

UC Irvine

UC Irvine Previously Published Works

Title

Publishing unbinned differential cross section results

Permalink

<https://escholarship.org/uc/item/75g8r5w1>

Journal

Journal of Instrumentation, 17(01)

ISSN

1748-0221

Authors

Arratia, Miguel
Butter, Anja
Campanelli, Mario
[et al.](#)

Publication Date

2022

DOI

10.1088/1748-0221/17/01/p01024

Peer reviewed

Presenting Unbinned Differential Cross Section Results

Miguel Arratia,^{a,b} Anja Butter,^c Mario Campanelli,^d Vincent Croft,^e Dag Gillberg,^f Aishik Ghosh,^{g,h} Kristin Lohwasser,ⁱ Bogdan Malaescu,^j Vinicius Mikuni,^k Benjamin Nachman,^{h,l} Juan Rojo,^{m,n} Jesse Thaler,^{o,p} Ramon Winterhalder^q

^a*Department of Physics and Astronomy, University of California, Riverside, CA 92521, USA*

^b*Thomas Jefferson National Accelerator Facility, Newport News, VA 23606, USA*

^c*Institut für Theoretische Physik, Universität Heidelberg, Heidelberg, Germany*

^d*University College London, London WC1E 6BT, UK*

^e*Department of Physics and Astronomy, Tufts University, Boston, MA 02155, USA*

^f*Department of Physics, Carleton University, Ottawa ON K1S 5B6, Canada*

^g*Department of Physics and Astronomy, University of California, Irvine, CA 92697, USA*

^h*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

ⁱ*University of Sheffield, Sheffield, S10 2TN, UK*

^j*LPNHE, Sorbonne Université, Université de Paris, CNRS/IN2P3, Paris, France*

^k*National Energy Research Scientific Computing Center, Berkeley, CA 94720, USA*

^l*Berkeley Institute for Data Science, University of California, Berkeley, CA 94720, USA*

^m*Nikhef Theory Group, Science Park 105, 1098 XG Amsterdam, The Netherlands*

ⁿ*Department of Physics and Astronomy, Vrije Universiteit Amsterdam, NL-1081 HV Amsterdam, The Netherlands*

^o*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

^p*The NSF AI Institute for Artificial Intelligence and Fundamental Interactions*

^q*Centre for Cosmology, Particle Physics and Phenomenology (CP3), Université catholique de Louvain, 1348 Louvain-la-Neuve, Belgium*

E-mail: bpnachman@lbl.gov

ABSTRACT: Machine learning tools have empowered a qualitatively new way to perform differential cross section measurements whereby the data are unbinned, possibly in many dimensions. Unbinned measurements can enable, improve, or at least simplify comparisons between experiments and with theoretical predictions. Furthermore, many-dimensional measurements can be used to define observables after the measurement instead of before. There is currently no community standard for publishing unbinned data. While there are also essentially no measurements of this type public, unbinned measurements are expected in the near future given recent methodological advances. The purpose of this paper is to propose a scheme for presenting and using unbinned results, which can hopefully form the basis for a community standard to allow for integration into analysis workflows. This is foreseen to be the start of an evolving community dialogue, in order to accommodate future developments in this field that is rapidly evolving.

Contents

1	Introduction	2
2	Motivation	3
2.1	Inference-Aware Binning	3
2.2	Derivative Measurements	4
2.3	Extension to Higher Dimensions	4
3	Overview of Unbinned Methods	5
3.1	Experimental Measurements	5
3.1.1	Density-Based Models	5
3.1.2	Classifier-Based Models	6
3.2	Theoretical Calculations	6
4	Statistics of Unbinned Spectra	7
4.1	Acceptance Effects	7
4.2	Background Subtraction	7
4.3	Local Statistical Uncertainty	7
5	Systematic Uncertainties	9
6	Proposed Format and Storage	9
6.1	Structure	9
6.2	Uncertainties	11
6.3	Storage	11
6.4	Common Tools	11
7	Example	12
8	Conclusions and Outlook	17

1 Introduction

One of the main components of high energy particle and nuclear physics research is the prediction and measurement of (differential) cross sections. Such spectra are sensitive to a wide range of phenomena and can be used to extract fundamental parameters of the Standard Model (SM), explore the quantum properties of the strong and electroweak forces, tune Monte Carlo (MC) event generators, and search for physics beyond the SM. In order to be comparable to data, predictions must be differential in the measurement observables; in order to be comparable to other experimental data and theoretical predictions, data must be corrected for detector effects. This is usually done by first discretizing the target phase space into a finite number of bins. Then, the theoretical and experimental results can be represented as histograms. These histograms can then be readily visualized and their contents can be encoded in an array for easy storage (e.g. in HEPDATA [1]) and statistical analysis including MC tuning [2], Parton Distribution Function (PDF) fits for protons [3] and heavy nuclei [4], Wilson coefficient extractions from higher-dimensional operators in the Effective Field Theories (EFTs) [5–14], and global fits of the parameters of UV-complete theories [15, 16] such as supersymmetry. Correcting for detector effects is known as *unfolding* (see Ref. [17–20] for recent reviews¹).

While the current paradigm has resulted in many exciting physics results, there are also significant limitations. For example, the observables and binning must be chosen ahead of time. This makes it difficult and in some cases impossible to compare different measurements with each other and with theoretical calculations. Furthermore, the optimal binning depends on the application. While there have been some proposals for unbinned measurements in the past [22–24], until now, they were not used for data analysis. Recent innovations in machine learning have empowered new algorithms that can process unbinned and high- (or even variable-)dimensional inputs/outputs (e.g. the momenta of all jets in an event) [25–35] which are already being used for the analysis of high energy physics data [36]. Representing unbinned spectra is very different than representing binned distributions and therefore it is critically important to discuss common formats and tools to maximize the science potential of the data and the corresponding theoretical predictions. Our goal is to propose an algorithm-agnostic framework that could become a standard across experiments and theory. In order to be most useful, widespread consensus is required and so input and feedback is most welcome.

This document is organized as follows. First, Sec 2 motivates the use of unbinned measurements. Then, Sec. 3 briefly describes existing proposals for unbinned unfolding methods. The goal of the proposed format is that it should be method-agnostic and so it needs to be able to accommodate all approaches. Section 4 discusses important statistical properties of unbinned spectra, which will be relevant for ensuring that all necessary information is stored in the proposed format. The new format itself is described in Sec. 6. An example is provided in Sec. 7. The paper ends with conclusions and outlook in Sec. 8.

¹Unfolding is not the only way to publish reusable information about data – a complementary approach is to publish the full statistical model of the data from parametric fits. See e.g. Ref. [21].

Our proposed format is meant to be the start of an evolving community dialogue, in order to accommodate future developments in this field that is currently rapidly evolving.

2 Motivation

There are generally three inter-related motivations for unbinned measurements, which are discussed below and combined with motivating physics examples.

2.1 Inference-Aware Binning

The first and foremost advantage is that measurements which are unbinned can be binned at inference time using the optimal choice of bin boundaries (‘inference-aware binning’). Currently, bins are chosen *a priori* without a particular statistical analysis in mind and use heuristic approaches (e.g. bin purity $\geq 50\%$) to achieve a useful binning for a wide range of *post hoc* analyses. For a given statistical analysis, there will be an optimal choice of binning that minimizes the uncertainty. For global fits combining multiple analyses, the optimal binning for a particular input measurement could even change with time as other results are updated or added. Multiple methods have been proposed to determine the preferred binning such as Bayesian reweighting [37] and Hessian profiling [38] for PDF fits and information geometry techniques for Higgs and EFT physics [39]. One can always attempt a brute-force approach whereby a global analysis is repeated several times for various choices of potential binning, but this option is usually disfavoured due to the heavy computational requirements.

The choice of an optimal binning at the time of a particular statistical analysis is facilitated by an unbinned result. This is particularly true for measurements of multiple moments of distributions (see e.g. Ref. [40, 41]) and one could even perform an unbinned or neural network-based analysis [42–48] without resorting to bins in the first place. With the unbinned data, one could study the loss in sensitivity from binning. In many cases, the information loss will be small (as the binning is adapted to the detector resolution), but the unbinned data can be used to ensure that there is never a significant loss in sensitivity from binning. This is particularly important for probing the tails of distributions, where any information about the distribution within a bin can improve the sensitivity. For example, such tails are critical for assessing the interplay between EFT and PDF fits [49]. Ultimately, the detector resolution and its uncertainty set a limit on the amount of information that can be extracted using unbinned spectra. In some cases, it may not be helpful to consider finer bins, but it could be useful for a variety of reasons to shift bin boundaries.

The covariance matrices released together with binned unfolded measurements sometimes have very small eigenvalues, translating into instabilities at the level of the χ^2 computations. The availability of unbinned observables would allow a more detailed investigation and possibly a solution for such problem since it would be possible to rebin without necessarily preserving the previous bin edges (in binned measurements, one can only merge, but not shift bins).

2.2 Derivative Measurements

The second advantage of unbinned unfolding is that it enables measurements of derivative observables. Suppose that one has measured a multi-differential cross section of x_0, \dots, x_n . Unbinned methods allow for the *post hoc* measurement of $f(x_0, \dots, x_n)$ for $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$. In the binned case, only a crude measurement of f is possible, using the values of f given by the combinations of bin centers over values of the x_i .

A common challenge in global analyses is establishing which observables are maximally sensitive to the parameters one aims to extract from the data. In some cases, one can exploit physical arguments in favor of specific options, e.g. in EFT fits it is often advantageous to bin data in terms of some invariant-mass variable to benefit from the energy-growth induced by the EFT operators. Another example is provided by PDF fits, where in general observables that are more closely related to the underlying partonic kinematics (such as lepton rapidity distributions in Drell-Yan production) exhibit a higher sensitivity to the PDFs. However, beyond these rather generic considerations, it is difficult to quantify *a priori* which choice of binning or even of differential distribution enjoys the maximal sensitivity. Even worse, the answer to this question may vary with time, since as the global fit evolves the relative weight of specific directions in the parameter space vary.

Another important example is the measurement of observables in different reference frames. For example, in ep scattering, it can be useful to change between the laboratory frame and the Breit frame (γ^*p center-of-mass frame). A seamless change of reference frame would benefit comparisons between different theoretical frameworks, which sometimes are used to study the same final state (e.g. Ref. [50] and Ref. [51]).

As with binned measurements, it is essential that unbinned measurements include a well-defined phase space (sometimes called the ‘fiducial volume’) for making comparisons to other measurements and to theory predictions. A benefit of unbinned unfolding is that one can more readily restrict the phase space after the measurement to match another measurement or prediction that has a smaller fiducial volume.

2.3 Extension to Higher Dimensions

For many of the machine learning-based methods, it is straightforward to increase the dimensionality of the unfolding. While not strictly a benefit of unbinned unfolding directly, there is a close connection between the unbinned nature of machine learning techniques and their ability to accommodate more observables. This is further enhanced by the considerations in the previous section. Highly differential binned cross sections are rare since it is cumbersome/impractical to construct and utilize measurements with so many bins. Unbinned measurements provide a practical alternative, but also come with smoothness assumptions related to the interpolations between different phase-space regions, implicitly done in the parameterisation of the detector response used at the unfolding step.

It is possible to take n different one-dimensional measurements and correlate them through bootstrapping for statistical uncertainties² (see e.g. Ref. [52] and Ref. [53] for an

²The nuisance parameters for systematic uncertainties of the same type are assumed fully correlated between measurements.

example using the inclusive jet and dijet double-differential cross-section measurements in ATLAS), but this requires access to the event numbers in data (and to a lesser extent, Monte Carlo) and the exact analysis selection code. These are usually only available within a collaboration and if this is not done soon after an analysis is done, it often becomes impossible or at least impractical later. In contrast, if all relevant dimensions are included in the initial unfolding, then performing a fit to multiple spectra is straightforward. See e.g. the discussion of Refs. [54–56] in the case of PDF analyses using top quark pair production measurements.

High-dimensional measurements are particularly important for probing hadronic final states, where it is important to simultaneously probe multiple objects. For example, studies of azimuthal modulations in semi-inclusive deep inelastic scattering (SIDIS) and related processes such as hadron-in-jet measurements need to include both electron and hadron variables [57]. Such measurements often require unfolding $\mathcal{O}(10)$ dimensions at once, which is challenging with binned approaches. It is known that limited geometrical acceptance of detectors introduce coupling between different harmonics, and those effects are compounded when performing binned measurements, especially when integrating over kinematic variables, see e.g. Ref. [58]. Furthermore, the detector resolution can depend on many quantities, so simultaneously including those observables will be necessary to achieve the ultimate precision.

All in all, there are strong reasons why unbinned measurements would represent a breakthrough in the way particle and nuclear physics measurements at the LHC and elsewhere are both presented and used in theoretical interpretations, which would be specially timely with the upcoming era of the high-statistics High Luminosity LHC [59, 60] and the electron ion collider (EIC) [61].

3 Overview of Unbinned Methods

As the goal of the proposed format is to be unfolding-procedure agnostic, it is important to briefly introduce current proposals so that the resulting format can be universally accommodating. This section is not meant to be comprehensive – please see the original papers for further details. Additional examples of generative models or classification models in high energy physics can be found in Ref. [62].

3.1 Experimental Measurements

For the remainder of this section, each data event will be represented by $x \in \mathbb{R}^n$, where n could be one for a single target observable or $n > 1$ for a multidimensional unfolding. The dimensionality of the detector-level and particle-level observables need not be the same and unless specified, x will refer to the particle-level features.

3.1.1 Density-Based Models

Approaches of this type learn an implicit or explicit function $\hat{p}(x) : \mathbb{R}^n \rightarrow \mathbb{R}$ that represents the probability density of the unfolded result that is designed to estimate the true density $p(x)$. Low-dimensional approaches without neural networks have been proposed based on

non-parametric density estimators [63]. Deep generative models like Generative Adversarial Networks (GAN) [64, 65] and Variational Autoencoders (VAE) [66, 67] produce implicit models for \hat{p} , i.e. they can sample from \hat{p} but do not provide a method for evaluating it explicitly. In contrast, deep models based on Normalizing Flows (NF) [68, 69] allow for both sampling and density estimation. GAN-based proposals include Ref. [27, 30], Ref. [34] relies on a VAE, and Ref. [29, 35] proposed using NFs.

3.1.2 Classifier-Based Models

Approaches of this type learn a function $\hat{w}(x)$ that approximates the likelihood ratio $w(x) = p(x)/p_{\text{MC}}(x)$, where $p(x)$ is the true probability density of the unfolded data as discussed in the previous subsection, and p_{MC} is the equivalent quantity predicted by a Monte Carlo algorithm. With this function one can sample events from the nominal MC and weight each event by $\hat{w}(x_i)$, where x_i are the features of the event. The produced sample has now been reweighted to match the unfolded data and constitutes the (central value) of the unbinned measurement. Any expectation values computed using this sample can be used as the central value of a measurement of this quantity, and if we include a normalization factor as part of $\hat{w}(x)$ that account for the integrated luminosity of the dataset, the sum of weights of any kinematic region defined from x will correspond to the associated fiducial cross section. Learning \hat{w} is achieved by making use of a well-known [70, 71] fact that one can directly learn likelihood ratios using classifiers, e.g. $\hat{w}(x) = f(x)/(1 - f(x))$ for classifier $f(x)$ trained with the cross-entropy loss function. This turns the problem from density estimation (estimating $p(x)$ and $p_{\text{MC}}(x)$ separately) into classification. Classifier-based proposals include Ref. [28, 31–33].

3.2 Theoretical Calculations

While the main motivation for the format proposed in this report is experimental measurements, it can also be used for theory predictions. In fact, we strongly encourage the publication of theoretical predictions alongside experimental measurements. If the latter are presented unbinned, then the predictions should also be presented unbinned³. Unbinned predictions could also be provided before a measurement is done, without the need to iterate between theorists and experimentalists on the exact analysis binning. Related ideas for unbinned theory predictions have been discussed in the context of (next-to)next-to leading order n -tuples (see Ref. [72] and references therein) whereby all or part of a calculation can be stored for later reuse. Just like for experimental results, theoretical predictions come with uncertainties. Some of these uncertainties can be computed ‘on the fly’ (e.g. in parton showers [73–75] and parton densities [76]) if certain information like the colliding particle types and energies are stored. This information is not observable (in machine learning, this is called ‘latent’) and so would not have a corresponding entry in an unbinned measurement. While it could be interesting to store such information for on-the-fly calculation of uncertainties, we propose to focus on fully prepared results that do

³Sometimes, predictions are complemented by weights that represent corrections for higher-orders (k -factors’) or for non-perturbative effects. These could also be unbinned (and maybe could use machine learning for parameterization in a given kinematic feature like jet p_T).

not need any additional machinery for interpretation. This could be revisited in the future. However, it is necessary to store the output of the weight variations so that uncertainties can be computed, even if the weights cannot be rederived from the available (observable) information.

4 Statistics of Unbinned Spectra

4.1 Acceptance Effects

An essential aspect of unfolded spectra is that they should not depend on detector-level quantities. This means that in addition to correcting for resolution effects, unfolding algorithms must also account for finite acceptance effects. One way to do this unbinned was proposed in Ref. [31]: simply assign an ‘empty’ event symbol to events that pass the detector-level or particle-level event selection (but not both). Particle-level empty events are discarded from the final result. One should also be cautious of matching objects between detector-level and particle-level because this can introduce a reliance on the detector-level objects. For example, if one measures properties about the leading jet in an event, it is important to not use a p_T -matching to pick the particle-level jet – often it is better⁴ to use the leading particle-level jet independent of how close it might be to the leading detector-level jet.

4.2 Background Subtraction

Many measurements have contamination from background processes that must be statistically subtracted to produce a measurement in the desired region of particle-level phase space. For example, a measurement involving charged leptons may need to estimate and subtract the contribution from events with *fake* charged leptons. When measuring binned spectra, this can be done on a bin-by-bin basis. There are multiple unbinned generalizations. For density-based methods (Sec. 3.1.1), one can learn a subtracted density [77]. For classifier-based approaches (Sec. 3.1.2), one can assign a negative weight to events that are to be statistically subtracted. If one wants to avoid negative weights, unbinned positive reweighting is possible [78]. The resulting weighted events could also be turned into unweighted events, possibly based also on neural networks [79, 80].

Many backgrounds are estimated using data-driven methods like the fake factor and ABCD methods. These approaches are usually constructed with bins, but unbinned generalizations are possible.

4.3 Local Statistical Uncertainty

To compute the expectation value of an observable from N weighted events, one needs to average that observable over the event weights:

$$\hat{\mu}_{\mathcal{O}} = \sum_{i=1}^N w_i \mathcal{O}(x_i), \quad (4.1)$$

⁴It may instead be better to use a geometric matching. That renders the measurement dependent on the detector-level observables, but given the superb angular resolution of modern detectors, this may be a negligible effect. Another approach is to use a different scheme altogether (e.g. rapidity ordering [41]).

where w_i are the individual event weights and \mathcal{O} is the observable on interest that depends on the event features x_i . For example, in the case of a binned differential cross section, the observable would be the indicator function for a histogram bin (0 if the event lands outside of the bin of interest, 1 if it lands inside). The functions \mathcal{O} and/or the weights w may be dimensionful depending on the application. For example, in the case of a binned differential cross section, $\hat{\mu}$ will have units of cross section divided by the units of the differential quantity. When measuring a summary statistic like the average of a quantity, the weights w_i are normalized to unity.

In the large N limit, the standard way to compute statistical uncertainties on $\hat{\mu}_{\mathcal{O}}$ is to sum over the squared weights:

$$\delta\hat{\mu}_{\mathcal{O}} = \sqrt{\sum_{i=1}^N w_i^2 \mathcal{O}(x_i)^2}. \quad (4.2)$$

This corresponds to estimating the standard error of the mean and dropping terms that are higher-order in N . For equal weight events and a histogram-based observable, this yields the standard result that $\delta\hat{\mu}_{\mathcal{O}}/\hat{\mu}_{\mathcal{O}} = 1/\sqrt{N_{\mathcal{O}}}$, where $N_{\mathcal{O}}$ is the number of events in the histogram bin.

Because Eqs. (4.1) and (4.2) depend on the same w_i , one can derive the standard statistical uncertainty by making sure each event is stored with its own weight. There are contexts, however, where the standard formula does not apply. For example, if there is a large variability in the event weights or if there are negative weight events, it may be desirable to reweight the events in a local phase space patch while preserving the cross section of that patch [81]:

$$w_{\text{patch}} = \frac{1}{N_{\text{patch}}} \sum_{\text{events } i \text{ in patch}} w_i. \quad (4.3)$$

For small enough patches, this replacement preserves the expectation value in Eq. (4.1). As emphasized in Ref. [78], however, it does not preserve the uncertainties:

$$w_{\text{patch}}^2 \neq \frac{1}{N_{\text{patch}}} \sum_{\text{events } i \text{ in patch}} w_i^2. \quad (4.4)$$

The solution advocated in Ref. [78] was to resample events, such that a fraction of events are removed and the remaining events are upweighted such that both Eqs. (4.1) and (4.2) are preserved. This solution might not be desirable in all circumstances, though, so we advocate for separately storing

$$w_i, \quad \tilde{w}_i^2, \quad (4.5)$$

where \tilde{w}_i^2 has the interpretation as the local average squared weight. In the reweighting example above, one would simply store the right hand sides of Eqs. (4.3) and (4.4). In situations where this extra information is not needed (i.e. there is no resampling), a value of $\tilde{w}_i^2 = w_i^2$ should be used.

5 Systematic Uncertainties

Typically, systematic uncertainties are estimated by varying some aspect of the simulation and repeating the unfolding procedure. Such uncertainties can be classified into two-point and replica-type uncertainties (discussed more in Sec. 6.2). This section briefly discusses uncertainties related to the unfolding procedure itself, which may be more involved when using machine learning.

Potential biases from the unfolding procedure are typically evaluated through a set of non-closure tests. In particular, the unfolding procedure is repeated for a case where the answer is known and the difference between the unfolded result and the known answer can be used to set a systematic uncertainty on the method bias. Examples of this type include unfolding one generator (labeled ‘data’) with another generator (labeled ‘simulation’) as well as stress-tests whereby one generator (labeled ‘data’) is deformed in some way and it is then unfolded with the same, non-deformed generator (labeled ‘simulation’). For example, a stress test may be created by weighting each event so that the ratio between the weighted and unweighted distribution is linear with a particular slope. As with all systematic uncertainties, there is no unique way to determine the non-closure uncertainty and ideally the data should be used in some way to bound the range of reasonable variations studied. Indeed, the method proposed in Ref. [82] allows to perform a data-driven closure test, introducing weights at particle-level so that the corresponding detector-level distribution is a better match to the data. This approach has the advantage of emulating the conditions of the actual data unfolding, hence its reliability. In general, several non-closure tests have to be performed and their output synthesized; and one has to make sure that the tests probe complementary effects, in order to avoid a possible double counting of some effects between such tests.

The same non-closure uncertainties discussed above can be used in the unbinned case. No additional uncertainties are required simply because neural networks are used. However, the architecture and training of neural networks may lead to sub-optimal performance and it may be useful to estimate the potential deviation from optimal performance (referred to as ‘optimally uncertainties’ in Ref. [48]) to ensure that they are small. In particular, density-based methods (Sec. 3.1.1) assume that the learned density is the true density and classifier-methods (Sec. 3.1.2) assume that the learned function is a monotonic transformation of the true likelihood-ratio. If ‘optimally uncertainties’ are not small, then the architecture/learning could be improved or additional steps could be added to enhance accuracy (such as calibration [83]). There are multiple ways to probe these uncertainties, including reinitializing with different starting parameters, bootstrapping, etc. Bayesian neural networks may also be able to help quantify some of these ‘uncertainties’ [84].

6 Proposed Format and Storage

6.1 Structure

There are generically two ways to represent unbinned data. If the results can be encoded by a parameterized probability density, then one could simply publish the functional form

and fitted parameters. A second approach would be to publish data sampled from the unbinned result. This second approach is more general as not all methods produce a fitted function (see Sec. 3) and therefore it is the baseline for our proposal⁵. However, we strongly encourage publishing any other representation as well, including neural network architectures (with fitted weights and biases) which can be used to either generate new samples (Sec. 3.1.1) or reweight new samples (Sec. 3.1.2). Neural networks would ideally be converted to a standard format like the Open Neural Network Exchange (ONNX) [85] or the native format if a given architecture is not supported in ONNX or precision is lost. If converting to a non-native format, it is important to verify that there is no loss in precision, which may be especially important for likelihood ratio approximations.

The proposal is to have two files: (1) a `submission` file, which has the same format as the `submission` YAML file used for HEPDATA and (2) a data file for each measurement. We encourage analyzers to include a default binning with the corresponding histogram contents for easy comparison and validation. The location of the data files are in the `submission` file. Note that a given analysis may include multiple measurements; also a given measurement could have multiple dimensions and from it, multiple observables could be constructed. Suppose that a given measurement has k -dimensions per event. Further suppose that there are N events sampled from the unfolded result and there are M systematic uncertainties. Then, the data files for the measurement will be a matrix of size $(M + 1) \times (k + 2) \times N$. A common block is repeated for the nominal and all uncertainties. This block has N columns and $k + 2$ rows. The first k rows are simply the values of the observable $\vec{x}_i \in \mathbb{R}^k$. The next row encodes per-event weights and the last row encodes the per-event statistical uncertainty. In most cases, the per-event statistical uncertainty will simply be the per-event weight squared, but it need not be the case (see Sec. 4). The structure of this block is illustrated below, for an algorithm that does not use event weights (top) and one that does (bottom):

Example: density-based approach (no event-weights)

$\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N$
 $1, 1, \dots, 1$
 $1, 1, \dots, 1$

Example: classifier-based approach

$\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N$
 w_1, w_2, \dots, w_N
 $w_1^2, w_2^2, \dots, w_N^2$

⁵In the future, there may also be interest in unfolding *parameterized observables* such that the observables themselves depend on an input μ . This could be done by providing values sampled in μ , but this could result in some loss of information depending on how the unfolding is performed. Exploring this possibility is left for future work.

Note that the above tables are represented as one event per column to fit on the page, but we envision that actually each event will be represented as a row (the transpose of the above tables). Furthermore, for classifier-based methods, for systematic variations that use the same particle-level events, but with different detector-level events or different weights, it could be possible to use the same feature list with a different set of unfolded weights for each uncertainty. In this case, only $(k + 2 + 2M) \times N$ numbers are required (or even fewer if the local uncertainty is w^2).

6.2 Uncertainties

There are two types of uncertainties: two-point and replica. Two-point uncertainties result from modifying a discrete parameter in the simulation and then rerunning the unfolding. In this case, there will be an entry in the output that has the same structure as the nominal result, but with different values (either different \vec{x}_i , different weights, or both). When the unfolded data are assembled into a bin, the uncertainty on the bin content is computed by taking the difference in the bin content when using the nominal entry and the varied entry.

Replica uncertainties are constructed by sampling many unfolded results from a distribution. They are described by not one entry, but many entries and the uncertainty when binned is computed by taking the standard deviation across values. The most common replica uncertainties are due to statistical uncertainties where the data or simulation are bootstrapped and the unfolding is repeated a number of times. Replica uncertainties are also encountered for resolution uncertainties where additional smearing may be added to a reconstructed object many times to determine the impact of an increased resolution.

6.3 Storage

Unlike recording histograms, storing unbinned data in the proposed format may require significant storage space ($\mathcal{O}(\text{GB})$). We therefore propose to host the results on ZENODO (<https://zenodo.org>), but linked to HEPDATA to make it easier to search and integrate with Rivet routines [86].

The access to the content of these large files can be streamlined to allow exploration of the data on the server while downloading only a relevant sub-region of the data. This would allow a user to download only the relevant observables from multiple files in an automated fashion and perform a combined analysis. The HDF5 format [87] allows for querying, and the HDF Server (h5serv) [88] defines universal unique identifiers (UUID) for each group and dataset object, which means that the user does not require prior knowledge of the HDF5 path in order to access a group in the file. Alternate implementations such as with XRootD [89] if using root format files or with Paquet [90] and Arrow [91] are also possible. The additional effort to add server side support required for these implementations is relatively small and will become worthwhile as presenting unbinned data becomes widespread.

6.4 Common Tools

The integration with standard analysis statistical (e.g. ROOUNFOLD [92]) and preservation tools (e.g. RIVET [86]) will depend on and benefit from the solution chosen for storage and

presentation of the unfolded result. This would require extensions of common tools, such as a machine learning interface to ROOUNFOLD and a utility in RIVET to rebin data on the fly. However, given the widespread use of these tools for comparing with predictions, it will be important to add to these packages as unbinned measurements become popular. Additional harmonization would also be useful, such as the standardization of names for various nuisance parameters. While this is not strictly a requirement of unbinned unfolding, it would facilitate easier comparisons and combinations with existing and future measurements.

7 Example

As an example, we consider the $Z + \text{jets}$ production process. The main purpose of this section, is to show in practice the differences of density-based and classifier-based models. As a classifier-based model we utilize the OMNIFOLD [28] approach which is an unbinned and multi-dimensional version of the well-established IBU [93–95] method. As a density-based model we employ a conditional invertible neural network (cINN) as being used in Ref. [29] which directly learns a stochastic mapping from detector-level to particle-level. The code to reproduce both methods can be found at [this Github repository](#).

We use the same simulations as in Ref. [28], which are briefly summarized in the following. Proton-proton collisions are simulated at $\sqrt{s} = 14$ TeV using PYTHIA 8.243 [96–98] with Tune 26 [99]. Detector effects are emulated using the DELPHES 3.4.2 [100] fast simulation of the CMS detector, which uses particle flow reconstruction. The PYTHIA simulation will be called ‘data’ and is also the model for the nominal ‘simulation’. An alternative simulation, to illustrate a systematic uncertainty, is provided by HERWIG 7.1.5 [101–103] with its default tune. Jets are clustered using the anti- k_t algorithm [104] with radius parameter $R = 0.4$ as implemented in FASTJET 3.3.2 [105, 106], using either all particle flow objects (detector-level) or stable non-neutrino truth particles (particle-level). To reduce acceptance effects, the leading jets are studied in events with a Z boson with transverse momentum $p_T^Z > 200$ GeV. A set of 500,000 events are set aside as ‘data’ and 500,000 events from both PYTHIA and HERWIG are then used for the simulation. To keep this example focused, we consider a two-dimensional unfolding of the N -subjettiness variables τ_1 and τ_2 [107, 108].

Both neural network models are implemented in TENSORFLOW [109] and optimized using ADAM [110]. The OMNIFOLD network consists of three hidden layers with 50 nodes per layer using the rectified linear activation function (ReLU) for intermediate layers and a sigmoid for the final layer. Training proceeds for 20 epochs with a batch size of 10,000 and the algorithm is repeated for five iterations. The cINN stacks 12 affine coupling blocks, where each block consists of a fully connected multi-layer perceptron with 4 hidden layers, 16 nodes per layer and leaky ReLU as activation function. It is trained for 50 epochs and a batch size of 1000. None of these parameters were optimized. An schematic overview comparing both employed models is illustrated in Fig. 1.

In the case of the cINN, the nominal result is built by training on the particle-level and detector-level ‘simulation’ and then applied to the detector-level ‘data’. Ten particle-

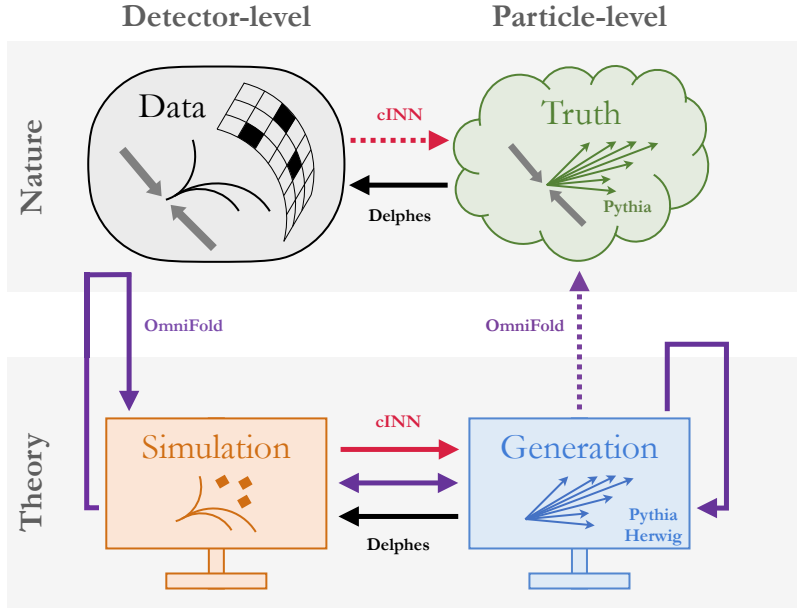


Figure 1. An schematic overview illustrating which parts are connected via several different models and tools. The black arrows indicate the usage of DELPHES. The red arrows show which parts are connected by the cINN method. In contrast, the iterative procedure of OMNIFOLD is highlighted by the purple arrows. Adapted from Ref. [28].

level events are sampled from the density for each detector-level ‘data’ event. In order to preserve the total cross section, each event is then given a weight of $w_i = 0.1$. The first five sampled events in the proposed format are shown below ($\vec{x}_i = (\tau_{1,i}, \tau_{2,i})$):

cINN: Nominal

(0.093,0.067), (0.138,0.079), (0.095,0.066), (0.056,0.054), (0.362,0.261) , ...
 0.1, 0.1, 0.1, 0.1, 0.1 , ...
 0.01, 0.01, 0.01, 0.01, 0.01 , ...

The systematic uncertainty due to the MC model is estimated similarly. In particular, the cINN is trained using the particle-level and detector-level samples from the alternative simulation. This model is then applied to the ‘data’, where ten events are sampled from each ‘data’ event. The first five samples are shown here:

cINN: Systematic Uncertainty (Alternative MC)

(0.059,0.042), (0.150,0.102), (0.117,0.086), (0.090,0.065), (0.091,0.050) , ...
 0.1, 0.1, 0.1, 0.1, 0.1 , ...
 0.01, 0.01, 0.01, 0.01, 0.01 , ...

For the nominal OmniFold result, both the ‘data’ and ‘simulation’ are used to derive

weights. Since the ‘data’ and ‘simulation’ are statistically identical, the weights are nearly constant and unity, as shown below:

OmniFold: Nominal

(0.070,0.071), (0.206,0.181), (0.117,0.053), (0.292,0.122), (0.029,0.027) , ...
 1.003, 1.003, 1.003, 1.003, 1.003 , ...
 1.007, 1.007, 1.007, 1.007, 1.007 , ...

Finally, the alternative MC OmniFold result is shown below. Now that the data and corresponding simulation are not statistically identical, the weights follow a non-trivial distribution:

OmniFold: Systematic Uncertainty (Alternative MC)

(0.185,0.086), (0.063,0.065), (0.400,0.156), (0.179,0.125), (0.050,0.047) , ...
 0.696, 1.626, 0.658, 0.777, 1.713 , ...
 0.485, 2.645, 0.433, 0.604, 2.934 , ...

A histogram of the weights for all methods is presented in Fig. 2 and the corresponding unfolded distributions, with a systematic uncertainty band is shown in Fig. 3. As advertised in Sec. 2, a benefit of unbinned unfolding is that it is easy to define new observables. This is demonstrated in Fig. 4, which shows unfolded results for $\tau_{21} = \tau_2/\tau_1$, which can be derived from the two-dimensional unfolded phase space. The OmniFold and cINN results are comparable; as stated earlier, neither method was optimized and it would be exciting to quantitatively compare the two approaches in the future.

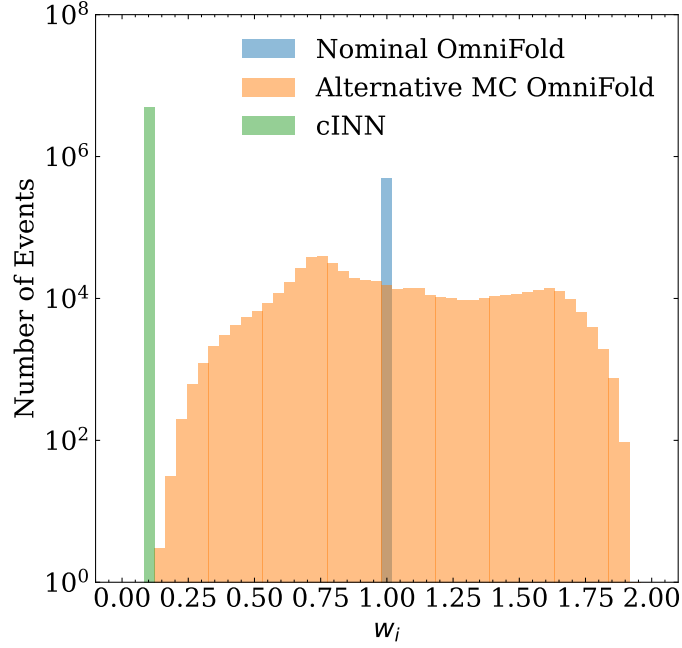


Figure 2. A histogram of the weights for each method. By definition, the weights for the cINN are 0.1, since 10 events are sampled from the density for each event in ‘data’. There are correspondingly ten times more events for the cINN than for OmniFold. For OmniFold, the nominal weights are nearly all one since the data and simulation are statistically identical.

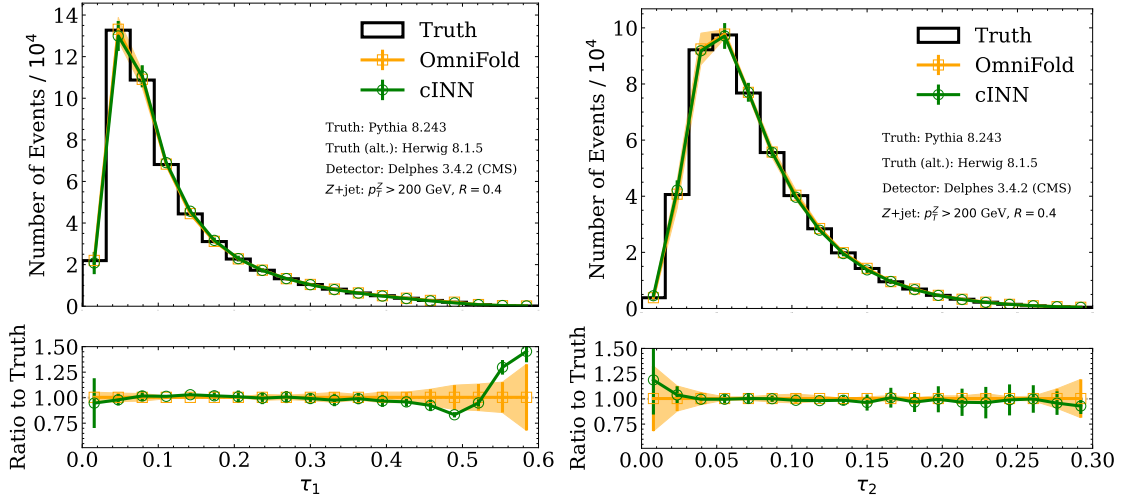


Figure 3. The measured histograms of τ_1 (left) and τ_2 (right) using OMNIFOLD and the cINN methods. The error band represents the modeling uncertainty from comparing PYTHIA and HERWIG.




Figure 4. The measured histogram of $\tau_{21} = \tau_2/\tau_1$ using OMNIFOLD and the cINN methods. The error band represents the modeling uncertainty from comparing PYTHIA and HERWIG.

8 Conclusions and Outlook

This report has presented a proposal for publishing unbinned differential cross section measurements and predictions. The methodology for unbinned measurements has been enhanced in part because of recent advances in machine learning. Unbinned measurements can always be rebinned and analyzed with traditional methods. In order to make best use of future measurements that use these tools, it is important to have some community standards and guidelines. We hope that this paper will serve the purpose of having a concrete proposal to discuss with all of the relevant experimental and theoretical communities. This is foreseen to be the start of an evolving community dialogue, in order to accommodate future developments in this field that is currently rapidly evolving.

Acknowledgments

AB is supported by the DFG (German Research Foundation) under grant 396021762 – TRR 257. BM gratefully acknowledges the continuous support from LPNHE, CNRS/IN2P3, Sorbonne Université and Université de Paris. BN was supported by the Department of Energy, Office of Science under contract number DE-AC02-05CH11231. JR is partially supported by NWO, the Dutch Science Council. JT is supported by the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, <http://iaifi.org/>), and by the U.S. DOE Office of High Energy Physics under grant number DE-SC0012567. RW is supported by FRS-FNRS (Belgian National Scientific Research Fund) IISN projects 4.4503.16. We would like to thank the organizers of the Les Houches PhysTeV workshop series (especially Emanuele Re), which hosted the online forum that began this document in June of 2021.

References

- [1] E. Maguire, L. Heinrich and G. Watt, *HEPData: a repository for high energy physics data*, *J. Phys. Conf. Ser.* **898** (2017) 102006, [[1704.05473](#)].
- [2] A. Buckley, H. Hoeth, H. Lacker, H. Schulz and J. E. von Seggern, *Systematic event generator tuning for the LHC*, *Eur. Phys. J. C* **65** (2010) 331–357, [[0907.2973](#)].
- [3] J. Gao, L. Harland-Lang and J. Rojo, *The Structure of the Proton in the LHC Precision Era*, *Phys. Rept.* **742** (2018) 1–121, [[1709.04922](#)].
- [4] R. Abdul Khalek, J. J. Ethier, J. Rojo and G. van Weelden, *nNNPDF2.0: quark flavor separation in nuclei from LHC data*, *JHEP* **09** (2020) 183, [[2006.14629](#)].
- [5] I. Brivio and M. Trott, *The Standard Model as an Effective Field Theory*, *Phys. Rept.* **793** (2019) 1–98, [[1706.08945](#)].
- [6] A. Biekötter, T. Corbett and T. Plehn, *The Gauge-Higgs Legacy of the LHC Run II*, *SciPost Phys.* **6** (2019) 064, [[1812.07587](#)].
- [7] E. da Silva Almeida, A. Alves, N. Rosa Agostinho, O. J. P. Éboli and M. C. Gonzalez-Garcia, *Electroweak Sector Under Scrutiny: A Combined Analysis of LHC and Electroweak Precision Data*, *Phys. Rev. D* **99** (2019) 033001, [[1812.01009](#)].

- [8] I. Brivio, S. Bruggisser, F. Maltoni, R. Moutafis, T. Plehn, E. Vryonidou et al., *O new physics, where art thou? A global search in the top sector*, *JHEP* **02** (2020) 131, [[1910.03606](#)].
- [9] S. Dawson, S. Homiller and S. D. Lane, *Putting standard model EFT fits to work*, *Phys. Rev. D* **102** (2020) 055012, [[2007.01296](#)].
- [10] J. Ellis, M. Madigan, K. Mimasu, V. Sanz and T. You, *Top, Higgs, Diboson and Electroweak Fit to the Standard Model Effective Field Theory*, *JHEP* **04** (2021) 279, [[2012.02779](#)].
- [11] J. J. Ethier, R. Gomez-Ambrosio, G. Magni and J. Rojo, *SMEFT analysis of vector boson scattering and diboson data from the LHC Run II*, *Eur. Phys. J. C* **81** (2021) 560, [[2101.03180](#)].
- [12] J. J. Ethier, G. Magni, F. Maltoni, L. Mantani, E. R. Nocera, J. Rojo et al., *Combined SMEFT interpretation of Higgs, diboson, and top quark data from the LHC*, [2105.00006](#).
- [13] I. Brivio, S. Bruggisser, E. Geoffray, W. Kilian, M. Krämer, M. Luchmann et al., *From Models to SMEFT and Back?*, [2108.01094](#).
- [14] E. d. S. Almeida, A. Alves, O. J. P. Éboli and M. C. Gonzalez-Garcia, *Electroweak legacy of the LHC Run II*, [2108.04828](#).
- [15] GAMBIT collaboration, P. Athron et al., *Global fits of GUT-scale SUSY models with GAMBIT*, *Eur. Phys. J. C* **77** (2017) 824, [[1705.07935](#)].
- [16] E. A. Bagnaschi et al., *Supersymmetric Dark Matter after LHC Run 1*, *Eur. Phys. J. C* **75** (2015) 500, [[1508.01173](#)].
- [17] G. Cowan, *A survey of unfolding methods for particle physics*, *Conf. Proc.* **C0203181** (2002) 248.
- [18] V. Blobel, *Unfolding Methods in Particle Physics*, *PHYSTAT2011 Proceedings* (2011) 240.
- [19] V. Blobel, *Unfolding*, *Data Analysis in High Energy Physics* (2013) 187.
- [20] R. Balasubramanian, L. Brenner, C. Burgard, G. Cowan, V. Croft, W. Verkerke et al., *Statistical method and comparison of different unfolding techniques using RooFit*, [1910.14654](#).
- [21] H. P. e. e. a. K. Cranmer, S. Kraml, *Publishing statistical models: Getting the most out of particle physics experiments*, [2109.04981](#).
- [22] L. Lindemann and G. Zech, *Unfolding by weighting Monte Carlo events*, *Nucl. Instrum. Meth.* **A354** (1995) 516–521.
- [23] G. Zech and B. Aslan, *Binning-Free Unfolding Based on Monte Carlo Migration*, *eConf* **C030908** (2003) TUGT001.
- [24] H. P. Dembinski and M. Roth, *An algorithm for automatic unfolding of one-dimensional data distributions*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **729** (2013) 410–416.
- [25] N. D. Gagunashvili, *Machine learning approach to inverse problem and unfolding procedure*, [1004.2006](#).
- [26] A. Glazov, *Machine learning as an instrument for data unfolding*, [1712.01814](#).
- [27] K. Datta, D. Kar and D. Roy, *Unfolding with Generative Adversarial Networks*, [1806.00433](#).

- [28] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman and J. Thaler, *OmniFold: A Method to Simultaneously Unfold All Observables*, *Phys. Rev. Lett.* **124** (2020) 182001, [[1911.09107](#)].
- [29] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, A. Rousselot, R. Winterhalder et al., *Invertible Networks or Partons to Detector and Back Again*, *SciPost Phys.* **9** (2020) 074, [[2006.06685](#)].
- [30] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn and R. Winterhalder, *How to GAN away Detector Effects*, [1912.00477](#).
- [31] 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 *9th International Conference on Learning Representations*, 5, 2021, [2105.04448](#).
- [32] M. Bunse, N. Piatkowski, T. Ruhe, W. Rhode and K. Morik, *Unification of deconvolution algorithms for Cherenkov astronomy*, in *5th International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 21–30, IEEE, 2018.
- [33] T. Ruhe, T. Voigt, M. Wornowizki, M. Börner, W. Rhode and K. Morik, *Mining for spectra - the dortmund spectrum estimation algorithm*, 2019.
- [34] J. N. Howard, S. Mandt, D. Whiteson and Y. Yang, *Foundations of a Fast, Data-Driven, Machine-Learned Simulator*, [2101.08944](#).
- [35] M. Vandegar, M. Kagan, A. Wehenkel and 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* (A. Banerjee and K. Fukumizu, eds.), vol. 130 of *Proceedings of Machine Learning Research*, pp. 2107–2115, PMLR, 11, 2021, [2011.05836](#), <https://proceedings.mlr.press/v130/vandegar21a.html>.
- [36] H1 Collaboration, *Measurement of lepton-jet correlation in deep-inelastic scattering with the H1 detector using machine learning for unfolding*, [2108.12376](#).
- [37] R. D. Ball, V. Bertone, F. Cerutti, L. Del Debbio, S. Forte, A. Guffanti et al., *Reweighting and Unweighting of Parton Distributions and the LHC W lepton asymmetry data*, *Nucl. Phys. B* **855** (2012) 608–638, [[1108.1758](#)].
- [38] H. Paukkunen and P. Zurita, *PDF reweighting in the Hessian matrix approach*, *JHEP* **12** (2014) 100, [[1402.6623](#)].
- [39] J. Brehmer, K. Cranmer, F. Kling and T. Plehn, *Better Higgs boson measurements through information geometry*, *Phys. Rev. D* **95** (2017) 073002, [[1612.05261](#)].
- [40] CLAS collaboration, M. Battaglieri et al., *Photoproduction of π^+ π^- meson pairs on the proton*, *Phys. Rev. D* **80** (2009) 072005, [[0907.1021](#)].
- [41] ATLAS Collaboration, *Measurement of jet charge in dijet events from $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*, *Phys. Rev. D* **93** (2016) 052003, [[1509.05190](#)].
- [42] S. Chen, A. Glioti, G. Panico and A. Wulzer, *Parametrized classifiers for optimal EFT sensitivity*, *JHEP* **05** (2021) 247, [[2007.10356](#)].
- [43] R. T. D’Agnolo and A. Wulzer, *Learning New Physics from a Machine*, *Phys. Rev. D* **99** (2019) 015014, [[1806.02350](#)].

- [44] R. T. D’Agnolo, G. Grosso, M. Pierini, A. Wulzer and M. Zanetti, *Learning multivariate new physics*, *Eur. Phys. J. C* **81** (2021) 89, [[1912.12155](#)].
- [45] J. Brehmer, G. Louppe, J. Pavez and K. Cranmer, *Mining gold from implicit models to improve likelihood-free inference*, [1805.12244](#).
- [46] J. Brehmer, K. Cranmer, G. Louppe and J. Pavez, *Constraining Effective Field Theories with Machine Learning*, *Phys. Rev. Lett.* **121** (2018) 111801, [[1805.00013](#)].
- [47] J. Brehmer, K. Cranmer, G. Louppe and J. Pavez, *A Guide to Constraining Effective Field Theories with Machine Learning*, *Phys. Rev.* **D98** (2018) 052004, [[1805.00020](#)].
- [48] B. Nachman, *A guide for deploying Deep Learning in LHC searches: How to achieve optimality and account for uncertainty*, [1909.03081](#).
- [49] A. Greljo, S. Iranipour, Z. Kassabov, M. Madigan, J. Moore, J. Rojo et al., *Parton distributions in the SMEFT from high-energy Drell-Yan tails*, *JHEP* **07** (2021) 122, [[2104.02723](#)].
- [50] X. Liu, F. Ringer, W. Vogelsang and F. Yuan, *Lepton-jet Correlations in Deep Inelastic Scattering at the Electron-Ion Collider*, *Phys. Rev. Lett.* **122** (2019) 192003, [[1812.08077](#)].
- [51] D. Gutierrez-Reyes, I. Scimemi, W. J. Waalewijn and L. Zoppi, *Transverse momentum dependent distributions with jets*, *Phys. Rev. Lett.* **121** (2018) 162001, [[1807.07573](#)].
- [52] ATLAS collaboration, *Evaluating statistical uncertainties and correlations using the bootstrap method*, Tech. Rep. ATL-PHYS-PUB-2021-011, CERN, Geneva, Apr, 2021.
- [53] ATLAS collaboration, M. Aaboud et al., *Measurement of inclusive jet and dijet cross-sections in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **05** (2018) 195, [[1711.02692](#)].
- [54] M. Czakon, N. P. Hartland, A. Mitov, E. R. Nocera and J. Rojo, *Pinning down the large- x gluon with NNLO top-quark pair differential distributions*, *JHEP* **04** (2017) 044, [[1611.08609](#)].
- [55] S. Bailey and L. Harland-Lang, *Differential Top Quark Pair Production at the LHC: Challenges for PDF Fits*, *Eur. Phys. J. C* **80** (2020) 60, [[1909.10541](#)].
- [56] T.-J. Hou et al., *New CTEQ global analysis of quantum chromodynamics with high-precision data from the LHC*, *Phys. Rev. D* **103** (2021) 014013, [[1912.10053](#)].
- [57] M. Arratia, Z.-B. Kang, A. Prokudin and F. Ringer, *Jet-based measurements of Sivers and Collins asymmetries at the future electron-ion collider*, *Phys. Rev. D* **102** (2020) 074015, [[2007.07281](#)].
- [58] T. B. Hayward et al., *Observation of Beam Spin Asymmetries in the Process $ep \rightarrow e' \pi^+ \pi^- X$ with CLAS12*, *Phys. Rev. Lett.* **126** (2021) 152501, [[2101.04842](#)].
- [59] P. Azzi et al., *Report from Working Group 1: Standard Model Physics at the HL-LHC and HE-LHC*, *CERN Yellow Rep. Monogr.* **7** (2019) 1–220, [[1902.04070](#)].
- [60] M. Cepeda et al., *Report from Working Group 2: Higgs Physics at the HL-LHC and HE-LHC*, *CERN Yellow Rep. Monogr.* **7** (2019) 221–584, [[1902.00134](#)].
- [61] R. Abdul Khalek et al., *Science Requirements and Detector Concepts for the Electron-Ion Collider: EIC Yellow Report*, [2103.05419](#).
- [62] M. Feickert and B. Nachman, *A Living Review of Machine Learning for Particle Physics*, [2102.02770](#).

- [63] H. P. Dembinski and M. Roth, *An algorithm for automatic unfolding of one-dimensional data distributions*, *Nucl. Instrum. Meth. A* **729** (2013) 410–416.
- [64] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair et al., *Generative adversarial nets*, in *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'14, (Cambridge, MA, USA), pp. 2672–2680, MIT Press, 2014, <http://dl.acm.org/citation.cfm?id=2969033.2969125>.
- [65] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta and A. A. Bharath, *Generative adversarial networks: An overview*, *IEEE Signal Processing Magazine* **35** (Jan, 2018) 53.
- [66] D. P. Kingma and M. Welling, *Auto-encoding variational bayes*, [1312.6114](#).
- [67] D. P. Kingma and M. Welling, *An Introduction to Variational Autoencoders*, *Foundations and Trends in Machine Learning* **12** (2019) 307.
- [68] D. J. Rezende and S. Mohamed, *Variational inference with normalizing flows*, *International Conference on Machine Learning* **37** (2015) 1530.
- [69] I. Kobyzev, S. Prince and M. Brubaker, *Normalizing Flows: An Introduction and Review of Current Methods*, *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020) 1.
- [70] T. Hastie, R. Tibshirani and J. Friedman, *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York Inc., New York, NY, USA, 2001.
- [71] M. Sugiyama, T. Suzuki and T. Kanamori, *Density Ratio Estimation in Machine Learning*. Cambridge University Press, 2012, [10.1017/CBO9781139035613](https://doi.org/10.1017/CBO9781139035613).
- [72] D. Maitre, G. Heinrich and M. Johnson, *N(N)LO event files: applications and prospects*, *PoS* **LL2016** (2016) 016, [[1607.06259](#)].
- [73] S. Mrenna and P. Skands, *Automated Parton-Shower Variations in Pythia 8*, *Phys. Rev. D* **94** (2016) 074005, [[1605.08352](#)].
- [74] J. Bellm, S. Plätzer, P. Richardson, A. Siódmok and S. Webster, *Reweighting Parton Showers*, *Phys. Rev. D* **94** (2016) 034028, [[1605.08256](#)].
- [75] E. Bothmann, M. Schönherr and S. Schumann, *Reweighting QCD matrix-element and parton-shower calculations*, *Eur. Phys. J. C* **76** (2016) 590, [[1606.08753](#)].
- [76] A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht et al., *LHAPDF6: parton density access in the LHC precision era*, *Eur. Phys. J. C* **75** (2015) 132, [[1412.7420](#)].
- [77] A. Butter, T. Plehn and R. Winterhalder, *How to GAN Event Subtraction*, [1912.08824](#).
- [78] B. Nachman and J. Thaler, *A Neural Resampler for Monte Carlo Reweighting with Preserved Uncertainties*, [2007.11586](#).
- [79] M. Backes, A. Butter, T. Plehn and R. Winterhalder, *How to GAN Event Unweighting*, *SciPost Phys.* **10** (2021) 089, [[2012.07873](#)].
- [80] B. Stienen and R. Verheyen, *Phase space sampling and inference from weighted events with autoregressive flows*, *SciPost Phys.* **10** (2021) 038, [[2011.13445](#)].
- [81] J. R. Andersen, C. Gütschow, A. Maier and S. Prestel, *A Positive Resampler for Monte Carlo events with negative weights*, *Eur. Phys. J. C* **80** (2020) 1007, [[2005.09375](#)].
- [82] B. Malaescu, *An Iterative, dynamically stabilized method of data unfolding*, [0907.3791](#).

- [83] K. Cranmer, J. Pavez and G. Louppe, *Approximating Likelihood Ratios with Calibrated Discriminative Classifiers*, [1506.02169](#).
- [84] M. Bellagente, M. Haußmann, M. Luchmann and T. Plehn, *Understanding Event-Generation Networks via Uncertainties*, [2104.04543](#).
- [85] J. Bai, F. Lu, K. Zhang et al., “Onnx: Open neural network exchange.”
- [86] A. Buckley, J. Butterworth, D. Grellscheid, H. Hoeth, L. Lonnblad, J. Monk et al., *Rivet user manual*, *Comput. Phys. Commun.* **184** (2013) 2803–2819, [[1003.0694](#)].
- [87] The HDF Group, “Hierarchical data format version 5.”
- [88] The HDF Group, “h5serv - REST-based service for HDF5 data.”
- [89] The XRootD Collaboration, “XRootD.”
- [90] The Apache Software Foundation, “Apache Parquet.”
- [91] The Apache Software Foundation, “Apache Arrow.”
- [92] T. Adye, *Unfolding algorithms and tests using rooufold*, 2011.
- [93] G. D’Agostini, *A multidimensional unfolding method based on bayes’ theorem*, *Nucl. Instrum. Meth.* **A362** (1995) 487–498.
- [94] L. B. Lucy, *An iterative technique for the rectification of observed distributions*, *Astronomical Journal* **79** (June, 1974) 745.
- [95] W. H. Richardson, *Bayesian-based iterative method of image restoration**, *J. Opt. Soc. Am.* **62** (Jan, 1972) 55–59.
- [96] T. Sjöstrand, S. Mrenna and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852–867, [[0710.3820](#)].
- [97] T. Sjöstrand, S. Mrenna and P. Z. Skands, *PYTHIA 6.4 Physics and Manual*, *JHEP* **05** (2006) 026, [[hep-ph/0603175](#)].
- [98] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159–177, [[1410.3012](#)].
- [99] *ATLAS Run 1 Pythia8 tunes*, Tech. Rep. ATL-PHYS-PUB-2014-021, CERN, Geneva, Nov, 2014.
- [100] DELPHES 3 collaboration, J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaitre, A. Mertens et al., *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057, [[1307.6346](#)].
- [101] M. Bahr et al., *Herwig++ Physics and Manual*, *Eur. Phys. J.* **C58** (2008) 639–707, [[0803.0883](#)].
- [102] J. Bellm et al., *Herwig 7.0/Herwig++ 3.0 release note*, *Eur. Phys. J.* **C76** (2016) 196, [[1512.01178](#)].
- [103] J. Bellm et al., *Herwig 7.1 Release Note*, [1705.06919](#).
- [104] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063, [[0802.1189](#)].
- [105] M. Cacciari, G. P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J.* **C72** (2012) 1896, [[1111.6097](#)].

- [106] M. Cacciari and G. P. Salam, *Dispelling the N^3 myth for the k_t jet-finder*, *Phys. Lett.* **B641** (2006) 57, [[hep-ph/0512210](#)].
- [107] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N -subjettiness*, *JHEP* **03** (2011) 015, [[1011.2268](#)].
- [108] J. Thaler and K. Van Tilburg, *Maximizing Boosted Top Identification by Minimizing N -subjettiness*, *JHEP* **02** (2012) 093, [[1108.2701](#)].
- [109] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean et al., *Tensorflow: A system for large-scale machine learning.*, in *OSDI*, vol. 16, pp. 265–283, 2016.
- [110] D. Kingma and J. Ba, *Adam: A method for stochastic optimization*, [1412.6980](#).