Neural resampler for Monte Carlo reweighting with preserved uncertainties

Benjamin Nachman^{1,*} and Jesse Thaler^{2,†}

¹*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

²*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

(Received 29 July 2020; accepted 17 September 2020; published 13 October 2020)

Monte Carlo event generators are an essential tool for data analysis in collider physics. To include subleading quantum corrections, these generators often need to produce negative weight events, which leads to statistical dilution of the datasets and downstream computational costs for detector simulation. Building on the recent proposal of a positive resampler method to rebalance weights within histogram bins, we introduce neural resampling: an unbinned approach to Monte Carlo reweighting based on neural networks that scales well to high-dimensional and variable-dimensional phase space. We pay particular attention to preserving the statistical properties of the event sample, such that neural resampling not only maintains the mean value of any observable but also its Monte Carlo uncertainty. This uncertainty preservation scheme is general and can also be applied to binned (non-neural network) resampling. To illustrate our neural resampling approach, we present a case study from the Large Hadron Collider of top quark pair production at next-to-leading order matched to a parton shower.

DOI: [10.1103/PhysRevD.102.076004](https://doi.org/10.1103/PhysRevD.102.076004)

I. INTRODUCTION

Data analysis in collider physics relies heavily on simulations to achieve precision. The state-of-the-art setup includes general purpose event generators like PYTHIA [1], HERWIG [2], and SHERPA [3], interfaced with next-to-leading order (NLO) corrections in the strong and/or electroweak couplings. These higher-order corrections are based on methods such as MC@NLO [4] and POWHEG [5] built into the general purpose generators or available as standalone packages like MG5_aMC@NLO [6] and POWHEG-BOX [7]. Outputs from these Monte Carlo generators are then fed into sophisticated detector simulation frameworks, such as those based on GEANT4 [8]. Event generation and simulation is becoming an increasingly relevant computational bottleneck for collider data analysis [9,10], particularly for the upcoming High Luminosity Large Hadron Collider (HL-LHC).

Part of the computational burden of Monte Carlo event production is related to the appearance of negative weight events. Each collision event produced by the pipeline above has an associated event weight w_i . The weights represent

the fact that a single simulated collision does not necessarily correspond to one real event. With the introduction of NLO corrections from quantum loops, some of these weights can even be negative, and the spread in event weights can get increasingly worse with higher-order calculations. As discussed more below, these negative weights complicate simulation-based inference and increase the computational demands.

In this paper, we present a *neural resampler* for Monte Carlo reweighting, which removes negative weights while preserving the statistical properties of the event sample. Our method builds upon the positive resampler approach introduced in Ref. [11], which uses histograms to determine bin-by-bin reweighting factors. The new feature of neural resampling is the use of neural networks to determine the reweighting factors, which allows us to work with the full unbinned, high-dimensional (and variable-dimensional) phase space. Another difference with Ref. [11] is that we pay particular attention to preserving the statistical uncertainties of the event sample, which we accomplish by pairing a local reweighting factor with a local resampling rate.

Let us review why weighted Monte Carlo events pose a computational burden for event generation and simulation. Here, we use “generation” to refer to particle-level event generation (e.g., PYTHIA) and “simulation” to refer to detector simulation (e.g., GEANT4). When a particle-level phase space point has a nontrivial spectrum of weights (e.g., positive and negative), this implies that a large number of generated events will be needed to achieve

*bpnachman@lbl.gov
†jthaler@mit.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

statistical accuracy. The reason is that only the average weight at a given phase space point is physically relevant, so if the local phase space weights have a large variance, then many generated events will be needed to accurately estimate the mean. In some cases, it might be possible to reduce the occurrence of negative weight events directly in the generation step [12], though this depends on the specific generator implementation. Barring that, a large number of particle-level events will need to be passed to the computationally expensive detector simulation step to achieve the desired accuracy.

The idea behind positive resampling [11] is to remove negative weights using a quasilocal weight rebalancing scheme. This method does not depend on how the particle-level events are generated and it is relevant anytime there is a nontrivial spectrum of weights (i.e., even if all the weights are positive but different). Though this method requires choosing a set of observables for histogram binning, the performance is very good when using a relatively small number of physically motivated observables. Note that positive resampling does not reduce the number of particle-level events that must be generated in order to achieve the desired statistical accuracy. Still, by rebalancing the event weights, positive resampling does decrease the number of particle-level events that are subsequently fed through the detector simulation. This can significantly reduce the computational cost of the full event production pipeline.

With neural resampling, we take advantage of the fact that neural networks, when paired with a suitable training algorithm, are excellent likelihood ratio estimators. This fact has been exploited in a variety of recent studies in high-energy physics [13–23]. The technique presented here is most closely related to the OMNIFOLD unfolding algorithm [20], which is based on the DCTR technique for full phase space reweighting [18]. Essentially, one can think of Monte Carlo reweighting as computing the likelihood ratio between the desired distribution and an unweighted baseline. Neural networks are particularly well suited for this task because they can naturally process the high-dimensional phase spaces encountered in Monte Carlo event generation. To handle variable-dimensional phase spaces, we take advantage of the particle flow network (PFN) architecture for point cloud learning [24,25].

After positive resampling or neural resampling, all of the event weights will be non-negative (assuming non-negative cross section), though they may not be uniform. The positive resampling method has a partial unweighting hyperparameter that determines what fraction of events are given a uniform weight [11]. In general, though, (partial) unweighting does not preserve the statistical properties of the event sample. Specifically, unweighting does preserve the expected mean of any observable, but it does not preserve the sample variance. As reviewed below, sample variance is the most common way to estimate

Monte Carlo uncertainties.¹ Crucially, neural resampling preserves the sample variance via a local resampling rate, without needing to choose any hyperparameters.

After neural resampling, one could optionally use standard unweighting techniques to make all of the weights uniform. While unweighting decreases the overall statistical power of the generated dataset by removing positive-weight events, it has the benefit of maximizing the per-event statistical power of the surviving subset. For this reason, there is a computational tradeoff between having weighted and unweighted events, with full unweighting being the best strategy when detector simulation is significantly more costly than event generation. The presence of negative-weight events, or any nontrivial distribution of weights at a given phase space point, is entirely deleterious and only slows statistical convergence. For this reason, neural resampling is always beneficial from the computational perspective, whether or not there is a subsequent unweighting step.

As an alternative use of neural networks, Monte Carlo unweighting could be achieved using generative models such as generative adversarial networks [26] and variational autoencoders [27,28]. Generative models have been extensively studied to accelerate or augment many aspects of high-energy physics simulations [29–60] and lattice field theory [61–64]. These approaches, however, require learning the full phase space density, instead of just the likelihood ratio, which is significantly more complicated than the neural resampling approach presented here. It is also worth mentioning that neural networks and other machine learning techniques have been studied to improve other components of event generation, including parton density modeling [65,66], phase space generation [67–71], matrix element calculations [72,73], and more [74–76].

The rest of this paper is organized as follows. In Sec. II, we review the statistics of event weights and introduce a method to preserve statistical uncertainties after local phase space reweighting. Then, Sec. III introduces our neural network-based resampling technique and shows how it can remove negative weights while preserving uncertainties. We demonstrate the performance of neural resampling for two case studies in Sec. IV: a simple example of two Gaussians and a realistic collider example of matched NLO top quark pair production. The paper ends with brief conclusions and outlook in Sec. V.

II. THE STATISTICS OF EVENT WEIGHTS

A. Review of weights and uncertainties

The output of a generic Monte Carlo event generator is a set of N simulated phase space points $\{x_i\}$ along with their associated weights $\{w_i\}$, sampled from random

¹The true variance, which is preserved by unweighting, is typically not accessible.

variables X and W , respectively. One can think of $X \in \mathbb{R}^d$ as representing the event features in a d -dimensional phase space. For simplicity, we work with normalized weights for this discussion:

$$\sum_{i=1}^N w_i = 1, \quad (1)$$

such that $w_i = 1/N$ for an unweighted event sample. It is straightforward to adapt the equations below to alternative weight conventions, and for the case studies in Sec. IV, we work with samples with initial weights $w_i = \pm 1$.

The expectation value of an observable $\mathcal{O}$ can be estimated through a weighted sum:

$$\langle \mathcal{O} \rangle \approx \hat{\mathcal{O}} \equiv \sum_{i=1}^N w_i \mathcal{O}(x_i), \quad (2)$$

where $\mathcal{O}(x)$ is the value of the observable at phase space point x . For example, $\mathcal{O}(x_i) = \delta(x_i \text{ in bin})$ for the contents of a histogram bin. In this discussion, hats always indicate estimates obtained through finite sampling.

The true value $\langle \mathcal{O} \rangle$ can only be obtained in the $N \rightarrow \infty$ limit, so the estimate in Eq. (2) is subject to uncertainties. The variance of Eq. (2) is N times the variance of each independent term in the sum:

$$\text{Var}[\hat{\mathcal{O}}] = N \text{Var}[W\mathcal{O}(X)], \quad (3)$$

where W is a random variable of which w_i is one realization. In general, though, the true value of $\text{Var}[W\mathcal{O}(X)]$ is not known *a priori*. Therefore, we have to estimate the variance through the sample variance:

$$\text{Var}[W\mathcal{O}(X)] \approx \frac{1}{N} \sum_{i=1}^N (w_i \mathcal{O}(x_i))^2 - \left(\frac{1}{N} \sum_{i=1}^N w_i \mathcal{O}(x_i) \right)^2. \quad (4)$$

In many cases of interest, the second term is subdominant, such as when $\mathcal{O}(x_i)$ is zero for most phase space points and one otherwise. Under that assumption, we obtain the standard Monte Carlo uncertainty estimate:

$$\text{Var}[\hat{\mathcal{O}}] \approx \sum_{i=1}^N (w_i \mathcal{O}(x_i))^2. \quad (5)$$

B. Monte Carlo reweighting

The idea behind reweighting is to select a subset of phase space points x_j with modified weights $\tilde{w}_j$ such that Eq. (2) [and later Eq. (5)] are preserved in the large N limit. Consider a small patch of phase space such that $\mathcal{O}$ is nearly constant. In such a region with N_{patch} events,

$$\hat{\mathcal{O}}_{\text{patch}} \approx \mathcal{O}(x_{\text{patch}}) \sum_{i=1}^{N_{\text{patch}}} w_i. \quad (6)$$

The key observation is that we can replace the N_{patch} original events with N_{patch}/K events of equal weight $\tilde{w}_j$,

$$\tilde{w}_j = \frac{K}{N_{\text{patch}}} \sum_{i=1}^{N_{\text{patch}}} w_i \approx K \langle W \rangle_{\text{patch}}. \quad (7)$$

This replacement preserves the estimate of any observable because

$$\sum_{j=1}^{N_{\text{patch}}/K} \tilde{w}_j = \sum_{i=1}^{N_{\text{patch}}} w_i, \quad (8)$$

where we are assuming N_{patch}/K is an integer. Since $\hat{\mathcal{O}}_{\text{patch}}$ is unchanged by this reweighting, the true value of $\text{Var}[\hat{\mathcal{O}}_{\text{patch}}]$ is also unchanged. Said another way, one can replace K events in a small patch with a single representative event and achieve approximately the same statistical properties.

As a result, one can reduce the number of events required for expensive detector simulations by picking $K > 1$. This reduction in statistics is possible for any K , though in general one has to be careful when accounting for the uncertainty, which we now discuss.

C. Preserving uncertainties

The naive uncertainty estimate in Eq. (5) is badly biased if one blindly uses the weights in Eq. (7). While it is possible to keep track of the local variance separately from the expected value, this is cumbersome and error prone. Instead, we advocate choosing K such that

$$\sum_{j=1}^{N_{\text{patch}}/K} \tilde{w}_j^2 = \sum_{i=1}^{N_{\text{patch}}} w_i^2. \quad (9)$$

Formally, Eqs. (8) and (9) require N/K to be an integer, but below we will take the continuum limit where this distinction is unnecessary. With this choice of K , one can estimate the variance in the usual way using Eq. (5) without any additional overhead. This is the key observation that underpins our neural resampling method.

In general, each phase space patch will require a different value of K . To determine this K_{patch} , it is convenient to rewrite Eq. (9) as

$$\tilde{w}_j^2 = \frac{K_{\text{patch}}}{N_{\text{patch}}} \sum_{i=1}^{N_{\text{patch}}} w_i^2 \approx K_{\text{patch}} \langle W^2 \rangle_{\text{patch}}. \quad (10)$$

Combining Eqs. (10) and (7) provides a prescription for choosing K_{patch} :

$$K_{\text{patch}} \approx \frac{\langle W^2 \rangle_{\text{patch}}}{\langle W \rangle_{\text{patch}}^2}. \quad (11)$$

Taking the continuum limit, Eq. (11) becomes

$$K(X) = \frac{\langle W^2 | X \rangle}{\langle W | X \rangle^2}, \quad (12)$$

with expectation values conditioned on the phase space points in X .

The above discussion can be encoded in the following practical algorithm to reweight and resample Monte Carlo events while preserving uncertainties:

- (1) Estimate $\widehat{W}(X) \approx \langle W | X \rangle$.
- (2) Estimate $\widehat{W^2}(X) \approx \langle W^2 | X \rangle$.
- (3) Define $\hat{K}(X) = \widehat{W^2}(X) / \widehat{W}(X)^2$.
- (4) For each event i , keep it with probability $1/\hat{K}(x_i)$; otherwise discard the event. Because $\hat{K}(x_i) \geq 1$ by construction, no event will be repeated.
- (5) For each kept event, set the new event weight to be $w_i \mapsto \tilde{W}(x_i) \equiv \widehat{W}(x_i) \hat{K}(x_i)$, which is the continuum limit of $\tilde{w}$ from Eq. (7).

The computational benefit of using $\tilde{W}(x_i)$ over w_i is true even if there are no negative weights. As with any Monte Carlo method, the accept-reject procedure in step 4 can only preserve Eqs. (8) and (9) in expectation value. As long as a given phase space point has a nontrivial spectrum of weights, the above reduction will decrease the computational cost of subsequent detector simulation with the same asymptotic statistical properties as captured by the first and second moments. The procedure above works for any estimation of $\langle W | X \rangle$ and $\langle W^2 | X \rangle$, including with histograms. The next section shows how to estimate these quantities without binning using neural networks.

III. NEURAL RESAMPLING

As described above, a Monte Carlo generator draws a sample $\{x_i\}$ from X . Each phase space point x_i has an associated weight w_i , which can be positive or negative. Moreover, the weights need not be a function of X , meaning that the same phase space point can have different weights, as determined by the Monte Carlo sampling scheme. The goal of the positive resampler method of Ref. [11] is to rebalance the weights such that each value of x has a unique weight. Our neural resampler accomplishes this same goal through binary classification with neural networks.

A. Learning event weights

To learn new event weights, we train a neural network to distinguish between two samples: the original sample $\{x_i\}$ with weights $\{w_i\}$ and a uniformly weighted sample with the same phase space points $\{x_i\}$ but weights set to 1. For concreteness, we use the binary cross-entropy loss for this

discussion, though other loss functions with the same asymptotic behavior would also work, such as the mean squared error.²

The loss function to be minimized is

$$\mathcal{L}[g] = - \sum_{i=1}^N w_i \log g(x_i) - \sum_{i=1}^N \log(1 - g(x_i)), \quad (13)$$

where $g(x)$ is parametrized as a neural network with output range $[0, 1]$. We emphasize that the two sums in this loss function run over the *same* phase space points x_i , just with different weights. This setup is identical to the second step of the OMNIFOLD unfolding algorithm [20], where a generated dataset is morphed into a weighted version of itself.

Taking a functional derivative of Eq. (13) with respect to $g(x)$ and setting it equal to zero, one can show that the loss function minimum provides an estimate of $\langle W | X \rangle$:

$$\frac{g(x)}{1 - g(x)} = \widehat{W}(x) \approx \langle W | X \rangle. \quad (14)$$

This is just a manifestation of the standard result that asymptotically (i.e., with infinite training data, maximally expressive neural network architecture, and ideal training procedure) the output of a binary classifier approaches a monotonic rescaling of the likelihood ratio; see e.g., Refs. [13–23]. In our case, the original sample has asymptotic probability distribution

$$p_{\text{original}}(x) = \langle W | x \rangle p_{\text{uniform}}(x), \quad (15)$$

where $p_{\text{uniform}}(x)$ is the phase space prior. The sample with uniform weights is not a proper probability distribution, since it is not normalized, but corresponds to N times $p_{\text{uniform}}(x)$. In this way, we learn local event weights that preserve the estimate of any observable via Eq. (2).

B. Learning uncertainties

For the case studies in Sec. IV, we focus on situations where the initial weights are all $\pm c$, for a fixed value of c . In such cases, $\langle W^2 | X \rangle = c^2$ by construction and no additional training is needed to perform neural reweighting.

²One key difference between binary cross entropy and mean squared error is that the former cannot learn negative weights. There are situations, particularly when using fixed-order Monte Carlo generators, where one encounters phase space regions with genuinely negative cross sections. We performed a preliminary test of this in the context of fixed-order top quark pair production with a parton shower subtraction scheme where, unlike the matched results in Sec. IV B, there are negative phase space regions. Using the mean squared error loss and linear activation in the final layer, we found good performance in the presence of both positive and negative cross section regions.

If there is a nontrivial spectrum of weight norms, we can repeat the logic of Sec. III A to estimate $\langle W^2|X \rangle$. This can be achieved using the loss function:

$$\mathcal{L}[h] = -\sum_{i=1}^N w_i^2 \log g(x_i) - \sum_{i=1}^N \log(1 - g(x_i)), \quad (16)$$

whose asymptotic minimum satisfies³

$$\frac{h(x)}{1 - h(x)} = \widehat{W}^2(x) \approx \langle W^2|X \rangle. \quad (17)$$

In this way, we learn local variances that preserve the estimate of observable uncertainties via Eq. (5).

C. Implementation details

To implement neural resampling, the (learned) $\widehat{W}(x)$ and $\widehat{W}^2(x)$ functions from Eqs. (14) and (17) are simply inserted into the algorithm described in Sec. II C, one could further unweight the samples to make all of the weights uniform; see further discussion in Ref. [11]. We omit this optional step in our case studies, since full (or partial) unweighting does not preserve the original Monte Carlo uncertainties.

For the following results, neural networks are constructed using rectified linear unit (ReLU) activation functions for hidden layers and the sigmoid function for the last layer to be the output in $[0, 1]$. All models are implemented in KERAS [77] with the TENSORFLOW backend [78] and trained using the cross entropy loss with the ADAM [79] optimizer. For the one-dimensional example, the neural network is fully connected with three hidden layers, 128 nodes per layer, and trained for ten epochs. The higher-dimensional example uses PFNs [24], based on the deep sets architecture [25], using the default parameters⁴ and trained for 100 epochs. None of the hyperparameters have been optimized for these studies.

IV. CASE STUDIES

To illustrate the potential of neural resampling, we present two case studies. First, we consider an example involving two Gaussians that highlights the potential of our method in a simple case where the optimal results can be obtained analytically. Then, we consider a realistic collider example of top quark pair production ($t\bar{t}$) at NLO matched to a parton shower, which allows us to demonstrate the robustness of our method to multidimensional (and variable-dimensional) phase space.

³As a generalization of this result, consider any random variable $X \in \mathbb{R}^d$ and write it as $X = (Y, Z)$ with $Y \in \mathbb{R}^k$ and $Z \in \mathbb{R}^\ell$ with $d = k + \ell$. Then, one can learn $\langle f(Z)|Y \rangle$ for a fixed function f by treating $Z|Y$ as a set of weights.

⁴<https://energyflow.network/>

For each of the following examples, we work with initial event weights of $w_i = \pm 1$, since this is the typical output of unweighted Monte Carlo generators. To map onto the discussion in Secs. II and III, one would need to normalize these weights to satisfy Eq. (1). Note that $\langle W^2|X \rangle = 1$ by construction, so we can skip the neural network training step in Eq. (16).

A. Two Gaussians

Our first case study involves two Gaussian distributions, one with positive weight and one with negative weight. Let $X_0 \sim \mathcal{N}(0, 1)$ and $X_1 \sim \mathcal{N}(0, 0.5)$, where $\mathcal{N}(\mu, \sigma)$ represents a Gaussian distribution with mean μ and standard deviation σ . The weight for X_0 events is $+1$ while the weight for X_1 events is -1 . We consider the case where the X_0 events happen 3 times more often than the X_1 events. The following results are based on 4M samples from X , i.e., 3M positive weight events from X_0 and 1M negative weight events from X_1 .

The weighted histogram of X is shown by the blue solid distribution in Fig. 1(a), corresponding to a Gaussian with a dip in the middle. After neural resampling with $K = 1$, shown by the orange dotted curve, we obtain the same true distribution up to statistical fluctuations. As desired, subsampling with the optimal value of K from Eq. (12) does not change the probability density, as shown by the green dashed curve. This subsampling does change the uncertainties, though, as discussed more below. We therefore conclude that neural resampling has successfully preserved the cross section via Eq. (2).

A nice feature of this example is that the correct unbinned weights are computable analytically using the asymptotic formulas in Sec. II. In Fig. 1(b), we show the original weight distribution together with the positive neural resampling weights using both $K = 1$ and the optimal value of K . While the original weights are both positive and negative, neural resampling yields strictly positive weights. We see that the finite sampling matches the expected analytic weight distributions, which is a nontrivial cross-check of our neural reweighting code. Note that even though a binning and finite sampling are chosen to represent the data in Fig. 1, all of these distributions are fundamentally unbinned.

In Fig. 1(c), we show the distribution of uncertainties using the standard Monte Carlo estimate based on summing the squared weights. With $K = 1$, the uncertainties are substantially underestimated relative to the original distribution, especially in the vicinity of $x = 0$, where there are large cancellations between negative and positive weights. Subsampling with the optimal K restores the original uncertainties, as desired. We therefore conclude that neural resampling has successfully preserved the uncertainties via Eq. (5).

As shown in Fig. 1(d), subsampling yields a significant savings in terms of the number of events required to obtain the same statistical properties as the original sample. Only

FIG. 1. Demonstration of neural resampling in a simple one-dimensional example of two Gaussians. (a) The (rescaled) cross section for observable x , estimated by summing the event weights in each histogram bin. (b) The distribution of the event weights, including solid lines corresponding to the analytic expectations. (c) The (rescaled) uncertainties for observable x , estimated by summing the squared event weights in each histogram bin and taking the square root. (d) The number of events as a function of x . The curves correspond to (solid blue) the original event sample with positive and negative weights, (orange dotted) neural resampling with $K = 1$ such that the cross section is preserved, and (green dashed) neural resampling with the optimal K value in Eq. (12) such that the cross sections and uncertainties are preserved. Note that the ratio between the orange and green curves in part (d) is $K(x)$, which is the local factor describing how many fewer events are needed after resampling.

one third of the events are needed to capture the same behavior of the original events, with the savings greatest near $x = 0$ where there are larger relative uncertainties.

This two Gaussian example highlights the efficacy of neural resampling in a simple one-dimensional example. We now turn to a case of relevance to collider physics where the phase space is multidimensional.

B. Top quark pair production

Our realistic collider case study is based on top quark pair production at NLO in quantum chromodynamics.

At fixed order in an expansion in the strong coupling constant α_s , it is well known that cross sections can become negative due to the breakdown of perturbation theory in the vicinity of soft/collinear singularities (see footnote 2). These unphysical phase space regions can be regulated by matching to a parton shower, rendering the differential cross section to be positive. In addition, matching improves the accuracy of the cross section prediction by including resummation effects beyond NLO.

The analysis below is based on 2M total events generated as follows. Fixed-order $t\bar{t}$ production is generated using MG5_aMC@NLO 5.2.7.2 [6] interfaced with the NNPDF 2.3 NLO

FIG. 2. Demonstration of neural resampling in a realistic collider example of top quark pair production at NLO matched to a parton shower. (a) The (rescaled) cross section for the leading ISR jet p_T . (b) The distribution of the event weights. The optimal K spectrum steeply falls out to about 10, but the distribution is truncated at 4.5 to aid the comparison with the other histograms. (c) The (rescaled) uncertainties for p_T . (d) The number of events as a function of p_T . As in Fig. 1, the curves correspond to (solid blue) the original event sample with positive and negative weights, (orange dotted) neural resampling with $K=1$, and (green dashed) neural resampling with the optimal K value in Eq. (12).

parton density function set [80]. This fixed-order calculation uses MADLOOP 2.7.2 [6,81], NINJA 1.2.0 [82,83], CUTTOOLS 1.9.3 [84], and ONELOOP 3.6 [85,86]. In these samples, the top quarks are forced to decay leptonically via $t\bar{t} \rightarrow b\bar{b}\mu^+\mu^-\nu\bar{\nu}$. Using the default FKS subtraction scheme [87,88], the resulting events are matched with the PYTHIA 8.230 [1,89] parton shower, keeping the default shower settings. Initial state radiation (ISR) is generated both from the NLO calculation and from the subsequent parton shower matching.

The resulting events in the HepMC 2.06.09 [90] format are processed with the HepMC reader module of DELPHES 3.4.2 [91]. This setup is also used to cluster $R=0.4$ anti- k_t jets [92] with FastJet [93]. Jets originating from a b -quark are

identified as such using the flavor tagging module in DELPHES. Jets are labeled as “ISR” if they are not b jets. We suppress the overall cross section information and take the weights to be ± 1 . After neural resampling, it is straightforward to rescale all of the weights to have the proper dimensions of a cross section, but we elide this step for simplicity.

To implement neural resampling, we use PFNs [24], where each event is represented as a variable-length set of four vectors that correspond to the muons, neutrinos, and clustered jets. The PDGID [94] of the muons and neutrinos are used as a per-particle feature in the PFN, while the b jets (ISR jets) are given the per-particle feature of 1 (0). In principle, we could have applied neural resampling

FIG. 3. The same as Fig. 2, but now plotting the number of ISR jets with $p_T > 10$ GeV. Note that the weight distribution in (b) is identical to Fig. 2(b), since neural resampling acts on the full unbinned phase space.

directly to the final-state hadrons, but since we are only going to plot jet-related quantities, it makes sense to perform jet clustering before neural resampling.

Note that the phase space that is being reweighted here is variable dimensional, with at least 20 dimensions from the 3-momenta for six particles and two jet masses. These data are constrained to an 18-dimensional manifold after considering transverse momentum conservation and are approximately constrained to a 14-dimensional manifold after accounting for mass-shell conditions. Many events have far more dimensions due to additional jets. Thus, this is a highly nontrivial test of whether neural resampling can yield sensible results across a multidimensional and variable-dimensional phase space.

In Fig. 2(a), we plot the distribution of transverse momentum (p_T) of the leading ISR jet. As expected,

neural resampling has matched the bulk of this distribution up to statistical uncertainties. In the tails, where training data are more sparse, there are larger variations, but we emphasize that these results were obtained “out of the box” with no attempt to optimize training parameters.

The distribution of event weights is shown in Fig. 2(b). The original sample had both positive and negative weights, but neural resampling has yielded a positive weight distribution as desired. With the optimal choice of K , the event weights extend out to about 10, which reflects the initial inefficiency of event generation with positive and negative weights.

The uncertainties are shown in Fig. 2(c), again using the standard Monte Carlo estimate from Eq. (5). With $K = 1$, the uncertainties are underestimated throughout the phase space, but with the optimal value of K , they are captured

correctly. Correspondingly, the optimal value of K yields a substantially lower number of events, as shown in Fig. 2(d), indicating a large gain in downstream computational efficiency. In particular, from the original 2M events, only about 600 k remain after neural resampling, with no change to the statistical properties.

To highlight that neural resampling yields sensible results across the whole (unbinned) phase space, we plot the distribution of the number of ISR jets with $p_T > 10$ GeV in Fig. 3. We emphasize that no retraining is needed here, since these plots are based on the same weights already shown in Fig. 2(b). Obtaining results like this that work for any observable of interest would be challenging with a binned approach. This case study therefore suggests that neural reweighting will be a powerful tool for efficient use of Monte Carlo generators at colliders.

V. CONCLUSIONS AND OUTLOOK

This paper has introduced neural resampling, a neural-network-based extension of the binned positive resampler proposed in Ref. [11]. By exploiting the ability of neural networks to approximate likelihood ratios, our new approach is able to eliminate negative weights without binning and with access to potentially high-dimensional (and variable-sized) phase spaces. Furthermore, neural resampling preserves statistical uncertainties with minimal overhead and no adjustable parameters. This uncertainty preservation scheme is general and can also be applied to binned (non-neural network) resampling.

Given the growing availability of higher-order corrections and the computational demands of detector simulations, there is a need to make the best use of limited resources. The neural resampler is able to preserve the statistical power of these precision calculations while decreasing downstream computational demands. For the future, it would be interesting to study whether neural resampling could be incorporated more directly into the

Monte Carlo generation process. This could lead to further computational gains, particularly if used in concert with emerging neural-network-based phase space integration techniques.

Beyond just computational efficiency, there may be other advantages of neural resampling. Many analysis strategies that consider the full unbinned log likelihood over events require per event weights to be positive [95], which is guaranteed by neural resampling (assuming non-negative cross section). Monte Carlo generators increasingly have the ability to keep track of uncertainties via weight variations [96,97], which may be straightforward to incorporate into neural resampling via parametrized networks [18,98]. Finally, neural resampling ensures that the event weights are a true function of phase space and not multi-valued maps, even allowing the weight function to be differentiated. This feature of the learned weight function may turn out to be beneficial for other aspects of data processing and analysis.

VI. CODE AND DATA

The code for this paper is listed in Ref. [99]. The top quark datasets are available upon request.

ACKNOWLEDGMENTS

We would like to thank A. Andreassen, P. Komiske, and E. Metodiev for many helpful discussions about reweighting with neural networks. We thank J. Andersen, C. Gutschow, A. Maier, and S. Prestel for helpful conversations about Ref. [11]. B.N. was supported by the U.S. Department of Energy (DOE) Office of Science under Contract No. DE-AC02-05CH11231. J.T. was supported by the U.S. DOE Office of High Energy Physics under Contract No. DE-SC0012567. B.N. would also like to thank Nvidia for providing Volta graphics processing units for neural network training.

-
- [1] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, An introduction to PYTHIA 8.2, *Comput. Phys. Commun.* **191**, 159 (2015).
 - [2] J. Bellm *et al.*, HERWIG 7.0/HERWIG++ 3.0 release note, *Eur. Phys. J. C* **76**, 196 (2016).
 - [3] E. Bothmann *et al.* (SHERPA Collaboration), Event generation with SHERPA 2.2, *SciPost Phys.* **7**, 034 (2019).
 - [4] P. Nason and G. Ridolfi, A positive-weight next-to-leading-order Monte Carlo for Z pair hadroproduction, *J. High Energy Phys.* **08** (2006) 077.
 - [5] S. Frixione and B.R. Webber, Matching NLO QCD computations and parton shower simulations, *J. High Energy Phys.* **06** (2002) 029.
 - [6] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *J. High Energy Phys.* **07** (2014) 079.
 - [7] S. Alioli, P. Nason, C. Oleari, and E. Re, A general framework for implementing NLO calculations in shower Monte Carlo programs: The POWHEG BOX, *J. High Energy Phys.* **06** (2010) 043.

- [8] S. Agostinelli *et al.* (GEANT4 Collaboration), GEANT4: A simulation toolkit, *Nucl. Instrum. Methods Phys. Res., Sect. A* **506**, 250 (2003).
- [9] J. Albrecht *et al.* (HEP Software Foundation), A roadmap for HEP software and computing R&D for the 2020s, *Comput. Softw. Big Sci.* **3**, 7 (2019).
- [10] S. Amoroso *et al.* (HSF Physics Event Generator WG), Challenges in Monte Carlo event generator software for high-luminosity LHC, [arXiv:2004.13687](#).
- [11] J. R. Andersen, C. Gutschow, A. Maier, and S. Prestel, A positive resampler for Monte Carlo events with negative weights, [arXiv:2005.09375](#).
- [12] R. Frederix, S. Frixione, S. Prestel, and P. Torrielli, On the reduction of negative weights in MC@NLO-type matching procedures, *J. High Energy Phys.* **07** (2020) 238.
- [13] K. Cranmer, J. Pavez, and G. Louppe, Approximating likelihood ratios with calibrated discriminative classifiers, [arXiv:1506.02169](#).
- [14] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, A guide to constraining effective field theories with machine learning, *Phys. Rev. D* **98**, 052004 (2018).
- [15] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, Constraining Effective Field Theories with Machine Learning, *Phys. Rev. Lett.* **121**, 111801 (2018).
- [16] J. Brehmer, G. Louppe, J. Pavez, and K. Cranmer, Mining gold from implicit models to improve likelihood-free inference, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 5242 (2020).
- [17] M. Stoye, J. Brehmer, G. Louppe, J. Pavez, and K. Cranmer, Likelihood-free inference with an improved cross-entropy estimator, [arXiv:1808.00973](#).
- [18] A. Andreassen and B. Nachman, Neural networks for full phase-space reweighting and parameter tuning, *Phys. Rev. D* **101**, 091901 (2020).
- [19] J. Brehmer, F. Kling, I. Espejo, and K. Cranmer, MadMiner: Machine learning-based inference for particle physics, *Comput. Softw. Big Sci.* **4**, 3 (2020).
- [20] 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**, 182001 (2020).
- [21] A. Andreassen, B. Nachman, and D. Shih, Simulation assisted likelihood-free anomaly detection, *Phys. Rev. D* **101**, 095004 (2020).
- [22] J. Hollingsworth and D. Whiteson, Resonance searches with machine learned likelihood ratios, [arXiv:2002.04699](#).
- [23] C. Badiali, F. Di Bello, G. Frattari, E. Gross, V. Ippolito, M. Kado, and J. Shlomi, Efficiency parameterization with neural networks, [arXiv:2004.02665](#).
- [24] P. T. Komiske, E. M. Metodiev, and J. Thaler, Energy flow networks: Deep sets for particle jets, *J. High Energy Phys.* **01** (2019) 121.
- [25] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Poczos, R. R. Salakhutdinov, and A. J. Smola, Deep sets, in *Advances in Neural Information Processing Systems 30*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2017), pp. 3391–3401, <http://papers.nips.cc/paper/6931-deep-sets.pdf>.
- [26] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative adversarial networks, [arXiv:1406.2661](#).
- [27] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, [arXiv:1312.6114](#).
- [28] D. J. Rezende, S. Mohamed, and D. Wierstra, Stochastic backpropagation and approximate inference in deep generative models, [arXiv:1401.4082](#).
- [29] L. de Oliveira, M. Paganini, and B. Nachman, Learning particle physics by example: Location-aware generative adversarial networks for physics synthesis, *Comput. Softw. Big Sci.* **1**, 4 (2017).
- [30] M. Paganini, L. de Oliveira, and B. Nachman, Accelerating Science with Generative Adversarial Networks: An Application to 3D Particle Showers in Multilayer Calorimeters, *Phys. Rev. Lett.* **120**, 042003 (2018).
- [31] M. Paganini, L. de Oliveira, and B. Nachman, CaloGAN: Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks, *Phys. Rev. D* **97**, 014021 (2018).
- [32] S. Alonso-Monsalve and L. H. Whitehead, Image-based model parameter optimization using model-assisted generative adversarial networks, <https://doi.org/10.1109/TNNLS.2020.2969327> (2018).
- [33] A. Butter, T. Plehn, and R. Winterhalder, How to GAN event subtraction, [arXiv:1912.08824](#).
- [34] J. A. Martinez, T. Q. Nguyen, M. Pierini, M. Spiropulu, and J.-R. Vlimant, Particle generative adversarial networks for full-event simulation at the LHC and their application to pileup description, *J. Phys. Conf. Ser.* **1525**, 012081 (2020).
- [35] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, and R. Winterhalder, How to GAN away detector effects, *SciPost Phys.* **8**, 070 (2020).
- [36] S. Vallecorsa, F. Carminati, and G. Khattak, 3D convolutional GAN for fast simulation, *EPJ Web Conf.* **214**, 02010 (2019).
- [37] C. Ahdida *et al.* (SHiP Collaboration), Fast simulation of muons produced at the SHiP experiment using generative adversarial networks, *J. Instrum.* **14**, P11028 (2019).
- [38] S. Carrazza and F. A. Dreyer, Lund jet images from generative and cycle-consistent adversarial networks, *Eur. Phys. J. C* **79**, 979 (2019).
- [39] A. Butter, T. Plehn, and R. Winterhalder, How to GAN LHC events, *SciPost Phys.* **7**, 075 (2019).
- [40] J. Lin, W. Bhimji, and B. Nachman, Machine learning templates for QCD factorization in the search for physics beyond the standard model, *J. High Energy Phys.* **05** (2019) 181.
- [41] R. Di Sipio, M. F. Giannelli, S. K. Haghighat, and S. Palazzo, DijetGAN: A generative-adversarial network approach for the simulation of QCD dijet events at the LHC, *J. High Energy Phys.* **08** (2019) 110.
- [42] B. Hashemi, N. Amin, K. Datta, D. Olivito, and M. Pierini, LHC analysis-specific datasets with generative adversarial networks, [arXiv:1901.05282](#).
- [43] V. Chekalina, E. Orlova, F. Ratnikov, D. Ulyanov, A. Ustyuzhanin, and E. Zakharov, Generative models for fast calorimeter simulation. LHCb case, *EPJ Web Conf.* **214**, 02034 (2019).

- [44] ATLAS Collaboration, Deep generative models for fast shower simulation in ATLAS, ATL-SOFT-PUB-2018-001, 2018, <http://cds.cern.ch/record/2630433>.
- [45] K. Zhou, G. Endrodi, L.-G. Pang, and H. Stocker, Regressive and generative neural networks for scalar field theory, *Phys. Rev. D* **100**, 011501 (2019).
- [46] F. Carminati, A. Gheata, G. Khattak, P.M. Lorenzo, S. Sharan, and S. Vallecorsa, Three dimensional Generative Adversarial Networks for fast simulation, *J. Phys. Conf. Ser.* **1085**, 032016 (2018).
- [47] S. Vallecorsa, Generative models for fast simulation, *J. Phys. Conf. Ser.* **1085**, 022005 (2018).
- [48] K. Datta, D. Kar, and D. Roy, Unfolding with generative adversarial networks, [arXiv:1806.00433](https://arxiv.org/abs/1806.00433).
- [49] P. Musella and F. Pandolfi, Fast and accurate simulation of particle detectors using generative adversarial networks, *Comput. Softw. Big Sci.* **2**, 8 (2018).
- [50] M. Erdmann, L. Geiger, J. Glombitza, and D. Schmidt, Generating and refining particle detector simulations using the Wasserstein distance in adversarial networks, *Comput. Softw. Big Sci.* **2**, 4 (2018).
- [51] K. Deja, T. Trzcinski, and Ł. Graczykowski, Generative models for fast cluster simulations in the TPC for the ALICE experiment, *EPJ Web Conf.* **214**, 06003 (2019).
- [52] D. Derkach, N. Kazeev, F. Ratnikov, A. Ustyuzhanin, and A. Volokhova, Cherenkov detectors fast simulation using neural networks, *Nucl. Instrum. Methods Phys. Res., Sect. A* (2020) 161804, **952**.
- [53] H. Erbin and S. Krippendorf, GANs for generating EFT models, [arXiv:1809.02612](https://arxiv.org/abs/1809.02612).
- [54] M. Erdmann, J. Glombitza, and T. Quast, Precise simulation of electromagnetic calorimeter showers using a Wasserstein generative adversarial network, *Comput. Softw. Big Sci.* **3**, 4 (2019).
- [55] L. de Oliveira, M. Paganini, and B. Nachman, Tips and tricks for training GANs with physics constraints, in *Proceedings of the Deep Learning for Physical Sciences Workshop at NIPS* (2017), https://dl4physicsciences.github.io/files/nips_dlps_2017_26.pdf.
- [56] L. de Oliveira, M. Paganini, and B. Nachman, Controlling physical attributes in GAN-accelerated simulation of electromagnetic calorimeters, *J. Phys. Conf. Ser.* **1085**, 042017 (2018).
- [57] S. Farrell, W. Bhimji, T. Kurth, M. Mustafa, D. Bard, Z. Lukic, B. Nachman, and H. Patton, Next generation generative neural networks for HEP, *EPJ Web Conf.* **214**, 09005 (2019).
- [58] B. Hooberman, A. Farbin, G. Khattak, V. Pacela, M. Pierini, J.-R. Vlimant, M. Spiropulu, W. Wei, M. Zhang, and S. Vallecorsa, Calorimetry with deep learning: Particle classification, energy regression, and simulation for high-energy physics, in *Proceedings of the Deep Learning for Physical Sciences Workshop at NIPS* (2017), https://dl4physicsciences.github.io/files/nips_dlps_2017_15.pdf.
- [59] D. Belayneh *et al.*, Calorimetry with deep learning: Particle simulation and reconstruction for collider physics, *Eur. Phys. J. C* **80**, 688 (2020).
- [60] E. Buhmann, S. Diefenbacher, E. Eren, F. Gaede, G. Kasieczka, A. Korol, and K. Kruger, Getting high: High fidelity simulation of high granularity calorimeters with high speed, [arXiv:2005.05334](https://arxiv.org/abs/2005.05334).
- [61] M. S. Albergo, G. Kanwar, and P. E. Shanahan, Flow-based generative models for Markov chain Monte Carlo in lattice field theory, *Phys. Rev. D* **100**, 034515 (2019).
- [62] G. Kanwar, M. S. Albergo, D. Boyda, K. Cranmer, D. C. Hackett, S. Racaniere, D. J. Rezende, and P. E. Shanahan, Equivariant Flow-Based Sampling for Lattice Gauge Theory, *Phys. Rev. Lett.* **125**, 121601 (2020).
- [63] J. M. Urban and J. M. Pawłowski, Reducing autocorrelation times in lattice simulations with generative adversarial networks, [arXiv:1811.03533](https://arxiv.org/abs/1811.03533).
- [64] K. A. Nicoli, C. J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel, S. Nakajima, and P. Stornati, On estimation of thermodynamic observables in lattice field theories with deep generative models, [arXiv:2007.07115](https://arxiv.org/abs/2007.07115).
- [65] R. D. Ball *et al.* (NNPDF Collaboration), Parton distributions from high-precision collider data, *Eur. Phys. J. C* **77**, 663 (2017).
- [66] S. Carrazza and J. Cruz-Martinez, Towards a new generation of parton densities with deep learning models, *Eur. Phys. J. C* **79**, 676 (2019).
- [67] J. Bendavid, Efficient Monte Carlo integration using boosted decision trees and generative deep neural networks, [arXiv:1707.00028](https://arxiv.org/abs/1707.00028).
- [68] M. D. Klimek and M. Perelstein, Neural network-based approach to phase space integration, [arXiv:1810.11509](https://arxiv.org/abs/1810.11509).
- [69] E. Bothmann, T. Janen, M. Knobbe, T. Schmale, and S. Schumann, Exploring phase space with neural importance sampling, *SciPost Phys.* **8**, 069 (2020).
- [70] C. Gao, J. Isaacson, and C. Krause, i-flow: High-dimensional integration and sampling with normalizing flows, [arXiv:2001.05486](https://arxiv.org/abs/2001.05486).
- [71] C. Gao, S. Hche, J. Isaacson, C. Krause, and H. Schulz, Event generation with normalizing flows, *Phys. Rev. D* **101**, 076002 (2020).
- [72] F. Bishara and M. Montull, (Machine) learning amplitudes for faster event generation, [arXiv:1912.11055](https://arxiv.org/abs/1912.11055).
- [73] S. Badger and J. Bullock, Using neural networks for efficient evaluation of high multiplicity scattering amplitudes, *J. High Energy Phys.* **06** (2020) 114.
- [74] J. W. Monk, Deep learning as a parton shower, *J. High Energy Phys.* **12** (2018) 021.
- [75] A. Andreassen, I. Feige, C. Frye, and M. D. Schwartz, Binary JUNIPR: An Interpretable Probabilistic Model for Discrimination, *Phys. Rev. Lett.* **123**, 182001 (2019).
- [76] A. Andreassen, I. Feige, C. Frye, and M. D. Schwartz, JUNIPR: A framework for unsupervised machine learning in particle physics, *Eur. Phys. J. C* **79**, 102 (2019).
- [77] F. Chollet, Keras, <https://github.com/fchollet/keras> (2017).
- [78] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, Tensorflow: A system for large-scale machine learning, in *OSDI (USENIX Association, Berkeley, CA, 2016)*, Vol. 16, pp. 265–283, <https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf>.
- [79] D. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [80] R. D. Ball *et al.*, Parton distributions with LHC data, *Nucl. Phys. B* **867**, 244 (2013).

- [81] V. Hirschi, R. Frederix, S. Frixione, M. V. Garzelli, F. Maltoni, and R. Pittau, Automation of one-loop QCD corrections, *J. High Energy Phys.* **05** (2011) 044.
- [82] P. Mastrolia, E. Mirabella, and T. Peraro, Integrand reduction of one-loop scattering amplitudes through Laurent series expansion, *J. High Energy Phys.* **06** (2012) 095; Erratum, *J. High Energy Phys.* **11** (2012) 128.
- [83] T. Peraro, Ninja: Automated integrand reduction via Laurent expansion for one-loop amplitudes, *Comput. Phys. Commun.* **185**, 2771 (2014).
- [84] G. Ossola, C. G. Papadopoulos, and R. Pittau, CutTools: A program implementing the OPP reduction method to compute one-loop amplitudes, *J. High Energy Phys.* **03** (2008) 042.
- [85] A. van Hameren, OneLoop: For the evaluation of one-loop scalar functions, *Comput. Phys. Commun.* **182**, 2427 (2011).
- [86] A. van Hameren, C. Papadopoulos, and R. Pittau, Automated one-loop calculations: A proof of concept, *J. High Energy Phys.* **09** (2009) 106.
- [87] S. Frixione, Z. Kunszt, and A. Signer, Three jet cross-sections to next-to-leading order, *Nucl. Phys.* **B467**, 399 (1996).
- [88] S. Frixione, A general approach to jet cross-sections in QCD, *Nucl. Phys.* **B507**, 295 (1997).
- [89] T. Sjöstrand, S. Mrenna, and P.Z. Skands, PYTHIA 6.4 physics and manual, *J. High Energy Phys.* **05** (2006) 026.
- [90] M. Dobbs and J. B. Hansen, The HEPMC C++ Monte Carlo event record for high energy physics, CERN, Geneva, Technical Report No. ATL-SOFT-2000-001, 2000, <https://cds.cern.ch/record/684090>.
- [91] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lematre, A. Mertens, and M. Selvaggi (DELPHES 3 Collaboration), DELPHES 3, A modular framework for fast simulation of a generic collider experiment, *J. High Energy Phys.* **02** (2014) 057.
- [92] M. Cacciari, G. P. Salam, and G. Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [93] M. Cacciari and G. P. Salam, Dispelling the N^3 myth for the k_t jet-finder, *Phys. Lett. B* **641**, 57 (2006).
- [94] M. Tanabashi *et al.* (Particle Data Group), Review of particle physics, *Phys. Rev. D* **98**, 030001 (2018).
- [95] B. Nachman, A guide for deploying deep learning in LHC searches: How to achieve optimality and account for uncertainty, *SciPost Phys.* **8**, 090 (2020).
- [96] S. Mrenna and P. Skands, Automated parton-shower variations in PYTHIA 8, *Phys. Rev. D* **94**, 074005 (2016).
- [97] E. Bothmann, M. Schönherr, and S. Schumann, Reweighting QCD matrix-element and parton-shower calculations, *Eur. Phys. J. C* **76**, 590 (2016).
- [98] P. Baldi, K. Cranmer, T. Faucett, P. Sadowski, and D. Whiteson, Parametrized neural networks for high-energy physics, *Eur. Phys. J. C* **76**, 235 (2016).
- [99] <https://github.com/bnachman/NeuralPositiveResampler>.