DOI: 10.48550/arXiv.2112.09709. URL: <http://arxiv.org/abs/2112.09709> (visited on 07/17/2025).

- [272] C. Burton, S. Stubbs, and P. Onyisi. “Mixture Density Network Estimation of Continuous Variable Maximum Likelihood Using Discrete Training Samples”. In: *The European Physical Journal C* 81.7 (July 2021). arXiv:2103.13416 [physics], p. 662. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-021-09469-y. URL: <http://arxiv.org/abs/2103.13416> (visited on 07/15/2025).
- [273] S. Chen, W. Wang, B. Xia, X. You, Z. Cao, and W. Ding. *CDE-GAN: Cooperative Dual Evolution Based Generative Adversarial Network*. Mar. 2021. DOI: 10.48550/arXiv.2008.09388. arXiv: 2008.09388 [cs]. (Visited on 07/20/2025).
- [274] C. Clason. “Regularization of Inverse Problems”. In: (Feb. 2021). URL: <http://arxiv.org/abs/2001.00617>.
- [275] A. Collaboration. “Jet fragmentation transverse momentum distributions in pp and p-Pb collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV”. In: *Journal of High Energy Physics* 09.9 (Sept. 2021), p. 211. ISSN: 1029-8479. DOI: 10.1007/jhep09(2021)211. URL: <http://arxiv.org/abs/2011.05904> (visited on 07/19/2025).
- [276] D. Collaboration. “Study of the normalized transverse momentum distribution of W bosons produced in $p \bar{p}$ collisions at $\sqrt{s}=1.96$ TeV”. In: *Physical Review D* 103.1 (Jan. 2021), p. 012003. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.103.012003. URL: <http://arxiv.org/abs/2007.13504> (visited on 07/21/2025).

- [277] G. Cowan. “Statistics”. In: *Handbook of Particle Detection and Imaging*. Ed. by C. Grupen and I. Buvat. 2021, pp. 103–124. DOI: 10.1007/978-3-319-93785-4_5.
- [278] M. Czakon, A. Mitov, and R. Poncelet. “Next-to-Next-to-Leading Order Study of Three-Jet Production at the LHC”. In: *Physical Review Letters* 127.15 (Oct. 2021), p. 152001. DOI: 10.1103/PhysRevLett.127.152001. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.152001> (visited on 07/19/2025).
- [279] DeepMind $\times$ UCL. *DeepMind UCL: Deep Learning Lecture Series 2021 [YouTube playlist]*. Online lecture series hosted on YouTube. 2021. URL: <https://www.youtube.com/playlist?list=PLqYmG7hTraZDVH599EItlEWsUOsJbAodm> (visited on 08/14/2025).
- [280] J. Djolonga, J. Yung, M. Tschannen, R. Romijnders, L. Beyer, A. Kolesnikov, J. Puigcerver, M. Minderer, A. D’Amour, D. Moldovan, S. Gelly, N. Houlsby, X. Zhai, and M. Lucic. *On Robustness and Transferability of Convolutional Neural Networks*. Mar. 2021. DOI: 10.48550/arXiv.2007.08558. URL: <http://arxiv.org/abs/2007.08558> (visited on 07/22/2025).
- [281] *Evaluating statistical uncertainties and correlations using the bootstrap method*. Place: Geneva. 2021. URL: <https://cds.cern.ch/record/2759945> (visited on 07/22/2025).
- [282] “Evaluating statistical uncertainties and correlations using the bootstrap method”. In: (2021).

- [283] P. d. Haan, M. Weiler, T. Cohen, and M. Welling. *Gauge Equivariant Mesh CNNs: Anisotropic convolutions on geometric graphs*. arXiv:2003.05425 [cs]. Nov. 2021. DOI: 10.48550/arXiv.2003.05425. URL: <http://arxiv.org/abs/2003.05425> (visited on 07/24/2025).
- [284] S. Hermann and M. Schmidt. “Noether’s theorem in statistical mechanics”. en. In: *Communications Physics* 4.1 (Aug. 2021). Publisher: Nature Publishing Group, p. 176. ISSN: 2399-3650. DOI: 10.1038/s42005-021-00669-2. URL: <https://www.nature.com/articles/s42005-021-00669-2> (visited on 07/24/2025).
- [285] P. Izmailov, S. Vikram, M. D. Hoffman, and A. G. Wilson. *What Are Bayesian Neural Network Posteriors Really Like?* Apr. 29, 2021. DOI: 10.48550/arXiv.2104.14421. arXiv: 2104.14421 [cs]. URL: <http://arxiv.org/abs/2104.14421> (visited on 07/24/2025). Pre-published.
- [286] A. Jones. *Effective sample size*. Nov. 2021. URL: <https://andrewcharlesjones.github.io/journal/21-effective-sample-size.html> (visited on 07/23/2025).
- [287] G. Kasieczka, B. Nachman, D. Shih, O. Amram, A. Andreassen, K. Benkendorfer, B. Bortolato, G. Brooijmans, F. Canelli, J. H. Collins, B. Dai, F. F. D. Freitas, B. M. Dillon, I.-M. Dinu, Z. Dong, J. Donini, J. Duarte, D. A. Faroughy, J. Gonski, P. Harris, A. Kahn, J. F. Kamenik, C. K. Khosa, P. Komiske, L. L. Pottier, P. Martín-Ramiro, A. Matevc, E. Metodiev, V. Mikuni, I. Ochoa, S. E. Park, M. Pierini, D. Rankin, V. Sanz, N. Sarda, U. Seljak, A. Smolkovic, G. Stein, C. M. Suarez, M. Szewc, J. Thaler, S. Tsan, S.-M. Udrescu, L. Vaslin, J.-R. Vlimant, D. Williams, and M. Yunus. “The

- LHC Olympics 2020: A Community Challenge for Anomaly Detection in High Energy Physics”. In: *Reports on Progress in Physics* 84.12 (Dec. 2021), p. 124201. ISSN: 0034-4885, 1361-6633. DOI: 10.1088/1361-6633/ac36b9. URL: <http://arxiv.org/abs/2101.08320> (visited on 07/24/2025).
- [288] O. Knapp, O. Cerri, G. Dissertori, T. Q. Nguyen, M. Pierini, and J. R. Vlimant. “Adversarially Learned Anomaly Detection on CMS Open Data: re-discovering the top quark”. In: *Eur.Phys.J.Plus* 136.2 (Feb. 2021), p. 236. ISSN: 21905444. DOI: 10.1140/epjp/s13360-021-01109-4.
- [289] P. Komiske, W. P. McCormack, and B. Nachman. “Preserving new physics while simultaneously unfolding all observables”. In: *Phys. Rev. D* 104.7 (2021). _eprint: 2105.09923, p. 076027. DOI: 10.1103/PhysRevD.104.076027.
- [290] S. Kuehn. “Impact of the HL-LHC detector upgrades on the physics program of the ATLAS and CMS experiments”. In: (May 2021). DOI: 10.48550/arXiv.2105.11514. URL: <http://arxiv.org/abs/2105.11514> (visited on 07/21/2025).
- [291] A. Loy and J. Korobova. *Bootstrapping Clustered Data in R using lmeresampler*. June 2021. DOI: 10.48550/arXiv.2106.06568. URL: <http://arxiv.org/abs/2106.06568> (visited on 07/22/2025).
- [292] A. Mehra. “Uncertainty Quantification in Deep Neural Networks: Techniques and Applications in Autonomous Decision-Making Systems”. In: *World Journal of Advanced Research and Reviews* 11.3 (2021), pp. 482–490. ISSN: 2581-9615. DOI: 10.30574/wjarr.2021.11.3.0421. URL: <https://wjarr.>

- com/content/uncertainty-quantification-deep-neural-networks-techniques-and-applications-autonomous (visited on 07/24/2025).
- [293] V. Mikuni and F. Canelli. “Point Cloud Transformers applied to Collider Physics”. In: *Machine Learning: Science and Technology* 2.3 (Sept. 2021). arXiv:2102.05073 [physics], p. 035027. ISSN: 2632-2153. DOI: 10.1088/2632-2153/ac07f6. URL: <http://arxiv.org/abs/2102.05073> (visited on 07/17/2025).
- [294] A. A. Mishra, A. Edelen, A. Hanuka, and C. Mayes. “Uncertainty quantification for deep learning in particle accelerator applications”. In: *Physical Review Accelerators and Beams* 24.11 (Nov. 2021). Publisher: American Physical Society, p. 114601. DOI: 10.1103/PhysRevAccelBeams.24.114601. URL: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.24.114601> (visited on 07/15/2025).
- [295] B. Nachman and J. Thaler. “E Pluribus Unum Ex Machina: Learning from Many Collider Events at Once”. In: *Physical Review D* 103.11 (June 2021). DOI: 10.1103/PhysRevD.103.116013. URL: <http://arxiv.org/abs/2101.07263>
<http://dx.doi.org/10.1103/PhysRevD.103.116013>.
- [296] W. Qian, Y. Xu, W. Zuo, and H. Li. “Self Sparse Generative Adversarial Networks”. In: *CAAI Artificial Intelligence Research* 1.1 (Jan. 2021), pp. 68–78. DOI: 10.26599/air.2022.9150005. URL: <http://arxiv.org/abs/2101.10556>
<http://dx.doi.org/10.26599/AIR.2022.9150005>.

- [297] M. C. Romao, N. F. Castro, and R. Pedro. “Finding New Physics without learning about it: Anomaly Detection as a tool for Searches at Colliders”. In: *European Physical Journal C* 81.1 (Nov. 2021). DOI: 10.1140/epjc/s10052-020-08807-w. URL: <http://arxiv.org/abs/2006.05432><http://dx.doi.org/10.1140/epjc/s10052-020-08807-w>.
- [298] M. Rosca, T. Weber, A. Gretton, and S. Mohamed. *A case for new neural network smoothness constraints*. arXiv:2012.07969 [stat]. July 2021. DOI: 10.48550/arXiv.2012.07969. URL: <http://arxiv.org/abs/2012.07969> (visited on 07/20/2025).
- [299] A. Stakia, T. Dorigo, G. Banelli, D. Bortoletto, A. Casa, P. de Castro, C. Delaere, J. Donini, L. Finos, M. Gallinaro, A. Giammanco, A. Held, F. J. Morales, G. Kotkowski, S. P. Liew, F. Maltoni, G. Menardi, I. Papavergou, A. Saggio, B. Scarpa, G. C. Strong, C. Tosciri, J. Varela, P. Vischia, and A. Weiler. “Advances in Multi-Variate Analysis Methods for New Physics Searches at the Large Hadron Collider”. In: *Reviews in Physics* 7 (Dec. 2021), p. 100063. ISSN: 2405-4283. DOI: 10.1016/j.revip.2021.100063. URL: <https://www.sciencedirect.com/science/article/pii/S2405428321000095> (visited on 07/21/2025).
- [300] “Studies of $t\text{-}\bar{t}$ interference effects in $b\text{-}\bar{b}\ell^+\ell^-\nu\text{-}\bar{\nu}$ final states with Powheg and MG5_aMC@NLO setups”. In: (2021).
- [301] D. M. Thompson. *Power and sample size for cluster randomized and stepped wedge trials: Comparing estimates obtained by applying design effects or by direct estimation in GLMM*. arXiv:2009.05171 [stat]. Mar. 2021. DOI: 10.

- 48550/arXiv.2009.05171. URL: <http://arxiv.org/abs/2009.05171> (visited on 07/22/2025).
- [302] O. T. Unke, S. Chmiela, H. E. Sauceda, M. Gastegger, I. Poltavsky, K. T. Schütt, A. Tkatchenko, and K.-R. Müller. “Machine Learning Force Fields”. In: *Chemical Reviews* 121.16 (Aug. 2021). Publisher: American Chemical Society, pp. 10142–10186. ISSN: 0009-2665. DOI: 10.1021/acs.chemrev.0c01111. URL: <https://doi.org/10.1021/acs.chemrev.0c01111> (visited on 07/20/2025).
- [303] F. Vargas, P. Thodoroff, N. D. Lawrence, and A. Lamacraft. “Solving Schrödinger Bridges via Maximum Likelihood”. In: *Entropy* 23.9 (Aug. 31, 2021), p. 1134. ISSN: 1099-4300. DOI: 10.3390/e23091134. arXiv: 2106.02081 [stat]. URL: <http://arxiv.org/abs/2106.02081> (visited on 07/24/2025).
- [304] Z. Wu, J. Zhang, and C. Mou. *Dense Deep Unfolding Network with 3D-CNN Prior for Snapshot Compressive Imaging*. arXiv:2109.06548 [eess]. Sept. 2021. DOI: 10.48550/arXiv.2109.06548. URL: <http://arxiv.org/abs/2109.06548> (visited on 07/15/2025).
- [305] G. Aad, B. Abbott, D. C. Abbott, et al. “Measurement of the Lund Jet Plane Using Charged Particles in 13 TeV Proton-Proton Collisions with the ATLAS Detector”. In: *Physical Review Letters* 124.22 (June 2020), p. 222002. DOI: 10.1103/PhysRevLett.124.222002. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.222002> (visited on 07/19/2025).
- [306] S. Algeri, J. Aalbers, K. D. Morå, and J. Conrad. “Searching for new physics with profile likelihoods: Wilks and beyond”. In: *Nature Reviews Physics* 2.5

- (Apr. 2020), pp. 245–252. ISSN: 2522-5820. DOI: 10.1038/s42254-020-0169-5. URL: <http://arxiv.org/abs/1911.10237> (visited on 07/22/2025).
- [307] M. Alhroob. “Single top quark production cross section using the ATLAS detector at the LHC”. In: *PoS EPS-HEP2019* (2020), p. 658. DOI: 10.22323/1.364.0658.
- [308] A. Andreassen, P. Komiske, E. Metodiev, B. Nachman, and J. Thaler. “OmniFold: Simultaneously Unfolding All Observables”. In: *Proceedings of the ML4Jets2020 Workshop*. New York, NY, USA, Jan. 2020. URL: <https://indico.cern.ch/event/809820/contributions/3632647/>.
- [309] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, and J. Thaler. “OmniFold: A Method to Simultaneously Unfold All Observables”. In: *Physical Review Letters* 124.18 (May 2020). Publisher: American Physical Society, p. 182001. DOI: 10.1103/PhysRevLett.124.182001. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.182001> (visited on 07/13/2025).
- [310] P. Baroň and J. Kvita. “Extending the Fully Bayesian Unfolding with Regularization Using a Combined Sampling Method”. In: *Symmetry 2020, Vol. 12, Page 2100* 12.12 (Dec. 2020), p. 2100. ISSN: 2073-8994. DOI: 10.3390/sym12122100. URL: <https://www.mdpi.com/2073-8994/12/12/2100/htm%20https://www.mdpi.com/2073-8994/12/12/2100>.
- [311] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, A. Rousselot, R. Winterhalder, L. Ardizzone, and U. Köthe. “Invertible Networks or Partons to

- Detector and Back Again”. In: *SciPost Phys.* 9.5 (Nov. 2020), p. 074. ISSN: 25424653. DOI: 10.21468/scipostphys.9.5.074.
- [312] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, and R. Winterhalder. “How to GAN away detector effects”. en. In: *SciPost Physics* 8.4 (Apr. 2020), p. 070. ISSN: 2542-4653. DOI: 10.21468/SciPostPhys.8.4.070. URL: <https://scipost.org/10.21468/SciPostPhys.8.4.070> (visited on 07/17/2025).
- [313] A. Bogatskiy, B. Anderson, J. T. Offermann, M. Roussi, D. W. Miller, and R. Kondor. “Lorentz Group Equivariant Neural Network for Particle Physics”. In: *37th International Conference on Machine Learning, ICML 2020 PartF168147-2* (June 2020), pp. 969–979. URL: <https://arxiv.org/pdf/2006.04780>.
- [314] A. Bogatskiy, B. Anderson, J. T. Offermann, M. Roussi, D. W. Miller, and R. Kondor. “Lorentz Group Equivariant Neural Network for Particle Physics”. In: (June 2020). arXiv:2006.04780 [hep-ph]. DOI: 10.48550/arXiv.2006.04780. URL: <http://arxiv.org/abs/2006.04780> (visited on 07/23/2025).
- [315] J. Brehmer, K. Cranmer, I. Espejo, F. Kling, G. Louppe, and J. Pavez. “Effective LHC measurements with matrix elements and machine learning”. In: *Journal of Physics: Conference Series* 1525 (2020), p. 12118. DOI: 10.1088/1742-6596/1525/1/012118.
- [316] J. Brehmer and K. Cranmer. “Simulation-based inference methods for particle physics”. In: (Nov. 2020). DOI: 10.48550/arXiv.2010.06439. URL: <http://arxiv.org/abs/2010.06439> (visited on 07/21/2025).

- [317] J. Brehmer, G. Louppe, J. Pavez, and K. Cranmer. “Mining gold from implicit models to improve likelihood-free inference”. In: *Proc. Nat. Acad. Sci.* 117.10 (2020), pp. 5242–5249. DOI: 10.1073/pnas.1915980117. arXiv: 1805.12244 [stat.ML].
- [318] A. Butter, S. Diefenbacher, G. Kasieczka, B. Nachman, and T. Plehn. “GAN-plifying Event Samples”. In: *SciPost Physics* 10.6 (Aug. 2020). ISSN: 25424653. DOI: 10.21468/scipostphys.10.6.139. URL: <https://arxiv.org/abs/2008.06545v3>.
- [319] Y. C. J. Chen, C. W. Chiang, G. Cottin, and D. Shih. “Boosted W and Z tagging with jet charge and deep learning”. In: *Physical Review D* 101.5 (Mar. 2020), p. 053001. ISSN: 24700029. DOI: 10.1103/physrevd.101.053001/figures/20/medium. URL: <https://journals.aps.org/prd/abstract/10.1103/PhysRevD.101.053001>.
- [320] Y. Chen, I. Moustaki, and H. Zhang. “A Note on Likelihood Ratio Tests for Models with Latent Variables”. In: *Psychometrika* 85.4 (Nov. 2020), pp. 996–1012. DOI: 10.1007/s11336-020-09735-0. URL: <http://arxiv.org/abs/2008.03971> <http://dx.doi.org/10.1007/s11336-020-09735-0>.
- [321] A. J. K. Chua and M. Vallisneri. “Learning Bayesian posteriors with neural networks for gravitational-wave inference”. In: *Phys. Rev. Lett.* 124.4 (2020), p. 041102. DOI: 10.1103/PhysRevLett.124.041102. arXiv: 1909.05966 [gr-qc].

- [322] C. M. S. Collaboration. “Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques”. In: *Journal of Instrumentation* 15.06 (June 2020), P06005. ISSN: 1748-0221. DOI: 10.1088/1748-0221/15/06/p06005. URL: <http://arxiv.org/abs/2004.08262> (visited on 07/19/2025).
- [323] S. Cotogno, T. Kasemets, and M. Myska. “Confronting same-sign W-boson production with parton correlations”. In: *Journal of High Energy Physics* 10.10 (Oct. 2020), p. 214. ISSN: 1029-8479. DOI: 10.1007/JHEP10(2020)214. URL: <http://arxiv.org/abs/2003.03347> (visited on 07/23/2025).
- [324] K. D. Doan, S. Manchanda, F. Wang, S. Keerthi, A. Bhowmik, and C. K. Reddy. *Image Generation Via Minimizing Fréchet Distance in Discriminator Feature Space*. Mar. 2020. DOI: 10.48550/arXiv.2003.11774. URL: <http://arxiv.org/abs/2003.11774> (visited on 07/20/2025).
- [325] T. Dorigo and P. de Castro. “Dealing with Nuisance Parameters using Machine Learning in High Energy Physics: a Review”. In: (July 2020). URL: <https://arxiv.org/abs/2007.09121v2>.
- [326] H. W. Engl and R. Ramlau. “Regularization of Inverse Problems”. In: *Encyclopedia of Applied and Computational Mathematics* (Jan. 2020), pp. 1233–1241. DOI: 10.1007/978-3-540-70529-1_{_}52. URL: <https://arxiv.org/abs/2001.00617v2>.
- [327] J. Fajtl, V. Argyriou, D. Monekosso, and P. Remagnino. “Latent Bernoulli Autoencoder”. In: (2020). URL: <https://github.com/ok1zjf/lbae>.

- [328] J. L. Fernández-Martínez and Z. Fernández-Muñiz. “The curse of dimensionality in inverse problems”. In: *Journal of Computational and Applied Mathematics* 369 (May 2020), p. 112571. ISSN: 0377-0427. DOI: 10.1016/j.cam.2019.112571. URL: <https://www.sciencedirect.com/science/article/pii/S037704271930576X> (visited on 07/14/2025).
- [329] S. G. From and A. W. Swift. “Some New Bounds And Approximations On Tail Probabilities Of The Poisson And Other Discrete Distributions”. en. In: *Probability in the Engineering and Informational Sciences* 34.1 (Jan. 2020), pp. 53–71. ISSN: 0269-9648, 1469-8951. DOI: 10.1017/s0269964818000347. URL: <https://www.cambridge.org/core/journals/probability-in-the-engineering-and-informational-sciences/article/abs/some-new-bounds-and-approximations-on-tail-probabilities-of-the-poisson-and-other-discrete-distributions/5D150AD3C15A0F8219340C7437513C3D?utm%5C%5Fsource=chatgpt.com> (visited on 07/13/2025).
- [330] A. Gaponenko. “A practical way to regularize unfolding of sharply varying spectra with low data statistics”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 960 (Apr. 2020), p. 163612. ISSN: 0168-9002. DOI: 10.1016/j.nima.2020.163612. URL: <https://www.sciencedirect.com/science/article/pii/S0168900220301790> (visited on 07/13/2025).
- [331] C. J. Geyer. “Stat 8112 Lecture Notes Unbiased Estimating Equations”. In: (2020).

- [332] S. R. Green, C. Simpson, and J. Gair. “Gravitational-wave parameter estimation with autoregressive neural network flows”. In: *Phys.Rev.D* 102.10 (Nov. 2020). ISSN: 24700029. DOI: 10.1103/physrevd.102.104057.
- [333] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, et al. “Array programming with NumPy”. In: *Nature* 2020 585:7825 585.7825 (Sept. 2020), pp. 357–362. ISSN: 1476-4687. DOI: 10.1038/s41586-020-2649-2. URL: <https://www.nature.com/articles/s41586-020-2649-2>.
- [334] Z. B. Kang, X. Liu, S. Mantry, and D. Y. Shao. “Jet Charge: A Flavor Prism for Spin Asymmetries at the Electron-Ion Collider”. In: *Physical Review Letters* 125.24 (Dec. 2020). ISSN: 10797114. DOI: 10.1103/physrevlett.125.242003.
- [335] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila. *Analyzing and Improving the Image Quality of StyleGAN*. arXiv:1912.04958 [cs]. Mar. 2020. DOI: 10.48550/arXiv.1912.04958. URL: <http://arxiv.org/abs/1912.04958> (visited on 07/17/2025).
- [336] A. J. Larkoski, I. Moult, and B. Nachman. “Jet substructure at the Large Hadron Collider: A review of recent advances in theory and machine learning”. In: *Physics Reports* 841 (Jan. 2020), pp. 1–63. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2019.11.001.
- [337] H. T. Li and I. Vitev. “Jet charge modification in dense QCD matter”. In: *Phys. Rev. D* 101 (2020), p. 076020. DOI: 10.1103/PhysRevD.101.076020. arXiv: 1908.06979 [hep-ph].

- [338] H. T. Li and I. Vitev. “Jet charge modification in finite QCD matter”. In: *Physical Review D* 101.7 (Apr. 2020). ISSN: 24700029. DOI: 10.1103/physrevd.101.076020.
- [339] A. Lobanov. “Precision timing calorimetry with the CMS HGCAL”. In: *JINST* 15.07 (July 2020), p. C07003. DOI: 10.1088/1748-0221/15/07/C07003. URL: <https://arxiv.org/abs/2005.13324>.
- [340] E. Mathieu and M. Nickel. *Riemannian Continuous Normalizing Flows*. arXiv:2006.10605 [stat]. Dec. 2020. DOI: 10.48550/arXiv.2006.10605. URL: <http://arxiv.org/abs/2006.10605> (visited on 07/16/2025).
- [341] J. Melbourne. “Strongly Convex Divergences”. In: *Entropy* 22.11 (Nov. 2020). arXiv:2009.10838 [cs], p. 1327. ISSN: 1099-4300. DOI: 10.3390/e22111327. URL: <http://arxiv.org/abs/2009.10838> (visited on 07/14/2025).
- [342] F. Nielsen. “An Elementary Introduction to Information Geometry”. en. In: *Entropy* 22.10 (Oct. 2020). Number: 10, p. 1100. ISSN: 1099-4300. DOI: 10.3390/e22101100. URL: <https://www.mdpi.com/1099-4300/22/10/1100> (visited on 07/24/2025).
- [343] H. Qu and L. Gouskos. “ParticleNet: Jet Tagging via Particle Clouds”. In: *Physical Review D* 101.5 (Mar. 2020). DOI: 10.1103/PhysRevD.101.056019. URL: <http://arxiv.org/abs/1902.08570><http://dx.doi.org/10.1103/PhysRevD.101.056019>.
- [344] K. Reygers and R. Stamen. *Statistical Methods in Particle Physics*. Lecture and tutorial information (Winter Term 2020/2021), Heidelberg University.

2020. URL: <https://uebungen.physik.uni-heidelberg.de/vorlesung/20202/1225> (visited on 08/14/2025).
- [345] E. de la Rosa, D. Robben, D. M. Sima, J. S. Kirschke, and B. Menze. “Differentiable Deconvolution for Improved Stroke Perfusion Analysis”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 12267 Lncs (2020), pp. 593–602. ISSN: 16113349. DOI: 10.1007/978-3-030-59728-3{_}58/figures/3. URL: <https://link.springer.com/chapter/10.1007/978-3-030-59728-3%5C%5F58>.
- [346] H. Serviansky, N. Segol, J. Shlomi, K. Cranmer, E. Gross, H. Maron, and Y. Lipman. “Set2Graph: Learning Graphs From Sets”. In: (Nov. 2020). DOI: 10.48550/arXiv.2002.08772. URL: <http://arxiv.org/abs/2002.08772> (visited on 07/24/2025).
- [347] J. Shlomi, P. Battaglia, and J.-R. Vlimant. “Graph Neural Networks in Particle Physics”. In: *Machine Learning: Science and Technology* 2.2 (Oct. 2020), p. 021001. DOI: 10.1088/2632-2153/abbf9a. URL: <http://arxiv.org/abs/2007.13681%20http://dx.doi.org/10.1088/2632-2153/abbf9a>.
- [348] A. Sirunyan, A. Tumasyan, W. Adam, F. Ambrogio, T. Bergauer, M. Dragicevic, J. Erö, A. E. D. Valle, M. Flechl, R. Frühwirth, et al. “Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques”. In: *Journal of Instrumentation* 15.06 (June 2020), P06005. DOI: 10.1088/1748-0221/15/06/p06005. URL: <https://dx.doi.org/10.1088/1748-0221/15/06/P06005>.

- [349] A. M. Sirunyan et al. “Measurement of quark- and gluon-like jet fractions using jet charge in PbPb and pp collisions at 5.02 TeV”. In: *Jhep* 07 (2020), p. 115. DOI: 10.1007/jhep07(2020)115. arXiv: 2004.00602 [hep-ex].
- [350] A. M. Sirunyan et al. “Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *JINST* 15.10 (Oct. 2020), P10017. DOI: 10.1088/1748-0221/15/10/P10017. URL: <https://arxiv.org/abs/2006.10165>.
- [351] A. M. Sirunyan, A. Tumasyan, W. Adam, et al. “Measurement of the Jet Mass Distribution and Top Quark Mass in Hadronic Decays of Boosted Top Quarks in pp Collisions at $\sqrt{s}=13\text{ TeV}$ ”. In: *Physical Review Letters* 124.20 (May 2020), p. 202001. DOI: 10.1103/PhysRevLett.124.202001. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.202001> (visited on 07/19/2025).
- [352] P. Vischia. “New approaches to the problem of unfolding”. In: (2020).
- [353] F. Wenzel, K. Roth, B. S. Veeling, J. Świątkowski, L. Tran, S. Mandt, J. Snoek, T. Salimans, R. Jenatton, and S. Nowozin. *How Good Is the Bayes Posterior in Deep Neural Networks Really?* July 2, 2020. DOI: 10.48550/arXiv.2002.02405. arXiv: 2002.02405 [stat]. URL: <http://arxiv.org/abs/2002.02405> (visited on 07/24/2025). Pre-published.
- [354] M. Wiatrak, S. V. Albrecht, and A. Nystrom. *Stabilizing Generative Adversarial Networks: A Survey*. Mar. 2020. DOI: 10.48550/arXiv.1910.00927. URL: <http://arxiv.org/abs/1910.00927> (visited on 07/20/2025).

- [355] R. A. Wilson. “A group-theorist’s perspective on symmetry groups in physics”. In: (Dec. 2020). arXiv:2009.14613 [math.GR]. DOI: 10.48550/arXiv.2009.14613. URL: <http://arxiv.org/abs/2009.14613> (visited on 07/24/2025).
- [356] P. Zyla, R. Barnett, J. Beringer, et al. “Review of Particle Physics”. In: *Progress of Theoretical and Experimental Physics* 2020.8 (Aug. 2020), p. 083c01. ISSN: 2050-3911. DOI: 10.1093/ptep/ptaa104. URL: <https://doi.org/10.1093/ptep/ptaa104> (visited on 07/19/2025).
- [357] M. S. Albergo, G. Kanwar, and P. E. Shanahan. “Flow-based generative models for Markov chain Monte Carlo in lattice field theory FLOW-BASED GENERATIVE MODELS for MARKOV CHAIN ... M. S. ALBERGO, G. KANWAR, and P. E. SHANAHAN”. In: *Physical Review D* 100.3 (Aug. 2019). ISSN: 24700029. DOI: 10.1103/physrevd.100.034515.
- [358] J. Albrecht et al. “A Roadmap for HEP Software and Computing R&D for the 2020s”. In: *Comput. Softw. Big Sci.* 3.1 (2019), p. 7. DOI: 10.1007/s41781-018-0018-8. arXiv: 1712.06982 [physics.comp-ph].
- [359] B. Bhattacharjee, S. Mukherjee, and R. Sengupta. “Study of energy deposition patterns in hadron calorimeter for prompt and displaced jets using convolutional neural network”. en. In: *Journal of High Energy Physics* 2019.11 (Nov. 2019), p. 156. ISSN: 1029-8479. DOI: 10.1007/jhep11(2019)156. URL: [https://doi.org/10.1007/JHEP11\(2019\)156](https://doi.org/10.1007/JHEP11(2019)156) (visited on 07/15/2025).
- [360] A. Brock, J. Donahue, and K. Simonyan. *Large Scale GAN Training for High Fidelity Natural Image Synthesis*. arXiv:1809.11096 [cs]. Feb. 2019. DOI:

- 10.48550/arXiv.1809.11096. URL: <http://arxiv.org/abs/1809.11096> (visited on 07/17/2025).
- [361] A. Butter, T. Plehn, and R. Winterhalder. “How to GAN LHC events”. en. In: *SciPost Physics* 7.6 (Dec. 2019), p. 075. ISSN: 2542-4653. DOI: 10.21468/SciPostPhys.7.6.075. URL: <https://scipost.org/10.21468/SciPostPhys.7.6.075> (visited on 07/17/2025).
- [362] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová. “Machine learning and the physical sciences”. In: *Rev.Mod.Phys.* 91.4 (Dec. 2019). ISSN: 15390756. DOI: 10.1103/revmodphys.91.045002.
- [363] S. Carrazza and F. A. Dreyer. “Lund jet images from generative and cycle-consistent adversarial networks”. en. In: *The European Physical Journal C* 79.11 (Nov. 2019), p. 979. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-019-7501-1. URL: <http://link.springer.com/10.1140/epjc/s10052-019-7501-1> (visited on 07/17/2025).
- [364] F. Di Mattia, P. Galeone, M. De Simoni, and E. Ghelfi. “A Survey on GANs for Anomaly Detection”. In: (June 2019). URL: <https://arxiv.org/abs/1906.11632v2>.
- [365] R. Di Sipio, M. F. Giannelli, S. K. Haghighat, and S. Palazzo. “DijetGAN: a Generative-Adversarial Network approach for the simulation of QCD dijet events at the LHC”. en. In: *Journal of High Energy Physics* 2019.8 (Aug.

- 2019), p. 110. ISSN: 1029-8479. DOI: 10.1007/jhep08(2019)110. URL: [https://doi.org/10.1007/JHEP08\(2019\)110](https://doi.org/10.1007/JHEP08(2019)110) (visited on 07/17/2025).
- [366] M. Erdmann, J. Glombitza, and T. Quast. “Precise Simulation of Electromagnetic Calorimeter Showers Using a Wasserstein Generative Adversarial Network”. In: *Computing and Software for Big Science* 3.1 (Jan. 2019). ISSN: 2510-2044. DOI: 10.1007/s41781-018-0019-7. URL: <http://dx.doi.org/10.1007/s41781-018-0019-7>.
- [367] S. Greydanus, M. Dzamba, and J. Yosinski. *Hamiltonian Neural Networks*. arXiv:1906.01563 [cs]. Sept. 2019. DOI: 10.48550/arXiv.1906.01563. URL: <http://arxiv.org/abs/1906.01563> (visited on 07/24/2025).
- [368] R. Haake and C. Loizides. “Machine Learning based jet momentum reconstruction in heavy-ion collisions”. In: *Phys. Rev. C* 99.6 (2019), p. 064904. DOI: 10.1103/PhysRevC.99.064904. arXiv: 1810.06324 [nucl-ex].
- [369] R. Haake and C. Loizides. “Machine-learning-based jet momentum reconstruction in heavy-ion collisions”. In: *Physical Review C* 99.6 (June 2019), p. 064904. ISSN: 24699993. DOI: 10.1103/physrevc.99.064904/figures/10/thumbnail. URL: <https://journals.aps.org/prc/abstract/10.1103/PhysRevC.99.064904>.
- [370] W. Handley and M. Millea. “Maximum-Entropy Priors with Derived Parameters in a Specified Distribution”. In: *Entropy 2019, Vol. 21, Page 272* 21.3 (Mar. 2019), p. 272. ISSN: 1099-4300. DOI: 10.3390/e21030272. URL:

- <https://www.mdpi.com/1099-4300/21/3/272/htm%20https://www.mdpi.com/1099-4300/21/3/272>.
- [371] N. Houlsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. d. Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly. *Parameter-Efficient Transfer Learning for NLP*. June 2019. DOI: 10.48550/arXiv.1902.00751. URL: <http://arxiv.org/abs/1902.00751> (visited on 07/21/2025).
- [372] J. Jia, F. Xie, and L. Xu. “Sparse Poisson regression with penalized weighted score function”. In: *Electronic Journal of Statistics* 13.2 (Jan. 2019). Publisher: Institute of Mathematical Statistics and Bernoulli Society, pp. 2898–2920. ISSN: 1935-7524, 1935-7524. DOI: 10.1214/19-ejs1580. URL: <https://projecteuclid.org/journals/electronic-journal-of-statistics/volume-13/issue-2/Sparse-Poisson-regression-with-penalized-weighted-score-function/10.1214/19-EJS1580.full> (visited on 07/14/2025).
- [373] T. Karras, S. Laine, and T. Aila. *A Style-Based Generator Architecture for Generative Adversarial Networks*. arXiv:1812.04948 [cs]. Mar. 2019. DOI: 10.48550/arXiv.1812.04948. URL: <http://arxiv.org/abs/1812.04948> (visited on 07/17/2025).
- [374] G. Kasieczka, B. Nachman, and D. Shih. *R&D Dataset for LHC Olympics 2020 Anomaly Detection Challenge*. Apr. 2019. URL: <https://zenodo.org/records/4536377> (visited on 07/24/2025).
- [375] D. P. Kingma and M. Welling. “An Introduction to Variational Autoencoders”. In: *Foundations and Trends in Machine Learning* 12.4 (June 2019), pp. 307–

392. ISSN: 19358245. DOI: 10.1561/22000000056. URL: <https://arxiv.org/pdf/1906.02691>.
- [376] I. Kobyzev, S. J. Prince, and M. A. Brubaker. “Normalizing Flows: An Introduction and Review of Current Methods”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43.11 (Aug. 2019), pp. 3964–3979. ISSN: 19393539. DOI: 10.1109/tpami.2020.2992934. URL: <https://arxiv.org/pdf/1908.09257>.
- [377] R. Kogler, B. Nachman, A. Schmidt, L. Asquith, E. Winkels, M. Campanelli, C. Delitzsch, P. Harris, A. Hinzmann, D. Kar, et al. “Jet substructure at the Large Hadron Collider”. In: *Reviews of Modern Physics* 91.4 (Dec. 2019). Publisher: American Physical Society, p. 045003. DOI: 10.1103/RevModPhys.91.045003. URL: <https://link.aps.org/doi/10.1103/RevModPhys.91.045003> (visited on 07/13/2025).
- [378] P. Komiske, R. Mastandrea, E. Metodiev, P. Naik, and J. Thaler. *CMS 2011A Open Data / Jet Primary Dataset / $pT > 375$ GeV / MOD HDF5 Format*. Aug. 2019. URL: <https://zenodo.org/records/3340205> (visited on 07/24/2025).
- [379] P. T. Komiske, E. M. Metodiev, and J. Thaler. “Energy Flow Networks: Deep Sets for Particle Jets”. In: *Journal of High Energy Physics* 2019.1 (Jan. 2019). DOI: 10.1007/jhep01(2019)121. URL: <http://arxiv.org/abs/1810.05165>
[http://dx.doi.org/10.1007/JHEP01\(2019\)121](http://dx.doi.org/10.1007/JHEP01(2019)121).
- [380] J. S. H. Lee, I. Park, I. J. Watson, and S. Yang. “Quark-Gluon Jet Discrimination Using Convolutional Neural Networks”. In: *Journal of the Ko-*

- rean Physical Society* 74.3 (Feb. 2019), pp. 219–223. ISSN: 19768524. DOI: 10.3938/jkps.74.219/metrics. URL: <https://link.springer.com/article/10.3938/jkps.74.219>.
- [381] W. Mulzer. “Five Proofs of Chernoff’s Bound with Applications *”. In: (2019).
- [382] F. Nielsen. “On the Jensen-Shannon Symmetrization of Distances Relying on Abstract Means”. en. In: *Entropy* 21.5 (May 2019). Number: 5 Publisher: Multidisciplinary Digital Publishing Institute, p. 485. ISSN: 1099-4300. DOI: 10.3390/e21050485. URL: <https://www.mdpi.com/1099-4300/21/5/485> (visited on 07/17/2025).
- [383] M. Raissi, P. Perdikaris, and G. E. Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational Physics* 378 (Feb. 2019), pp. 686–707. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2018.10.045. URL: <https://www.sciencedirect.com/science/article/pii/S0021999118307125> (visited on 07/15/2025).
- [384] M. Rieger and I. u. N. RWTH Aachen University (Germany). Fakultät für Mathematik. “Search for Higgs boson production in association with top quarks and decaying into bottom quarks using deep learning techniques with the CMS experiment”. In: (June 2019). DOI: 10.18154/rwth-2019-06415.
- [385] K. Scaman and A. Virmaux. *Lipschitz regularity of deep neural networks: analysis and efficient estimation*. arXiv:1805.10965 [stat]. Oct. 2019. DOI:

- 10.48550/arXiv.1805.10965. URL: <http://arxiv.org/abs/1805.10965> (visited on 07/20/2025).
- [386] D. Scheirich and V. Pleskot. *Statistical Methods in High Energy Physics*. Lecture notes, Institute of Particle and Nuclear Physics, Charles University. Taught as NJSF143 in 2019–2020; script updated January 2021. 2019. URL: <https://ipnp.cz/wp-content/uploads/2019/10/notes-2.pdf>.
- [387] S. Shi, X. Wang, and H. Li. *PointRCNN: 3D Object Proposal Generation and Detection from Point Cloud*. arXiv:1812.04244 [cs]. May 2019. DOI: 10.48550/arXiv.1812.04244. URL: <http://arxiv.org/abs/1812.04244> (visited on 07/15/2025).
- [388] A. M. Sirunyan et al. “Search for nonresonant Higgs boson pair production in the $b\bar{b}b\bar{b}$ final state at $\sqrt{s} = 13$ TeV”. In: *JHEP* 04 (2019), p. 112. DOI: 10.1007/JHEP04(2019)112. arXiv: 1810.11854 [hep-ex].
- [389] D. Terjék, Terjék, and Dávid. “Adversarial Lipschitz Regularization”. In: *arXiv* (Jan. 2019), arXiv:1907.05681. DOI: 10.48550/arxiv.1907.05681. URL: <https://ui.adsabs.harvard.edu/abs/2019arXiv190705681T/abstract>.
- [390] C. Winkler, D. Worrall, E. Hoozeboom, and M. Welling. “Learning Likelihoods with Conditional Normalizing Flows”. In: (Nov. 2019). URL: <https://arxiv.org/pdf/1912.00042>.
- [391] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena. *Self-Attention Generative Adversarial Networks*. arXiv:1805.08318 [stat]. June 2019. DOI: 10.48550/

- arXiv.1805.08318. URL: <http://arxiv.org/abs/1805.08318> (visited on 07/17/2025).
- [392] A. Bozson, G. Cowan, and F. Spanò. “Unfolding with Gaussian Processes”. In: (Nov. 2018). _eprint: 1811.01242.
- [393] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez. “A guide to constraining effective field theories with machine learning”. In: *Physical Review D* 98.5 (Sept. 2018). Publisher: American Physical Society, p. 052004. DOI: 10.1103/PhysRevD.98.052004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.052004> (visited on 07/15/2025).
- [394] K. Datta, D. Kar, and D. Roy. *Unfolding with Generative Adversarial Networks*. arXiv:1806.00433 [physics]. Aug. 2018. DOI: 10.48550/arXiv.1806.00433. URL: <http://arxiv.org/abs/1806.00433> (visited on 07/17/2025).
- [395] S. Depeweg, J. M. Hernández-Lobato, F. Doshi-Velez, and S. Udluft. *Decomposition of Uncertainty in Bayesian Deep Learning for Efficient and Risk-sensitive Learning*. June 15, 2018. DOI: 10.48550/arXiv.1710.07283. arXiv: 1710.07283 [stat]. URL: <http://arxiv.org/abs/1710.07283> (visited on 07/24/2025). Pre-published.
- [396] F. A. Dreyer, G. P. Salam, and G. Soyez. “The Lund Jet Plane”. In: *Journal of High Energy Physics* 12.12 (Dec. 2018), p. 064. ISSN: 1029-8479. DOI: 10.1007/JHEP12(2018)064. URL: <http://arxiv.org/abs/1807.04758> (visited on 07/24/2025).

- [397] M. Erdmann, L. Geiger, J. Glombitza, and D. Schmidt. *Generating and refining particle detector simulations using the Wasserstein distance in adversarial networks*. arXiv:1802.03325 [astro-ph]. Feb. 2018. DOI: 10.48550/arXiv.1802.03325. URL: <http://arxiv.org/abs/1802.03325> (visited on 07/17/2025).
- [398] C. K. Fisher, A. M. Smith, and J. R. Walsh. *Boltzmann Encoded Adversarial Machines*. arXiv:1804.08682 [stat]. Apr. 2018. DOI: 10.48550/arXiv.1804.08682. URL: <http://arxiv.org/abs/1804.08682> (visited on 07/24/2025).
- [399] P. Glaysher. “BDTs in the Search for $t\bar{t}H$ Production with Higgs Decays to $b\bar{b}$ at ATLAS”. In: *PoS EPS-HEP2017* (2018). Ed. by P. Checchia et al., p. 698. DOI: 10.22323/1.314.0698.
- [400] P. Henderson, J. Romoff, and J. Pineau. *Where Did My Optimum Go?: An Empirical Analysis of Gradient Descent Optimization in Policy Gradient Methods*. Oct. 2018. DOI: 10.48550/arXiv.1810.02525. URL: <http://arxiv.org/abs/1810.02525> (visited on 07/21/2025).
- [401] T. Karras, T. Aila, S. Laine, and J. Lehtinen. *Progressive Growing of GANs for Improved Quality, Stability, and Variation*. arXiv:1710.10196 [cs]. Feb. 2018. DOI: 10.48550/arXiv.1710.10196. URL: <http://arxiv.org/abs/1710.10196> (visited on 07/17/2025).
- [402] P. T. Komiske, E. M. Metodiev, B. Nachman, and M. D. Schwartz. “Pileup Mitigation with Machine Learning (PUMML)”. In: *Journal of High Energy Physics* 2017.12 (Jan. 2018). DOI: 10.1007/jhep12(2017)051. URL: <http://arxiv.org/abs/1707.08600>
[http://dx.doi.org/10.1007/JHEP12\(2017\)051](http://dx.doi.org/10.1007/JHEP12(2017)051).

- [403] P. T. Komiske, E. M. Metodiev, and M. D. Schwartz. “Deep learning in color: towards automated quark/gluon jet discrimination”. In: *Journal of High Energy Physics* 2017.1 (Sept. 2018). DOI: 10.1007/jhep01(2017)110. URL: <http://arxiv.org/abs/1612.01551>[http://dx.doi.org/10.1007/JHEP01\(2017\)110](http://dx.doi.org/10.1007/JHEP01(2017)110).
- [404] P. T. Komiske, E. M. Metodiev, and J. Thaler. “Energy flow polynomials: A complete linear basis for jet substructure”. In: *Journal of High Energy Physics* 2018.4 (Apr. 2018). DOI: 10.1007/jhep04(2018)013. URL: <http://arxiv.org/abs/1712.07124>[http://dx.doi.org/10.1007/JHEP04\(2018\)013](http://dx.doi.org/10.1007/JHEP04(2018)013).
- [405] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. *Spectral Normalization for Generative Adversarial Networks*. arXiv:1802.05957 [cs]. Feb. 2018. DOI: 10.48550/arXiv.1802.05957. URL: <http://arxiv.org/abs/1802.05957> (visited on 07/17/2025).
- [406] S.-M. Moosavi-Dezfooli, A. Fawzi, J. Uesato, and P. Frossard. *Robustness via curvature regularization, and vice versa*. arXiv:1811.09716 [cs]. Nov. 2018. DOI: 10.48550/arXiv.1811.09716. URL: <http://arxiv.org/abs/1811.09716> (visited on 07/14/2025).
- [407] M. Paganini, L. De Oliveira, and B. Nachman. “CaloGAN : Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks”. In: *Phys.Rev.D* 97.1 (Jan. 2018). ISSN: 24700029. DOI: 10.1103/physrevd.97.014021.

- [408] G. Papamakarios, T. Pavlakou, and I. Murray. *Masked Autoregressive Flow for Density Estimation*. arXiv:1705.07057 [stat]. June 2018. DOI: 10.48550/arXiv.1705.07057. URL: <http://arxiv.org/abs/1705.07057> (visited on 07/16/2025).
- [409] P. Petersen and F. Voigtlaender. *Optimal approximation of piecewise smooth functions using deep ReLU neural networks*. arXiv:1709.05289 [math]. May 2018. DOI: 10.48550/arXiv.1709.05289. URL: <http://arxiv.org/abs/1709.05289> (visited on 07/15/2025).
- [410] T. J. Sejnowski. “The Deep Learning Revolution: Machine Intelligence Meets Human Intelligence”. In: *MIT Press* (2018).
- [411] M. Stoye, J. Brehmer, G. Louppe, J. Pavez, and K. Cranmer. “Likelihood-free inference with an improved cross-entropy estimator”. In: (Aug. 2018). URL: <https://arxiv.org/pdf/1808.00973>.
- [412] M. Tanabashi, K. Hagiwara, K. Hikasa, et al. “Review of Particle Physics”. In: *Physical Review D* 98.3 (Aug. 2018). Publisher: American Physical Society, p. 030001. DOI: 10.1103/PhysRevD.98.030001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.030001> (visited on 07/21/2025).
- [413] A. Vogt, S.-O. Moch, B. Ruijl, T. Ueda, and J. A. M. Vermaseren. “Four-Loop Results on Anomalous Dimensions and Splitting Functions in QCD”. In: *PoS Radcor2017* (Jan. 2018). Ed. by A. Hoang and C. Schneider, p. 046. DOI: 10.22323/1.290.0046. arXiv: 1801.06085 [hep-ph].

- [414] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. “Dynamic Graph CNN for Learning on Point Clouds”. In: (Jan. 2018). arXiv: 1801.07829 [cs.CV].
- [415] Y. Wu, Y. Qin, and M. Zhu. *High-dimensional covariance matrix estimation using a low-rank and diagonal decomposition*. arXiv:1802.06048 [stat]. Feb. 2018. DOI: 10.48550/arXiv.1802.06048. URL: <http://arxiv.org/abs/1802.06048> (visited on 07/23/2025).
- [416] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Poczos, R. Salakhutdinov, and A. Smola. “Deep Sets”. In: (Apr. 2018). DOI: 10.48550/arXiv.1703.06114. URL: <http://arxiv.org/abs/1703.06114> (visited on 07/24/2025).
- [417] M. Aaboud, G. Aad, B. Abbott, O. Abdinov, B. Abeloos, S. H. Abidi, O. S. AbouZeid, N. L. Abraham, H. Abramowicz, H. Abreu, et al. “Measurement of Lepton Differential Distributions and the Top Quark Mass in $t\bar{t}$ production in Pp Collisions at $\sqrt{s}=8$ TeV with the ATLAS Detector”. In: *The European Physical Journal C* 77.11 (Nov. 2017), p. 804. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-5349-9. arXiv: 1709.09407 [hep-ex]. (Visited on 07/19/2025).
- [418] T. Adye. “RooUnfold a history, scope, and status DIANA Meeting-Unfolding in HEP 13 th”. In: (2017). URL: <http://hepunix.rl.ac.uk/~adye/software/unfold/RooUnfold.html>.
- [419] M. Arjovsky and L. Bottou. *Towards Principled Methods for Training Generative Adversarial Networks*. arXiv:1701.04862 [stat]. Jan. 2017. DOI: 10.48550/

- arXiv:1701.04862. URL: <http://arxiv.org/abs/1701.04862> (visited on 07/17/2025).
- [420] M. Arjovsky, S. Chintala, and L. Bottou. *Wasserstein GAN*. arXiv:1701.07875 [stat]. Dec. 2017. DOI: 10.48550/arXiv.1701.07875. URL: <http://arxiv.org/abs/1701.07875> (visited on 07/17/2025).
- [421] J. Bellm, S. Gieseke, D. Grellscheid, P. Kirchgaesser, F. Loshaj, G. Nail, A. Papaefstathiou, S. Plätzer, R. Podskubka, M. Rauch, et al. “Herwig 7.1 Release Note”. In: (May 2017). URL: <http://arxiv.org/abs/1705.06919>.
- [422] N. Berger. *Lecture on Statistical analysis methods by ATLAS*. Sept. 2017. URL: <https://cds.cern.ch/record/2285058>.
- [423] M. Bessner. “Measurement of differential di-photon plus jet cross sections using the ATLAS detector”. PhD thesis. Hamburg U., 2017. DOI: 10.3204/pubdb-2017-10063.
- [424] A. Bevan, R. G. Goñi, T. Stevenson, and T. Stevenson. “Support vector machines and generalisation in HEP”. In: *J. Phys. Conf. Ser.* 898.7 (2017). Ed. by R. Mount and C. Tull. _eprint: 1702.04686, p. 072021. DOI: 10.1088/1742-6596/898/7/072021.
- [425] V. Cheung and R. Vogt. “Polarized Heavy Quarkonium Production in the Color Evaporation Model”. In: *Physical Review D* 95.7 (Apr. 2017), p. 074021. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.95.074021. URL: <http://arxiv.org/abs/1702.07809> (visited on 07/23/2025).

- [426] C. M. S. Collaboration. “Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV”. en. In: *Journal of Instrumentation* 12.02 (Feb. 2017), P02014. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/02/P02014. URL: <http://arxiv.org/abs/1607.03663> (visited on 07/21/2025).
- [427] M. Dashti and A. M. Stuart. “The bayesian approach to inverse problems”. In: *Handbook of Uncertainty Quantification* (June 2017), pp. 311–428. DOI: 10.1007/978-3-319-12385-1_{_}7.
- [428] A. Di Meglio, M. Girone, F. Rademakers, and A. Purcell. “submitter : CERN openlab white paper on future ICT challenges in scientific research”. In: (2017). DOI: 10.5281/zenodo.998694. URL: <https://cds.cern.ch/record/2301895>.
- [429] L. Dinh, J. Sohl-Dickstein, and S. Bengio. *Density estimation using Real NVP*. arXiv:1605.08803 [cs]. Feb. 2017. DOI: 10.48550/arXiv.1605.08803. URL: <http://arxiv.org/abs/1605.08803> (visited on 07/16/2025).
- [430] M. Frate, K. Cranmer, S. Kalia, A. Vandenberg-Rodes, and D. Whiteson. “Modeling Smooth Backgrounds and Generic Localized Signals with Gaussian Processes”. In: (Sept. 2017). DOI: 10.48550/arXiv.1709.05681. URL: <http://arxiv.org/abs/1709.05681> (visited on 07/21/2025).
- [431] T. Gehrmann, G. Luisoni, and P. F. Monni. “Resummation of Jet Rates and Event-Shape Distributions in E^+E^- ”. In: *Parton Radiation and Fragmentation from LHC to FCC-ee*. 2017, pp. 97–102.

- [432] T. Gehrmann, G. Luisoni, and P. F. Monni. “Resummation of jet rates and event-shape distributions in e^+e^- ”. In: *Parton radiation and fragmentation from LHC to FCC-ee*. 2017, pp. 97–102.
- [433] I. Goodfellow. *NIPS 2016 Tutorial: Generative Adversarial Networks*. arXiv:1701.00160 [cs]. Apr. 2017. DOI: 10.48550/arXiv.1701.00160. URL: <http://arxiv.org/abs/1701.00160> (visited on 07/17/2025).
- [434] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville. *Improved Training of Wasserstein GANs*. arXiv:1704.00028 [cs]. Dec. 2017. DOI: 10.48550/arXiv.1704.00028. URL: <http://arxiv.org/abs/1704.00028> (visited on 07/17/2025).
- [435] A. Kendall and Y. Gal. “What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?” In: *arXiv:1703.04977 [cs.CV]* (Oct. 5, 2017). DOI: 10.48550/arXiv.1703.04977. arXiv: 1703.04977 [cs.CV]. URL: <http://arxiv.org/abs/1703.04977> (visited on 07/24/2025).
- [436] D. P. Kingma and J. Ba. *Adam: A Method for Stochastic Optimization*. arXiv:1412.6980 [cs]. Jan. 2017. DOI: 10.48550/arXiv.1412.6980. URL: <http://arxiv.org/abs/1412.6980> (visited on 07/15/2025).
- [437] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. *Improving Variational Inference with Inverse Autoregressive Flow*. arXiv:1606.04934 [cs]. Jan. 2017. DOI: 10.48550/arXiv.1606.04934. URL: <http://arxiv.org/abs/1606.04934> (visited on 07/16/2025).

- [438] S. Kolouri, S. R. Park, M. Thorpe, D. Slepcev, and G. K. Rohde. “Optimal Mass Transport: Signal processing and machine-learning applications”. In: *IEEE Signal Processing Magazine* 34.4 (2017), pp. 43–59. DOI: 10.1109/MSP.2017.2695801.
- [439] M. Kuusela and P. B. Stark. “Shape-constrained uncertainty quantification in unfolding steeply falling elementary particle spectra”. In: *The Annals of Applied Statistics* 11.3 (Sept. 2017). Publisher: Institute of Mathematical Statistics, pp. 1671–1710. ISSN: 1932-6157, 1941-7330. DOI: 10.1214/17-aoas1053. URL: <https://projecteuclid.org/journals/annals-of-applied-statistics/volume-11/issue-3/Shape-constrained-uncertainty-quantification-in-unfolding-steeply-falling-elementary-particle/10.1214/17-AOAS1053.full> (visited on 07/13/2025).
- [440] Z. Lu, H. Pu, F. Wang, Z. Hu, and L. Wang. *The Expressive Power of Neural Networks: A View from the Width*. arXiv:1709.02540 [cs]. Nov. 2017. DOI: 10.48550/arXiv.1709.02540. URL: <http://arxiv.org/abs/1709.02540> (visited on 07/15/2025).
- [441] A. Ly, M. Marsman, J. Verhagen, R. Grasman, and E.-J. Wagenmakers. *A Tutorial on Fisher Information*. Oct. 2017. DOI: 10.48550/arXiv.1705.01064. URL: <http://arxiv.org/abs/1705.01064> (visited on 07/21/2025).
- [442] E. Maguire, L. Heinrich, and G. Watt. “HEPData: a repository for high energy physics data”. In: *J. Phys. Conf. Ser.* 898.10 (2017). Ed. by R. Mount and C. Tull, p. 102006. DOI: 10.1088/1742-6596/898/10/102006. arXiv: 1704.05473 [hep-ex].

- [443] M. U. Mozer. “Jet reconstruction and substructure measurements in ATLAS and CMS with first Run-2 data”. en. In: *Proceedings of Fourth Annual Large Hadron Collider Physics – PoS(LHCP2016)*. Vol. 276. Conference Name: Fourth Annual Large Hadron Collider Physics. SISSA Medialab, Mar. 2017, p. 090. DOI: 10.22323/1.276.0090. URL: <https://pos.sissa.it/276/090> (visited on 07/13/2025).
- [444] Y. Mroueh, T. Sercu, and V. Goel. *McGan: Mean and Covariance Feature Matching GAN*. arXiv:1702.08398 [cs]. June 2017. DOI: 10.48550/arXiv.1702.08398. URL: <http://arxiv.org/abs/1702.08398> (visited on 07/17/2025).
- [445] J. Pearkes, W. Fedorko, A. Lister, and C. Gay. “Jet Constituents for Deep Neural Network Based Top Quark Tagging”. In: (Apr. 2017). URL: <https://arxiv.org/pdf/1704.02124>.
- [446] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. C. Courville. “FiLM: Visual Reasoning with a General Conditioning Layer”. In: *CoRR* abs/1709.07871 (2017). arXiv: 1709.07871. URL: <http://arxiv.org/abs/1709.07871>.
- [447] P. Ramachandran, B. Zoph, and Q. V. Le Google Brain. “Searching for Activation Functions”. In: *6th International Conference on Learning Representations, ICLR 2018 - Workshop Track Proceedings* (Oct. 2017). URL: <https://arxiv.org/pdf/1710.05941>.
- [448] S. Ruder. *An overview of gradient descent optimization algorithms*. arXiv:1609.04747 [cs]. June 2017. DOI: 10.48550/arXiv.1609.04747. URL: <http://arxiv.org/abs/1609.04747> (visited on 07/15/2025).

- [449] S. Schmitt. “Data Unfolding Methods in High Energy Physics”. In: *EPJ Web Conf.* 137 (2017). Ed. by Y. Foka, N. Brambilla, and V. Kovalenko. _eprint: 1611.01927, p. 11008. DOI: 10.1051/epjconf/201713711008.
- [450] S. Schmitt. “Data Unfolding Methods in High Energy Physics”. In: *EPJ Web of Conferences* 137 (Mar. 2017), p. 11008. ISSN: 2100-014x. DOI: 10.1051/epjconf/201713711008. URL: <https://www.epj-conferences.org/articles/epjconf/abs/2017/06/epjconf%5C%5Fconf2017%5C%5F11008/epjconf%5C%5Fconf2017%5C%5F11008.html>.
- [451] A. M. Sirunyan et al. “Particle-flow reconstruction and global event description with the CMS detector”. In: *Jinst* 12.10 (2017), P10003. DOI: 10.1088/1748-0221/12/10/p10003. arXiv: 1706.04965 [physics.ins-det].
- [452] A. M. Sirunyan et al. “Measurements of jet charge with dijet events in pp collisions at $\sqrt{s} = 8$ TeV”. In: *Jhep* 10 (2017), p. 131. DOI: 10.1007/jhep10(2017)131. arXiv: 1706.05868 [hep-ex].
- [453] W. Tang, X. Li, X. Qian, H. Wei, and C. Zhang. “Data Unfolding with Wiener-SVD Method”. In: (2017).
- [454] W. Tang, X. Li, X. Qian, H. Wei, and C. Zhang. “Data Unfolding with Wiener-SVD Method”. en. In: *Journal of Instrumentation* 12.10 (Oct. 2017), P10002. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/10/p10002. URL: <https://dx.doi.org/10.1088/1748-0221/12/10/P10002> (visited on 07/13/2025).

- [455] A. Vaswani, G. Brain, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. “Attention Is All You Need”. In: (June 2017), p. 1. URL: <https://arxiv.org/pdf/1706.03762>.
- [456] N. Zardoshti. “Investigating the Role of Coherence Effects on Jet Quenching in Pb-Pb Collisions at $\sqrt{s_{NN}}=2.76$ TeV using Jet Substructure”. In: *Nucl. Phys. A* 967 (2017). Ed. by U. Heinz, O. Evdokimov, and P. Jacobs. _eprint: 1705.03383, pp. 560–563. DOI: 10.1016/j.nuclphysa.2017.05.055.
- [457] “6.1 Metric entropy and its uses”. In: (2016).
- [458] G. Aad et al. “Measurement of fiducial differential cross sections of gluon-fusion production of Higgs bosons decaying to $WW^* \rightarrow e\nu\mu\nu$ with the ATLAS detector at $\sqrt{s} = 8$ TeV”. In: *Jhep* 08 (2016), p. 104. DOI: 10.1007/jhep08(2016)104. arXiv: 1604.02997 [hep-ex].
- [459] G. Aad et al. “Measurement of jet charge in dijet events from $\sqrt{s}=8$ TeV pp collisions with the ATLAS detector”. In: *Phys. Rev. D* 93.5 (2016), p. 052003. DOI: 10.1103/PhysRevD.93.052003. arXiv: 1509.05190 [hep-ex].
- [460] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al. “TensorFlow: A system for large-scale machine learning”. In: *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016* (May 2016), pp. 265–283. URL: <https://arxiv.org/pdf/1605.08695>.
- [461] J. Allison, K. Amako, J. Apostolakis, P. Arce, M. Asai, T. Aso, E. Bagli, A. Bagulya, S. Banerjee, G. Barrand, et al. “Recent developments in Geant4”. In:

- Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 835 (Nov. 2016), pp. 186–225. ISSN: 0168-9002. DOI: 10.1016/j.nima.2016.06.125. URL: <https://www.sciencedirect.com/science/article/pii/S0168900216306957> (visited on 07/13/2025).
- [462] T. Chen and C. Guestrin. “XGBoost: A Scalable Tree Boosting System”. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Aug. 2016, pp. 785–794. DOI: 10.1145/2939672.2939785. URL: <https://doi.org/10.1145/2939672.2939785> (visited on 08/11/2025).
- [463] D.-A. Clevert, T. Unterthiner, and S. Hochreiter. *Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)*. arXiv:1511.07289 [cs]. Feb. 2016. DOI: 10.48550/arXiv.1511.07289. URL: <http://arxiv.org/abs/1511.07289> (visited on 07/15/2025).
- [464] T. S. Cohen and M. Welling. “Group Equivariant Convolutional Networks”. In: (June 2016). DOI: 10.48550/arXiv.1602.07576. URL: <http://arxiv.org/abs/1602.07576> (visited on 07/24/2025).
- [465] R. Contino, A. Falkowski, F. Goertz, C. Grojean, and F. Riva. “On the validity of the effective field theory approach to SM precision tests”. en. In: *Journal of High Energy Physics* 2016.7 (July 2016), p. 144. ISSN: 1029-8479. DOI: 10.1007/jhep07(2016)144. URL: [https://doi.org/10.1007/JHEP07\(2016\)144](https://doi.org/10.1007/JHEP07(2016)144) (visited on 07/13/2025).

- [466] A. Deur, S. J. Brodsky, and G. F. de T ramond. “The QCD running coupling”. In: *Progress in Particle and Nuclear Physics* 90 (Sept. 2016), pp. 1–74. ISSN: 01466410. DOI: 10.1016/j.pnnp.2016.04.003.
- [467] M. J. Kuusela. “Uncertainty quantification in unfolding elementary particle spectra at the Large Hadron Collider”. PhD Thesis. Ecole Polytechnique, Lausanne, 2016. DOI: 10.5075/epfl-thesis-7118.
- [468] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, and A. Schwartzman. “Jet-images – deep learning edition”. In: *Journal of High Energy Physics* 2016.7 (July 2016), pp. 1–32. ISSN: 10298479. DOI: 10.1007/jhep07(2016)069/metrics. URL: [https://link.springer.com/article/10.1007/JHEP07\(2016\)069](https://link.springer.com/article/10.1007/JHEP07(2016)069).
- [469] G. Papamakarios and I. Murray. “Fast ϵ -free Inference of Simulation Models with Bayesian Conditional Density Estimation”. In: *Advances in Neural Information Processing Systems* (May 2016), pp. 1036–1044. ISSN: 10495258. URL: <https://arxiv.org/pdf/1605.06376>.
- [470] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. *Improved Techniques for Training GANs*. arXiv:1606.03498 [cs]. June 2016. DOI: 10.48550/arXiv.1606.03498. URL: <http://arxiv.org/abs/1606.03498> (visited on 07/17/2025).
- [471] P. Shah, U. K. Khankhoje, and M. Moghaddam. “Inverse Scattering Using a Joint L1-L2 Norm-Based Regularization”. In: *IEEE Transactions on Antennas and Propagation* 64.4 (Apr. 2016), pp. 1373–1384. ISSN: 1558-2221. DOI: 10.

- 1109/tap.2016.2529641. URL: <https://ieeexplore.ieee.org/document/7406678> (visited on 07/15/2025).
- [472] I. O. Tolstikhin, B. K. Sriperumbudur, and B. Schölkopf. “Minimax Estimation of Maximum Mean Discrepancy with Radial Kernels”. In: *Advances in Neural Information Processing Systems*. Vol. 29. Curran Associates, Inc., 2016. URL: https://papers.nips.cc/paper_files/paper/2016/hash/5055cbf43fac3f7e2336b27310f0b9ef-Abstract.html (visited on 07/24/2025).
- [473] C. Weisser and M. Williams. “Machine learning and multivariate goodness of fit”. In: (Dec. 2016). DOI: 10.48550/arXiv.1612.07186. URL: <http://arxiv.org/abs/1612.07186> (visited on 07/21/2025).
- [474] L. Yan, S. Pal, and J.-Y. Ollitrault. “Nonlinear hydrodynamic response confronts LHC data”. In: *Nuclear Physics A* 956 (Dec. 2016), pp. 340–343. ISSN: 03759474. DOI: 10.1016/j.nuclphysa.2016.01.010. URL: <http://arxiv.org/abs/1601.00040> (visited on 07/24/2025).
- [475] G. Zech. “Analysis of distorted measurements – parameter estimation and unfolding”. In: (July 2016). _eprint: 1607.06910.
- [476] G. Aad, T. Abajyan, B. Abbott, J. Abdallah, S. Abdel Khalek, O. Abdinov, R. Aben, B. Abi, M. Abolins, O. S. AbouZeid, et al. “Jet energy measurement and its systematic uncertainty in proton–proton collisions at $\sqrt{s}=7$ TeV with the ATLAS detector”. en. In: *The European Physical Journal C* 75.1 (Jan. 2015), p. 17. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-014-3190-y.

- URL: <https://doi.org/10.1140/epjc/s10052-014-3190-y> (visited on 07/19/2025).
- [477] C. Adam-Bourdarios, G. Cowan, C. Germain-Renaud, I. Guyon, B. Kégl, and D. Rousseau. “The Higgs Machine Learning Challenge”. In: *Journal of Physics: Conference Series* 664.7 (Dec. 2015), p. 072015. ISSN: 1742-6596. DOI: 10.1088/1742-6596/664/7/072015. URL: <https://iopscience.iop.org/article/10.1088/1742-6596/664/7/072015%20https://iopscience.iop.org/article/10.1088/1742-6596/664/7/072015/meta>.
- [478] J. A. Aguilar-Saavedra, D. Amidei, A. Juste, and M. Perez-Victoria. “Asymmetries in top quark pair production at hadron colliders”. In: *Rev. Mod. Phys.* 87 (2015), pp. 421–455. DOI: 10.1103/RevModPhys.87.421. arXiv: 1406.1798 [hep-ph].
- [479] M. Baak, S. Gadatsch, R. Harrington, and W. Verkerke. “Interpolation between multi-dimensional histograms using a new non-linear moment morphing method”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 771 (Jan. 2015), pp. 39–48. ISSN: 0168-9002. DOI: 10.1016/j.nima.2014.10.033. URL: <https://www.sciencedirect.com/science/article/pii/S0168900214011814> (visited on 07/14/2025).
- [480] P. Baldi, P. Sadowski, and D. Whiteson. “Enhanced Higgs boson to $\tau+\tau$ - search with deep learning”. In: *Physical Review Letters* 114.11 (Mar. 2015), p. 111801. ISSN: 10797114. DOI: 10.1103/physrevlett.114.111801/figures/3/

- thumbnail. URL: <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.114.111801>.
- [481] M. Cacciari, G. P. Salam, and G. Soyez. “SoftKiller, a particle-level pileup removal method”. en. In: *The European Physical Journal C* 75.2 (Feb. 2015), p. 59. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-015-3267-2. URL: <https://doi.org/10.1140/epjc/s10052-015-3267-2> (visited on 07/21/2025).
- [482] T. Carli, K. Rabbertz, and S. Schumann. “Studies of Quantum Chromodynamics at the LHC”. In: *The Large Hadron Collider. Harvest of Run 1*. Ed. by T. Schörner-Sadenius. 2015, pp. 139–194. DOI: 10.1007/978-3-319-15001-7_5. arXiv: 1506.03239 [hep-ex].
- [483] F. Chollet et al. *Keras*. <https://keras.io>. 2015.
- [484] J. R. Christiansen and P. Z. Skands. “String Formation Beyond Leading Colour”. In: *Jhep* 08 (2015), p. 003. DOI: 10.1007/jhep08(2015)003. arXiv: 1505.01681 [hep-ph].
- [485] G. Cowan. “Statistics for Searches at the LHC”. In: *LHC Phenomenology*. Ed. by E. Gardi, N. Glover, and A. Robson. Cham: Springer International Publishing, 2015, pp. 321–355. ISBN: 978-3-319-05362-2. DOI: 10.1007/978-3-319-05362-2_9. URL: http://link.springer.com/10.1007/978-3-319-05362-2_9 (visited on 07/13/2025).
- [486] E. Denton, S. Chintala, A. Szlam, and R. Fergus. *Deep Generative Image Models using a Laplacian Pyramid of Adversarial Networks*. arXiv:1506.05751

- [cs]. June 2015. DOI: 10.48550/arXiv.1506.05751. URL: <http://arxiv.org/abs/1506.05751> (visited on 07/17/2025).
- [487] L. Dinh, D. Krueger, and Y. Bengio. *NICE: Non-linear Independent Components Estimation*. arXiv:1410.8516 [cs]. Apr. 2015. DOI: 10.48550/arXiv.1410.8516. URL: <http://arxiv.org/abs/1410.8516> (visited on 07/16/2025).
- [488] B. Hall. “Matrix Lie Groups”. en. In: *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*. Ed. by B. C. Hall. Cham: Springer International Publishing, 2015, pp. 3–30. ISBN: 978-3-319-13467-3. DOI: 10.1007/978-3-319-13467-3_1. URL: https://doi.org/10.1007/978-3-319-13467-3_1 (visited on 07/24/2025).
- [489] K. He, X. Zhang, S. Ren, and J. Sun. *Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification*. arXiv:1502.01852 [cs] version: 1. Feb. 2015. DOI: 10.48550/arXiv.1502.01852. URL: <http://arxiv.org/abs/1502.01852> (visited on 07/15/2025).
- [490] K. He, X. Zhang, S. Ren, and J. Sun. “Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification”. In: *2015 IEEE International Conference on Computer Vision (ICCV)*. 2015, pp. 1026–1034. DOI: 10.1109/iccv.2015.123.
- [491] S. Ioffe and C. Szegedy. “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift”. In: *32nd International Conference on Machine Learning, ICML 2015 1* (Feb. 2015), pp. 448–456. URL: <https://arxiv.org/pdf/1502.03167>.

- [492] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu. “Spatial Transformer Networks”. In: *Advances in Neural Information Processing Systems* 2015-January (June 2015), pp. 2017–2025. ISSN: 10495258. URL: <https://arxiv.org/abs/1506.02025v3>.
- [493] V. Khachatryan et al. “Search for quark contact interactions and extra spatial dimensions using dijet angular distributions in proton–proton collisions at $\sqrt{s} = 8$ TeV”. In: *Phys. Lett. B* 746 (2015), pp. 79–99. DOI: 10.1016/j.physletb.2015.04.042. arXiv: 1411.2646 [hep-ex].
- [494] M. Kuusela and V. M. Panaretos. “Statistical unfolding of elementary particle spectra: Empirical Bayes estimation and bias-corrected uncertainty quantification”. In: *Annals of Applied Statistics* 9.3 (Nov. 2015), pp. 1671–1705. DOI: 10.1214/15-AOAS857. URL: <http://arxiv.org/abs/1505.04768><http://dx.doi.org/10.1214/15-AOAS857>.
- [495] A. J. Larkoski, S. Marzani, and J. Thaler. “Sudakov Safety in Perturbative QCD”. In: *Physical Review D* 91.11 (June 2015), p. 111501. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.91.111501. URL: <http://arxiv.org/abs/1502.01719> (visited on 07/19/2025).
- [496] A. Mertens. “New features in Delphes 3”. In: *J. Phys. Conf. Ser.* 608.1 (2015). Ed. by L. Fiala, M. Lokajicek, and N. Tumova, p. 012045. DOI: 10.1088/1742-6596/608/1/012045.

- [497] M. A. Nielsen. “Neural Networks and Deep Learning”. en. In: (2015). Publisher: Determination Press. URL: <http://neuralnetworksanddeeplearning.com> (visited on 07/15/2025).
- [498] M. H. Pesaran. *Time Series and Panel Data Econometrics*. en. Google-Books-ID: 7RokCwAAQBAJ. Oxford University Press, Oct. 2015. ISBN: 978-0-19-105847-9.
- [499] M.-T. Puth, M. Neuhäuser, and G. D. Ruxton. “On the variety of methods for calculating confidence intervals by bootstrapping”. en. In: *Journal of Animal Ecology* 84.4 (2015). _eprint: <https://besjournals.onlinelibrary.wiley.com/doi/pdf/10.1111/1365-2656.12382>, pp. 892–897. ISSN: 1365-2656. DOI: 10.1111/1365-2656.12382. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1365-2656.12382> (visited on 07/22/2025).
- [500] D. J. Rezende and S. Mohamed. “Variational Inference with Normalizing Flows”. In: *32nd International Conference on Machine Learning, ICML 2015* 2 (May 2015), pp. 1530–1538. URL: <https://arxiv.org/pdf/1505.05770>.
- [501] K. Sharp and F. Matschinsky. “Translation of Ludwig Boltzmann’s Paper “On the Relationship between the Second Fundamental Theorem of the Mechanical Theory of Heat and Probability Calculations Regarding the Conditions for Thermal Equilibrium” Sitzungberichte der Kaiserlichen Akademie der Wissenschaften. Mathematisch-Naturwissen Classe. Abt. II, LXXVI 1877, pp 373-435 (Wien. Ber. 1877, 76:373-435). Reprinted in Wiss. Abhandlungen, Vol. II, reprint 42, p. 164-223, Barth, Leipzig, 1909”. In: *Entropy 2015, Vol. 17, Pages 1971-2009* 17.4 (Apr. 2015), pp. 1971–2009. ISSN: 1099-4300. DOI:

- 10.3390/e17041971. URL: <https://www.mdpi.com/1099-4300/17/4/1971/htm%20https://www.mdpi.com/1099-4300/17/4/1971>.
- [502] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands. “An introduction to PYTHIA 8.2”. In: *Comput. Phys. Commun.* 191 (2015), pp. 159–177. DOI: 10.1016/j.cpc.2015.01.024. arXiv: 1410.3012 [hep-ph].
- [503] B. Szabó, A. W. v. d. Vaart, and J. H. v. Zanten. “Frequentist coverage of adaptive nonparametric Bayesian credible sets”. In: *The Annals of Statistics* 43.4 (Aug. 2015). arXiv:1310.4489 [math]. ISSN: 0090-5364. DOI: 10.1214/14-aos1270. URL: <http://arxiv.org/abs/1310.4489> (visited on 07/14/2025).
- [504] G. Aad et al. “Measurement of the cross-section of high transverse momentum vector bosons reconstructed as single jets and studies of jet substructure in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”. In: *New J. Phys.* 16.11 (2014), p. 113013. DOI: 10.1088/1367-2630/16/11/113013. arXiv: 1407.0800 [hep-ex].
- [505] G. Aad, T. Abajyan, B. Abbott, et al. “Measurements of normalized differential cross sections for $\overline{t}t$ production in pp collisions at $\sqrt{s}=7\text{ }\mathrm{TeV}$ using the ATLAS detector”. In: *Physical Review D* 90.7 (Oct. 2014), p. 072004. DOI: 10.1103/PhysRevD.90.072004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.90.072004> (visited on 07/23/2025).
- [506] “ATLAS Pythia 8 tunes to 7 TeV data”. In: (Nov. 2014).

- [507] P. Baldi, P. Sadowski, and D. Whiteson. “Searching for exotic particles in high-energy physics with deep learning”. In: *Nature Communications* 5.1 (July 2014), pp. 1–9. ISSN: 20411723. DOI: 10.1038/ncomms5308;techmeta=129,141;subjmeta=25,419,483,639,640,766;kwr=applied+physics,particle+physics,theoretical+physics. URL: <https://www.nature.com/articles/ncomms5308>.
- [508] G. Bohm and G. Zech. “Statistics of weighted Poisson events and its applications”. In: *Nucl. Instrum. Meth. A* 748 (2014), pp. 1–6. DOI: 10.1016/j.nima.2014.02.021. arXiv: 1309.1287 [physics.data-an].
- [509] T. Chen and T. He. “Higgs Boson Discovery with Boosted Trees”. In: *Hepmlnips* (2014).
- [510] Y.-T. Chien and I. Vitev. “Jet Shape Resummation Using Soft-Collinear Effective Theory”. In: *Journal of High Energy Physics* 2014.12 (Dec. 2014). DOI: 10.1007/jhep12(2014)061. URL: <http://arxiv.org/abs/1405.4293> 20[http://dx.doi.org/10.1007/JHEP12\(2014\)061](http://dx.doi.org/10.1007/JHEP12(2014)061).
- [511] A. Collaboration. “Measurement of differential production cross-sections for a Z boson in association with b -jets in 7 TeV proton-proton collisions with the ATLAS detector”. In: *Journal of High Energy Physics* 10.10 (Oct. 2014), p. 141. ISSN: 1029-8479. DOI: 10.1007/JHEP10(2014)141. URL: <http://arxiv.org/abs/1407.3643> (visited on 07/23/2025).

- [512] K. Cranmer. “Practical Statistics for the LHC”. In: *2011 European School of High-Energy Physics*. 2014, pp. 267–308. DOI: 10.5170/CERN-2014-003.267. arXiv: 1503.07622 [physics.data-an].
- [513] J. De Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi. “DELPHES 3: A modular framework for fast simulation of a generic collider experiment”. In: *Journal of High Energy Physics* 2014.2 (Feb. 2014), pp. 1–26. ISSN: 10298479. DOI: 10.1007/jhep02(2014)057/metrics. URL: [https://link.springer.com/article/10.1007/JHEP02\(2014\)057](https://link.springer.com/article/10.1007/JHEP02(2014)057).
- [514] J. L. Fernández-Martínez, J. L. G. Pallero, Z. Fernández-Muñiz, and L. M. Pedruelo-González. “The effect of noise and Tikhonov’s regularization in inverse problems. Part I: The linear case”. In: *Journal of Applied Geophysics* 108 (Sept. 2014), pp. 176–185. ISSN: 0926-9851. DOI: 10.1016/j.jappgeo.2014.05.006. URL: <https://www.sciencedirect.com/science/article/pii/S0926985114001402> (visited on 07/13/2025).
- [515] S. Fiorendi. “Measurement of the branching fraction of $B_c \rightarrow J/\psi \pi^+ \pi^+ \pi^-$ relative to $B_c \rightarrow J/\psi \pi^+$ and of the production cross-section of $B_c \rightarrow J/\psi \pi^+$ relative to $B^+ \rightarrow J/\psi K^+$ with the CMS experiment”. PhD thesis. Milan Bicocca U., 2014.
- [516] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. “Generative Adversarial Networks”. In: *Science Robotics* 3.January (June 2014), pp. 2672–2680. ISSN: 10495258. URL: <https://arxiv.org/pdf/1406.2661>.

- [517] W. Hollik. “Quantum Field Theory and the Standard ModelQuantum Field Theory and the Standard Model , Matthew D. Schwartz, Cambridge U. Press, 2014. \$90.00 (850 pp.). ISBN 978-1-107-03473-0 ”. In: *Physics Today* 67.12 (2014), pp. 57–58. ISSN: 0031-9228.
- [518] D. Kotlorz and S. V. Mikhailov. “Cut Moments and a Generalization of DGLAP Equations”. In: *Journal of High Energy Physics* 06.6 (2014), p. 065. ISSN: 1029-8479. DOI: 10.1007/jhep06(2014)065. arXiv: 1404.5172 [hep-th]. (Visited on 07/18/2025).
- [519] A. J. Larkoski, S. Marzani, G. Soyez, and J. Thaler. “Soft drop”. In: *Journal of High Energy Physics* 2014 2014:5 2014.5 (May 2014), pp. 1–46. ISSN: 1029-8479. DOI: 10.1007/jhep05(2014)146. URL: [https://link.springer.com/article/10.1007/JHEP05\(2014\)146](https://link.springer.com/article/10.1007/JHEP05(2014)146).
- [520] X. Lin, C. Genest, D. L. Banks, G. Molenberghs, D. W. Scott, and J.-L. Wang, eds. *Past, Present, and Future of Statistical Science*. New York: Chapman and Hall/CRC, Mar. 2014. ISBN: 978-0-429-17128-4. DOI: 10.1201/b16720.
- [521] M. Perry, H. B. Prosper, and A. Meyer-Baese. “GPU Implementation of Bayesian Neural Network Construction for Data-Intensive Applications”. In: *J. Phys. Conf. Ser.* 513 (2014). Ed. by D. L. Groep and D. Bonacorsi, p. 022027. DOI: 10.1088/1742-6596/513/2/022027.
- [522] M. D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge: Cambridge University Press, 2014. ISBN: 978-1-107-03473-0. URL: <https://www>.

- cambridge.org/core/books/quantum-field-theory-and-the-standard-model/4AC1D88C3AD14C55A49B1A8E99CF85A0 (visited on 08/14/2025).
- [523] S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge: Cambridge University Press, 2014. ISBN: 978-1-107-05713-5. DOI: 10.1017/CB09781107298019. URL: <https://www.cambridge.org/core/books/understanding-machine-learning/3059695661405D25673058E> (visited on 07/24/2025).
- [524] P. Skands, S. Carrazza, and J. Rojo. “Tuning PYTHIA 8.1: the Monash 2013 Tune”. In: *Eur. Phys. J. C* 74.8 (2014), p. 3024. DOI: 10.1140/epjc/s10052-014-3024-y. arXiv: 1404.5630 [hep-ph].
- [525] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting”. In: *Journal of Machine Learning Research* 15.56 (2014), pp. 1929–1958. ISSN: 1533-7928. URL: <http://jmlr.org/papers/v15/srivastava14a.html> (visited on 07/20/2025).
- [526] D. Strózik-Kotlorz and A. Kotlorz. “DGLAP Evolution of the Truncated Moments of PDFs in Spin Physics”. In: *Physics of Particles and Nuclei* 45.1 (2014), pp. 52–53. ISSN: 1531-8559. DOI: 10.1134/s1063779614010985. (Visited on 07/18/2025).
- [527] G. Aad et al. “Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV”. In: *Eur. Phys. J. C* 73.3 (2013), p. 2304. DOI: 10.1140/epjc/s10052-013-2304-2. arXiv: 1112.6426 [hep-ex].

- [528] L. Apolinario, N. Armesto, and L. Cunqueiro. “An analysis of the influence of background subtraction and quenching on jet observables in heavy-ion collisions”. In: *Jhep* 02 (2013), p. 022. DOI: 10.1007/jhep02(2013)022. arXiv: 1211.1161 [hep-ph].
- [529] Belle Collaboration, M. Leitgab, R. Seidl, M. Grosse Perdekamp, A. Vossen, I. Adachi, H. Aihara, D. M. Asner, V. Aulchenko, T. Aushev, et al. “Precision Measurement of Charged Pion and Kaon Differential Cross Sections in e^+e^- Annihilation at $\sqrt{s} = 10.52 \text{ GeV}$ ”. In: *Physical Review Letters* 111.6 (Aug. 2013). Publisher: American Physical Society, p. 062002. DOI: 10.1103/PhysRevLett.111.062002. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.111.062002> (visited on 07/13/2025).
- [530] V. Blobel. “Unfolding”. In: *Data Analysis in High Energy Physics*. John Wiley & Sons, Ltd, June 2013, pp. 187–225. DOI: 10.1002/9783527653416.CH6. URL: <https://onlinelibrary.wiley.com/doi/full/10.1002/9783527653416.ch6>
<https://onlinelibrary.wiley.com/doi/abs/10.1002/9783527653416.ch6>
<https://onlinelibrary.wiley.com/doi/10.1002/9783527653416.ch6>.
- [531] S. J. Brodsky, F. Fleuret, C. Hadjidakis, and J. P. Lansberg. “Physics opportunities of a fixed-target experiment using LHC beams”. In: *Physics Reports* 522.4 (Jan. 2013), pp. 239–255. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2012.10.001.

- [532] C. M. S. Collaboration. “Measurement of the differential and double-differential Drell-Yan cross sections in proton-proton collisions at $\sqrt{s} = 7$ TeV”. In: *Journal of High Energy Physics* 12.12 (Dec. 2013), p. 030. ISSN: 1029-8479. DOI: 10.1007/JHEP12(2013)030. URL: <http://arxiv.org/abs/1310.7291> (visited on 07/23/2025).
- [533] M. Dasgupta, A. Fregoso, S. Marzani, and G. P. Salam. “Towards an understanding of jet substructure”. In: *Journal of High Energy Physics* 2013 2013:9 2013.9 (Sept. 2013), pp. 1–57. ISSN: 1029-8479. DOI: 10.1007/jhep09(2013)029. URL: [https://link.springer.com/article/10.1007/JHEP09\(2013\)029](https://link.springer.com/article/10.1007/JHEP09(2013)029).
- [534] C. J. Geyer. “5601 Notes: The Sandwich Estimator”. In: (2013).
- [535] H. Gordon. “Measurement of production asymmetries”. In: *6th International Workshop on Charm Physics*. Nov. 2013. arXiv: 1311.5745 [hep-ex].
- [536] D. P. Kingma and M. Welling. “Auto-Encoding Variational Bayes”. In: *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings* (Dec. 2013). DOI: 10.61603/ceas.v2i1.33. URL: <https://arxiv.org/pdf/1312.6114>.
- [537] D. Krohn, M. D. Schwartz, T. Lin, and W. J. Waalewijn. “Jet charge at the LHC”. In: *Physical Review Letters* 110.21 (May 2013). ISSN: 00319007. DOI: 10.1103/physrevlett.110.212001.
- [538] A. L. Maas, A. Y. Hannun, and A. Y. Ng. “Rectifier Nonlinearities Improve Neural Network Acoustic Models”. In: (2013).

- [539] S. Maeda and T. Iguchi. “A new unfolding code combining maximum entropy and maximum likelihood for neutron spectrum measurement”. In: *Journal of Nuclear Science and Technology* 50.4 (Apr. 2013). Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/00223131.2013.773162>, pp. 381–386. ISSN: 0022-3131. DOI: 10.1080/00223131.2013.773162. URL: <https://doi.org/10.1080/00223131.2013.773162> (visited on 07/13/2025).
- [540] F. Sforza and V. Lippi. “Support vector machine classification on a biased training set: Multi-jet background rejection at hadron colliders”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 722 (Sept. 2013), pp. 11–19. ISSN: 0168-9002. DOI: 10.1016/j.nima.2013.04.046. URL: <https://www.sciencedirect.com/science/article/pii/S0168900213004671> (visited on 07/17/2025).
- [541] P. Skands. “Introduction to QCD”. In: *Searching for New Physics at Small and Large Scales*. 2013, pp. 341–420. DOI: 10.1142/9789814525220_0008. arXiv: 1207.2389 [hep-ph]. (Visited on 07/18/2025).
- [542] M. Thomson. *Modern particle physics*. eng. Oclc: 840462909. Cambridge, United Kingdom: Cambridge University Press, 2013. ISBN: 978-1-107-03426-6. URL: <http://digitale-objekte.hbz-nrw.de/storage/2013/11/04/file%5C%5F5/5371837.pdf> (visited on 07/18/2025).
- [543] G. Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Phys. Lett.*

- B* 716 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].
- [544] R. Abbate, M. Fickinger, A. H. Hoang, V. Mateu, and I. W. Stewart. “Precision Thrust Cumulant Moments at $N^3\text{LL}$ ”. In: *Phys. Rev. D* 86 (2012), p. 094002. DOI: 10.1103/PhysRevD.86.094002. arXiv: 1204.5746 [hep-ph].
- [545] A. Altheimer, S. Arora, L. Asquith, G. Brooijmans, J. Butterworth, M. Campanelli, B. Chapleau, A. E. Cholakian, J. P. Chou, M. Dasgupta, et al. “Jet substructure at the Tevatron and LHC: New results, new tools, new benchmarks”. In: *Journal of Physics G: Nuclear and Particle Physics* 39.6 (June 2012). ISSN: 09543899. DOI: 10.1088/0954-3899/39/6/063001.
- [546] I. Bionis and P. S. Koutsourelakis. “Free energy computations by minimization of Kullback–Leibler divergence: An efficient adaptive biasing potential method for sparse representations”. In: *Journal of Computational Physics* 231.9 (May 2012), pp. 3849–3870. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2012.01.033. URL: <https://www.sciencedirect.com/science/article/pii/S0021999112000630> (visited on 07/16/2025).
- [547] M. Cacciari, G. P. Salam, and G. Soyez. “FastJet User Manual”. In: *Eur. Phys. J. C* 72.3 (Mar. 2012), p. 1896. ISSN: 14346052. DOI: 10.1140/epjc/s10052-012-1896-2.
- [548] G. Choudalakis. *Fully Bayesian Unfolding*. arXiv:1201.4612 [physics]. May 2012. DOI: 10.48550/arXiv.1201.4612. URL: <http://arxiv.org/abs/1201.4612> (visited on 07/14/2025).

- [549] A. Collaboration. “Measurement of event background fluctuations for charged particle jet reconstruction in Pb-Pb collisions at $\sqrt{s_{\rm NN}} = 2.76$ TeV”. en. In: *Journal of High Energy Physics* 03.3 (2012), p. 053. ISSN: 1029-8479. DOI: 10.1007/JHEP03(2012)053. URL: <http://arxiv.org/abs/1201.2423> (visited on 07/21/2025).
- [550] T. C. Collaboration. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Physics Letters B* 716.1 (Sept. 2012). arXiv:1207.7235 [hep-ex], pp. 30–61. ISSN: 03702693. DOI: 10.1016/j.physletb.2012.08.021. URL: <http://arxiv.org/abs/1207.7235> (visited on 07/13/2025).
- [551] K. Cranmer, G. Lewis, L. Moneta, A. Shibata, and W. Verkerke. *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*. New York, 2012. DOI: 10.17181/cern-open-2012-016. URL: <https://cds.cern.ch/record/1456844> (visited on 07/13/2025).
- [552] G. E. Hinton, N. Srivastava, and K. Swersky. *Neural Networks for Machine Learning (2012 Coursera Lectures)*. Online lecture notes. Includes lecture slides with multiple sections (e.g. lectures 1a 16e). 2012. URL: https://www.cs.toronto.edu/~hinton/coursera_lectures.html (visited on 08/14/2025).
- [553] B. T. Knapik, A. W. van der Vaart, and J. H. van Zanten. “Bayesian inverse problems with Gaussian priors”. In: *Annals of Statistics* 39.5 (Feb. 2012), pp. 2626–2657. DOI: 10.1214/11-aos920. URL: <http://arxiv.org/abs/1103.2692>
<http://dx.doi.org/10.1214/11-AOS920>.

- [554] C. Kreutz, A. Raue, and J. Timmer. “Likelihood based observability analysis and confidence intervals for predictions of dynamic models”. In: *BMC Systems Biology* 6.1 (2012), p. 120. ISSN: 1752-0509. DOI: 10.1186/1752-0509-6-120. URL: <http://arxiv.org/abs/1107.0013> (visited on 07/22/2025).
- [555] LHC Higgs Cross Section Working Group, S. Dittmaier, C. Mariotti, G. Passarino, R. Tanaka, S. Alekhin, J. Alwall, E. A. Bagnaschi, A. Banfi, J. Blumlein, et al. “Handbook of LHC Higgs Cross Sections: 2. Differential Distributions”. In: (Jan. 2012). DOI: 10.5170/cern-2012-002. URL: <http://arxiv.org/abs/1201.3084><http://dx.doi.org/10.5170/CERN-2012-002>.
- [556] E. Marchand and W. E. Strawderman. *On Bayesian credible sets in restricted parameter space problems and lower bounds for frequentist coverage*. arXiv:1208.0028 [math]. Dec. 2012. DOI: 10.48550/arXiv.1208.0028. URL: <http://arxiv.org/abs/1208.0028> (visited on 07/14/2025).
- [557] S. Schmitt. “TUnfold: an algorithm for correcting migration effects in high energy physics”. In: *Jinst* 7 (2012). _eprint: 1205.6201, T10003. DOI: 10.1088/1748-0221/7/10/t10003.
- [558] J. Thaler and K. Van Tilburg. “Maximizing boosted top identification by minimizing N-subjettiness”. In: *Journal of High Energy Physics* 2012.2 (Feb. 2012), pp. 1–33. ISSN: 11266708. DOI: 10.1007/jhep02(2012)093/metrics. URL: [https://link.springer.com/article/10.1007/JHEP02\(2012\)093](https://link.springer.com/article/10.1007/JHEP02(2012)093).

- [559] W. J. Waalewijn. “Calculating the charge of a jet”. In: *Physical Review D - Particles, Fields, Gravitation and Cosmology* 86.9 (Nov. 2012). ISSN: 15507998. DOI: 10.1103/physrevd.86.094030.
- [560] G. Aad et al. “Charged-particle multiplicities in pp interactions measured with the ATLAS detector at the LHC”. In: *New J. Phys.* 13 (2011), p. 053033. DOI: 10.1088/1367-2630/13/5/053033. arXiv: 1012.5104 [hep-ex].
- [561] G. Aad, B. Abbott, J. Abdallah, A. A. Abdelalim, A. Abdesselam, O. Abdinov, B. Abi, M. Abolins, H. Abramowicz, H. Abreu, et al. “Measurement of inclusive jet and dijet cross sections in proton-proton collisions at 7 TeV centre-of-mass energy with the ATLAS detector”. en. In: *The European Physical Journal C* 71.2 (Feb. 2011), p. 1512. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-010-1512-2. URL: <https://doi.org/10.1140/epjc/s10052-010-1512-2> (visited on 07/13/2025).
- [562] T. Adye. “Unfolding algorithms and tests using RooUnfold”. In: *Phystat 2011*. _eprint: 1105.1160. Geneva: Cern, 2011, pp. 313–318. DOI: 10.5170/cern-2011-006.313.
- [563] A. Ali and G. Kramer. “JETS and QCD: a historical review of the discovery of the quark and gluon jets and its impact on QCD”. In: *The European Physical Journal H* 2011 36:2 36.2 (Sept. 2011), pp. 245–326. ISSN: 2102-6467. DOI: 10.1140/epjh/e2011-10047-1. URL: <https://link.springer.com/article/10.1140/epjh/e2011-10047-1>.

- [564] V. Blobel. “Unfolding Methods in Particle Physics”. In: *Phystat 2011*. Geneva: Cern, 2011, pp. 240–251. DOI: 10.5170/cern-2011-006.240.
- [565] V. Blobel. “Unfolding Methods in Particle Physics”. In: (2011), pp. 240–251. DOI: 10.5170/CERN-2011-006.240. URL: <https://cds.cern.ch/record/2203257> (visited on 07/22/2025).
- [566] CMS Collaboration, S. Chatrchyan, V. Khachatryan, A. M. Sirunyan, A. Tumasyan, W. Adam, T. Bergauer, M. Dragicevic, J. Erö, C. Fabjan, et al. “Measurement of the differential cross section for isolated prompt photon production in $p\bar{p}$ collisions at 7 TeV”. In: *Physical Review D* 84.5 (Sept. 2011). Publisher: American Physical Society, p. 052011. DOI: 10.1103/PhysRevD.84.052011. URL: <https://link.aps.org/doi/10.1103/PhysRevD.84.052011> (visited on 07/13/2025).
- [567] R. D. Cousins. *Negatively Biased Relevant Subsets Induced by the Most-Powerful One-Sided Upper Confidence Limits for a Bounded Physical Parameter*. arXiv:1109.2023 [physics]. Sept. 2011. DOI: 10.48550/arXiv.1109.2023. URL: <http://arxiv.org/abs/1109.2023> (visited on 07/22/2025).
- [568] G. Cowan. “Highlights from PHYSTAT 2011”. In: *Phystat 2011*. Geneva: Cern, 2011, pp. 215–224. DOI: 10.5170/cern-2011-006.215.
- [569] G. Cowan. *Statistical Data Analysis for High Energy Physics: Lecture Notes (Freiburg, 2011)*. Lecture notes for the Graduiertenkolleg "Physik an Hadron-Beschleunigern", Universität Freiburg. Available online. 2011. URL: https://www.pp.rhul.ac.uk/~cowan/stat_freiburg.html.

- [570] G. Cowan, K. Cranmer, E. Gross, and O. Vitells. “Asymptotic formulae for likelihood-based tests of new physics”. In: *Eur.Phys.J.C* 71.2 (Feb. 2011), p. 1554. ISSN: 14346052. DOI: 10.1140/epjc/s10052-011-1554-0.
- [571] G. Cowan, K. Cranmer, E. Gross, and O. Vitells. “Asymptotic formulae for likelihood-based tests of new physics”. In: *The European Physical Journal C* 71.2 (2011), p. 1554. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-011-1554-0. URL: <http://arxiv.org/abs/1007.1727> (visited on 07/22/2025).
- [572] J. Duchi, E. Hazan, and Y. Singer. “Adaptive Subgradient Methods for Online Learning and Stochastic Optimization”. In: *Journal of Machine Learning Research* 12.61 (2011), pp. 2121–2159. URL: <http://jmlr.org/papers/v12/duchi11a.html>.
- [573] R. K. Ellis, W. J. Stirling, and B. R. Webber. *QCD and Collider Physics*. Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol. 8. Cambridge: Cambridge University Press, Feb. 2011. ISBN: 978-0-521-54589-1. DOI: 10.1017/cbo9780511628788. (Visited on 07/18/2025).
- [574] C. J. Geyer. “Introduction to Markov Chain Monte Carlo”. In: *Handbook of Markov Chain Monte Carlo*. CRC Press, May 2011, pp. 3–48. ISBN: 978-1-4200-7941-8. DOI: 10.1201/b10905-2. URL: <http://www.scopus.com/inward/record.url?scp=84918818472&partnerID=8YFLogxK> (visited on 07/23/2025).

- [575] X. Glorot, A. Bordes, and Y. Bengio. “Deep Sparse Rectifier Neural Networks”. en. In: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. Issn: 1938-7228. JMLR Workshop and Conference Proceedings, June 2011, pp. 315–323. URL: <https://proceedings.mlr.press/v15/glorot11a.html> (visited on 07/21/2025).
- [576] A. Harel. “Statistical methods in CMS searches”. In: (2011), pp. 88–93. DOI: 10.5170/CERN-2011-006.88. URL: <https://cds.cern.ch/record/2203243> (visited on 07/22/2025).
- [577] V. A. Kostelecký and J. D. Tasson. “Matter-gravity couplings and Lorentz violation”. In: *Physical Review D* 83.1 (2011). arXiv:1006.4106 [gr-qc], p. 016013. DOI: 10.1103/PhysRevD.83.016013. URL: <https://link.aps.org/doi/10.1103/PhysRevD.83.016013> (visited on 07/23/2025).
- [578] L. Lyons. “Unfolding: Introduction”. In: *Phystat 2011*. Geneva: Cern, 2011, pp. 225–228. DOI: 10.5170/cern-2011-006.225.
- [579] “Measurement of the Charge Asymmetry in Top Quark Pair Production”. In: (2011).
- [580] G. P. Salam. “Elements of QCD for hadron colliders”. In: (Jan. 2011). arXiv:1011.5131 [hep-ph]. DOI: 10.48550/arXiv.1011.5131. URL: <http://arxiv.org/abs/1011.5131> (visited on 07/18/2025).
- [581] J. Thaler and K. Van Tilburg. “Identifying boosted objects with N-subjettiness”. In: *Journal of High Energy Physics* 2011.3 (Mar. 2011), pp. 1–28. ISSN:

10298479. DOI: 10.1007/jhep03(2011)015/metrics. URL: [https://link.springer.com/article/10.1007/JHEP03\(2011\)015](https://link.springer.com/article/10.1007/JHEP03(2011)015).
- [582] C. Varin, N. Reid, and D. Firth. “An Overview of Composite Likelihood Methods”. In: *Statistica Sinica* 21 (Jan. 2011).
- [583] Y. Yang, E. F. Remmers, C. B. Ogunwole, D. L. Kastner, P. K. Gregersen, and W. Li. “Effective Sample Size: Quick Estimation of the Effect of Related Samples in Genetic Case-Control Association Analyses”. In: *Computational Biology and Chemistry* 35.1 (May 2011), pp. 40–49. DOI: 10.1016/j.compbiolchem.2010.12.006. URL: <http://arxiv.org/abs/q-bio/0611093%20http://dx.doi.org/10.1016/j.compbiolchem.2010.12.006>.
- [584] G. Zech. “Regularization and error assignment to unfolded distributions”. In: *Phystat 2011*. Geneva: Cern, 2011, pp. 252–259. DOI: 10.5170/cern-2011-006.252.
- [585] J.-H. Zhong, R.-S. Huang, S.-C. Lee, R.-S. Huang, and S.-C. Lee. “A program for the Bayesian Neural Network in the ROOT framework”. In: *Comput. Phys. Commun.* 182 (2011), pp. 2655–2660. DOI: 10.1016/j.cpc.2011.07.019. arXiv: 1103.2854 [physics.data-an].
- [586] G. Cowan. *Topics in statistical data analysis for high-energy physics*. en. 2010. DOI: 10.5170/cern-2010-002.197. URL: <http://cds.cern.ch/record/1281954> (visited on 07/13/2025).

- [587] G. D’Agostini. *Improved iterative Bayesian unfolding*. arXiv:1010.0632 [physics]. Oct. 2010. DOI: 10.48550/arXiv.1010.0632. URL: <http://arxiv.org/abs/1010.0632> (visited on 07/13/2025).
- [588] S. D. Ellis, C. K. Vermilion, J. R. Walsh, A. Hornig, and C. Lee. “Jet shapes and jet algorithms in SCET”. In: *Journal of High Energy Physics* 2010.11 (Nov. 2010), pp. 1–83. ISSN: 11266708. DOI: 10.1007/jhep11(2010)101/metrics. URL: [https://link.springer.com/article/10.1007/JHEP11\(2010\)101](https://link.springer.com/article/10.1007/JHEP11(2010)101).
- [589] P. Giovanni and C. Zheng. “Multi-Dimensional Gaussian Fluctuations on the Poisson Space”. In: *Electronic Journal of Probability* 15.none (Jan. 2010). Publisher: Institute of Mathematical Statistics and Bernoulli Society, pp. 1487–1527. ISSN: 1083-6489, 1083-6489. DOI: 10.1214/EJP.v15-813. URL: <https://projecteuclid.org/journals/electronic-journal-of-probability/volume-15/issue-none/Multi-Dimensional-Gaussian-Fluctuations-on-the-Poisson-Space/10.1214/EJP.v15-813.full> (visited on 07/13/2025).
- [590] J. F. Grosse-Oetringhaus and K. Reygers. “Charged-Particle Multiplicity in Proton-Proton Collisions”. In: *J. Phys. G* 37 (2010), p. 083001. DOI: 10.1088/0954-3899/37/8/083001. arXiv: 0912.0023 [hep-ex].
- [591] D. W. Hogg, J. Bovy, and D. Lang. “Data analysis recipes: Fitting a model to data”. In: (Aug. 2010). DOI: 10.48550/arXiv.1008.4686. URL: <http://arxiv.org/abs/1008.4686> (visited on 07/22/2025).
- [592] *Statistics and Discoveries at the LHC (1/4)*. en. June 2010. URL: <https://indico.cern.ch/event/77830/> (visited on 07/17/2025).

- [593] L. Verde. “Statistical methods in cosmology”. In: *Lect. Notes Phys.* Vol. 800. 2010, pp. 147–177. DOI: 10.1007/978-3-642-10598-2_4. URL: https://doi.org/10.1007/978-3-642-10598-2_4 (visited on 07/22/2025).
- [594] B. J. Brewer and M. J. Francis. “Entropic Priors and Bayesian Model Selection”. In: *AIP Conference Proceedings* 1193.1 (Dec. 2009), pp. 179–186. ISSN: 0094-243x. DOI: 10.1063/1.3275612. URL: <https://doi.org/10.1063/1.3275612> (visited on 07/14/2025).
- [595] G. Cowan. *Statistical Methods for Particle Physics, Lecture 4: Unfolding*. Statistical Inference for Astro and Particle Physics Workshop, Weizmann Institute, Rehovot. 2009. URL: <https://indico.weizmann.ac.il/conferenceDisplay.py?confId=52>.
- [596] G. Cowan. “Topics in statistical data analysis for HEP”. In: *65th Scottish Universities Summer School in Physics: LHC Physics*. Aug. 2009, pp. 341–369. DOI: 10.1201/b11865-15.
- [597] D. Froidevaux and V. A. Mitsou. “Experimental prospects at the Large Hadron Collider”. In: *Journal of Physics: Conference Series* 171 (May 2009). DOI: 10.1088/1742-6596/171/1/012021. URL: <http://arxiv.org/abs/0905.0258>
<http://dx.doi.org/10.1088/1742-6596/171/1/012021>.
- [598] T. Hastie, S. Rosset, J. Zhu, and H. Zou. *Multi-class AdaBoost*. Vol. 2. 3. International Press of Boston, 2009, pp. 349–360. DOI: 10.4310/SII.2009.V2.N3.A8.

- [599] T. Hastie, R. Tibshirani, and J. Friedman. “Boosting and Additive Trees”. en. In: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Ed. by T. Hastie, R. Tibshirani, and J. Friedman. New York, NY: Springer, 2009, pp. 337–387. ISBN: 978-0-387-84858-7. DOI: 10.1007/978-0-387-84858-7_10. URL: https://doi.org/10.1007/978-0-387-84858-7_10 (visited on 08/11/2025).
- [600] C. Pahl, S. Bethke, S. Kluth, and J. Schieck. “Study of Moments of Event Shapes and a Determination of α_s Using e^+e^- Annihilation Data from $\sqrt{s}=10.58$ GeV”. In: *The European Physical Journal C* 60.2 (Mar. 2009), p. 181. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-009-0930-5. arXiv: 0810.2933 [hep-ex]. URL: <http://arxiv.org/abs/0810.2933> (visited on 07/24/2025).
- [601] A. G.-D. Ridder, T. Gehrmann, E. W. N. Glover, and G. Heinrich. “NNLO moments of event shapes in e^+e^- annihilation”. In: *Journal of High Energy Physics* 2009.5 (May 2009). DOI: 10.1088/1126-6708/2009/05/106. URL: <http://arxiv.org/abs/0903.4658><http://dx.doi.org/10.1088/1126-6708/2009/05/106>.
- [602] C. Villani. “The Wasserstein distances”. en. In: *Optimal Transport: Old and New*. Ed. by C. Villani. Berlin, Heidelberg: Springer, 2009, pp. 93–111. ISBN: 978-3-540-71050-9. DOI: 10.1007/978-3-540-71050-9_6. URL: https://doi.org/10.1007/978-3-540-71050-9_6 (visited on 07/17/2025).
- [603] Q. Zhou. “Information Matrices in Estimating Function Approach: Tests for Model Misspecification and Model Selection”. Master’s thesis. University of

- Waterloo, 2009. URL: <http://hdl.handle.net/10012/4614> (visited on 08/14/2025).
- [604] G. Altarelli. *Elementary Particles*. 1st ed. Vol. 4. 1. Berlin ; Heidelberg: Springer, 2008. ISBN: 3-540-74202-6.
- [605] C. Amsler, M. Doser, M. Antonelli, D. M. Asner, K. S. Babu, H. Baer, H. R. Band, R. M. Barnett, E. Bergren, J. Beringer, et al. “Monte Carlo techniques”. In: *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* 667.1-5 (Sept. 2008), pp. 1–6. ISSN: 03702693. DOI: 10.1016/j.physletb.2008.07.018.
- [606] M. Bahr et al. “Herwig++ Physics and Manual”. In: *Eur. Phys. J. C* 58 (2008), pp. 639–707. DOI: 10.1140/epjc/s10052-008-0798-9. arXiv: 0803.0883 [hep-ph].
- [607] M. Cacciari, G. P. Salam, and G. Soyez. “The anti- k_t jet clustering algorithm”. In: *Jhep* 04 (2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph].
- [608] M. Cacciari, G. P. Salam, and G. Soyez. “The Anti- $k(t)$ jet clustering algorithm”. In: *Jhep* 04.4 (Apr. 2008), p. 063. ISSN: 11266708. DOI: 10.1088/1126-6708/2008/04/063.
- [609] A. C. Cameron, J. B. Gelbach, and D. L. Miller. “Bootstrap-Based Improvements for Inference with Clustered Errors”. In: *The Review of Economics and Statistics* 90.3 (2008). Publisher: The MIT Press, pp. 414–427. ISSN: 0034-6535.

- DOI: 10.1162/rest.90.3.414. URL: <https://www.jstor.org/stable/40043157> (visited on 07/22/2025).
- [610] A. Carpio, O. Dorn, M. Moscoso, F. Natterer, G. C. Papanicolaou, M. L. Rapún, and A. Teta. *Inverse Problems and Imaging*. Ed. by L. L. Bonilla, J. .-. Morel, F. Takens, and B. Teissier. Vol. 1943. Lecture Notes in Mathematics. Berlin, Heidelberg: Springer, 2008. ISBN: 978-3-540-78547-7. DOI: 10.1007/978-3-540-78547-7. URL: <http://link.springer.com/10.1007/978-3-540-78547-7> (visited on 07/13/2025).
- [611] J. Schmalian and C. D. Batista. “Emergent symmetry and dimensional reduction at a quantum critical point”. In: *Physical Review B* 77.9 (Mar. 2008). Publisher: American Physical Society, p. 094406. DOI: 10.1103/PhysRevB.77.094406. URL: <https://link.aps.org/doi/10.1103/PhysRevB.77.094406> (visited on 07/24/2025).
- [612] T. Sjostrand, S. Mrenna, and P. Z. Skands. “A Brief Introduction to PYTHIA 8.1”. In: *Comput. Phys. Commun.* 178 (2008), pp. 852–867. DOI: 10.1016/j.cpc.2008.01.036. arXiv: 0710.3820 [hep-ph].
- [613] L. Teodorescu. “Artificial neural networks in high-energy physics”. In: (2008). DOI: 10.5170/cern-2008-002.13. URL: <https://cds.cern.ch/record/1100521>.
- [614] W. Verkerke and D. Kirkby. “RooFit Users Manual version 2.91”. In: 14 (2008).

- [615] G. Cowan. “Bayesian Statistical Methods In Particle Physics”. In: *42nd Rencontres de Moriond on QCD and High Energy Hadronic Interactions*. Hanoi, Vietnam: Gioi Publ., 2007, pp. 7–10.
- [616] C. A. Field and A. H. Welsh. “Bootstrapping Clustered Data”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 69.3 (June 2007), pp. 369–390. ISSN: 1369-7412. DOI: 10.1111/j.1467-9868.2007.00593.x. URL: <https://doi.org/10.1111/j.1467-9868.2007.00593.x> (visited on 07/22/2025).
- [617] J. D. Hunter. “Matplotlib: A 2D graphics environment”. In: *Computing in Science and Engineering* 9.3 (2007), pp. 90–95. ISSN: 15219615. DOI: 10.1109/mcse.2007.55.
- [618] D. Kotlorz and A. Kotlorz. “Evolution Equations for Truncated Moments of the Parton Distributions”. In: *Physics Letters B* 644.4 (2007), pp. 284–287. ISSN: 03702693. DOI: 10.1016/j.physletb.2006.11.054. arXiv: hep-ph/0610282. (Visited on 07/18/2025).
- [619] H. Voss, A. Höcker, J. Stelzer, and F. Tegenfeldt. “TMVA, the Toolkit for Multivariate Data Analysis with ROOT”. en. In: *PoS*. Vol. ACAT. Nikhef, Amsterdam, the Netherlands: Sissa Medialab, 2007, p. 040. DOI: 10.22323/1.050.0040. URL: <https://doi.org/10.22323/1.050.0040> (visited on 08/11/2025).
- [620] *All of Nonparametric Statistics*. Springer New York, 2006. DOI: 10.1007/0-387-30623-4.

- [621] J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce Dubois, M. Asai, G. Barrand, R. Capra, S. Chauvie, R. Chytrcek, G. Cirrone, G. Cooperman, G. Cosmo, G. Cuttone, G. Daquino, M. Donszelmann, M. Dressel, G. Folger, F. Foppiano, J. Generowicz, V. Grichine, S. Guatelli, P. Gumplinger, A. Heikkinen, I. Hrivnacova, A. Howard, S. Incerti, V. Ivanchenko, T. Johnson, F. Jones, T. Koi, R. Kokoulin, M. Kossov, H. Kurashige, V. Lara, S. Larsson, F. Lei, O. Link, F. Longo, M. Maire, A. Mantero, B. Mascialino, I. McLaren, P. Mendez Lorenzo, K. Minamimoto, K. Murakami, P. Nieminen, L. Pandola, S. Parlati, L. Peralta, J. Perl, A. Pfeiffer, M. Pia, A. Ribon, P. Rodrigues, G. Russo, S. Sadilov, G. Santin, T. Sasaki, D. Smith, N. Starkov, S. Tanaka, E. Tcherniaev, B. Tome, A. Trindade, P. Truscott, L. Urban, M. Verderi, A. Walkden, J. Wellisch, D. Williams, D. Wright, and H. Yoshida. “Geant4 developments and applications”. en. In: *IEEE Transactions on Nuclear Science* 53.1 (2006), p. 270. ISSN: 0018-9499. DOI: 10.1109/TNS.2006.869826. URL: <http://ieeexplore.ieee.org/document/1610988/> (visited on 08/14/2025).
- [622] G. L. Bayatian et al. “CMS Physics: Technical Design Report Volume 1: Detector Performance and Software”. In: (2006). DOI: 10.2172/2510878.
- [623] G. Cowan. “Bayesian statistical methods for parton analyses”. In: *14th International Workshop on Deep Inelastic Scattering*. Apr. 2006, pp. 157–160.
- [624] M. Göckeler, R. Horsley, H. Perlt, P. E. L. Rakow, G. Schierholz, and A. Schiller. “Operator Product Expansion on the Lattice: Analytic Wilson Coefficients”. In: *Lat2006* (2006), p. 119. DOI: 10.22323/1.032.0119. arXiv: [hep-lat/0610064](https://arxiv.org/abs/hep-lat/0610064). (Visited on 07/18/2025).

- [625] P. C. Hansen, J. G. Nagy, and D. P. O’Leary. “6. Regularization by Spectral Filtering”. In: *Deblurring Images. Fundamentals of Algorithms*. Society for Industrial and Applied Mathematics, Jan. 2006, pp. 71–86. ISBN: 978-0-89871-618-4. DOI: 10.1137/1.9780898718874.ch6. URL: <https://epubs.siam.org/doi/10.1137/1.9780898718874.ch6> (visited on 07/17/2025).
- [626] C. Lee and G. Sterman. “Momentum Flow Correlations from Event Shapes: Factorized Soft Gluons and Soft-Collinear Effective Theory”. In: *Physical Review D - Particles, Fields, Gravitation and Cosmology* 75.1 (Nov. 2006). DOI: 10.1103/PhysRevD.75.014022. URL: <http://arxiv.org/abs/hep-ph/0611061><http://dx.doi.org/10.1103/PhysRevD.75.014022>.
- [627] S. Luzzatto. “Chapter 3 - Stochastic-Like Behaviour in Nonuniformly Expanding Maps”. In: *Handbook of Dynamical Systems*. Ed. by B. Hasselblatt and A. Katok. Vol. 1. Handbook of Dynamical Systems. Elsevier Science, Jan. 2006, pp. 265–326. DOI: 10.1016/S1874-575X(06)80028-7. URL: <https://www.sciencedirect.com/science/article/pii/S1874575X06800287> (visited on 07/24/2025).
- [628] “Precision electroweak measurements on the Z resonance”. In: *Physics Reports* 427.5 (May 2006), pp. 257–454. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2005.12.006. URL: <https://www.sciencedirect.com/science/article/pii/S0370157305005119> (visited on 07/13/2025).
- [629] C. Van Den Broeck. “Binary black hole detection rates in inspiral gravitational wave searches”. In: *Class.Quant.Grav.* 23.13 (July 2006), pp. L51–L58. ISSN: 0264-9381. DOI: 10.1088/0264-9381/23/13/L01.

- [630] T. Zhang. “A Modification for Bayesian Credible Intervals”. In: *Communications in Statistics - Theory and Methods* 35.9 (Sept. 2006), pp. 1703–1711. ISSN: 03610926. DOI: 10.1080/03610920600683838. URL: <https://www.tandfonline.com/doi/abs/10.1080/03610920600683838>.
- [631] G. Abbiendi et al. “Measurement of event shape distributions and moments in $e^+e^- \rightarrow \text{hadrons}$ at 91-GeV - 209-GeV and a determination of $\alpha(s)$ ”. In: *Eur. Phys. J. C* 40 (2005), pp. 287–316. DOI: 10.1140/epjc/s2005-02120-6. arXiv: hep-ex/0503051.
- [632] D. A. Belsley, E. Kuh, and R. E. Welsch. *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity*. en. Google-Books-ID: GECBEU-JVNe0C. John Wiley & Sons, Feb. 2005. ISBN: 978-0-471-72514-5.
- [633] J. Charles, A. Hocker, H. Lacker, S. Laplace, F. R. L. Diberder, J. Malcles, J. Ocariz, M. Pivk, and L. Roos. “CP Violation and the CKM Matrix: Assessing the Impact of the Asymmetric B Factories”. In: *European Physical Journal C* 41.1 (Mar. 2005), pp. 1–131. DOI: 10.1140/epjc/s2005-02169-1. URL: <http://arxiv.org/abs/hep-ph/0406184><http://dx.doi.org/10.1140/epjc/s2005-02169-1>.
- [634] “Inverse Problems and Interpretation of Measurements”. In: *Applied Mathematical Sciences (Switzerland)* 160 (2005), pp. 1–5. ISSN: 2196968x. DOI: 10.1007/0-387-27132-5{_}1.

- [635] J. P. Kaipio and E. Somersalo. “Statistical and computational inverse problems”. In: *Applied Mathematical Sciences (Switzerland)* 160 (2005), pp. i–339. ISSN: 2196968x. DOI: 10.1007/b138659/cover.
- [636] W. C. Karl. “Regularization in Image Restoration and Reconstruction”. In: *Handbook of Image and Video Processing, Second Edition* (Jan. 2005), pp. 183–202. DOI: 10.1016/b978-012119792-6/50075-9.
- [637] “Measurement of event shape distributions and moments in $e^+e^- \rightarrow$ hadrons at 91-209 GeV and a determination of α_s ”. In: *The European Physical Journal C - Particles and Fields 2005 40:3* 40.3 (Apr. 2005), pp. 287–316. ISSN: 1434-6052. DOI: 10.1140/epjc/s2005-02120-6. URL: <https://link.springer.com/article/10.1140/epjc/s2005-02120-6>.
- [638] “Statistical Inversion Theory”. In: *Applied Mathematical Sciences (Switzerland)* 160 (2005), pp. 49–114. ISSN: 2196968x. DOI: 10.1007/0-387-27132-5{_}3.
- [639] S. Weinberg. *The Quantum Theory of Fields: Volume 1: Foundations*. Vol. 1. Cambridge: Cambridge University Press, June 2005. ISBN: 978-0-521-67053-1. DOI: 10.1017/CB09781139644167. URL: <https://doi.org/10.1017/CB09781139644167> (visited on 07/23/2025).
- [640] G. Abbiendi, T. Schorner-Sadenius, C. Geich-Gimbel, Y. Rozen, M. Boutemur, H. Burckhart, D. Axen, E. Wolf, P. Mattig, K. Sachs, et al. “Measurement of event shape distributions and moments in $e^+e^- \rightarrow$ hadrons at 91-209 GeV and a determination of α_s ”. In: *Eur. Phys. J. C* 40.3 (Aug. 2004),

- pp. 287–316. ISSN: 14346044. DOI: 10.1140/epjc/s2005-02120-6. URL: <https://cds.cern.ch/record/802086>.
- [641] J. Abdallah et al. “The Measurement of $\alpha(s)$ from event shapes with the DELPHI detector at the highest LEP energies”. In: *Eur. Phys. J. C* 37 (2004), pp. 1–23. DOI: 10.1140/epjc/s2004-01889-x. arXiv: hep-ex/0406011.
- [642] P. Achard et al. “Studies of hadronic event structure in e^+e^- annihilation from 30-GeV to 209-GeV with the L3 detector”. In: *Phys. Rept.* 399 (2004), pp. 71–174. DOI: 10.1016/j.physrep.2004.07.002. arXiv: hep-ex/0406049.
- [643] D. Anguita, S. Ridella, F. Riveccio, R. Zunino, S. Amerio, and I. Lazzizzera. “Model selection in top quark tagging with a support vector classifier”. In: *IEEE Conf.Proc.* 3.3 (2004), pp. 2059–2064. ISSN: 10987576. DOI: 10.1109/ijcnn.2004.1380934.
- [644] A. Caticha and R. Preuss. “Entropic Priors”. In: *AIP Conference Proceedings* 707.1 (Apr. 2004), pp. 371–380. ISSN: 0094-243x. DOI: 10.1063/1.1751380. URL: <https://doi.org/10.1063/1.1751380> (visited on 07/14/2025).
- [645] M. Dasgupta and G. P. Salam. “Event shapes in e^+e^- annihilation and deep inelastic scattering”. In: *Journal of Physics G: Nuclear and Particle Physics* 30.5 (May 2004). ISSN: 09543899. DOI: 10.1088/0954-3899/30/5/r01.
- [646] M. Feindt. “A Neural Bayesian Estimator for Conditional Probability Densities”. In: (Feb. 2004). arXiv: physics/0402093.
- [647] D. A. S. Fraser, N. Reid, and A. C. M. Wong. “Inference for bounded parameters”. In: *Physical Review D* 69.3 (2004), p. 033002. DOI: 10.1103/PhysRevD.

- 69.033002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.69.033002> (visited on 07/22/2025).
- [648] B. Fuglede and F. Topsoe. “Jensen-Shannon divergence and Hilbert space embedding”. In: *International Symposium on Information Theory, 2004. ISIT 2004. Proceedings*. June 2004, pp. 31–. DOI: 10.1109/isit.2004.1365067. URL: <https://ieeexplore.ieee.org/document/1365067> (visited on 07/17/2025).
- [649] A. Heister, S. Schael, R. Barate, R. Brunelière, I. De Bonis, D. Decamp, C. Goy, S. Jézéquel, J. P. Lees, F. Martin, et al. “Studies of QCD at e+e- centre-of-mass energies between 91 and 209 GeV”. In: *European Physical Journal C* 35.4 (2004), pp. 457–486. ISSN: 14346044. DOI: 10.1140/epjc/s2004-01891-4.
- [650] F. James, R. Cousins, and G. Cowan. “Statistics”. In: *Phys. Lett. B* 592 (2004), pp. 279–288.
- [651] S. Moch, J. A. M. Vermaseren, and A. Vogt. “The Three-Loop Splitting Functions in QCD: The Non-Singlet Case”. In: *Nuclear Physics B* 688.1-2 (2004), pp. 101–134. ISSN: 05503213. DOI: 10.1016/j.nuclphysb.2004.03.030. arXiv: [hep-ph/0403192](https://arxiv.org/abs/hep-ph/0403192). (Visited on 07/18/2025).
- [652] L. Rosasco, E. D. Vito, A. Caponnetto, M. Piana, and A. Verri. “Are Loss Functions All the Same?” In: *Neural Computation* 16.5 (May 1, 2004), pp. 1063–1076. ISSN: 0899-7667. DOI: 10.1162/089976604773135104. URL: <https://doi.org/10.1162/089976604773135104> (visited on 08/11/2025).

- [653] The DELPHI Collaboration and J. Abdallah. “The measurement of α_s from event shapes with the DELPHI detector at the highest LEP energies”. In: *European Physical Journal C* 37.1 (June 2004), pp. 1–23. DOI: 10.1140/epjc/s2004-01889-x. URL: <http://arxiv.org/abs/hep-ex/0406011><http://dx.doi.org/10.1140/epjc/s2004-01889-x>.
- [654] J. Abdallah et al. “A Study of the energy evolution of event shape distributions and their means with the DELPHI detector at LEP”. In: *Eur. Phys. J. C* 29 (2003), pp. 285–312. DOI: 10.1140/epjc/s2003-01198-0. arXiv: [hep-ex/0307048](http://arxiv.org/abs/hep-ex/0307048).
- [655] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(03)01368-8. URL: <https://www.sciencedirect.com/science/article/pii/S0168900203013688> (visited on 08/14/2025).
- [656] G. D. Agostini. *Bayesian Reasoning In Data Analysis: A Critical Introduction*. English. River Edge, NJ: World Scientific Publishing Company, 2003. ISBN: 978-981-238-356-3.
- [657] D. Asner. “Charm Dalitz plot analysis formalism and results: Expanded RPP-2004 version”. In: (Nov. 2003). URL: <https://arxiv.org/abs/hep-ex/0410014>.
- [658] J. Conrad, O. Botner, A. Hallgren, and C. Pérez de los Heros. “Including systematic uncertainties in confidence interval construction for Poisson statistics”.

- In: *Physical Review D* 67.1 (Jan. 2003). Publisher: American Physical Society, p. 012002. DOI: 10.1103/PhysRevD.67.012002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.67.012002> (visited on 07/13/2025).
- [659] G. Dissertori, I. G. Knowles, and a. M. Schmelling. *Quantum Chromodynamics: High Energy Experiments and Theory*. International Series of Monographs on Physics. Oxford, New York: Oxford University Press, Apr. 2003. ISBN: 978-0-19-850572-3.
- [660] L. E. Eberly and G. Casella. “Estimating Bayesian credible intervals”. In: *Journal of Statistical Planning and Inference*. Special issue II: Model Selection, Model Diagnostics, Empirical Bayes and Hierarchical Bayes 112.1 (Mar. 2003), pp. 115–132. ISSN: 0378-3758. DOI: 10.1016/S0378-3758(02)00327-0. URL: <https://www.sciencedirect.com/science/article/pii/S0378375802003270> (visited on 07/14/2025).
- [661] D. Endres and J. Schindelin. “A new metric for probability distributions”. In: *IEEE Transactions on Information Theory* 49.7 (July 2003), pp. 1858–1860. ISSN: 1557-9654. DOI: 10.1109/tit.2003.813506. URL: <https://ieeexplore.ieee.org/document/1207388> (visited on 07/17/2025).
- [662] P. Achard, O. Adriani, M. Aguilar-Benitez, J. Alcaraz, G. Alemanni, J. Allaby, A. Aloisio, M. G. Alviggi, H. Anderhub, V. P. Andreev, et al. “Determination of α_s from hadronic event shapes in e^+e^- annihilation at $192 \leq \sqrt{s} \leq 208$ GeV”. In: *Physics Letters B* 536.3-4 (June 2002), pp. 217–228. ISSN: 0370-2693. DOI: 10.1016/S0370-2693(02)01814-2. URL: <https://www.sciencedirect.com/science/article/abs/pii/S0370269302018142?via%3Dihub>.

- [663] P. Achard, O. Adriani, M. Aguilar-Benitez, J. Alcaraz, G. Alemanni, J. Allaby, A. Aloisio, M. G. Alviggi, H. Anderhub, V. P. Andreev, et al. “Determination of α_s from hadronic event shapes in e^+e^- annihilation at $192 \leq \sqrt{s} \leq 208$ GeV”. In: *Phys. Lett. B* 536.3-4 (June 2002), p. 217. ISSN: 03702693. DOI: 10.1016/S0370-2693(02)01814-2.
- [664] V. Blobel. *An Unfolding Method for High Energy Physics Experiments*. arXiv:hep-ex/0208022. Aug. 2002. DOI: 10.48550/arXiv.hep-ex/0208022. URL: <http://arxiv.org/abs/hep-ex/0208022> (visited on 07/17/2025).
- [665] V. Blobel. “An Unfolding method for high-energy physics experiments”. In: *Conference on Advanced Statistical Techniques in Particle Physics*. Aug. 2002, pp. 258–267. arXiv: hep-ex/0208022.
- [666] G. Cowan. “A survey of unfolding methods for particle physics”. In: *Conf. Proc. C* 0203181 (2002). Ed. by M. R. Whalley and L. Lyons, pp. 248–257.
- [667] M. Mandelkern. “Setting confidence intervals for bounded parameters”. In: *Statistical Science* 17.2 (May 2002), pp. 149–172. ISSN: 0883-4237, 2168-8745. DOI: 10.1214/ss/1030550859. URL: <https://projecteuclid.org/journals/statistical-science/volume-17/issue-2/Setting-confidence-intervals-for-bounded-parameters/10.1214/ss/1030550859.full> (visited on 07/22/2025).
- [668] A. L. Read. “Presentation of search results: the CLs technique”. en. In: *Journal of Physics G: Nuclear and Particle Physics* 28.10 (2002), pp. 2693–2704. ISSN:

- 0954-3899. DOI: 10.1088/0954-3899/28/10/313. URL: <https://dx.doi.org/10.1088/0954-3899/28/10/313> (visited on 07/22/2025).
- [669] C. C. Rodriguez. “Entropic priors for discrete probabilistic networks and for mixtures of Gaussians models”. In: *AIP Conference Proceedings* 617.1 (May 2002), pp. 410–432. ISSN: 0094-243x. DOI: 10.1063/1.1477063. URL: <https://doi.org/10.1063/1.1477063> (visited on 07/14/2025).
- [670] J. Back, C.-H. Cheng, U. Egede, D. Kirkby, G. Raven, C. Roat, A. Sarti, M. Serra, A. Soffer, J. Stark, M. Turri, W. Verkerke, and C. Voena. “A User’s Guide to the RooFitTools Package for Unbinned Maximum Likelihood Fitting”. In: (2001).
- [671] A. Böhm, A. Ewers, H. Fesefeldt, D. Käfer, W. Krenz, W. Lübelmeyer, M. vonÉderÉMey, J. Mnich, T. Niessen, D. Pandoulas, et al. “Measurement of the tau branching fractions into leptons”. In: *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* 507.1-4 (May 2001), pp. 47–60. ISSN: 03702693. DOI: 10.1016/s0370-2693(01)00294-5.
- [672] G. P. Salam and D. Wicke. “Hadron masses and power corrections to event shapes”. In: *Journal of High Energy Physics* 5.5 (June 2001). DOI: 10.1088/1126-6708/2001/05/061. URL: <http://arxiv.org/abs/hep-ph/0102343> 20[http://dx.doi.org/10.1088/1126-6708/2001/05/061](https://dx.doi.org/10.1088/1126-6708/2001/05/061).
- [673] R. Solomon. “A brief history of the classification of the finite simple groups”. en. In: *Bulletin of the American Mathematical Society* 38.3 (Mar. 2001), pp. 315–352. ISSN: 0273-0979, 1088-9485. DOI: 10.1090/s0273-0979-01-00909-0.

- URL: <https://www.ams.org/bull/2001-38-03/S0273-0979-01-00909-0/>
(visited on 07/23/2025).
- [674] M. Acciarri, P. Achard, O. Adriani, M. Aguilar-Benitez, J. Alcaraz, G. Alemani, J. Allaby, A. Aloisio, M. G. Alvigi, G. Ambrosi, et al. “QCD studies in e^+e^- annihilation from 30 GeV to 189 GeV”. In: *Phys. Lett. B* 489.1-2 (Sept. 2000), p. 65. ISSN: 03702693. DOI: 10.1016/S0370-2693(00)00891-1.
- [675] C. Dugas, Y. Bengio, F. Bélisle, C. Nadeau, and R. Garcia. “Incorporating Second-Order Functional Knowledge for Better Option Pricing”. In: *Advances in Neural Information Processing Systems*. Vol. 13. MIT Press, 2000. URL: <https://papers.nips.cc/paper%5C%5Ffiles/paper/2000/hash/44968aece94f667e4095002d140b5896-Abstract.html> (visited on 07/21/2025).
- [676] D. H. Fremlin. *Measure Theory*. en. Google-Books-ID: 2_lHYXeEd7YC. Torres Fremlin, 2000. ISBN: 978-0-9538129-0-5.
- [677] C. D. Manning and H. Schütze. *Foundations of Statistical Natural Language Processing*. English. Cambridge, Mass: The MIT Press, 1999. ISBN: 978-0-262-13360-9.
- [678] L. Schoeffel. “An Elegant and Fast Method to Solve QCD Evolution Equations. Application to the Determination of the Gluon Content of the Pomeron”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 423.2 (1999), pp. 439–445. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(98)01316-3. (Visited on 07/18/2025).

- [679] G. Cowan. *Statistical data analysis*. 1998. ISBN: 978-0-19-850156-5.
- [680] A. Kosteletsky. “The Status of CPT”. In: (Oct. 1998). DOI: 10.48550/arXiv.hep-ph/9810365. URL: <http://arxiv.org/abs/hep-ph/9810365> (visited on 07/23/2025).
- [681] A. Neumaier. “Solving Ill-Conditioned and Singular Linear Systems: A Tutorial on Regularization”. In: *SIAM Review* 40.3 (Jan. 1998). Publisher: Society for Industrial and Applied Mathematics, pp. 636–666. ISSN: 0036-1445. DOI: 10.1137/S0036144597321909. URL: <https://epubs.siam.org/doi/10.1137/S0036144597321909> (visited on 07/13/2025).
- [682] Y. Rubner, C. Tomasi, and L. Guibas. “A metric for distributions with applications to image databases”. In: *Sixth International Conference on Computer Vision*. Jan. 1998, pp. 59–66. DOI: 10.1109/iccv.1998.710701. URL: <https://ieeexplore.ieee.org/document/710701> (visited on 07/14/2025).
- [683] M. Acciarri, O. Adriani, M. Aguilar-Benitez, S. Ahlen, J. Alcaraz, G. Alemanni, J. Allaby, A. Aloisio, G. Alverson, M. G. Alviggi, et al. “QCD studies and determination of α_s in e+e- collisions at $\sqrt{s} = 161$ GeV and 172 GeV L3 Collaboration”. In: *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* 404.3-4 (July 1997), pp. 390–402. ISSN: 03702693. DOI: 10.1016/S0370-2693(97)00647-3.
- [684] M. Acciarri, O. Adriani, M. Aguilar-Benitez, S. Ahlen, J. Alcaraz, G. Alemanni, J. Allaby, A. Aloisio, G. Alverson, M. G. Alviggi, et al. “Study of hadronic events and measurements of α_s between 30 and 91 GeV”. In: *Physics Letters*,

- Section B: Nuclear, Elementary Particle and High-Energy Physics* 411.3-4 (Oct. 1997), pp. 339–353. ISSN: 03702693. DOI: 10.1016/s0370-2693(97)01000-9.
- [685] K. Ackerstaff, G. Alexander, J. Allison, N. Altekamp, K. Ametewee, K. J. Anderson, S. Anderson, S. Arcelli, S. Asai, D. Axen, et al. “QCD studies with e+e- annihilation data at 161 GeV”. In: *Z. Phys. C* 75.2 (June 1997), p. 193. ISSN: 01709739. DOI: 10.1007/s002880050462.
- [686] M. S. Bartlett and R. H. Fowler. “Properties of sufficiency and statistical tests”. In: *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences* 160.901 (Jan. 1997), pp. 268–282. DOI: 10.1098/rspa.1937.0109. URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.1937.0109> (visited on 07/22/2025).
- [687] D. Buskulic, I. De Bonis, D. Decamp, P. Ghez, C. Goy, J. P. Lees, A. Lucotte, M. N. Minard, P. Odier, B. Pietrzyk, et al. “Studies of QCD in e+e- → hadrons at $E_{\text{cm}} = 130$ and 136 GeV”. In: *Z. Phys. C* 73.3 (1997), p. 409. ISSN: 01709739. DOI: 10.1007/s002880050330.
- [688] D. Wolpert and W. Macready. “No free lunch theorems for optimization”. In: *IEEE Transactions on Evolutionary Computation* 1.1 (1997), pp. 67–82. DOI: 10.1109/4235.585893.
- [689] M. Acciarri, A. Adam, O. Adriani, M. Aguilar-Benitez, S. Ahlen, B. Alpat, J. Alcaraz, J. Allaby, A. Aloisio, G. Alverson, et al. “Study of the structure of hadronic events and determination of α_s at $\sqrt{s} = 130$ GeV and 136 GeV”. In: *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy*

- Physics* 371.1-2 (Mar. 1996), pp. 137–148. ISSN: 03702693. DOI: 10.1016/0370-2693(96)00096-2.
- [690] G. Alexander, J. Allison, N. Altekamp, K. Ametewee, K. J. Anderson, S. Anderson, S. Arcelli, S. Asai, D. Axen, G. Azuelos, et al. “QCD studies with $e+e-$ annihilation data at 130 and 136 GeV”. In: *Z. Phys. C* 72.2 (1996), p. 191. ISSN: 01709739. DOI: 10.1007/s002880050237.
- [691] Y. L. Dokshitzer, G. Marchesini, and B. R. Webber. “Dispersive approach to power-behaved contributions in QCD hard processes”. In: *Nuclear Physics B* 469.1-2 (June 1996), pp. 93–142. ISSN: 05503213. DOI: 10.1016/0550-3213(96)00155-1. URL: <http://arxiv.org/abs/hep-ph/9512336>[http://dx.doi.org/10.1016/0550-3213\(96\)00155-1](http://dx.doi.org/10.1016/0550-3213(96)00155-1).
- [692] D. J. Gross. “The role of symmetry in fundamental physics”. en. In: *Proceedings of the National Academy of Sciences* 93.25 (Dec. 1996), p. 14256. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.93.25.14256. URL: <https://pnas.org/doi/full/10.1073/pnas.93.25.14256> (visited on 07/24/2025).
- [693] A. Höcker and V. Kartvelishvili. “SVD approach to data unfolding”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 372.3 (Apr. 1996), pp. 469–481. ISSN: 0168-9002. DOI: 10.1016/0168-9002(95)01478-0. URL: <https://www.sciencedirect.com/science/article/pii/0168900295014780> (visited on 07/13/2025).

- [694] M. Neubert. “B DECAYS AND CP VIOLATION”. en. In: *International Journal of Modern Physics A* 11.23 (1996). arXiv:hep-ph/9604412, pp. 4173–4240. ISSN: 0217-751X, 1793-656X. DOI: 10.1142/S0217751X96001966. URL: <https://www.worldscientific.com/doi/abs/10.1142/S0217751X96001966> (visited on 07/23/2025).
- [695] D. Buskulic et al. “Performance of the ALEPH detector at LEP”. In: *Nucl. Instrum. Meth. A* 360 (1995), pp. 481–506. DOI: 10.1016/0168-9002(95)00138-7.
- [696] D. A. Fish, A. M. Brinicombe, E. R. Pike, and J. G. Walker. “Blind deconvolution by means of the Richardson–Lucy algorithm”. En. In: *Josa A* 12.1 (Jan. 1995). Publisher: Optica Publishing Group, pp. 58–65. ISSN: 1520-8532. DOI: 10.1364/josaa.12.000058. URL: <https://opg.optica.org/josaa/abstract.cfm?uri=josaa-12-1-58> (visited on 07/13/2025).
- [697] Y. Freund and R. E. Schapire. “A decision-theoretic generalization of on-line learning and an application to boosting”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 904 (1995), pp. 23–37. ISSN: 1611-3349. DOI: 10.1007/3-540-59119-2_{_}166. URL: https://link.springer.com/chapter/10.1007/3-540-59119-2_166.
- [698] F. Girosi, M. Jones, and T. Poggio. “Regularization Theory and Neural Networks Architectures”. In: *Neural Computation* 7.2 (Mar. 1995), pp. 219–269. ISSN: 0899-7667. DOI: 10.1162/neco.1995.7.2.219. URL: <https://dx.doi.org/10.1162/neco.1995.7.2.219>.

- [699] D. J. Miller, A. Rao, K. Rose, and A. Gersho. “An Information-theoretic Learning Algorithm for Neural Network Classification”. In: *Advances in Neural Information Processing Systems*. Vol. 8. MIT Press, 1995. URL: https://proceedings.neurips.cc/paper_files/paper/1995/hash/38ca89564b2259401518960f7a06Abstract.html (visited on 07/22/2025).
- [700] M. E. Peskin and D. V. Schroeder. *An Introduction to quantum field theory*. Reading, USA: Addison-Wesley, 1995. ISBN: 978-0-429-49417-8. DOI: 10.1201/9780429503559.
- [701] S. Zaroubi. “Wiener reconstruction, SVD and ’optimal’ functional bases: Application for redshift galaxy catalogs”. In: *30th Rencontres de Moriond: Euroconferences: Clustering in the Universe*. _eprint: astro-ph/9505103. 1995, pp. 135–142.
- [702] G. D’Agostini. *A multidimensional unfolding method based on Bayes Theorem*. Tech. rep. Hamburg: Desy, 1994. URL: <https://cds.cern.ch/record/265717>.
- [703] W. R. Leo. *Techniques for Nuclear and Particle Physics Experiments: A How-to Approach*. en. Berlin, Heidelberg: Springer, 1994. ISBN: 978-3-642-57920-2. DOI: 10.1007/978-3-642-57920-2. URL: <https://link.springer.com/10.1007/978-3-642-57920-2> (visited on 07/13/2025).
- [704] M. Matzke. *Unfolding of pulse height spectra*. Tech. rep. Braunschweig: Braunschweig Univ. Phys.-Tech. Bundesanst, 1994. URL: <https://cds.cern.ch/record/275528>.

- [705] C. R. Rao. *Selected papers of C.R. Rao / 1. eng.* OCLC: 174244259. New York: Wiley, 1994. ISBN: 978-0-470-22091-7.
- [706] M. Schmelling. “The Method of reduced cross entropy: A General approach to unfold probability distributions”. In: *Nucl. Instrum. Meth. A* 340 (1994), pp. 400–412. DOI: 10.1016/0168-9002(94)90119-8.
- [707] N. B. Skachkov Jinr and R. Dubna. *On the determination of alpha s from the QCD analysis of the scaling violation in the fragmentation functions of the process $e+e- \rightarrow h + X$ studied by DELPHI, TASSO and other collaborations.* June 1994. URL: <https://cds.cern.ch/record/2627790>.
- [708] M. P. Wand and M. C. Jones. *Kernel Smoothing.* New York: Chapman and Hall/CRC, Nov. 1994. ISBN: 978-0-429-17059-1. DOI: 10.1201/b14876.
- [709] S. Catani, L. Trentadue, G. Turnock, and B. R. Webber. “Resummation of Large Logarithms in e^+e^- Event Shape Distributions”. In: *Nucl. Phys. B* 407 (1993), pp. 3–42. DOI: 10.1016/0550-3213(93)90271-p.
- [710] Y. Ohnishi, I. Adachi, J. Fujimoto, R. Itoh, A. Sugiyama, M. Yamauchi, K. Abe, T. Abe, M. Aoki, S. Awa, et al. “Measurements of α_s in e^+e^- annihilation at TRISTAN”. In: *Phys.Lett.B* 313.3-4 (Sept. 1993), pp. 475–482. ISSN: 03702693. DOI: 10.1016/0370-2693(93)90022-a.
- [711] J. Praestgaard and J. A. Wellner. “Exchangeably Weighted Bootstraps of the General Empirical Process”. In: *The Annals of Probability* 21.4 (Oct. 1993), pp. 2053–2086. ISSN: 0091-1798, 2168-894X. DOI: 10.1214/aop/1176989011. URL: <https://projecteuclid.org/journals/annals-of-probability/>

- volume - 21 / issue - 4 / Exchangeably - Weighted - Bootstraps - of - the - General - Empirical - Process / 10 . 1214 / aop / 1176989011 . full (visited on 07/22/2025).
- [712] D. Decamp, B. Deschizeaux, C. Goy, J. P. Lees, M. N. Minard, R. Alemany, F. Ariztizabal, P. Comas, J. M. Crespo, M. Delfino, et al. “Measurement of α_s in hadronic Z decays using all-orders resummed predictions”. In: *Physics Letters B* 284.1-2 (June 1992), pp. 163–176. ISSN: 03702693. DOI: 10.1016/0370-2693(92)91942-3.
- [713] A. Gelman and D. B. Rubin. “Inference from Iterative Simulation Using Multiple Sequences”. In: *Statist. Sci.* 7 (1992), pp. 457–472. DOI: 10.1214/ss/1177011136.
- [714] S. Catani, G. Turnock, and B. R. Webber. “Heavy jet mass distribution in e^+e^- annihilation”. In: *Physics Letters B* 272.3-4 (Dec. 1991), pp. 368–372. ISSN: 03702693. DOI: 10.1016/0370-2693(91)91845-m.
- [715] S. Catani, G. Turnock, B. R. Webber, and L. Trentadue. “Thrust distribution in e^+e^- annihilation”. In: *Physics Letters B* 263.3-4 (July 1991), pp. 491–497. ISSN: 03702693. DOI: 10.1016/0370-2693(91)90494-b.
- [716] E. b. V. P. Godambe, ed. *Estimating Functions*. Oxford Statistical Science Series. Oxford, New York: Oxford University Press, Sept. 1991. ISBN: 978-0-19-852228-7.
- [717] J. Lin. “Divergence measures based on the Shannon entropy”. In: *IEEE Transactions on Information Theory* 37.1 (Jan. 1991), pp. 145–151. ISSN: 1557-9654.

- DOI: 10.1109/18.61115. URL: <https://ieeexplore.ieee.org/document/61115> (visited on 07/17/2025).
- [718] K. Abe, K. Amako, Y. Arai, Y. Asano, M. Chiba, Y. Chiba, M. Daigo, T. Emura, M. Fukawa, T. Fukui, et al. “Determination of the QCD scale parameter Λ_{MS} with QCD cascade on the basis of the next-to-leading logarithmic approximation”. In: *Physics Letters B* 240.1-2 (Apr. 1990), pp. 232–236. ISSN: 03702693. DOI: 10.1016/0370-2693(90)90440-h.
- [719] B. Adeva, M. Aguilar-Benitez, H. Akbari, J. Alcaraz, A. Aloisio, J. Alvarez-Taviel, G. Alverson, M. G. Alviggi, H. Anderhub, A. L. Anderson, et al. “The construction of the L3 experiment”. In: *Nuclear Inst. and Methods in Physics Research, A* 289.1-2 (Apr. 1990), pp. 35–102. ISSN: 01689002. DOI: 10.1016/0168-9002(90)90250-a.
- [720] G. J. H. Burgers, C. Fuglesang, R. Hagedorn, and V. Kuvshinov. “Multiplicity Distributions in Hadron Interactions Derived From the Statistical Bootstrap Model”. In: *Z. Phys. C* 46 (1990), pp. 465–480. DOI: 10.1007/BF01621036.
- [721] G. Cybenko. “Approximation by superpositions of a sigmoidal function”. en. In: *Mathematics of Control, Signals and Systems* 2.4 (Dec. 1989), pp. 303–314. ISSN: 1435-568x. DOI: 10.1007/bf02551274. URL: <https://doi.org/10.1007/BF02551274> (visited on 07/15/2025).
- [722] K. Hornik, M. Stinchcombe, and H. White. “Multilayer feedforward networks are universal approximators”. In: *Neural Networks* 2.5 (Jan. 1989), pp. 359–366. ISSN: 0893-6080. DOI: 10.1016/0893-6080(89)90020-8. URL: [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8).

- [//www.sciencedirect.com/science/article/pii/S08933608089900208](https://www.sciencedirect.com/science/article/pii/S08933608089900208)
(visited on 07/15/2025).
- [723] K. Kato and T. Munehisa. “Double-cascade scheme for QCD jets in $e+e-$ annihilation”. In: *Physical Review D* 39.1 (1989), pp. 156–162. ISSN: 05562821. DOI: 10.1103/physrevd.39.156.
- [724] H. R. Kunsch. “The Jackknife and the Bootstrap for General Stationary Observations”. In: *The Annals of Statistics* 17.3 (Sept. 1989), pp. 1217–1241. ISSN: 0090-5364, 2168-8966. DOI: 10.1214/aos/1176347265. URL: <https://projecteuclid.org/journals/annals-of-statistics/volume-17/issue-3/The-Jackknife-and-the-Bootstrap-for-General-Stationary-Observations/10.1214/aos/1176347265.full> (visited on 07/22/2025).
- [725] R. F. Streater and A. S. Wightman. *PCT, spin and statistics, and all that*. 1989. ISBN: 978-0-691-07062-9.
- [726] G. Zech. “Upper limits in experiments with background or measurement errors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 277.2 (1989), p. 608. ISSN: 0168-9002. DOI: 10.1016/0168-9002(89)90795-X. URL: <https://www.sciencedirect.com/science/article/pii/S016890028990795X> (visited on 07/22/2025).
- [727] H. N. Multhei and B. Schorr. “On an Iterative Method for the Unfolding of Spectra”. In: *Nucl. Instrum. Meth. A* 257 (1987), p. 371. DOI: 10.1016/0168-9002(87)90759-5.

- [728] L. L. Cam. “Asymptotic Methods in Statistical Decision Theory”. In: Springer Series in Statistics (1986). DOI: 10.1007/978-1-4612-4946-7. URL: <http://link.springer.com/10.1007/978-1-4612-4946-7>.
- [729] R. Narayan and R. Nityananda. “Maximum entropy image restoration in astronomy”. In: *Ann. Rev. Astron. Astrophys.* 24 (1986), pp. 127–170. DOI: 10.1146/annurev.aa.24.090186.001015.
- [730] M. Althoff, W. Braunschweig, F. J. Kirschfink, K. Lübelmeyer, H. U. Martyn, P. Roskamp, H. G. Sander, D. Schmitz, H. Siebke, W. Wallraff, et al. “Determination of α_s in first and second order QCD from $e+e^-$ annihilation into hadrons”. In: *Zeitschrift für Physik C Particles and Fields* 26.2 (1984), pp. 157–174. ISSN: 14346052. DOI: 10.1007/bf01421751/metrics. URL: <https://link.springer.com/article/10.1007/BF01421751>.
- [731] E. D’Hoker and E. Farhi. “Decoupling a fermion in the standard electro-weak theory”. In: *Nuclear Physics, Section B* 248.1 (1984). ISSN: 05503213. DOI: 10.1016/0550-3213(84)90587-x.
- [732] B. Adeva, D. P. Barber, U. Becker, J. Berdugo, G. Berghoff, A. Böhm, J. G. Branson, J. D. Burger, M. Capell, M. Cerrada, et al. “Model-independent second-order determination of the strong-coupling constant α_s ”. In: *Phys. Rev. Lett.* 50.26 (1983), p. 2051. ISSN: 00319007. DOI: 10.1103/physrevlett.50.2051.

- [733] G. Altarelli. “Partons in Quantum Chromodynamics”. In: *Physics Reports* 81.1 (1982), p. 1. ISSN: 0370-1573. DOI: 10.1016/0370-1573(82)90127-2. (Visited on 07/18/2025).
- [734] G. Arnison, A. Astbury, B. Aubert, C. Bacci, R. Bernabei, A. Bézaguët, R. Böck, T. J. V. Bowcock, M. Calvetti, T. Carroll, et al. “Transverse momentum spectra for charged particles at the CERN proton-antiproton collider”. In: *Physics Letters B* 118.1 (Dec. 1982), pp. 167–172. ISSN: 0370-2693. DOI: 10.1016/0370-2693(82)90623-2. URL: <https://www.sciencedirect.com/science/article/pii/0370269382906232> (visited on 07/14/2025).
- [735] L. A. Shepp and Y. Vardi. “Maximum Likelihood Reconstruction for Emission Tomography”. In: *IEEE Transactions on Medical Imaging* 1.2 (Oct. 1982), pp. 113–122. ISSN: 1558-254x. DOI: 10.1109/tmi.1982.4307558. URL: <https://ieeexplore.ieee.org/document/4307558> (visited on 07/13/2025).
- [736] H. White. “Maximum Likelihood Estimation of Misspecified Models”. In: *Econometrica* 50.1 (1982), pp. 1–25. ISSN: 0012-9682. DOI: 10.2307/1912526. URL: <https://www.jstor.org/stable/1912526> (visited on 07/22/2025).
- [737] T. Chandramohan and L. Clavelli. “Consequences of 2nd order QCD for jet structure in e^+e^- annihilation”. In: *Nuclear Physics, Section B* 184.C (1981), pp. 365–380. ISSN: 05503213. DOI: 10.1016/0550-3213(81)90224-8.
- [738] D. B. Rubin. “The Bayesian Bootstrap”. In: *The Annals of Statistics* 9.1 (Jan. 1981), pp. 130–134. ISSN: 0090-5364, 2168-8966. DOI: 10.1214/aos/1176345338. URL: <https://projecteuclid.org/journals/annals-of->

- statistics/volume-9/issue-1/The-Bayesian-Bootstrap/10.1214/aos/1176345338.full (visited on 07/22/2025).
- [739] I. Vincze. “On the Concept and Measure of Information Contained in an Observation”. In: *Contributions to Probability*. Ed. by J. Gani and V. K. Rohatgi. Academic Press, Jan. 1981, pp. 207–214. ISBN: 978-0-12-274460-0. DOI: 10.1016/b978-0-12-274460-0.50023-0. URL: <https://www.sciencedirect.com/science/article/pii/B9780122744600500230> (visited on 07/14/2025).
- [740] A. Salam. “Gauge Unification of Fundamental Forces”. In: *Rev. Mod. Phys.* 52 (1980). Ed. by A. Ali, C. Isham, T. Kibble, and Riazuddin, pp. 525–538. DOI: 10.1103/RevModPhys.52.525.
- [741] *The Classical Theory of Fields*. en-US. Jan. 1980. ISBN: 978-0-08-050349-3. URL: <https://shop.elsevier.com/books/the-classical-theory-of-fields/landau/978-0-08-050349-3> (visited on 07/24/2025).
- [742] S. Weinberg. “Conceptual Foundations of the Unified Theory of Weak and Electromagnetic Interactions”. In: *Rev. Mod. Phys.* 52 (1980). Ed. by S. Lundqvist, pp. 515–523. DOI: 10.1103/RevModPhys.52.515.
- [743] H. White. “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity”. In: *Econometrica* 48.4 (1980), pp. 817–838. ISSN: 0012-9682. DOI: 10.2307/1912934. URL: <https://www.jstor.org/stable/1912934> (visited on 07/22/2025).

- [744] P. Hoyer, P. Osland, H. G. Sander, T. F. Walsh, and P. M. Zerwas. “Quantum chromodynamics and jets in $e+e-$ ”. In: *Nucl. Phys. B* 161.2-3 (1979), p. 349. ISSN: 05503213. DOI: 10.1016/0550-3213(79)90217-7.
- [745] W. A. Bardeen, A. J. Buras, D. W. Duke, and T. Muta. “Deep-Inelastic Scattering beyond the Leading Order in Asymptotically Free Gauge Theories”. In: *Physical Review D* 18.11 (1978), p. 3998. DOI: 10.1103/PhysRevD.18.3998. (Visited on 07/18/2025).
- [746] L. Boltzmann. “Ableitung des Stefanschen Gesetzes, 1) betreffend die Abhängigkeit der Wärmestrahlung von der Temperatur aus der elektromagnetischen Lichttheorie”. In: *Von Kirchhoff bis Planck* (1978), pp. 152–156. DOI: 10.1007/978-3-663-13885-3_{11}. URL: <https://link.springer.com/chapter/10.1007/978-3-663-13885-3%5C%5F11>.
- [747] J. Kodaira and T. Uematsu. “How Can We Define Properly the Q^2 Dependent Parton Distribution Function in QCD?” In: *Nucl.Phys.B* 141.4 (Sept. 1978), pp. 497–506. ISSN: 05503213. DOI: 10.1016/0550-3213(78)90042-1.
- [748] G. Altarelli and G. Parisi. “Asymptotic Freedom in Parton Language”. In: *Nucl. Phys. B* 126 (1977), pp. 298–318. DOI: 10.1016/0550-3213(77)90384-4.
- [749] A. Birnbaum. “The Neyman-Pearson theory as decision theory, and as inference theory; with a criticism of the Lindley-savage argument for Bayesian theory”. en. In: *Synthese* 36.1 (Sept. 1977), pp. 19–49. ISSN: 1573-0964. DOI: 10.1007/BF00485690. URL: <https://doi.org/10.1007/BF00485690> (visited on 07/22/2025).

- [750] Y. L. Dokshitzer. “Calculation of the Structure Functions for Deep Inelastic Scattering and $e^+ e^-$ Annihilation by Perturbation Theory in Quantum Chromodynamics.” In: *Sov. Phys. JETP* 46 (1977), pp. 641–653.
- [751] E. Farhi. “Quantum Chromodynamics Test for Jets”. In: *Physical Review Letters* 39.25 (1977), pp. 1587–1588. ISSN: 00319007. DOI: 10.1103/physrevlett.39.1587.
- [752] R. N. Giere. “Allan Birnbaum’s conception of statistical evidence”. en. In: *Synthese* 36.1 (Sept. 1977), pp. 5–13. ISSN: 1573-0964. DOI: 10.1007/BF00485688. URL: <https://doi.org/10.1007/BF00485688> (visited on 07/22/2025).
- [753] J. L. Ibanez and Y. Pomeau. “A Simple Case of Strange Attractor”. In: (Mar. 1976).
- [754] L. B. Lucy. “An iterative technique for the rectification of observed distributions”. en. In: *The Astronomical Journal* 79 (June 1974), p. 745. ISSN: 0004-6256. DOI: 10.1086/111605. URL: <https://ui.adsabs.harvard.edu/abs/1974AJ.....79..745L/abstract> (visited on 07/13/2025).
- [755] D. J. Gross and F. Wilczek. “Asymptotically Free Gauge Theories. I”. In: *Physical Review D* 8.10 (Nov. 1973), pp. 3633–3652. DOI: 10.1103/PhysRevD.8.3633. (Visited on 07/18/2025).
- [756] D. J. Gross and F. Wilczek. “Ultraviolet Behavior of Non-Abelian Gauge Theories”. In: *Physical Review Letters* 30.26 (June 1973), pp. 1343–1346. DOI: 10.1103/PhysRevLett.30.1343. (Visited on 07/18/2025).

- [757] H. D. Politzer. “Reliable Perturbative Results for Strong Interactions?” In: *Physical Review Letters* 30.26 (1973), pp. 1346–1349. DOI: 10.1103/PhysRevLett.30.1346. (Visited on 07/18/2025).
- [758] V. N. Gribov and L. N. Lipatov. “Deep inelastic e p scattering in perturbation theory”. In: *Sov. J. Nucl. Phys.* 15 (1972), pp. 438–450.
- [759] W. H. Richardson. “Bayesian-Based Iterative Method of Image Restoration”. En. In: *Josa* 62.1 (Jan. 1972). Publisher: Optica Publishing Group, pp. 55–59. DOI: 10.1364/josa.62.000055. URL: <https://opg.optica.org/josa/abstract.cfm?uri=josa-62-1-55> (visited on 07/13/2025).
- [760] V. N. Vapnik and A. Y. Chervonenkis. “On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities”. In: *Theory of Probability & Its Applications* 16.2 (Jan. 1971), pp. 264–280. ISSN: 0040-585X. DOI: 10.1137/1116025. URL: <https://epubs.siam.org/doi/10.1137/1116025> (visited on 07/24/2025).
- [761] W. E. Deming, H. Tamura, F. W. Bennett, I. H. LaValle, E. Dayhoff, W. Knight, I. J. Good, L. R. Shenton, C. E. Stephenson, J. D. Williams, S. N. Dalal, M.-T. Chao, and M. N. Das. “Letters to the Editor”. In: *The American Statistician* 24.2 (1970). Publisher: [American Statistical Association, Taylor & Francis, Ltd.], pp. 35–39. ISSN: 0003-1305. DOI: 10.1080/00031305.1970.10478892. URL: <https://www.jstor.org/stable/2681931> (visited on 07/24/2025).
- [762] L. N. Vaserstein. “Markov Processes over Denumerable Products of Spaces, Describing Large Systems of Automata”. In: *Problemy Peredachi Informatsii*

- 5.3 (1969). English translation available, pp. 64–72. URL: <https://m.mathnet.ru/ppi1811>.
- [763] F. Eicker. “Limit theorems for regressions with unequal and dependent errors”. In: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*. Vol. 5.1. University of California Press, Jan. 1967, pp. 59–83. URL: <https://projecteuclid.org/ebooks/berkeley-symposium-on-mathematical-statistics-and-probability/Proceedings-of-the-Fifth-Berkeley-Symposium-on-Mathematical-Statistics-and/chapter/Limit-theorems-for-regressions-with-unequal-and-dependent-errors/bsmsp/1200512981> (visited on 07/22/2025).
- [764] P. J. Huber. “The behavior of maximum likelihood estimates under nonstandard conditions”. In: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*. Vol. 5.1. University of California Press, Jan. 1967, pp. 221–234. URL: <https://projecteuclid.org/ebooks/berkeley-symposium-on-mathematical-statistics-and-probability/Proceedings-of-the-Fifth-Berkeley-Symposium-on-Mathematical-Statistics-and/chapter/The-behavior-of-maximum-likelihood-estimates-under-nonstandard-conditions/bsmsp/1200512988> (visited on 07/22/2025).
- [765] P. W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Phys. Rev. Lett.* 13 (16 Oct. 1964), pp. 508–509. DOI: 10.1103/PhysRevLett.13.508. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.508>.

- [766] P. J. Huber. “Robust Estimation of a Location Parameter”. In: *The Annals of Mathematical Statistics* 35.1 (Mar. 1964), pp. 73–101. ISSN: 0003-4851, 2168-8990. DOI: 10.1214/aoms/1177703732. (Visited on 08/11/2025).
- [767] A. N. Tikhonov. “Solution of incorrectly formulated problems and the regularization method”. English. In: *Sov. Math., Dokl.* 5 (1963), pp. 1035–1038. ISSN: 0197-6788.
- [768] S. Weinberg. “Elementary Particle Theory of Composite Particles”. In: *Physical Review* 130.2 (Apr. 1963). Publisher: American Physical Society, pp. 776–783. DOI: 10.1103/PhysRev.130.776. URL: <https://link.aps.org/doi/10.1103/PhysRev.130.776> (visited on 07/13/2025).
- [769] A. Birnbaum. “On the Foundations of Statistical Inference”. In: *Journal of the American Statistical Association* 57.298 (June 1962). Publisher: ASA Website _eprint: <https://www.tandfonline.com/doi/pdf/10.1080/01621459.1962.10480660>, pp. 269–306. ISSN: 0162-1459. DOI: 10.1080/01621459.1962.10480660. URL: <https://www.tandfonline.com/doi/abs/10.1080/01621459.1962.10480660> (visited on 07/22/2025).
- [770] S. L. Glashow. “Partial Symmetries of Weak Interactions”. In: *Nucl. Phys.* 22 (1961), pp. 579–588. DOI: 10.1016/0029-5582(61)90469-2.
- [771] A. Kolmogorov. “Théorie générale des systèmes dynamiques de la mécanique classique”. fra. In: *Séminaire Janet. Mécanique analytique et mécanique céleste* 1 (1957). Publisher: Secrétariat mathématique, pp. 1–20. URL: <https://eudml.org/doc/112607> (visited on 07/24/2025).

- [772] G. Luders and B. Zumino. “Some Consequences of TCP-Invariance”. In: *Phys. Rev.* 106 (1957), pp. 385–386. DOI: 10.1103/PhysRev.106.385.
- [773] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. “Experimental Test of Parity Conservation in Beta Decay”. In: *Physical Review* 105.4 (Feb. 1957), pp. 1413–1414. DOI: 10.1103/PhysRev.105.1413. URL: <https://link.aps.org/doi/10.1103/PhysRev.105.1413> (visited on 07/23/2025).
- [774] D. N. LAWLEY. “TESTS OF SIGNIFICANCE FOR THE LATENT ROOTS OF COVARIANCE AND CORRELATION MATRICES”. In: *Biometrika* 43.1-2 (June 1956), pp. 128–136. ISSN: 0006-3444. DOI: 10.1093/biomet/43.1-2.128. URL: <https://doi.org/10.1093/biomet/43.1-2.128> (visited on 07/22/2025).
- [775] D. N. Lawley. “A General Method for Approximating to the Distribution of Likelihood Ratio Criteria”. In: *Biometrika* 43.3/4 (1956), pp. 295–303. ISSN: 0006-3444. DOI: 10.2307/2332908. URL: <https://www.jstor.org/stable/2332908> (visited on 07/22/2025).
- [776] T. D. Lee and C.-N. Yang. “Question of Parity Conservation in Weak Interactions”. In: *Phys. Rev.* 104 (1956), pp. 254–258. DOI: 10.1103/PhysRev.104.254.
- [777] T. W. Anderson and D. A. Darling. “Asymptotic Theory of Certain "Goodness of Fit" Criteria Based on Stochastic Processes”. In: *The Annals of Mathematical Statistics* 23.2 (June 1952). Publisher: Institute of Mathematical Statistics,

- pp. 193–212. ISSN: 0003-4851, 2168-8990. DOI: 10.1214/aoms/1177729437. URL: <https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-23/issue-2/Asymptotic-Theory-of-Certain-Goodness-of-Fit-Criteria-Based-on/10.1214/aoms/1177729437.full> (visited on 07/24/2025).
- [778] H. Chernoff. “A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the sum of Observations”. In: *The Annals of Mathematical Statistics* 23.4 (Dec. 1952), pp. 493–507. ISSN: 0003-4851, 2168-8990. DOI: 10.1214/aoms/1177729330. URL: <https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-23/issue-4/A-Measure-of-Asymptotic-Efficiency-for-Tests-of-a-Hypothesis/10.1214/aoms/1177729330.full> (visited on 07/22/2025).
- [779] S. Kullback and R. A. Leibler. “On Information and Sufficiency”. In: *The Annals of Mathematical Statistics* 22.1 (Mar. 1951). Publisher: Institute of Mathematical Statistics, pp. 79–86. ISSN: 0003-4851, 2168-8990. DOI: 10.1214/aoms/1177729694. URL: <https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-22/issue-1/On-Information-and-Sufficiency/10.1214/aoms/1177729694.full> (visited on 07/14/2025).
- [780] G. Goertzel. *Quota Sampling and Importance Functions in Stochastic Solution of Particle Problems*. English. Tech. rep. AECD-2793. Originating Research Org. not identified, Nov. 1949. URL: <https://www.osti.gov/biblio/4405957> (visited on 07/24/2025).

- [781] K. Iwasawa. “On Some Types of Topological Groups”. In: *Annals of Mathematics* 50.3 (1949). Publisher: [Annals of Mathematics, Trustees of Princeton University on Behalf of the Annals of Mathematics, Mathematics Department, Princeton University], pp. 507–558. ISSN: 0003-486X. DOI: 10.2307/1969548. URL: <https://www.jstor.org/stable/1969548> (visited on 07/24/2025).
- [782] H. Cramér. *Mathematical Methods of Statistics*. Princeton University Press, 1946. ISBN: 9781400883868.
- [783] G. Darmais. “Sur les limites de la dispersion de certaines estimations”. In: *Revue de l’Institut International de Statistique / Review of the International Statistical Institute* 13.1/4 (1945). Publisher: [International Statistical Institute (ISI), Wiley], pp. 9–15. ISSN: 0373-1138. DOI: 10.2307/1400974. URL: <https://www.jstor.org/stable/1400974> (visited on 07/24/2025).
- [784] C. Radhakrishna Rao. “Information and accuracy attainable in the estimation of statistical parameters”. In: (1945). ISSN: 0008-0659. URL: <http://bulletin.calmathsoc.org/article.php?ID=B.1945.37.14>.
- [785] A. C. Aitken and H. Silverstone. “XV.—On the Estimation of Statistical Parameters”. en. In: *Proceedings of the Royal Society of Edinburgh Section A: Mathematics* 61.2 (Jan. 1942), pp. 186–194. ISSN: 2053-5902, 0080-4541. DOI: 10.1017/S008045410000618X. URL: <https://www.cambridge.org/core/journals/proceedings-of-the-royal-society-of-edinburgh-section-a-mathematics/article/abs/xvon-the-estimation-of-statistical-parameters/74C2B968275EED1A954DB316FBC3635E> (visited on 07/24/2025).

- [786] A. S. Householder. “A theory of steady-state activity in nerve-fiber networks: I. Definitions and preliminary lemmas”. In: *The Bulletin of Mathematical Biophysics* 3.2 (June 1941), pp. 63–69. ISSN: 00074985. DOI: 10.1007/bf02478220/metrics. URL: <https://link.springer.com/article/10.1007/BF02478220>.
- [787] W. Pauli. “The Connection Between Spin and Statistics”. In: *Phys. Rev.* 58 (1940), pp. 716–722. DOI: 10.1103/PhysRev.58.716.
- [788] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. In: <https://doi.org/10.1214/aoms/1177732360> 9.1 (Mar. 1938), pp. 60–62. ISSN: 0003-4851. DOI: 10.1214/AOMS/1177732360. URL: <https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-9/issue-1/The-Large-Sample-Distribution-of-the-Likelihood-Ratio-for-Testing/10.1214/aoms/1177732360.full%20https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-9/issue-1/The-Large-Sample-Distribution-of-the-Likelihood-Ratio-for-Testing/10.1214/aoms/1177732360.short>.
- [789] “Properties of sufficiency and statistical tests”. In: *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences* 160.901 (May 1937), pp. 268–282. ISSN: 0080-4630. DOI: 10.1098/RSPA.1937.0109.
- [790] O. Nikodym. “Sur une généralisation des intégrales de M. J. Radon”. In: *Fundamenta Mathematicae* 15.1 (1930), pp. 131–179. ISSN: 0016-2736.

- [791] H. Cramér. “On the composition of elementary errors: First paper: Mathematical deductions”. In: *Scandinavian Actuarial Journal* 1928.1 (Jan. 1928). Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/03461238.1928.10416862>, pp. 13–74. ISSN: 0346-1238. DOI: 10.1080/03461238.1928.10416862. URL: <https://doi.org/10.1080/03461238.1928.10416862> (visited on 07/24/2025).
- [792] R. Von Mises. *Wahrscheinlichkeit Statistik und Wahrheit*. de. Berlin, Heidelberg: Springer Berlin Heidelberg, 1928. ISBN: 978-3-662-36230-3. DOI: 10.1007/978-3-662-36230-3. URL: <http://link.springer.com/10.1007/978-3-662-36230-3> (visited on 07/24/2025).
- [793] E. Slutsky. *Über stochastische Asymptoten und Grenzwerte*. German. Metron 5, Nr. 3, 3-89 (1925). zbMATH: 2592216, JFM: 51.0380.03. 1925.
- [794] F. Hausdorff. “Summationsmethoden und Momentfolgen I.” ger. In: *Mathematische Zeitschrift* 9 (1921), pp. 74–109. URL: <http://eudml.org/doc/167613>.
- [795] A. Einstein. “The foundation of the general theory of relativity.” In: *Annalen Phys.* 49.7 (1916). Ed. by J.-P. Hsu and D. Fine, pp. 769–822. DOI: 10.1002/andp.19163540702.
- [796] J. Radon. *Theorie und anwendungen der absolut additiven mengenfunktionen ...* Wien: Hölder, 1913. URL: <https://catalog.hathitrust.org/Record/008892713> (visited on 07/20/2025).
- [797] P. E. R. F.R.S. “LXXIX. The scattering of α and β particles by matter and the structure of the atom”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 21.125 (May 1911), pp. 669–688. ISSN: 1941-

5982. DOI: 10.1080/14786440508637080. URL: <https://www.tandfonline.com/doi/abs/10.1080/14786440508637080>.
- [798] I. Fredholm. “Sur une classe d’équations fonctionnelles”. In: *Acta Mathematica* 27.none (Jan. 1903). Publisher: Institut Mittag-Leffler, pp. 365–390. ISSN: 0001-5962, 1871-2509. DOI: 10.1007/bf02421317. URL: <https://projecteuclid.org/journals/acta-mathematica/volume-27/issue-none/Sur-une-classe-d%3a9equations-fonctionnelles/10.1007/BF02421317.full> (visited on 07/13/2025).
- [799] O. Hölder. “Die einfachen Gruppen im ersten und zweiten Hundert der Ordnungszahlen”. In: *Mathematische Annalen* 40 (1892), pp. 55–88. ISSN: 0025-5831.
- [800] K. Weierstraß. *Über die analytische Darstellbarkeit sogenannter willkürlicher Functionen einer reellen Veränderlichen*. ger. Sitzungsberichte der Preussischen Akademie der Wissenschaften Jg. 1885, 34 ; Jg. 1885, 38. Berlin: : 1885.
- [801] C. G. J. Jacobi. “De Determinantibus functionalibus.” lat. In: *Journal für die reine und angewandte Mathematik* 22 (1841), pp. 319–359. ISSN: 0075-4102; 1435-5345/e. URL: <https://eudml.org/doc/147138> (visited on 07/24/2025).
- [802] B. Taylor. *Methodus incrementorum directa & inversa*. Auctore Brook Taylor, LL. D. & Regiae Societatis Secretario. lat. Londini: typis Pearsonianis: prostant apud Gul. Innys ad Insignia Principis in Coemeterio Paulino, 1715.
- [803] *Philosophiae Naturalis Principia Mathematica*. Auctore Js. Newton ... en. 1687. URL: <http://www.e-rara.ch/zut/129833> (visited on 07/24/2025).

- [804] N. R. Council. “Elementary-Particle Physics: Revealing the Secrets of Energy and Matter”. In: (Apr. 1). DOI: 10.17226/6045. URL: <http://www.nap.edu/catalog/6045>.

Appendix A

Glossary

Angular distance The *angular distance metric* is defined as

$$\Delta R^2 = \Delta \eta^2 + \Delta \varphi^2, \quad (\text{A.1})$$

where φ is the azimuthal angle around the beam axis. This metric approximates the geometric angle between particles in the detector while remaining approximately longitudinally boost invariant.

Cross section For collisions between discrete particles, the *cross section* is defined as the area transverse to their relative motion. If the particles were to interact via contact forces (e.g., hard spheres), the cross section corresponds to their geometric size. For long range forces however, the cross section is larger than the physical dimensions of the particles due to action-at-a-distance effects.

Differential cross section The *differential cross section* ($\frac{d\sigma}{d\Omega}$) provides additional granularity by describing how the probability of scattering depends on specific

final state variables, such as scattering angle (θ) or energy transfer. It is defined as

$$\frac{d\sigma}{d\Omega} = \frac{\text{Number of events scattered into } d\Omega}{\text{Incident flux} \times \text{Target density}}. \quad (\text{A.2})$$

Integrated luminosity The *integrated luminosity* $\mathcal{L}_{\text{int}}$ is defined as

$$\mathcal{L}_{\text{int}} = \int_{t_0}^{t_f} \mathcal{L}(t) dt, \quad \sigma = \frac{N_{\text{obs}} - N_{\text{bkg}}}{\mathcal{L}_{\text{int}} \epsilon A}. \quad (\text{A.3})$$

Here N_{obs} is the number of selected events, N_{bkg} an estimate of background contaminations, ϵ the detector and selection efficiency and A the geometric-plus-kinematic acceptance of the analysis.

Luminosity The *luminosity* $\mathcal{L}(t)$ is defined such that the interaction rate for a process with cross section σ is $dN/dt = \mathcal{L}(t) \sigma$.

Monte Carlo tuning *Monte Carlo tuning* is the iterative process of adjusting simulation parameters to match measured cross sections, ensuring that detector simulations accurately reproduce real experimental data.

Natural units *Natural Units* are a system of units where $\hbar = c = 1$, effectively setting the speed of light and quantum of action as fundamental measuring sticks. This choice transforms all physical quantities into powers of energy.

Parton distribution functions *Parton distribution functions* (PDFs) are probability distributions that describe the probability of finding a parton (quark or gluon) in a hadron at a given momentum fraction x and scale Q^2 . PDFs are determined from global fits to a wide range of hard scattering processes, including deep inelastic scattering, Drell–Yan production, and jet production.

Parton luminosity *Parton luminosity* is the effective luminosity for specific parton–parton interactions, calculated by convolving the total luminosity with parton distribution functions.

Poisson distribution. The *Poisson distribution* is the fundamental distribution governing counting experiments. For n observed events with expected value λ ,

$$\mathbb{P}(n \mid \lambda) = \frac{\lambda^n e^{-\lambda}}{n!} . \quad (\text{A.4})$$

Profile likelihood ratio The *profile likelihood ratio*

$$\mathcal{L}(\mu) = \frac{L(\mu, \hat{\theta}_\mu)}{L(\hat{\mu}, \hat{\theta})} \quad (\text{A.5})$$

where $\hat{\theta}_\mu$ maximises L for fixed signal strength μ is the optimal test statistics for hypothesis testing under the Neyman–Pearson lemma.

Pseudorapidity The *pseudorapidity* η , is defined as

$$\eta = -\ln \tan \frac{\theta}{2} \quad (\text{A.6})$$

where θ is the polar angle from the beam axis [252]. Under a longitudinal boost with rapidity $\beta = \text{artanh } v/c$, a massless particle’s η transforms as $\eta \mapsto \eta + \beta$

Correlation matrix. A useful, related measure is the *correlation matrix*, $\boldsymbol{\rho}$, defined as the normalised covariance matrix,

$$\rho_{ij} = \frac{\Sigma_{ij}}{\sqrt{\Sigma_{ii}\Sigma_{jj}}}, \quad (\text{A.7})$$

where Σ is the correlation matrix.

Covariance matrix. The *covariance matrix* Σ is defined as

$$\Sigma_{ij} := \mathbb{E}[(\hat{t}_i - \mathbb{E}[\hat{t}_i])(\hat{t}_j - \mathbb{E}[\hat{t}_j])]. \quad (\text{A.8})$$

KL divergence

$$D_{\text{KL}}(p||q) = \int p(x) \ln \frac{p(x)}{q(x)} dx, \quad (\text{A.9})$$

where p is the true distribution and q is the unfolded distribution.

Pull distribution. For a given unfolded bin or parameter θ , the *pull* is defined as

$$\text{Pull } \theta = \frac{\hat{\theta} - \theta_{\text{true}}}{\sigma_{\hat{\theta}}} \quad (\text{A.10})$$

where $\hat{\theta}$ is the unfolded estimate, θ_{true} is the true value, and $\sigma_{\hat{\theta}}$ is the reported uncertainty.

Residuals For binned methods, the *bin by bin residual* is defined as

$$\delta_i = \hat{t}_i - t_i \quad (\text{A.11})$$

where $\hat{t}_i$ is the unfolded count in bin i and t_i is the true count.

VLC Divergence

$$\Delta_{\text{VLC}}(p, q) = \frac{1}{2} \int \frac{(p(\lambda) - q(\lambda))^2}{p(\lambda) + q(\lambda)} d\lambda. \quad (\text{A.12})$$

Wasserstein α -metric

$$W_{\alpha}(p, q) = \left(\inf_{\gamma \in \Gamma(p, q)} \iint |x - y|^{\alpha} d\gamma(x, y) \right)^{\frac{1}{\alpha}}, \quad (\text{A.13})$$

where $\Gamma(p, q)$ is the set of all joint distributions with marginals p and q .

Adaptive methods. *Adaptive methods* are optimisation techniques that vary the learning rate η over the course of the training rather than keeping it constant, and allow different learning rates for different dimensions of the input, rather than the traditional scalar learning rate that was constant across dimensions.

Batch A *mini-batch*, often simply called a *batch*, is a subset of the training data used to compute the gradient of the loss function. In practice, neural networks are trained using mini-batch gradient descent, where parameter updates are performed using gradients computed on subsets (batches) of the training data.

Deep neural network A *deep neural network* is a neural network with depth, $d > 2$. That is to say standard deep feed forward neural network consists of an input layer, *more than one* hidden layers, and an output layer.

Depth of a neural network The *depth*, $d = \ell_{\max}$, of a neural network is defined as the depth of its deepest layer.

Depth of a neuron The *depth* of a neuron is defined as the length of the longest path between the input and the neuron.

Divergence The *divergence* $\text{Div}(f(x; W), g(x))$ is a continuous valued proxy for the error between $f(x; W)$ and $g(x)$.

Empirical risk The *empirical* risk as an approximation of the risk function defined as

$$\hat{R}(f_\theta) = \frac{1}{N} \sum_{i=1}^n L(y_i, f_\theta(x_i)) , \quad (\text{A.14})$$

where L is the loss function, the empirical expected divergence.

Forward propagation *Forward propagation* is the process of computing the network’s output given an input.

Hidden layer Any layer of a neural network except the input layer (often labelled layer 0) and the output layer ($\ell_{\max}$) is called a hidden layer.

Layer of a neuron The set of all neurons at depth ℓ is called *layer* ℓ of the network.

Learning rate scheduling *Learning rate scheduling* is the process of dynamically adjusting the learning rate during training.

Loss function The *loss function* L is defined as the unbiased empirical average divergence between the neural network output and the target output over all training instances.

$$L(W) = \frac{1}{|X|} \sum_{i=1}^{|X|} \text{Div} \left(f(x_i; W), g(x_i) \right). \quad (\text{A.15})$$

Neuron The fundamental unit of a neural network, a *neuron* or *node*, computes a weighted sum of its inputs followed by a non-linear transformation:

$$y = \sigma(w^T x + b), \quad (\text{A.16})$$

where $x \in \mathbb{R}^n$ is the input vector, $y \in \mathbb{R}^m$ is the output vector, $w \in \mathbb{R}^{n \times m}$ is a weight matrix, $b \in \mathbb{R}^m$ is a bias term, and $\sigma(\cdot)$ is a (typically non-linear) activation function.

Physics informed neural network A *physics informed neural network* is a neural network that explicitly incorporates domain knowledge about physical laws, conservation principles, and symmetries.

Risk function. For a space of functions f_θ parametrised by θ the *risk function* is defined as

$$R(f_\theta) = \mathbb{E}_{(x,y) \sim p(X,Y)} [\text{Div}(y, f_\theta(x))] , \quad (\text{A.17})$$

where Div is the divergence between the predicted output $f_\theta(x)$ and the true output y , and $P(X, Y)$ is the joint probability distribution of inputs and outputs.

Spectral normalisation *Spectral normalisation* is the process of controlling the Lipschitz constant of a discriminator by normalising its weights,

$$W_{\text{SN}} = \frac{W}{\sigma(W)} , \quad (\text{A.18})$$

where $\sigma(W)$ is the spectral norm of weight matrix W , provides additional tools to avoid training instabilities in GANs [389, 405].

Stochastic gradient descent (SGD) *Stochastic gradient descent* is an adaptation of gradient descent that updates parameters using gradients computed on mini-batches, introducing noise that can help escape local minima.

$$W_{t+1} = W_t - \eta \nabla L_{\text{batch}}(W_t) . \quad (\text{A.19})$$

Width of a layer The *width* of a layer is the number of neurons in that layer.

Binned unfolding A *binned unfolding method* is one in which both the particle level and detector level distributions are discretised into histograms, transforming the continuous convolution of Equation (I.16) into the matrix equation Equation (II.13).

Event shape observables. *Event shape observables*, are observables which characterise the geometrical properties of energy flow in collisions, are often studied through their moments.

MLC loss Consider the asymptotic *maximum likelihood classifier* (MLC) loss

$$L(d, g) = - \int p(z) \ln d(z) dz - \int \tilde{q}(z)(1 - d(z)) dz, \quad (\text{A.20})$$

where $d : \mathbb{R} \rightarrow (0, 1]$ is a discriminator.

Moments of PDFs. For a parton distribution function $f(x, Q^2)$, where x is the momentum fraction and Q^2 is the energy scale, the k^{th} moment is defined as

$$M_k(Q^2) = \int_0^1 x^{k-1} f(x, Q^2) dx. \quad (\text{A.21})$$

The DGLAP equations for these moments take the form

$$\frac{dM_k(Q^2)}{d \ln Q^2} = \frac{\alpha_s(Q^2)}{2\pi} P_k M_k(Q^2). \quad (\text{A.22})$$

where P_k are the moments of the splitting functions [139, 526, 573, 580, 678]. This formulation converts the complex integro-differential evolution equations into a set of simple ordinary differential equations, one for each moment independently.

Statistical moment. The k^{th} *moment* of random variable Z a probability density $p(z)$ is defined as

$$\mathbb{E}[Z^k] = \int_{-\infty}^{\infty} z^k p(z) dz. \quad (\text{A.23})$$

Moment generating function (MGF) The *moment generating function (MGF)* of a random variable Z , $M_Z(t) = \mathbb{E}[e^{tZ}]$, is defined as

$$M_Z(t) = \sum_{n=0}^{\infty} \frac{t^n}{n!} \mathbb{E}[Z^n] . \quad (\text{A.24})$$

The MGF encapsulates the entire moment sequence via its Taylor expansion, and serves as a proxy for the full distribution.

Effective sample size The *effective size* of a sample is the effective number of independent points, $N_{\text{eff}} = N/[1 + (N - 1)\rho]$ [25, 260, 301, 583] in that sample accounting for correlations.

Event independence in HEP In high energy physics cross section measurements, standard statistical treatments for unbinned inference rely on a set of core assumptions collectively referred to as *event independence*. These assumptions posit that interaction events are generated and observed as independent trials of a stochastic process, which greatly simplifies the construction of likelihood functions and inference procedures.

Pileup At particle physics experiments multiple events can occur in the same detector read-out window and be recorded together. This phenomenon is known as *pileup*. Pileup entails that what is recorded as a single ‘event’ may in fact be a composite of several independent interactions superimposed in the detector readout.

Poisson bootstrap The *Poisson bootstrap* is a variant of non-parametric bootstrap resampling in which instead of drawing exactly N samples each time, each

original event, x_i , is assigned a Poisson(1) random weight, w_{ij} in each replica, j [508, 578]. x_i is then included in replica j w_{ij} times.

i.i.d. *Independence and identical distribution* (i.i.d) is the assumption that each event can be treated as an independent draw from the same underlying probability density function $p(x \mid \theta)$, where x denotes the measured observables (kinematic variables, detector signals, *et cetera*) for that event. In other words, conditional on the physics model parameters θ , all events are statistically independent and governed by an identical distribution [263, 295, 299, 679].

Symmetry discovery *Symmetry discovery* is the process of identifying invariant transformations of a dataset, that is to say, operations under which the probability distribution of the dataset remains unchanged in a statistical sense [76, 138, 229].

Diffeomorphism invariant probability densities. Let $X \subseteq \mathbb{R}^n$ be an open set and Φ_G be a C^1 action of a group G on X . Let dx denote the Lebesgue measure on X . A probability density $p : X \rightarrow [0, \infty)$ is diffeomorphism invariant under the action of Φ_G if

$$\forall \text{ measurable } A \subseteq X, \forall g \in G, \mu_p(g \cdot A) = \mu_p(A), \quad (\text{A.25})$$

where

$$\mu_p(A) = \int_A p(x) dx. \quad (\text{A.26})$$

This definition implies that, for almost every $x \in X$ and for every $g \in G$,

$$p(g(x)) |\det Dg(x)| = p(x), \quad (\text{A.27})$$

where $Dg(x)$ is the Jacobian of the diffeomorphism $x \mapsto g(x)$.

Equivariant neural network. A neural network f_θ is said to be *equivariant* under the action group G on its input space X if

$$\forall g \in G \quad \forall x \in X \quad f_\theta(g \cdot x) = g \cdot f_\theta(x) . \quad (\text{A.28})$$

Group invariant function A function $f : X \rightarrow \mathbb{R}$ is *G-invariant* if there exists a function $\tilde{f} : X/G \rightarrow \mathbb{R}$ with $f = \tilde{f} \circ \pi$ where $\pi : X \rightarrow X/G$ is the canonical projection.

Invariant measure For every measure space (X, Σ, μ) and for every group G acting on X through a Σ -measurable action Φ , μ is called an *invariant measure under G* if $\forall (g, A) \in G \times \Sigma$

1. $g(A) \in \Sigma$.
2. $\mu(g(A)) = \mu(A)$

Equivalently, using the pushforward $g_\#(\mu(A)) := \mu(g^{-1}(A))$, μ is defined to be invariant under G if

$$\forall (A, g) \in \Sigma \times G \quad g_\#(\mu(A)) = \mu(A) . \quad (\text{A.29})$$

Reference density A *reference density* $\rho(x)$ is a probability density function that establishes a canonical measure for comparing probability distributions by enabling the definition of relative entropy

$$H_\rho[X] = -\mathbb{E}_\mu \left[\ln \frac{d\mu}{d\rho} \right] . \quad (\text{A.30})$$

Relative symmetry. A diffeomorphism g is defined to be a symmetry of a density p *relative* to p_I if both p and p_I are preserved.

$$p(g(x)) |\det Dg(x)| = p(x), \quad p_I(g(x)) |\det Dg(x)| = p_I(x). \quad (\text{A.31})$$

Symmetric action of a group [627]. Let $(\Omega, \mathcal{F}, \mu)$ be a probability space and G a group. G is defined to have a *symmetric action* on the measure space $(\Omega, \mathcal{F})$ if there exists a function

$$\phi : G \times \Omega \rightarrow \Omega, \quad (g, \omega) \mapsto g \cdot \omega \quad (\text{A.32})$$

such that

1.

$$\forall \omega \in \Omega \quad \phi(1, \omega) = \omega \quad (\text{A.33})$$

2.

$$\forall \omega \in \Omega \quad \forall g_1, g_2 \in G \quad \phi(g_1 g_2, \omega) = \phi(g_1(\phi(g_2, \omega))) \quad (\text{A.34})$$

3.

$$\forall g \in G \quad \forall A \in \mathcal{F} \quad \mu(g \cdot A) = \mu(A), \quad (\text{A.35})$$

Symmetry of a function In group theoretic terms, a *function* $f : X \rightarrow Y$ is *symmetric* under a group G there exists a group action of G on X

$$\begin{aligned} \Phi : G \times X &\rightarrow X \\ (g, x) &\mapsto g(x) \end{aligned} \quad (\text{A.36})$$

such that for every $x \in X$ and every $g \in G$, the value of the function $f(x) = f(g(x))$.

Symmetry of a probability density function If $p(x)$ denotes a probability density on X , a group G is a *symmetry of p* if

$$\forall g \in G \quad p(x) \, dx = p(g(x)) \, d(g(x)). \quad (\text{A.37})$$

Symmetry under a group. Let G be a group of transformations (continuous or discrete) acting on a space of states or observations X . A physical system or probability distribution is *symmetric under G* if applying any transformation $g \in G$ leaves the relevant observables unchanged.

G -Equivariant Fisher Information For a likelihood $\mathcal{L}(x \mid \theta)$ whose parameter space carries an action of a finite or Lie group G , the *G -equivariant Fisher information matrix* is

$$[\mathcal{I}_G(\theta)]_{ij} = \mathbb{E}\left[\frac{\partial \ln \mathcal{L}}{\partial \theta_i} \frac{\partial \ln \mathcal{L}}{\partial \theta_j}\right] - \sum_{g \in G} w_g \mathbb{E}\left[\frac{\partial \ln \mathcal{L}}{\partial \theta_i} \frac{\partial \ln \mathcal{L}(g \cdot \theta)}{\partial \theta_j}\right], \quad (\text{A.38})$$

where $\{w_g\}_{g \in G}$ are averaging weights satisfying $\sum_{g \in G} w_g = 1$. Equivalently, $\mathcal{I}_G(\theta) = P_G \mathcal{I}(\theta) P_G^T$ where P_G is the orthogonal projection operator onto the G -invariant subspace [342].

Appendix B

Theorems and proofs.

Theorem 1 (Scaling of relative uncertainty). Consider a d dimensional measurement space partitioned into histogram bins of widths $\Delta z_1, \dots, \Delta z_d$, and let N_{events} denote the total number of events uniformly distributed over the full measurement range. Then the relative uncertainty in the estimated differential event density in any given bin obeys

$$\text{Rel. uncertainty} = \frac{\sigma_Q}{Q} \propto \frac{1}{\sqrt{N_{\text{events}}} \prod_{i=1}^d \Delta z_i}, \quad (\text{B.1})$$

where Q is the differential quantity (events per unit hypervolume).

Proof of Theorem 1.

A single histogram bin occupies hypervolume

$$V = \prod_{i=1}^d \Delta z_i \quad (\text{B.2})$$

in the measurement space. Under the uniform–density assumption, the expected number of events in that bin scales as

$$N_{\text{bin}} \propto N_{\text{events}} V. \quad (\text{B.3})$$

By Poisson statistics the absolute uncertainty on the bin count is

$$\sigma_{N_{\text{bin}}} = \sqrt{N_{\text{bin}}}, \quad (\text{B.4})$$

and hence the relative uncertainty on N_{bin} is

$$\frac{\sigma_{N_{\text{bin}}}}{N_{\text{bin}}} = \frac{1}{\sqrt{N_{\text{bin}}}} \propto \frac{1}{\sqrt{N_{\text{events}} V}}. \quad (\text{B.5})$$

Since the differential measurement is $Q = N_{\text{bin}}/V$, its relative uncertainty coincides with that of the bin count.

$$\frac{\sigma_Q}{Q} = \frac{(\sigma_{N_{\text{bin}}}/V)}{(N_{\text{bin}}/V)} \quad (\text{B.6})$$

$$= \frac{\sigma_{N_{\text{bin}}}}{N_{\text{bin}}} \quad (\text{B.7})$$

$$\propto \frac{1}{\sqrt{N_{\text{events}} \prod_{i=1}^d \Delta z_i}}, \quad (\text{B.8})$$

which establishes the claimed scaling. $\square$

Theorem 2 (Asymptotic standard normality of pulls). Let $\hat{\theta}_n$ be an estimator of a scalar parameter θ based on n independent events, and let σ_n denote the reported (1σ) uncertainty for $\hat{\theta}_n$. Assume:

Asymptotic unbiasedness: $\mathbb{E}[\hat{\theta}_n] - \theta \rightarrow 0$ as $n \rightarrow \infty$.

Variance calibration: $\sigma_n^2 / \text{Var}(\hat{\theta}_n) \xrightarrow{p} 1$.

Asymptotic normality: $\sqrt{n}(\hat{\theta}_n - \theta) \Rightarrow \mathcal{N}(0, \tau^2)$ for some $\tau^2 > 0$.

For the *pull*, P_n defined according to Definition 17,

$$\lim_{n \rightarrow \infty} P_n \sim \mathcal{N}(0, 1). \quad (\text{B.9})$$

That is to say, across pseudo-experiments, the empirical distribution of pulls converges to the standard normal law.

Proof of Theorem 2. By asymptotic normality,

$$\lim_{n \rightarrow \infty} \frac{\hat{\theta}_n - \theta}{\sqrt{\text{Var}(\hat{\theta}_n)}} \sim \mathcal{N}(0, 1). \quad (\text{B.10})$$

By variance calibration,

$$\lim_{n \rightarrow \infty} \frac{\sqrt{\text{Var}(\hat{\theta}_n)}}{\sigma_n} = 1. \quad (\text{B.11})$$

Hence, by Slutsky's theorem [793],

$$\lim_{n \rightarrow \infty} P_n = \lim_{n \rightarrow \infty} \frac{\hat{\theta}_n - \theta}{\sqrt{\text{Var}(\hat{\theta}_n)}} \cdot \lim_{n \rightarrow \infty} \frac{\sqrt{\text{Var}(\hat{\theta}_n)}}{\sigma_n} \sim \mathcal{N}(0, 1) \cdot 1 = \mathcal{N}(0, 1) \quad (\text{B.12})$$

□

Theorem 3 (Effective sample size with correlated event weights). Let $w_1, \dots, w_N$ be event weights, and let $\rho(w_i, w_j)$ denote the correlation coefficient between the contributions of events i and j to a given observable. The *effective sample size* accounting for correlations defined in Equation (II.52) computes the ratio of the squared total weight to the variance of the weighted sum of N correlated unit-variance contributions.

Proof of Theorem 3. Consider the weighted sum $S := \sum_{i=1}^N w_i X_i$ with mean μ and variance 1 for all i . The variance of S is

$$\sigma_S^2 = \sum_{i=1}^N w_i^2 \sigma_i^2 + 2 \sum_{1 \leq i < j \leq N} w_i w_j \Sigma_{ij} \quad (\text{B.13})$$

$$= \sum_{i=1}^N w_i^2 + 2 \sum_{1 \leq i < j \leq N} w_i w_j \rho(w_i, w_j). \quad (\text{B.14})$$

The definition of effective sample size equates σ_S^2 to the variance of N_{eff} independent, equally weighted observations of variance 1.

$$\sigma_S^2 = \frac{\left(\sum_{i=1}^N w_i \right)^2}{N_{\text{eff}}}. \quad (\text{B.15})$$

Solving for N_{eff} yields Equation (II.52). $\square$

Theorem 4 (Universal Approximation Theorem).

$$\forall (\sigma, g, \epsilon) \in C^0(\mathbb{R}) \times C^0([0, 1]^n) \times (0, \infty) \exists M \in (0, \infty) \quad (\text{B.16})$$

$$\left[\forall x \in \mathbb{R} \ |\sigma(x)| < M \implies \exists N \in \mathbb{N} \exists (c, b, W) \in \mathbb{R}^N \times \mathbb{R}^N \times \mathbb{R}^{N \times n} \right. \\ \left. \{ \forall z \in [0, 1]^n \ \|c^T \sigma(Wz + b) - g(z)\| < \epsilon \} \vee \{ \exists A \in \mathbb{R} \ \forall x \in \mathbb{R} \ \sigma(x) = A \} \right].$$

Theorem 5 (Change of variables for pushforward densities). For all probability spaces $(\Omega, \mathcal{F}, \mathbb{P})$, for all random variables $Z : \Omega \rightarrow \mathbb{R}^d$ with measure law $\mu_Z \ll \lambda^d$ and density $p_Z = \frac{d\mu_Z}{d\lambda^d}$, for all open sets $U, V \subseteq \mathbb{R}^d$ and for all C^1 bijections with C^1 inverses (i.e. C^1 diffeomorphisms) $f : U \rightarrow V, X = f(Z)$ has measure law $\mu_X = f_{\#} \mu_Z$ only if $\mu_X \ll \lambda^d$ and its Radon–Nikodým derivative [10, 676] is, for almost every

$x \in \mathbb{R}^d$,

$$\frac{d\mu_X}{d\lambda^d}(x) = \begin{cases} p_Z(f^{-1}(x)) |\det Df^{-1}(x)| & x \in V, \\ 0 & x \notin V. \end{cases} \quad (\text{B.17})$$

Equivalently, if $x = f(z)$ then for almost every $z \in U$,

$$\frac{d\mu_X}{d\lambda^d}(f(z)) = \frac{p_Z(z)}{|\det Df(z)|} \quad (\text{B.18})$$

Proof of Theorem 5. Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be an arbitrary Borel function. By the definition of the pushforward,

$$\int_{\mathbb{R}^d} g(x) d\mu_X(x) = \int_{\Omega} g \circ f \circ Z(\omega) d\mathbb{P}(\omega) = \int_{\mathbb{R}^d} g \circ f(z) p_Z(z) d\lambda^d(z). \quad (\text{B.19})$$

Apply the classical change-of-variables theorem for C^1 diffeomorphisms (equivalently, the area formula for diffeomorphisms) to the integrand $\psi(z) := g(f(z)) p_Z(z)$.

$$\int_U \psi(z) d\lambda^d(z) = \int_V \psi \circ f^{-1}(x) |\det Df^{-1}(x)| d\lambda^d(x). \quad (\text{B.20})$$

Since p_Z is zero almost everywhere outside U , one may restrict the domain of integration to U . This yields

$$\int_{\mathbb{R}^d} g \circ (f(z) p_Z(z) d\lambda^d(z) = \int_V g(x) p_Z \circ f^{-1}(x) |\det Df^{-1}(x)| d\lambda^d(x). \quad (\text{B.21})$$

Combining Equations (III.42) and (III.44) yields

$$\int_{\mathbb{R}^d} g(x) d\mu_X(x) = \int_{\mathbb{R}^d} g(x) [\mathbb{1}_V(x) p_Z \circ f^{-1}(x) |\det Df^{-1}(x)|] d\lambda^d(x). \quad (\text{B.22})$$

Since the above equality holds for every bounded Borel function g , for almost every x the Radon–Nikodým theorem yields

$$\frac{d\mu_X}{d\lambda^d}(x) = \mathbb{1}_V(x) p_Z \circ f^{-1}(x) |\det Df^{-1}(x)|. \quad (\text{B.23})$$

as claimed. The equivalent form follows from

$$|\det Df^{-1}(f(z))| = \frac{1}{|\det Df(z)|}. \quad (\text{B.24})$$

□

Theorem 6 (Entropy projection with moment constraints). Under the assumptions above, the minimiser of $D_{\text{KL}}(\ell \| q)$ over densities all ℓ subject to $\mathbb{E}_\ell[Z^a] = \mathbb{E}_p[Z^a]$, $a = 0, 1, \dots, n$ (where $a = 0$ enforces normalisation) exists in the exponential family

$$\ell^*(z) = q(z) \exp\left(-1 - \sum_{a=0}^n \beta_a z^a\right), \quad (\text{B.25})$$

for some Lagrange multipliers $\{\beta_a\}_{a=0}^n$.

Proof of Theorem 6. Form the Lagrangian

$$\mathcal{L}(\ell, \beta) = D_{\text{KL}}(\ell \| q) + \sum_{a=0}^n \beta_a \left(\int z^a \ell(z) dz - \mathbb{E}_p[Z^a] \right). \quad (\text{B.26})$$

A standard variational calculation gives the Gâteaux derivative

$$\frac{\delta \mathcal{L}}{\delta \ell(z)} = \ln \frac{\ell(z)}{q(z)} + 1 + \sum_{a=0}^n \beta_a z^a. \quad (\text{B.27})$$

Setting this to zero yields Equation (V.16). Feasibility (i.e., choice of β) enforces the constraints. □

Theorem 7 (Optimal discriminator and reduced objective). For fixed g , the unique maximiser of $d \mapsto -L(d, g)$ is $d^*(z) = \frac{p(z)}{\tilde{q}(z)}$. At d^* ,

$$\tilde{L}(g) = L(d^*, g) = -D_{\text{KL}}(p \| \tilde{q}). \quad (\text{B.28})$$

Proof of Theorem 7. Varying d in Equation (V.20) gives

$$\frac{\delta L}{\delta d(z)} = -\frac{p(z)}{d(z) + \tilde{q}(z)}. \quad (\text{B.29})$$

Hence $d^*(z) = p(z)/\tilde{q}(z)$. Substituting and using $\int p = \int \tilde{q} = 1$ gives Equation (V.21). $\square$

Theorem 8 (Closure under perfect response). Any stationary point of $g \mapsto \tilde{L}(g)$ of the form Equation (V.19) satisfies

$$\forall a \in \{1, \dots, n\} \quad \mathbb{E}_{\tilde{q}}[Z^a] = \mathbb{E}_p[Z^a]. \quad (\text{B.30})$$

Proof of Theorem 8. From Equation (V.21),

$$\tilde{L}(g) = - \int p(z) \ln \tilde{q}(z) \, dz + \text{const}. \quad (\text{B.31})$$

Hence, the functional derivative with respect to g is

$$\frac{\delta \tilde{L}}{\delta g(z)} = \frac{p(z)}{g(z)}. \quad (\text{B.32})$$

For the parametrisation Equation (V.19),

$$\forall a \in \{1, \dots, n\} \quad \frac{\partial g(z)}{\partial \beta_a} = -g(z)(z^a - \mathbb{E}_{\tilde{q}}[Z^a]). \quad (\text{B.33})$$

Since $\frac{\partial \ln P}{\partial \beta_a} = -\mathbb{E}_{\tilde{q}}[Z^a]$, thus

$$\frac{\partial \tilde{L}}{\partial \beta_a} = \int \frac{\delta \tilde{L}}{\delta g(z)} \frac{\partial g(z)}{\partial \beta_a} \, dz \quad (\text{B.34})$$

$$= - \int p(z) \left(z^a - \mathbb{E}_{\tilde{q}}[Z^a] \right) \, dz \quad (\text{B.35})$$

$$= - \left(\mathbb{E}_p[Z^a] - \mathbb{E}_{\tilde{q}}[Z^a] \right). \quad (\text{B.36})$$

At any stationary point, this vanishes for $a = 1, \dots, n$, proving Equation (V.23). $\square$

Theorem 9 (Reduced objective at detector level). For the discriminator acting on x , the discriminator optimised MLC loss is

$$\tilde{L}(g) = -D_{\text{KL}}(p(x) \parallel \tilde{q}(x)). \quad (\text{B.37})$$

Proof of Theorem 9. Identical to Theorem 7 after replacing z by x and using Equation (V.30). $\square$

Theorem 10 (Moment matching). Under the assumptions above, any stationary point of $g \mapsto \tilde{L}(g)$ in the class Equation (V.19) satisfies

$$\forall a \in \{1, \dots, n\} \quad \mathbb{E}_{\tilde{q}}[Z^a] = \mathbb{E}_{p_{\text{mod}}}[Z^a], \quad (\text{B.38})$$

where the modified particle level density p_{mod} is

$$p_{\text{mod}}(z) = \int \frac{p(x) r(x \mid z) \tilde{q}(z)}{\tilde{q}(x)} dx. \quad (\text{B.39})$$

Proof of Theorem 10. By Equation (V.31) and the chain rule for functional derivatives,

$$\frac{\partial \tilde{L}}{\partial \beta_a} = \iint \frac{\delta \tilde{L}}{\delta \tilde{q}(x)} \frac{\delta \tilde{q}(x)}{\delta g(z)} \frac{\partial g(z)}{\partial \beta_a} dz dx. \quad (\text{B.40})$$

$$\frac{\delta \tilde{L}}{\delta \tilde{q}(x)} = \frac{p(x)}{\tilde{q}(x)} \quad (\text{B.41})$$

and

$$\frac{\delta \tilde{q}(x)}{\delta g(z)} = r(x \mid z) q(z). \quad (\text{B.42})$$

Using Equation (V.26) and $\tilde{q}(z) = q(z) g(z)$ gives

$$\frac{\partial \tilde{L}}{\partial \beta_a} = - \int \left(z^a - \mathbb{E}_{\tilde{q}}[Z^a] \right) \left(\int \frac{p(x)}{\tilde{q}(x)} r(x \mid z) dx \right) \tilde{q}(z) dz \quad (\text{B.43})$$

$$= - \left(\mathbb{E}_{p_{\text{mod}}}[Z^a] - \mathbb{E}_{\tilde{q}}[Z^a] \right), \quad (\text{B.44})$$

where p_{mod} is as in Equation (V.33). Stationarity enforces Equation (V.32). $\square$

Theorem 11 (Transfer operator representation). Let

$$\tilde{q}(z | x) := \frac{r(x | z) \tilde{q}(z)}{\tilde{q}(x)}. \quad (\text{B.45})$$

Define the transfer kernel

$$f(z | z') := \int \tilde{q}(z | x) r(x | z') dx. \quad (\text{B.46})$$

Then

$$p_{\text{mod}}(z) = \int f(z | z') p(z') dz'. \quad (\text{B.47})$$

Moreover, if $r(x | z) = \delta(x - z)$ (perfect detector), then $f(z | z') = \delta(z - z')$ and $p_{\text{mod}} = p$.

Proof of Theorem 11. Substitute $p(x) = \int r(x | z') p(z') dz'$ into Equation (V.33) and exchange integrals

$$p_{\text{mod}}(z) = \iint \frac{r(x | z) \tilde{q}(z)}{\tilde{q}(x)} r(x | z') dx p(z') dz' \quad (\text{B.48})$$

$$= \int \left(\int \tilde{q}(z | x) r(x | z') dx \right) p(z') dz'. \quad (\text{B.49})$$

If $r(x | z) = \delta(x - z)$, then $\tilde{q}(z | x) = \delta(z - x)$ and Equation (V.40) gives $f(z | z') = \delta(z - z')$. $\square$

Theorem 12 (Moment representation of a random variable [604, 722, 794, 800, 802]).

Let λ denote Lebesgue measure on R and $\mathcal{P}(\mathbb{R})$ the set of Borel probability measures on $\mathbb{R}$. Then

$$\begin{aligned}
 & \forall m : \mathbb{N} \rightarrow \mathbb{R} \left[m(0) = 1 \implies \forall \mu \in \mathcal{P}(\mathbb{R}) \right. \\
 & \quad \left. \left\{ \mu \ll \lambda \wedge \left(\forall n \in \mathbb{N} \int_{\mathbb{R}} z^n d\mu(z) = m(n) \right) \right. \right. \\
 & \quad \wedge \left(\exists \text{ open } I \subseteq \mathbb{R} \wedge 0 \in I \forall t \in I \int_{\mathbb{R}} e^{tz} d\mu(z) < \infty \right) \\
 & \quad \left. \implies \forall (\nu, n) \in \mathcal{P}(\mathbb{R}) \times \mathbb{N} \int_{\mathbb{R}} z^n d\nu(z) = m(n) \implies \nu = \mu \right\} \Bigg].
 \end{aligned} \tag{B.50}$$

Equivalently, writing $d\mu = f dz$ and $d\nu = g dz$ when they exist,

$$\begin{aligned}
 & \forall n \in \mathbb{N} \int_{\mathbb{R}} z^n f(z) dz = \int_{\mathbb{R}} z^n g(z) dz \wedge \exists \text{ open } I \ni 0 \forall t \in I \\
 & \quad \int_{\mathbb{R}} e^{tz} f(z) dz < \infty \implies f = g \text{ almost everywhere}
 \end{aligned} \tag{B.51}$$

Theorem 13 (Equicorrelated sampling error and effective sample size). For every $N \in \mathbb{N}_{\geq 2}$, and for every set of real valued random variables $\{X_i\}_{i=1}^N$, if:

$$\begin{aligned}
 & \exists (\mu, \sigma, \rho) \in \mathbb{R} \times (0, \infty) \times \left[-\frac{1}{N-1}, 1 \right] \forall i \in \mathbb{N}_{\leq N} \mathbb{E}(X_i) = \mu \\
 & \quad \wedge \forall i, j \in \mathbb{N}_{\leq N}^2 \Sigma_{ij} = \sigma^2[(1 - \rho)\delta_{ij} + \rho]
 \end{aligned} \tag{B.52}$$

i.e.

$$\Sigma_{ij} = \begin{cases} \sigma^2, & i = j, \\ \rho \sigma^2, & i \neq j. \end{cases} \tag{B.53}$$

Then:

1. $\Sigma \succeq 0$ iff $\Sigma \in \left[-\frac{1}{N-1}, 1 \right]$ and $R(\rho) \succ 0$ iff $\rho \in \left(-\frac{1}{N-1}, 1 \right)$.

2. For the sample mean $\bar{X} \equiv \frac{1}{N} \sum_{i=1}^N X_i$,

$$\text{Err}(\bar{X}) \equiv \sqrt{\text{Var}(\bar{X})} = \sigma \sqrt{\frac{1}{N} + \left(1 - \frac{1}{N}\right)\rho}. \quad (\text{B.54})$$

Proof of Theorem 13. (1) The matrix Σ has eigenvalue $\sigma^2(1 + (N - 1)\rho)$ with eigenvector $\begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix}^T$ and eigenvalue $\sigma^2(1 - \rho)$ with multiplicity $N - 1$. Hence $\Sigma \succeq 0$ iff $\sigma^2(1 - \rho \geq 0)$ and $\sigma^2(1 + (N - 1)\rho \geq 0)$, i.e. $\rho \in [-1/(N - 1), 1]$, and $R(\rho) \succ 0$ iff both inequalities are strict.

(2) Using $\bar{X} = N^{-1} \sum_i X_i$,

$$\text{Var}(\bar{X}) = \frac{1}{N^2} \sum_{i,j=1}^N \Sigma_{ij} = \frac{N\sigma^2}{N^2} (1 + (N - 1)\rho) = \sigma^2 \left(\frac{1}{N} + \left(1 - \frac{1}{N}\right)\rho \right), \quad (\text{B.55})$$

which yields the stated $\text{Err}(\bar{X})$. $\square$

Theorem 14 (Asymptotic Normality of the MLE [39, 83, 84, 194, 277, 441, 570]).

Consider a set of i.i.d. observations $\{X_j\}_{j=1}^N$ and a model parametrised by parameter $\theta \in \Omega \subseteq \mathbb{R}^d$. The likelihood is

$$\mathcal{L}(\theta) = \prod_{i=1}^N p(X_j | \theta), \quad (\text{B.56})$$

and the Fisher information matrix is

$$I(\theta) = \mathbb{E}_\theta \left[\nabla_\theta^2 (-\ln \mathcal{L}(\theta)) \right]. \quad (\text{B.57})$$

Under regularity conditions ensuring consistency and differentiability, the maximum likelihood estimator is distributed as

$$\lim_{N \rightarrow \infty} \hat{\theta} \sim \mathcal{N}(\theta_{\text{true}}, I(\theta_{\text{true}})^{-1}) \quad (\text{B.58})$$

Theorem 15 (Wilks [788]). Let $\{X_j\}_{j=1}^N$ be a set of i.i.d measurements with likelihood

$$L(\theta) = \prod_{i=1}^N p(X_i \mid \theta), \quad \theta \in \Omega \subseteq \mathbb{R}^d. \quad (\text{B.59})$$

Let the null hypothesis be

$$H_0 : \theta \in \Theta_0, \quad \dim \Theta_0 = r, \quad (\text{B.60})$$

and the alternative is $H_1 : \theta \in \Theta_1$, with $\Omega = \Theta_0 \cup \Theta_1$. The generalised likelihood ratio statistic is

$$\mathcal{L} = \frac{\sup_{\theta \in \Omega} L(\theta)}{\sup_{\theta \in \Theta_0} L(\theta)}. \quad (\text{B.61})$$

Under standard regularity conditions and assuming H_0 is true,

$$\lim_{N \rightarrow \infty} -2 \ln \mathcal{L} \sim \chi_{d-r}^2 \quad (\text{B.62})$$

Theorem 16 (Godambe Information [716]). Let $\{X_j\}_{j=1}^N$ be i.i.d. observations modelled by $p(x \mid \theta)$, $\theta \in \Omega \subseteq \mathbb{R}^d$. Consider an estimating function

$$U(\theta) = \sum_{j=1}^N u(X_j \mid \theta), \quad (\text{B.63})$$

where

$$\forall \theta \in \Omega \quad \mathbb{E}_\theta [u(X_j \mid \theta)] = 0 \quad (\text{B.64})$$

such that standard regularity conditions hold ensuring consistency and asymptotic normality of the root $\hat{\theta}$ solving $U(\theta) = 0$. Let

$$H(\theta) = -\mathbb{E}_\theta \left[\frac{\partial u(X_j \mid \theta)}{\partial \theta} \right] \quad (\text{the sensitivity matrix}), \text{ and} \quad (\text{B.65})$$

$$J(\theta) = \text{Var}_\theta [u(X_j \mid \theta)] \quad (\text{the variability matrix}). \quad (\text{B.66})$$

Then the *Godambe information matrix* can be computed as,

$$G(\theta) = H(\theta)^T J(\theta)^{-1} H(\theta). \quad (\text{B.67})$$

and

$$\lim_{N \rightarrow \infty} \hat{\theta} \sim \mathcal{N} \left(\theta, \frac{1}{N} G(\theta)^{-1} \right). \quad (\text{B.68})$$

That is to say, the asymptotic covariance of $\hat{\theta}$ is $N^{-1}G(\theta)^{-1}$.

Theorem 17 (Extension of Theorem 16 [381, 520, 582, 778]). If the estimating function is the score of a correctly specified likelihood,

$$u(X_j \mid \theta) = -\nabla_{\theta} \ln p(X_j \mid \theta), \quad (\text{B.69})$$

then

$$H = -\mathbb{E}[\nabla^2 \ln \mathcal{L}], \quad (\text{B.70})$$

$$J = \text{Var}(\nabla \ln \mathcal{L}), \quad (\text{B.71})$$

and the sandwich (or Godambe) information matrix is given by

$$G(\hat{\theta}) = H(\hat{\theta}) J(\hat{\theta})^{-1} H(\hat{\theta}), \quad (\text{B.72})$$

Theorem 18 (Bartlett correction [686, 775]). Let $\{X_j\}_{j=1}^N$ be i.i.d. according a model $p(x \mid \theta)$ and let the hypotheses be

$$H_0 : \theta \in \Theta_0 \quad H_1 : \theta \in \Theta_1 \quad \Omega = \Theta_0 \cup \Theta_1, \quad (\text{B.73})$$

with $\dim \Omega = d$, $\dim \Theta_0 = r$, and the usual likelihood-ratio statistic

$$\mathcal{L} = \frac{\sup_{\theta \in \Theta_0} L(\theta)}{\sup_{\theta \in \Omega} L(\theta)}. \quad (\text{B.74})$$

Under standard regularity conditions and assuming H_0 is true,

$$\exists c \in \mathbb{R} \quad c \frac{\cdot \mathbb{E}_{\Theta_0}[-\ln \mathcal{L}]}{d - r} = 1 + \frac{B}{N} + \mathcal{O}(N^{-2}), \quad (\text{B.75})$$

with B determined by lower order cumulants of the log-likelihood.¹

¹B can be worked out via the third and fourth derivatives of the log-likelihood [774].

Theorem 19 (Asymptotics of the Bartlett correction). The Bartlett corrected statistic

$$\tilde{\mathcal{L}} = c \cdot \mathcal{L} \quad (\text{B.76})$$

satisfies

$$\mathbb{E}_{\Theta_0}[\tilde{\mathcal{L}}] = d - r \quad \text{and} \quad \lim_{N \rightarrow \infty} \tilde{\mathcal{L}} \sim \chi_{d-r}^2. \quad (\text{B.77})$$

Theorem 20 (Iwasawa Decomposition of the Affine Group [781]). Let

$$\text{Aff}(n) = GL(n, \mathbb{R}) \ltimes \mathbb{R}^n \quad (\text{B.78})$$

be the real affine group. Define the subgroups

$$\begin{aligned} Z &= \{\pm I_n\} \cong \mathbb{Z}_2, \\ K &= SO(n), \\ A &= \{D(\mathbf{r}) = \text{diag}(r_1, \dots, r_n) \mid r_i \in (0, \infty)\}, \\ N &= \{S(u) \mid S_{ii} = 1, S_{i < j} \in \mathbb{R}, S_{i > j} = 0\}, \\ V &= \{\mathbf{v} \mid v_i \in \mathbb{R}\} \cong \mathbb{R}^n. \end{aligned} \quad (\text{B.79})$$

Then every element $M \in \text{Aff}_n(\mathbb{R})$ admits a unique factorisation

$$M = \sqrt{|\det M|} (\pm I_n)^{\frac{1 - \text{sgn}(\det M)}{2}} R(\theta) D(\mathbf{r}) S(u) + \mathbf{v}, \quad (\text{B.80})$$

where $\pm I_n \in Z$, $R(\theta) \in K$, $D(\mathbf{r}) \in A$, $S(u) \in N$, $\mathbf{v} \in V$.

Theorem 21 (Symmetry-Adapted Cramér–Rao Bound [705, 761, 782–785]). Let

$\hat{\theta}$ be any unbiased estimator of $\theta \in \mathbb{R}^d$ based on N i.i.d. samples, and assume the statistical model $p(x \mid \theta)$ is G -equivariant. Then the covariance matrix obeys the

lower bound

$$\Sigma(\hat{\theta}) \succeq \frac{1}{N} \mathcal{I}_G^{-1}(\theta), \quad (\text{B.81})$$

where $\succeq$ denotes the Loewner partial order. The lower bound is saturated if and only if the score vectors projected onto the invariant subspace are linear in $(\hat{\theta} - \theta)$.

Proof of Theorem 21. (Sketch) The classical information inequality states that $\Sigma(\hat{\theta}) \succeq N^{-1} \mathcal{I}^{-1}(\theta)$. Because the model and the estimator respect the group action, both sides commute with P_G . Projecting the score function onto the invariant subspace and using $P_G^2 = P_G$ yields $\text{Cov}(\hat{\theta}) \succeq N^{-1} (P_G \mathcal{I} P_G)^{-1} = N^{-1} \mathcal{I}_G^{-1}$. $\square$

Theorem 22 (Variance Reduction for Finite Symmetries). For a discrete group G of order $|G|$ if

$$\rho = \frac{\mathbb{E} \left[\frac{\partial \ln \mathcal{L}}{\partial \theta} \frac{\partial \ln \mathcal{L} \circ g}{\partial \theta} \right]}{\mathbb{E} \left[\left(\frac{\partial \ln \mathcal{L}}{\partial \theta} \right)^2 \right]}; \quad (\text{B.82})$$

for the average pairwise correlation of score functions across group orbits, then

$$\frac{\text{Var}(\hat{\theta}_{\text{sym}})}{\text{Var}(\hat{\theta}_{\text{unconstrained}})}, = \frac{1}{|G|} + \frac{|G| - 1}{|G|} \rho. \quad (\text{B.83})$$

This means that symmetry is maximally beneficial ($\rho \rightarrow 0$) when orbit scores are uncorrelated, recovering the $1/|G|$ ideal reduction factor.