



Unbinned inference with correlated events

Krish Desai^{1,2,a} , Owen Long^{3,b} , Benjamin Nachman^{1,4,c}

¹ Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

² Department of Physics, University of California, Berkeley, CA 94720, USA

³ Department of Physics and Astronomy, University of California, Riverside, CA 92521, USA

⁴ Berkeley Institute for Data Science, University of California, Berkeley, CA 94720, USA

Received: 9 May 2025 / Accepted: 19 September 2025

© The Author(s) 2025

Abstract Modern machine learning has enabled parameter inference from event-level data without the need to first summarize all events with a histogram. All of these unbinned inference methods make use of the fact that the events are statistically independent so that the log likelihood is a sum over events. However, this assumption is not valid for unbinned inference on unfolded data, where the deconvolution process induces a correlation between events. We explore the impact of event correlations on downstream inference tasks in the context of the OmniFold unbinned unfolding method. We find that uncertainties may be significantly underestimated when event correlations are excluded from uncertainty quantification.

Contents

1	Introduction	
2	Methodology	
2.1	Unbinned unfolding	
2.2	Unbinned inference	
2.3	Implementation	
3	Setup	
3.1	Data	
3.2	Neural networks	
4	Numerical results	
4.1	Binned baseline	
4.1.1	Binned unfolding with IBU and inference	
4.2	Event-by-event correlations	
4.3	Parameter estimation in 1D with a KDE	
4.4	Parameter estimation in 1D with a NN	

4.5	Parameter estimation beyond 1D	
5	Conclusions and outlook	
Appendix A: Parameter values for the 2D, 4D, and 6D studies		
Appendix B: Activation function dependence of event correlations		
Appendix C: Ensembling and inference resolution from NN model parameter initialization		
References		

1 Introduction

Accurate and precise parameter estimation is a central task in high-energy physics data analysis. Traditional approaches usually involve summarizing high-dimensional, event-level data into a small number of observables, whose differential cross sections are approximated with histograms. Best fit parameters and confidence intervals are then extracted using first-principles-based or Monte Carlo simulation-based templates using binned likelihood approaches. While effective, these methods can suffer from significant information loss and potential biases due to the dimensionality reduction and the binning. Furthermore, classical methods cannot be employed to study the multidifferential cross section of many observables simultaneously due to the increasing difficulty of binning effectively as the number of dimensions increases. Even if a simultaneous binned measurement was performed, it is not feasible to extract the differential cross section of lower-dimensional observables computed from the higher-dimensional phase space.

The advent of modern machine learning techniques has enabled unbinned inference methods that operate directly on event-level data without the need for dimensionality reduction or histogramming [1]. These methods exploit the full

^a e-mail: krish.desai@berkeley.edu

^b e-mail: owenl@ucr.edu (corresponding author)

^c e-mail: bpnachman@lbl.gov

information content of the data, leading to improved sensitivity and reduced biases in parameter estimation. A key assumption underlying these unbinned methods is the statistical independence of events, which allows the joint likelihood to be expressed as a product over individual event likelihoods, or equivalently, the log-likelihood as a sum over events:

$$\ln \mathcal{L}(\theta) = \sum_{i=1}^N \ln p(x_i | \theta), \quad (1.1)$$

where x_i denotes the observed data for event i , θ represents the parameters of interest, and $p(x_i | \theta)$ is the probability density function for event i given θ . The widely-used asymptotic formulae [2] for likelihood-based statistics apply also to the unbinned case and can significantly accelerate the calculation of confidence intervals and p -values.

Such methods have been effectively applied to detector-level data [3], where the assumption of the statistical independence of events holds. However, this assumption of event independence does not hold in the context of unfolded data. Unfolding is the process of statistically removing detector distortions to produce cross-section measurements at the particle-level. These results are independent of the detector and can thus be analyzed outside of experimental collaborations. Unfolding is an ill-posed inverse problem that requires regularization techniques [4–7] and introduces correlations between events due to the deconvolution process. These correlations violate the independence assumption. Nonetheless, unfolded data have been used extensively to measure fundamental parameters, fit parton distribution functions [8], tune parton shower Monte Carlo simulators [9], and search for physics beyond the Standard Model [10]. All of these studies use binned data, where correlations can be properly accounted with a covariance matrix. The usual Gaussian limit of Eq. (1.1) results in a χ^2 test statistic. With bin-by-bin correlations, the statistic most often used is

$$\chi_{\text{full}}^2(\theta) = (D - P(\theta))^T \Sigma^{-1} (D - P(\theta)), \quad (1.2)$$

where D is the vector of unfolded data counts, $P(\theta)$ is the vector of predicted counts from the model at parameters θ , and Σ is the covariance matrix encoding statistical and systematic uncertainties. As with Eq. (1.1), asymptotic formulae apply to Eq. (1.2) for deriving confidence intervals and p -values.

Recent innovations from machine learning have enabled unbinned unfolding [11–13]. These methods are based on discriminative [14–16] or generative methods [17–28] and there are now also experimental results with some of these approaches [29–38]. One can bin the unbinned results and use Eq. (1.2) for parameter estimation, but we are not aware of an analog to Eq. (1.1) that accounts for correlations. The goal

of this paper is explore how unbinned unfolded results can be used for parameter estimation. While there are a growing number of unbinned experimental results, the methods and publications are sufficiently new that there are currently no parameter extractions using these data. Our study will establish some best practices for using unbinned unfolded data for parameter estimation as well as new research questions for improving these protocols in the future.

In particular, we are interested in a number of questions: (1) is there an advantage in precision for binning after the unfolding instead of before? (2) is there any effect on the inference precision if correlations are ignored when creating the test statistic? and (3) if we ignore correlations and use Eq. (1.1) for the unbinned data, are the asymptotic uncertainties still valid? For (1), we will compare classical binned unfolding with an unbinned approach that is binned post-hoc for inference. To address (2), we have to use binned inference since we do not know the unbinned analog of Eq. (1.2). Lastly, we will compare asymptotic and numerical uncertainties to answer (3). Reference [14] found that binning before or after unfolding gave similar accuracy, but a study of precision for downstream parameter estimation has not yet been addressed. Since the detector distortions are treated unbinned, we hypothesized that unbinned unfolding could be more precise than traditional, binned unfolding. Furthermore, we expected that using correlations in the test statistic would lead to more precise inference than ignoring correlations in the test statistic. Lastly, we hypothesized that the asymptotic formulae would not give correct uncertainties when correlations are ignored (and correct ones when they are included when possible). For a series of Gaussian examples, we will confront each of these hypotheses with empirical evidence in order to address our main questions above.

The remainder of this paper is organized as follows. In Sect. 2, we detail the methodology of unbinned unfolding and inference. Section 3 outlines the experimental setup, including the construction of simulated datasets. Numerical results for one and many dimensions are presented in Sect. 4. The paper ends with conclusions and outlook in Sect. 5.

2 Methodology

2.1 Unbinned unfolding

We studied unbinned unfolding using the OmniFold [14, 15] method, as it is the one that has been used in most recent unbinned measurements. The task is to find the true distribution given observed data from a detector that includes non-negligible resolution or smearing effects. An essential requirement is to have an accurate detector simulation that maps the generated particle level observables (gen particle) to

the output of the detector (sim particle). An a-priori estimate of the true distribution is used to generate the gen-particle distribution in the Monte Carlo. The a-priori estimate is necessary because the true distribution is what we would like to measure using the data and it is either not known in advance, or the goal is to improve upon prior measurements, which are the basis for the a-priori estimate, with the data that are the input to the unfolding. The OmniFold method finds the Monte Carlo event weights necessary to reweight the gen-particle distribution, from the a-priori estimate, to the true distribution using an iterative process.

Each OmniFold iteration consists of two steps. The first step is to find the Monte Carlo event (“pull”) weights that will make the sim-particle distribution match the detector data. The desired unfolding output is a function that gives an event weight from gen-particle inputs, not sim-particle inputs. The second step of OmniFold provides this by finding Monte Carlo event (“push”) weights that will make the gen-particle distribution match the pull-weighted gen-particle distribution. The push-weighted Monte Carlo sim-particle distribution can then be compared to the detector data distribution. If they agree, the task is complete. If they disagree, the next iteration repeats the procedure starting with the push-weighted Monte Carlo sample, where the product of the push weights of all previous iterations are retained as event weights. The step weights in each iteration become closer to one as the process converges, usually after about five iterations. The final output is the weighted generation, i.e. particle level Monte Carlo sample, where each event is weighted by the product of the push weights for all of the iterations.

A detailed discussion of unbinned unfolding and the optimization of hyperparameters, such as the number of iterations, can be found in reference [39].

2.2 Unbinned inference

Inference of model parameters from unbinned data is commonly done using a maximum likelihood fit. Asymptotic uncertainties on the model parameters are determined from finding the parameter values that give a change in log likelihood of 0.5, corresponding to one standard deviation in the Gaussian approximation (e.g. the likelihood is locally quadratic near the maximum). Methods exist [40] for taking into account event weights in the evaluation of the asymptotic uncertainties if the weights are statistically independent. Reference [40] suggests a method for taking into account weight correlations if they can be parameterized. For our case, there is no simple parameterization, but we leave further explorations of related strategies to future work. In our study, we fit the output of the OmniFold unfolding, which is the weighted Monte Carlo gen-particle distribution. The OmniFold event weights are not statistically independent. We study how these

weight correlations impact the evaluation of the asymptotic uncertainties from the maximum likelihood fit.

Statistically valid inference uncertainties, even when analyzing unfolding output events with weight correlations, can be obtained numerically, either by using the bootstrap method [41], or from statistically independent pseudo datasets. In the work presented here, we use the numerical evaluation of the inference uncertainties as a valid reference for our studies of the asymptotic inference uncertainties. The numerical uncertainty is defined as the root-mean-square (RMS) of the distribution of inference results from either a collection of statistically independent pseudo datasets or a set of bootstrap samples derived from the dataset being unfolded and analyzed. We use “RMS” as a useful shorthand for the numerical uncertainty in the text and figures below.

2.3 Implementation

We consider both a one dimensional and multi-dimensional inference task described in the next section. For the one-dimensional study, we use kernel density estimation (KDE) with a uniform kernel over a fixed radius to estimate the probability densities that are used within OmniFold. The estimated probability density function (PDF) p at point x in the feature space is proportional to the number of points within δx of x . The pull weight in step 1 for a Monte Carlo event at x_s is $w_1 = p_d(x_s)/p_s(x_s)$, where p_d and p_s are the PDFs for the data and sim-particle Monte Carlo, respectively. Likewise, the push weight for step 2 is $w_2 = p_g(x_g, w_1)/p_g(x_g)$, where $p_g(x_g, w_1)$ is the gen-particle PDF evaluated with the weights w_1 from step 1.

In higher dimensions, we instead train neural network binary classifiers to find the weight functions used in OmniFold. In step 1, a classifier is trained to distinguish data from sim-particle Monte Carlo and the pull weight is $w_1 = c_1(x_s)/(1 - c_1(x_s))$, where $c_1(x_s)$ is the classifier output, which is the probability that an event at x_s is from the data distribution. In step 2, another classifier is trained to distinguish gen-particle Monte Carlo weighted by w_1 from gen-particle Monte Carlo and the push weight is $w_2 = c_2(x_g)/(1 - c_2(x_g))$, where $c_2(x_g)$ is the classifier output.

3 Setup

To demonstrate parameter estimation in addition to unfolding, we need data generated from a model with a known parametric family. In order to run enough (pseudo)experiments to test the coverage of the inference (in addition to research and development), we also need to be able to generate many examples from the model. Therefore, we focus on Gaussian examples. Many physics problems of interest share many fea-

tures with this setup; at the same time, we expect the sign and magnitude of the features we observe to be problem-specific.

3.1 Data

In our one-dimensional studies, the true model is a Gaussian with parameters $\mu_{\text{true}} = 0.2$ and $\sigma_{\text{true}}^2 = 0.81$. The parameter values used in the multidimensional studies are given in Appendix A. 10,000 events are generated for the “truth” dataset. A larger Monte Carlo (MC) sample of size 100,000 events is also generated from a Gaussian with mean 0.0 and variance 1.0. In traditional unfolding, this MC dataset would be the one to derive the response matrix. It also serves as the initialization for the unfolded result. Both the true and ‘particle-level’ MC sample are passed through the same detector distortions: a Gaussian smearing with mean zero and various resolutions (i.e. detector distortions) ranging from 0 to 0.75. For each parameter set, datasets are generated and binned into a fixed number of bins (15 bins spanning the range $[-3, 4]$). The true mean is slightly positive, so the range included in the binning is larger on the positive side. In the unbinned cases, the result is unbinned and presented as a histogram for illustration and comparison purposes.

For each parameter set that we study, we generate 500 statistically independent pseudo datasets where the number of events in each dataset is chosen from a Poisson distribution with a mean of 10,000. Figure 1 shows example distributions. In a real application, where there is only one dataset being analyzed, the bootstrapping method [41] could be used to determine the statistical uncertainties numerically. We have verified that bootstrapping and independent pseudo datasets give the same results for the statistical uncertainties, as evaluated by the RMS of the inference results of 500 datasets.

3.2 Neural networks

We use the Tensorflow 2.17.0 [42] and Keras 3.4.1 frameworks to implement and train our Neural Networks (NNs). The NN consists of three fully connected layers with 50 nodes per layer. In our nominal results, we use a modified version of the `gelu` [43] activation function, though we have also explored using `relu` activation. The `gelu` function is modified as

$$\text{GELU}(x) \approx 0.5 x \left(1 + \tanh \left[\beta \left(x + 0.044715 x^3 \right) \right] \right)$$

where β is a parameter that tunes the distance scale of the transition from flat to a slope of 1. With $\beta = 1$, the function is the original `gelu`, which gives a high degree of regularization and a larger distance scale for the output event weight correlations. With $\beta = 4$, the regularization is more modest and the correlations are closer to the output with `relu`, which is the asymptotic case for large β . More details on our choice

of activation function are given in Appendix B. The NNs are trained using the Adam optimizer [44] with a learning rate of 5×10^{-4} . The loss function is binary cross entropy. Data are equally divided into train and test samples for monitoring. The NNs are trained for 50 epochs with a batch size of 10^4 . We remove the output dependence on the random initial NN model weights, which are set with the `GolrotUniform` method [45], by averaging the final event weights of 10 runs of the full OmniFold procedure, each with different random initial model weights. The effects of this ensembling [46] are described in more detail in Appendix C.

These hyperparameter values were chosen to give stable results with reasonable convergence. When exploring the hyperparameter space, we monitored several aspects of the unfolding, including:

- The mean and RMS of the step 1 and step 2 weight distributions for each OmniFold iteration. Satisfactory convergence is indicated by a mean weight of 1 and a narrow weight distribution.
- The residuals of the unfolded distribution, compared to the known true distribution after each iteration for both the gen-particle and sim-particle distributions in all feature dimensions.
- The mean and width of the distribution of the difference between the OmniFold event weight and the known true event weight. Since the multivariate normal parameters are known in our study for both the Monte Carlo and true distribution, the true event weight can be calculated for each event. For the 1D study, this can be visualized as a simple graph. Example graphs of the event weight function can be found in Appendix B.

Examples of the monitoring output can be found in the code repository for this paper.

As an alternative to NNs in the one-dimensional studies, we also use KDEs with a fixed kernel of fixed radius 0.1.

4 Numerical results

We start with the one-dimensional example and first establish a binned baseline. Then, we consider the unbinned case (also in one dimension), by first characterizing the size of the correlations. Lastly, we explore the higher-dimensional inference task (which has no binned baseline).

4.1 Binned baseline

4.1.1 Binned unfolding with IBU and inference

For comparison with unbinned approaches, we establish a baseline in which the data are binned from the start. We per-




Fig. 1 Histograms showing the datasets from the one-dimensional study. The distributions on the left show the true distribution of the observable compared with the gen-particle distribution of the Monte Carlo. The center histograms show a sample detector-level distribution

of the observable compared with the sim-particle distribution of the Monte Carlo. The right histogram shows the resolution function. For these histograms, we used 31 bins in the range $[-5, 5]$

form Iterative Bayesian Unfolding (IBU) [47] (also known as Lucy-Richardson deconvolution [48, 49]) on the binned data, which is similar to OmniFold in the binned limit. Both IBU and OmniFold are iterative processes that converge to the maximum likelihood estimate of the particle-level distribution given the detector-level data (see the Appendix of [14]). The unfolded binned distribution is fit to a binned Gaussian using a χ^2 statistic. In our inference, two variants are considered: one that employs the full covariance matrix computed from 500 bootstrap replicas of the unfolded data, and another that uses only the diagonal elements of the covariance matrix. The covariance is estimated numerically from the ensemble of bootstrapped histograms. Parameter estimation is performed by minimizing the χ^2 , and the uncertainties are determined either by the asymptotic change in χ^2 (with $\Delta\chi^2 = 1$) or numerically from the spread (standard deviation or RMS) of the inferred parameters over the bootstrap ensemble.

This binned baseline procedure provides a benchmark against which the performance of the unbinned unfolding (followed by either post-hoc binning for inference, or unbinned inference) can be compared. For each parameter, two sets of asymptotic analyses are performed: one obtained from the fit using the full covariance matrix (“cov”) and one using only the diagonal approximation (“cov_diag”). We then calculate the mean asymptotic error and its standard error (SEM) over the 500 replicas. In parallel, the numerical uncertainty is obtained from the standard deviation of the best-fit parameter values across these replicas.

First, we study the uncertainties in Fig. 2a and b for the fitted mean (μ) and variance (σ^2), respectively. As the detector resolution deteriorates, the overall uncertainty increases, as expected. Notably, the numerical uncertainty matches the asymptotic error using the full covariance matrix but is consistently smaller than the asymptotic error computed using


the diagonal covariance matrix, particularly for larger smearing values. This indicates that the standard asymptotic error estimation using only the diagonal values of the covariance matrix overestimates the true confidence intervals when the unfolding-induced correlations become significant. The cause of the jump in the graphs of Fig. 2b is not known, but this does not affect our conclusions. These findings underscore the importance of incorporating the full covariance structure in the inference process when computing asymptotic uncertainties.

Next, we quantify the bias in Fig. 3a and b for the evolution of the mean best-fit values of μ and σ^2 , respectively. In these plots, the data points represent the average best-fit value at a given smearing, with vertical error bars corresponding to the SEM computed from the 500 replicas. In addition, a horizontal dashed red line marks the true parameter value. For both μ (Fig. 3a), and σ^2 (Fig. 3b), the estimates computed both with and without using off-diagonal elements of the covariance matrix are similar. In particular, all of the methods show nearly the same bias as a function of smearing, which is covered by the uncertainty. The truth spectra are the same for each smearing and thus the uncertainties are very correlated between points.


These results provide an important baseline: they demonstrate that with IBU and conventional binned inference, the choice of covariance treatment can affect the inferred uncertainties. In the subsequent sections, we compare these findings with those obtained from the unbinned unfolding approach.

4.2 Event-by-event correlations

For each parameter set, 500 statistically independent test datasets are unfolded without binning using a fixed Monte Carlo sample. Each Monte Carlo event i in the output of




(a)




(b)

Fig. 2 Mean asymptotic error versus detector smearing for the **a** Gaussian mean, μ , and **b** the variance σ^2 obtained from the full covariance analysis (green circles) and the diagonal covariance approximation (pink squares). The green stars represent the standard deviation com-

puted from the spread of the best-fit values over 500 bootstrap replicas. Note that the numerical uncertainty agrees well with the asymptotic error from the full covariance fit while the diagonal approximation consistently overestimates the uncertainty at larger smearing values



(a)



(b)

Fig. 3 The mean best-fit value of **a** μ and **b** σ^2 is shown as a function of the detector smearing, with error bars indicating the standard error of the mean (SEM) over 500 bootstrap test datasets. Horizontal dashed

red lines mark the true values of $\mu = 0.2$ and $\sigma^2 = 0.81$. Both fit methods (full covariance and diagonal approximation) yield similar central values




Fig. 4 Average weight correlation between two events as a function of the absolute distance between the events in the observable. The four curves show different values for the detector resolution. The top plots

show unfolding with the KDE approach within OmniFold, while the bottom plots show results from using NNs within OmniFold. The error bars show the RMS of the correlation values

unfolding dataset j is assigned a weight w_{ij} by the unfolding algorithm. This allows us to quantify the degree to which the event weights for a pair of Monte Carlo events are correlated in the unfolding.

Figure 4 shows the average weight correlation between pairs of events as a function of the distance between the pair of events in the observable for both the KDE and NN approaches. The closer the events are to each other, the stronger the weights are correlated, with a correlation of 1

at a distance of zero by construction. For perfect resolution, the correlation quickly dies off as the distance between the events grows. In principle, in this case, there should be no correlations between events and so any residual correlation is due to smoothing in the KDE or NN. The implicit regularization within the KDE and NN is different. For example, the KDE range parameter (δx) limits the length scale of correlations, in the absence of detector smearing, while the NN has no explicit length scale built into it. As the detector reso-




Fig. 5 Covariance matrices of a histogram of the unfolding output for four detector resolution values (0, 0.25, 0.50, 0.75). The top plots show unfolding with the KDE approach within OmniFold, while the bottom plots show results from using NNs within OmniFold. The histogram has 40 bins in the range $[-4, 4]$

lution gets worse, the range of the correlation increases. The correlations themselves show a wave-like pattern as nearby events are correlated and then anti-correlated beyond some point in order to preserve the total number of events.

Another way to visualize the correlations is to bin the data and construct covariance matrices. Figure 5 shows these matrices using histograms (40 bins from -4 to 4 in the observable) of the unfolding output, where the covariance is calculated numerically from the 500 test datasets for each resolution value. The results are consistent with the unbinned evaluation of the weight correlations, shown in Fig. 4, and additionally show that the distance scale of the correlations does not depend strongly on the observable itself.

4.3 Parameter estimation in 1D with a KDE

To investigate the impact of the event weight correlations in the unbinned unfolding output, we compare three ways of performing the model parameter inference:

- An unbinned ML fit of the unbinned unfolding output, which is the weighted Monte Carlo sample.
- A binned χ^2 fit of the unfolding output that uses the full covariance matrix for the histogram bins.
- A binned χ^2 fit of the unfolding output that uses only the diagonal elements of the covariance matrix for the histogram bins.

For the binned χ^2 fits, the covariance matrix of the histogram bins is evaluated numerically from a set of histograms made from the 500 test datasets. The inference uncertainty is determined for each model parameter from the RMS of the inference results for the 500 test datasets. We consider this to be the most correct estimate of the inference uncertainty. We compare this with the asymptotic uncertainties from the fits, determined from the change in the model parameter that gives a change of 0.5 in the log likelihood for the ML fit or a change of 1 in the χ^2 . We also compare the unbinned unfolding with the results of performing a binned unfolding using the Iterative Bayesian Unfolding technique. For all binned results, the Gaussian model predictions are integrated over the bins in order to have histogram templates that are properly parameterized.

Figure 6 shows the inference uncertainty and bias as a function of the detector resolution. We draw the following conclusions from the results shown:

- The asymptotic uncertainties in the unbinned ML fit of the unbinned unfolding are too low, do not depend on the detector resolution, and agree with the correct uncertainty only for the perfect detector resolution case. This can be seen in the top plots, noting that the orange circles do not depend on the smearing and they disagree with the orange stars (hidden under the red stars) for all levels of smearing except no smearing.




Fig. 6 Inference uncertainty (top) and bias (bottom) on the model mean (left) and variance (right) as a function of the detector resolution. Unbinned unfolding (UBU) is compared with the results of binned Iterative Bayesian Unfolding (IBU). The inference is done three ways: an unbinned ML fit, a binned χ^2 fit using the full covariance matrix for the

fit histogram, and a binned χ^2 fit using only the diagonal elements of the covariance matrix for the histogram. The blue band centered on the true value shows $\pm$ the RMS inference uncertainty from the unbinned unfolding unbinned ML fit to provide a scale reference

- The asymptotic uncertainties in the binned χ^2 fits agree with the correct uncertainty, as evaluated numerically from the RMS of the inference distribution, when the full covariance matrix is used. Pictorially, blue circles agree with blue stars (hidden under the red stars) and the green circles agree with the green stars.
- When the off-diagonal elements of the covariance matrix are set to zero in the χ^2 fit, the asymptotic uncertainties

are incorrect. This corresponds to the difference between the red circles and red stars, is the binned analog of the first point, and was observed in the fully binned case in Fig. 2).

- From the lower plots, we conclude that the bias in the unbinned ML fit is the same as the bias in the χ^2 fit.
- The inference uncertainty and bias for the unbinned unfolding is better than the for the binned unfolding with




Fig. 7 Inference uncertainty (top) and bias (bottom) on the model mean (left) and variance (right) as a function of the detector resolution. Unbinned unfolding (UBU) is done two ways: with the KDE and with Neural Networks (NN) in OmniFold. The inference is done two

ways: an unbinned ML fit and a binned χ^2 fit using the full covariance matrix for the fit histogram. The blue band centered on the true value shows $\pm$ the RMS inference uncertainty from the unbinned unfolding unbinned ML fit to provide a scale reference

IBU for the binning and method hyperparameter choices that produced the results shown in Fig. 6. This suggests that some useful information may be lost when the data are binned before unfolding.

The results from the χ^2 fits are as expected. When correlations from the unfolding are properly included by using

the full covariance matrix for the fit histogram, the asymptotic uncertainties are correct. When these correlations are ignored, the asymptotic uncertainties are incorrect. Correlations in the unbinned unfolding are negligible when the detector resolution is perfect and, in this case, the asymptotic uncertainties in the unbinned ML fit are correct because the event weights from the unfolding are uncorrelated.




Fig. 8 Scatter plots of the RMS vs average asymptotic inference uncertainty for unbinned unfolding using NNs in OmniFold in 1D, 2D, 4D, and 6D. Each model parameter is represented by a point (N means in N dimensions and $N(N + 1)/2$ covariance elements). The vertical coordinate is the RMS of the fitted model parameter from 500 experiments,

while the horizontal coordinate is the average asymptotic uncertainty from an unbinned ML fit. The dashed gray line has a slope of 1. The slope of the solid green line is set to the average of the RMS/asymptotic uncertainty ratio




Fig. 9 Scatter plots of the RMS vs average asymptotic inference uncertainty for unbinned unfolding using NNs in OmniFold in 6D with varying amounts of detector resolution smearing. The detector resolution from the 6D study shown in Fig. 8 is scaled by the following factors: 0, 1.0, 2.0, and 3.0. Each model parameter is represented by a point.

The vertical coordinate is the RMS of the fitted model parameter from 500 experiments, while the horizontal coordinate is the average asymptotic uncertainty from an unbinned ML fit. The dashed gray line has a slope of 1. The slope of the solid green line is set to the average of the RMS/asymptotic uncertainty ratio

4.4 Parameter estimation in 1D with a NN

Figure 7 has a similar setup as Fig. 6, but compares the KDE and NN approaches within OmniFold. The trends with the NN are similar.

We found that the activation function can affect the degree of regularization in the NN output, which also affects the strength and distance scale of the weight correlations in the OmniFold output. Interestingly, we found that the inference uncertainty is insensitive to these correlations. This is discussed in more detail in Appendix B.

4.5 Parameter estimation beyond 1D

In dimensions higher than 2 or 3, unbinned unfolding becomes the only practical method. We investigate higher-dimensional feature spaces using a multivariate normal dis-

tribution as our ‘physics’ model. The covariance parameters of the model are chosen such that the densities do not factorize in the variables spanning the feature space. The mean and sigma values were chosen to include a range of mean values within $[-1, 1]$, sigma values within $[0.6, 1.5]$, and non-trivial off-diagonal correlation coefficients within $[-0.7, 0.7]$. The Monte Carlo gen particle and pseudo data true values differ by small amounts, so that the unfolding task is non trivial. The chosen values give an invertible model multivariate Gaussian covariance matrix. We model the detector resolution with a Gaussian smearing that is separate and uncorrelated for each feature dimension. The model and resolution parameters used are given in Appendix A.

Figure 8 compares the inference precision as estimated with the RMS of the 500 pseudo datasets with the asymptotic uncertainty from an unbinned ML fit, using Eq. (1.1) for each model parameter for studies in 1D, 2D, 4D, and 6D feature

space. The detector resolution is set to nominal values that are kept the same when additional feature dimensions are added. In each study, the RMS uncertainty is higher than the asymptotic uncertainty. The RMS/asymptotic uncertainty ratio appears to be the roughly the same for all of the model parameters within each plot with the ratio ranging from 1.18 to 1.28.

Figure 9 compares four studies in the 6D feature space with varying detector resolution. The results are qualitatively similar to the 1D case. The RMS/asymptotic uncertainty ratio is consistent with 1 for perfect detector resolution and increases gradually as the detector smearing increases.

5 Conclusions and outlook

Simulation-based inference is enabling a new class of analyses that make use of all the available information without first (lossily) compressing the data into a small number of summary statistics that are discretized via histograms. New tools have enabled unbinned and high-dimensional inference of parameters and unbinned and high-dimensional inference of differential cross sections. In this paper, we explore the fusion of these two approaches: unbinned parameter estimation from unbinned unfolded results.

We used a Gaussian setup, where the simulation/forward model of the data is known analytically and pseudodata were unfolded using the OmniFold unbinned unfolding method. We identified a number of critical observations: (1) there can be an advantage of unbinned unfolding for parameter estimation even if the inference is performed using histograms. This matches with our intuition that it is best to reduce the information as late as possible in the statistical analysis. (2) Ignoring correlations – as is currently required in the unbinned inference case – seems to have little or no effect on the accuracy and precision of the fit as long as numerical methods are used for the statistical analysis. This does not seem to be a formal consequence of the setup and may be a feature of the Gaussian example. The motivating question of the entire study led us to (3). Asymptotic formulae are not valid when correlations are ignored and the reported uncertainties can be significantly underestimated without computing them numerically.

With a growing number of unbinned cross section measurements, it is critical to build a robust statistical and software suite for analyzing these data in the context of physics models. Our main recommendation is to avoid asymptotic formulae for unbinned inference of unbinned measurements until the statistical formalism can be extended to include event-to-event correlations. Valid inference uncertainties can still be obtained, when analyzing the output of unbinned unfolding with event weight correlations, if they are evaluated numerically using the bootstrap method. An important area of future research will be to directly study the interplay

between machine learning regularization and the impact on downstream inference. In this work, we have focused on the case where the ‘physics’ model of the data is known; in the most general case, this model is not known explicitly and one must also use machine learning to approximate it from simulations. Studying the end-to-end approach may reveal connections between the (implicit) regularization used by the machine learning models at both steps. While there is still significant research required to optimally employ these methods, current results can already be used with care. The scientific gains from these approaches hold immense potential for particle physics and beyond.

Acknowledgements We thank Jesse Thaler for useful discussions and feedback on the manuscript. This material is based upon work supported by the National Science Foundation under Grant No. 2311666.

Data Availability Statement My manuscript has no associated data. [Authors’ comment: The parameters necessary to generate the simulation samples for the work presented here are given in Appendix A.]

Code Availability Statement My manuscript has associated code/software in a data repository. [Authors’ comment: The code used to generate the samples and perform the studies is available in the code repository <https://github.com/owen234/unbinned-inference-paper>.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.
Funded by SCOAP³.

Appendix A: Parameter values for the 2D, 4D, and 6D studies

This appendix gives the parameter values used in the multi-dimensional studies. The model parameters are the mean μ , width σ , and correlation coefficients ρ for the multivariate normal distribution. Values are given for the Monte Carlo (MC) and the true values used to generate the pseudo data. The detector resolution is σ_{det} .
Parameters for the 2D study.

$$\begin{aligned} \mu_{\text{MC}} &= \begin{pmatrix} 0.0 \\ 1.0 \end{pmatrix} & \mu_{\text{true}} &= \begin{pmatrix} 0.2 \\ 0.8 \end{pmatrix} & \sigma_{\text{MC}} &= \begin{pmatrix} 1.0 \\ 1.5 \end{pmatrix} \\ \sigma_{\text{true}} &= \begin{pmatrix} 0.9 \\ 1.3 \end{pmatrix} & \sigma_{\text{det}} &= \begin{pmatrix} 0.5 \\ 0.8 \end{pmatrix} \end{aligned}$$

$$\rho_{\text{MC}} = \begin{pmatrix} 1.0 & -0.6 \\ -0.6 & 1.0 \end{pmatrix} \quad \rho_{\text{true}} = \begin{pmatrix} 1.0 & -0.6 \\ -0.6 & 1.0 \end{pmatrix}$$

Parameters for the 4D study.

$$\mu_{\text{MC}} = \begin{pmatrix} 1.0 \\ 0.0 \\ -0.5 \\ 0.5 \end{pmatrix} \quad \mu_{\text{true}} = \begin{pmatrix} 0.8 \\ 0.1 \\ -0.6 \\ 0.7 \end{pmatrix} \quad \sigma_{\text{MC}} = \begin{pmatrix} 1.0 \\ 0.7 \\ 1.1 \\ 0.8 \end{pmatrix}$$

$$\sigma_{\text{true}} = \begin{pmatrix} 0.8 \\ 0.6 \\ 1.0 \\ 0.6 \end{pmatrix} \quad \sigma_{\text{det}} = \begin{pmatrix} 0.4 \\ 0.5 \\ 0.6 \\ 0.3 \end{pmatrix}$$

$$\rho_{\text{MC}} = \begin{pmatrix} 1.0 & 0.1 & -0.2 & 0.3 \\ 0.1 & 1.0 & 0.0 & 0.1 \\ -0.2 & 0.0 & 1.0 & 0.7 \\ 0.3 & 0.1 & 0.7 & 1.0 \end{pmatrix}$$

$$\rho_{\text{true}} = \begin{pmatrix} 1.0 & 0.0 & -0.3 & 0.4 \\ 0.0 & 1.0 & 0.2 & 0.0 \\ -0.3 & 0.2 & 1.0 & 0.5 \\ 0.4 & 0.0 & 0.5 & 1.0 \end{pmatrix}$$

Parameters for the 6D study.

$$\mu_{\text{MC}} = \begin{pmatrix} 1.0 \\ 0.0 \\ -0.5 \\ 0.5 \\ -1.0 \\ 0.3 \end{pmatrix} \quad \mu_{\text{true}} = \begin{pmatrix} 0.8 \\ 0.1 \\ -0.6 \\ 0.7 \\ -0.8 \\ 0.1 \end{pmatrix} \quad \sigma_{\text{MC}} = \begin{pmatrix} 1.0 \\ 0.7 \\ 1.1 \\ 0.8 \\ 1.2 \\ 1.4 \end{pmatrix}$$

$$\sigma_{\text{true}} = \begin{pmatrix} 0.8 \\ 0.6 \\ 1.0 \\ 0.6 \\ 1.0 \\ 1.1 \end{pmatrix} \quad \sigma_{\text{det}} = \begin{pmatrix} 0.4 \\ 0.5 \\ 0.6 \\ 0.3 \\ 0.4 \\ 0.4 \end{pmatrix}$$

$$\rho_{\text{MC}} = \begin{pmatrix} 1.0 & 0.1 & 0.2 & -0.3 & 0.0 & 0.0 \\ 0.1 & 1.0 & 0.0 & -0.2 & 0.3 & 0.1 \\ 0.2 & 0.0 & 1.0 & 0.1 & -0.2 & 0.3 \\ -0.3 & -0.2 & 0.1 & 1.0 & 0.1 & 0.0 \\ 0.0 & 0.3 & -0.2 & 0.1 & 1.0 & 0.7 \\ 0.0 & 0.1 & 0.3 & 0.0 & 0.7 & 1.0 \end{pmatrix}$$

$$\rho_{\text{true}} = \begin{pmatrix} 1.0 & 0.0 & 0.2 & -0.2 & 0.1 & 0.0 \\ 0.0 & 1.0 & 0.0 & -0.1 & 0.2 & 0.0 \\ 0.2 & 0.0 & 1.0 & 0.0 & -0.3 & 0.4 \\ -0.2 & -0.1 & 0.0 & 1.0 & 0.2 & 0.0 \\ 0.0 & 0.2 & -0.3 & 0.2 & 1.0 & 0.5 \\ 0.0 & 0.0 & 0.4 & 0.0 & 0.5 & 1.0 \end{pmatrix}$$

Appendix B: Activation function dependence of event correlations

We have found that the event weight correlations in the unfolding output are sensitive to the activation functions used in the NN model. Figure 10 shows the average event weight correlation as a function of distance in feature space for our 1D study with a detector smearing of 0.5 for the KDE compared to three different choices for the NN activation func-




Fig. 10 Average weight correlation between two events as a function of the absolute distance between the events in the observable. The results for the KDE are compared with three different NN activation function choices: `relu`, `gelu` with $\beta = 1$, and `gelu` with $\beta = 4$




Fig. 11 Graphs of the event weight function from the unfolding compared with the true event weight function. The columns show different options within OmniFold: left is the KDE, center left is NNs with `relu` activation, center right is NNs with `gelu` activation with $\beta = 4$, and

right is NNs with `gelu` activation with $\beta = 1$. Each row is the results for a particular test sample and the columns within a row all use the same test sample

tion: `relu`, `gelu` with $\beta = 1$, and `gelu` with $\beta = 4$. We find that `gelu` with $\beta = 1$, which is the standard version of `gelu`, introduces the most regularization and has the longest distance scale for the correlations. If the transition from flat to a slope of 1 is shortened within `gelu` by setting $\beta = 4$, the pattern of event weight correlations is quite similar to `relu` and this is similar to the results from the KDE.

Figure 11 shows graphs of the event weight function from the unfolding compared with the true event weight function. The columns compare the KDE with the three NN activation function choices. Each row shows the results for unfolding a particular test dataset. The KDE has the least amount of regularization of the weight function. Statistical fluctuations in the tails are visible. The `relu` activation allows for sharper kinks in the weight function and tends to converge to a linear function on the right side, where the statistics are lower. The `gelu` with $\beta = 1$ is much more regularized than `gelu` with

$\beta = 4$, which is closest to a smoothed version of the KDE weight function.

The true weight function is smooth and slowly varying, so a high degree of regularization is not necessarily a problem for our particular toy physics model. The inference results from the unfolding are given for each plot in Fig. 11. Even though the weight functions differ substantially, they all give very similar values in the inference of the model parameters. After seeing the graphs for `gelu` with $\beta = 1$ (right column in Fig. 11), we expected it to have better precision, but that's not the case. The KDE and all three NN activation choices all give inference uncertainties of around 0.010 for the model μ and 0.014 for the model σ^2 . Note that the inference uncertainty scale is quite small compared to the horizontal axis range shown in Fig. 11.

Appendix C: Ensembling and inference resolution from NN model parameter initialization

The converged state of the unbinned unfolding when using NNs within OmniFold has some dependence on the random initialization of the NN model weights. This impacts the resolution of the unfolding and downstream inference. However, this effect can be removed by averaging the output of several

runs of the full unfolding, where each run starts with a different set of random initial NN model weights. This averaging is known as ensembling. Averaging the final unfolding weights is known as parallel ensembling [46]. Figure 12 shows the inference uncertainty on the model mean and model variance for our 1D study with detector smearing of 0.5 as a function of the number of runs used in the ensemble average. Our interpretation of these results is that an ensemble of 10 runs is




Fig. 12 Inference uncertainty on the model parameters for the 1D study as a function of the number of runs in the ensemble average. A detector smearing of 0.5 was used




Fig. 13 Inference uncertainty on the model parameters for our 1D study as a function of detector smearing. The orange circles are a measure of the contribution from the initialization of the random NN weights σ_p . The blue squares show the resolution without ensembling σ_1 . The

Green triangles show the resolution for a 10-run ensemble average σ_{10} . The quadrature difference $\sqrt{\sigma_1^2 - \sigma_p^2}$ agrees fairly well with the 10-run ensemble average

sufficient to remove the parameter initialization contribution to the unfolding resolution.

Figure 13 shows the components of the overall inference uncertainty as a function of the level of detector smearing. The orange circles show the resolution, from the RMS of the inference distribution, when the inference is done 200 times on the same dataset with different initial random NN model weights in a single run. This is a direct measure of the size of the parameter initialization resolution contribution σ_p . The blue squares show the results for unfolding 200 independent datasets without ensembling σ_1 . The green triangles show the results for unfolding the same 200 independent datasets, where the ensemble average of 10 runs is used σ_{10} . The red diamonds show the quadrature difference of the 1-run minus the parameter initialization resolution $\sqrt{\sigma_1^2 - \sigma_p^2}$, which agrees fairly well with the 10-run ensemble average resolution σ_{10} , as expected.

References

1. K. Cranmer, J. Brehmer, G. Louppe, The frontier of simulation-based inference. *Proc. Natl. Acad. Sci.* **117**(48), 30055–30062 (2020). <https://doi.org/10.1073/pnas.1912789117>
2. G. Cowan, K. Cranmer, E. Gross, O. Vitells, Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J. C* **71**, 1554 (2011) (Erratum: *Eur. Phys. J. C* **73**, 2501 (2013)). <https://doi.org/10.1140/epjc/s10052-011-1554-0>. arXiv:1007.1727
3. G. Aad et al., Measurement of off-shell Higgs boson production in the $H^* \rightarrow ZZ \rightarrow 4\ell$ decay channel using a neural simulation-based inference technique in 13 TeV pp collisions with the ATLAS detector. *Rep. Prog. Phys.* **88**(5), 057803 (2025). <https://doi.org/10.1088/1361-6633/adcd9a>. arXiv:2412.01548
4. G. Cowan, A survey of unfolding methods for particle physics. *Conf. Proc. C* **0203181**, 248–257 (2002)
5. V. Blobel, Unfolding methods in particle physics, in *PHYSTAT2011 Proceedings* (2011). <https://doi.org/10.5170/CERN-2011-006>
6. V. Blobel, Unfolding in data analysis, in *High Energy Physics*, ed. by O. Behnke, K. Kröninger, G. Schott, T. Schörner-Sadenius. Wiley-VCH, Weinheim (2013). <https://doi.org/10.1002/9783527653416.ch6>
7. L. Brenner, R. Balasubramanian, C. Burgard, W. Verkerke, G. Cowan, P. Verschuur, V. Croft, Comparison of unfolding methods using RooFitUnfold. *Int. J. Mod. Phys. A* **35**(24), 2050145 (2020). <https://doi.org/10.1142/S0217751X20501456>. arXiv:1910.14654
8. J.J. Ethier, E.R. Nocera, Parton distributions in nucleons and nuclei. *Ann. Rev. Nucl. Part. Sci.* **70**, 43–76 (2020). <https://doi.org/10.1146/annurev-nucl-011720-042725>. arXiv:2001.07722
9. A. Buckley, H. Hoeth, H. Lacker, H. Schulz, J.E. von Seggern, Systematic event generator tuning for the LHC. *Eur. Phys. J. C* **65**, 331–357 (2010). <https://doi.org/10.1140/epjc/s10052-009-1196-7>. arXiv:0907.2973
10. J.M. Butterworth, D. Grellscheid, M. Krämer, B. Sarrazin, D. Yallup, Constraining new physics with collider measurements of standard model signatures. *JHEP* **03**, 078 (2017). [https://doi.org/10.1007/JHEP03\(2017\)078](https://doi.org/10.1007/JHEP03(2017)078). arXiv:1606.05296
11. M. Arratia et al., Publishing unbinned differential cross section results. *JINST* **17**(01), P01024 (2022). <https://doi.org/10.1088/1748-0221/17/01/P01024>. arXiv:2109.13243
12. S. Badger et al., Machine learning and LHC event generation. *SciPost Phys.* **14**(4), 079 (2023). <https://doi.org/10.21468/SciPostPhys.14.4.079>. arXiv:2203.07460
13. N. Huetsch et al., The landscape of unfolding with machine learning. *SciPost Phys.* **18**(2), 070 (2025). <https://doi.org/10.21468/SciPostPhys.18.2.070>. arXiv:2404.18807
14. A. Andreassen, P.T. Komiske, E.M. Metodiev, B. Nachman, J. Thaler, Omnifold: a method to simultaneously unfold all observables. *Phys. Rev. Lett.* **124**, 182001 (2020). <https://doi.org/10.1103/PhysRevLett.124.182001>. arXiv:1911.09107
15. A. Andreassen, P.T. Komiske, E.M. Metodiev, B. Nachman, A. Suresh, J. Thaler, Scaffolding simulations with deep learning for high-dimensional deconvolution, in *9th International Conference on Learning Representations* (2021). arXiv:2105.04448
16. C.-C. Pan, X. Dong, Y.-C. Sun, A.-Y. Cheng, A.-B. Wang, Y.-X. Hu, H. Cai, SwdFold: a reweighting and unfolding method based on optimal transport theory. arXiv:2406.01635
17. K. Datta, D. Kar, D. Roy, Unfolding with generative adversarial networks. arXiv:1806.00433
18. J.N. Howard, S. Mandt, D. Whiteson, Y. Yang, Learning to simulate high energy particle collisions from unlabeled data. *Sci. Rep.* **12**, 7567 (2022). <https://doi.org/10.1038/s41598-022-10966-7>. arXiv:2101.08944
19. S. Diefenbacher, G.-H. Liu, V. Mikuni, B. Nachman, W. Nie, Improving generative model-based unfolding with Schrödinger bridges. *Phys. Rev. D* **109**(7), 076011 (2024). <https://doi.org/10.1103/PhysRevD.109.076011>. arXiv:2308.12351
20. A. Butter, T. Jezo, M. Klasen, M. Kuschick, S. Palacios Schweitzer, T. Plehn, Kicking it off(-shell) with direct diffusion. *SciPost Phys. Core* **7**(3), 064 (2024). <https://doi.org/10.21468/SciPostPhysCore.7.3.064>. arXiv:2311.17175
21. M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, R. Winterhalder, How to GAN away detector effects. *SciPost Phys.* **8**(4), 070 (2020). <https://doi.org/10.21468/SciPostPhys.8.4.070>. arXiv:1912.00477
22. M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, A. Rousselot, R. Winterhalder, L. Ardizzone, U. Köthe, Invertible networks or partons to detector and back again. *SciPost Phys.* **9**, 074 (2020). <https://doi.org/10.21468/SciPostPhys.9.5.074>. arXiv:2006.06685
23. M. Vandegar, M. Kagan, A. Wehenkel, G. Louppe, Neural empirical bayes: source distribution estimation and its applications to simulation-based inference, in *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics, Vol. 130 of Proceedings of Machine Learning Research* ed. by A. Banerjee, K. Fukumizu (PMLR, 2021), p. 2107–2115. <https://proceedings.mlr.press/v130/vandegar21a.html>. arXiv:2011.05836
24. M. Backes, A. Butter, M. Dunford, B. Malaescu, An unfolding method based on conditional invertible neural networks (cINN) using iterative training. *SciPost Phys. Core* **7**(1), 007 (2024). <https://doi.org/10.21468/scipostphyscore.7.1.007>. arXiv:2212.08674
25. M. Leigh, J.A. Raine, K. Zoch, T. Golling, ν -flows: conditional neutrino regression. *SciPost Phys.* **14**(6), 159 (2023). <https://doi.org/10.21468/SciPostPhys.14.6.159>. arXiv:2207.00664
26. J. Ackerschott, R.K. Barman, D. Gonçalves, T. Heimel, T. Plehn, Returning CP-observables to the frames they belong. *SciPost Phys.* **17**(1), 001 (2024). <https://doi.org/10.21468/SciPostPhys.17.1.001>. arXiv:2308.00027
27. A. Shmakov, K. Greif, M. Fenton, A. Ghosh, P. Baldi, D. Whiteson, End-to-end latent variational diffusion models for inverse problems in high energy physics. arXiv:2305.10399
28. A. Shmakov, K.T. Greif, M.J. Fenton, A. Ghosh, P. Baldi, D. Whiteson, Full event particle-level unfolding with variable-length latent variational diffusion. *SciPost Phys.* **18**(4), 117 (2025). <https://doi.org/10.21468/SciPostPhys.18.4.117>. arXiv:2404.14332
29. V. Andreev et al., Measurement of lepton-jet correlation in deep-inelastic scattering with the H1 detector using machine learning

- for unfolding. Phys. Rev. Lett. **128**(13), 132002 (2022). <https://doi.org/10.1103/PhysRevLett.128.132002>. arXiv:2108.12376
30. H1 Collaboration, Machine learning-assisted measurement of multi-differential lepton-jet correlations in deep-inelastic scattering with the H1 detector, H1prelim-22-031. <https://www-h1.desy.de/h1/www/publications/htmlsplit/H1prelim-22-031.long.html>
 31. V. Andreev et al., Unbinned deep learning jet substructure measurement in high q_2 ep collisions at hera. Phys. Lett. B **844**, 138101 (2023). <https://doi.org/10.1016/j.physletb.2023.138101>. arXiv:2303.13620
 32. V. Andreev, et al., Machine Learning-Assisted Measurement of Lepton-Jet Azimuthal Angular Asymmetries in Deep-Inelastic Scattering at HERA. arXiv:2412.14092
 33. R. Aaij et al., Multidifferential study of identified charged hadron distributions in z -tagged jets in proton–proton collisions at $\sqrt{s} = 13$ tev. Phys. Rev. D **108**, L031103 (2023). <https://doi.org/10.1103/PhysRevD.108.L031103>. arXiv:2208.11691
 34. G. Aad et al., Simultaneous unbinned differential cross-section measurement of twenty-four Z +jets kinematic observables with the ATLAS detector. Phys. Rev. Lett. **133**(26), 261803 (2024). <https://doi.org/10.1103/PhysRevLett.133.261803>. arXiv:2405.20041
 35. G. Aad et al., Measurement of jet track functions in pp collisions at $s=13$ tev with the atlas detector. Phys. Lett. B **868**, 139680 (2025). <https://doi.org/10.1016/j.physletb.2025.139680>. arXiv:2502.02062
 36. Measurement of event shapes in minimum bias events from pp collisions at 13 TeV, Tech. rep., CERN, Geneva (2024). <https://cds.cern.ch/record/2899591>
 37. 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
 38. T. Pani, Generalized angularities measurements from star at $\sqrt{s_{nn}} = 200$ gev. EPJ Web Conf. **296**, 11003 (2024). <https://doi.org/10.1051/epjconf/202429611003>. arXiv:2403.13921
 39. F. Canelli, et al., A practical guide to unbinned unfolding. arXiv:2507.09582
 40. C. Langenbruch, Parameter uncertainties in weighted unbinned maximum likelihood fits. Eur. Phys. J. C **82**(5), 393 (2022). <https://doi.org/10.1140/epjc/s10052-022-10254-8>. arXiv:1911.01303
 41. B. Efron, Bootstrap methods: another look at the jackknife. Ann. Stat. **7**(1), 1–26 (1979). <https://doi.org/10.1214/aos/1176344552>
 42. M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G.S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, X. Zheng, TensorFlow: large-scale machine learning on heterogeneous systems, software available from tensorflow.org (2015). <https://www.tensorflow.org/>
 43. D. Hendrycks, K. Gimpel, Gaussian error linear units (gelus) (2023). arXiv:1606.08415
 44. D. P. Kingma, J. Ba, Adam: a method for stochastic optimization (2017). arXiv:1412.6980
 45. X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, Vol. 9 of Proceedings of Machine Learning Research* ed. by Y. W. Teh, M. Titterton. PMLR, Chia Laguna Resort, Sardinia, Italy, 2010, p. 249–256. <https://proceedings.mlr.press/v9/glorot10a.html>
 46. F.T. Acosta, T. Wamorkar, V. Mikuni, B. Nachman, Stabilizing neural likelihood ratio estimation. arXiv:2503.20753
 47. G. D'Agostini, A multidimensional unfolding method based on Bayes' theorem. [https://doi.org/10.1016/0168-9002\(95\)00274-X](https://doi.org/10.1016/0168-9002(95)00274-X)
 48. W.H. Richardson, Bayesian-based iterative method of image restoration*. J. Opt. Soc. Am. **62**(1), 55–59 (1972). <https://doi.org/10.1364/JOSA.62.000055>
 49. L.B. Lucy, An iterative technique for the rectification of observed distributions. Astron. J. **79**, 745 (1974). <https://doi.org/10.1086/111605>