

Event generation with normalizing flowsChristina Gao¹, Stefan Höche¹, Joshua Isaacson¹, Claudius Krause^{1,*} and Holger Schulz²¹*Fermi National Accelerator Laboratory, Batavia, Illinois, 60510, USA*²*Department of Physics, University of Cincinnati, Cincinnati, Ohio 45219, USA*

(Received 18 February 2020; accepted 9 March 2020; published 6 April 2020)

We present a novel integrator based on normalizing flows which can be used to improve the unweighting efficiency of Monte Carlo event generators for collider physics simulations. In contrast to machine learning approaches based on surrogate models, our method generates the correct result even if the underlying neural networks are not optimally trained. We exemplify the new strategy using the example of Drell-Yan type processes at the LHC, both at leading and partially at next-to-leading order QCD.

DOI: [10.1103/PhysRevD.101.076002](https://doi.org/10.1103/PhysRevD.101.076002)**I. INTRODUCTION**

Numerical simulation programs are a cornerstone of collider physics. They are used for the planning of future experiments, analysis of current measurements and, finally, reinterpretation based on an improved theoretical understanding of nature. They employ Monte Carlo methods to link theory and experiment by generating virtual collider events, which can then be analyzed like actual events observed in detectors [1,2].

With more and more data available from the Large Hadron Collider (LHC) and the high-luminosity upgrade, the task of simulating collisions at high precision becomes a matter of concern for the high-energy physics community. The projected amount of computational resources falls far short of the needs for precision event generation [3]. Past studies of the scaling behavior of multijet simulations have shown that the computing needs are largely determined by the gradually decreasing unweighting efficiency [4,5]. Except for dedicated integrators, which require a detailed understanding of the physics problem at hand, adaptive Monte Carlo methods seem to be the only choice to address this problem [6–13].

With the rise of machine learning, this topic has seen a resurgence of interest recently. The possibility of using these techniques for integration in high-energy physics was first discussed in [14]. Boosted decision trees and generative adversarial networks (GANs) were investigated as possible general purpose integrators. This new technique

improved the integration of nonseparable high dimensional functions, for which traditional algorithms failed. The first true physics application was presented in [15]. The authors used dense neural networks (DNN) in order to perform a variable transformation and demonstrate that they obtain significantly larger efficiencies for three body decay integrals than standard approaches [16]. The major drawback of this method is its computational cost. Since the network acts as a variable transformation, its gradient must be computed for each inference point in order to determine the Jacobian. This becomes computationally heavy for high multiplicity processes.

A completely orthogonal approach utilizes machine learning techniques directly for amplitude evaluation [17] or event generation [18–24]. Training data for these approaches are obtained from traditional event generation techniques, and hence the problem of efficient event generation still remains. In addition, one needs to ensure that the neural networks are trained well in order to approximate the original integrand. If this is not the case, the resulting generator will not only be inefficient but may actually yield the wrong result [25].

In this publication we propose a novel idea to address the problem: We replace standard adaptive algorithms like VEGAS [6,7], by the extension [26,27] of a nonlinear independent components estimation (NICE) technique [28,29], also known as a normalizing flow. This algorithm is combined with a recursive multichannel [30,31] to form a generic integrator for collider event generation. We test its performance in Drell-Yan type processes at the LHC, computed both at leading and partially at next-to-leading order QCD. We focus our study on the event generation efficiency in comparison to the general-purpose matrix element generator COMIX [31]. While the training of neural networks during the adaptation stage of the normalizing flow integrator is a very time consuming operation, event generation is inexpensive, because no gradients need to be

*Corresponding author.
ckrause@fnal.gov

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. 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.

computed. This could make the technique a prime candidate for LHC event generation in the near future. Normalizing flows have also been combined with Markov chain Monte Carlo methods, showing promising results [32–34].

This manuscript is organized as follows. Section II briefly reviews the technique of Monte Carlo integration and introduces the concept of normalizing flows. Section III presents our new integrator. Section IV discusses its computing performance and presents some first applications to LHC event generation. Section V contains an outlook.

II. NORMALIZING FLOWS

Monte Carlo or quasi-Monte Carlo methods are known as the only viable option to tackle high-dimensional integration problems. The basic technique relies on approximating the integrand by randomly sampling points in the integration domain Ω and weighting each point, x , by the value of the integrand, $f(x)$. The value of the integral is then obtained as the statistical average of all points, and the uncertainty is determined by its variance:

$$I = \int_{\Omega} f(x) dx = \frac{\Omega}{N} \sum_{i=1}^N f(x_i) = \Omega \langle f \rangle_x, \quad (1)$$

$$\sigma_I = \Omega \sqrt{\frac{\langle f^2 \rangle_x - \langle f \rangle_x^2}{N-1}}.$$

In this context, $\langle \cdot \rangle_x$ indicates that the average is taken with respect to a uniform distribution in x . The variance of the integral can be reduced by importance sampling or stratified sampling [35]. In particular, using the transformation $dx = dG(x)/g(x)$, with $G(x)$ the primitive of $g(x)$, one obtains

$$I = \int_{\Omega} \frac{f(x)}{g(x)} dG(x) = \Omega \langle f/g \rangle_G, \quad (2)$$

$$\sigma_I = \Omega \sqrt{\frac{\langle (f/g)^2 \rangle_G - \langle f/g \rangle_G^2}{N-1}}.$$

The function $g(x)$ can now be chosen appropriately, such as to minimize the variance. In the limit $g(x) \rightarrow f(x)/I$, Eq. (2) would be estimated with vanishing uncertainty. The goal is thus to find a distribution $g(x)$ that resembles the shape of $f(x)$ most closely, while being integrable and invertible in order to allow for faster sampling.

For multidimensional integrals, where the variable transformation reads $d\vec{x} \rightarrow d\vec{x}' |d\vec{x}(\vec{x}')/d\vec{x}'|$, we can simply replace the Jacobian $g(x) \rightarrow |d\vec{x}'/d\vec{x}|$, and Eq. (2) remains valid. This forms the basis for the concept of a normalizing flow: For a bijective map, $G(\vec{x})$, of the random variable \vec{x} that is drawn from a flat probability distribution, the variable $\vec{x}' = G(\vec{x})$ follows the probability distribution

$$d\vec{x}' = dG^{-1}(\vec{x}') \left| \frac{\partial G^{-1}(\vec{x}')}{\partial \vec{x}'} \right|^{-1} = d\vec{x} |g(\vec{x})|. \quad (3)$$

Applying a series of transformations, G_k , where $k = 1, \dots, K$, one defines the normalizing flow as a bijective mapping between statistical distributions of the random variables \vec{x} and \vec{x}_K ,

$$d\vec{x}_K = dG_K(G_{K-1}(\dots G_2(G_1(\vec{x})))) \quad (4)$$

$$= d\vec{x} \prod_{k=1}^K |g_k(\vec{x}_{k-1})|, \quad \text{where } \vec{x}_0 = \vec{x}.$$

In practice, G_k is often limited to simple functions, in order to make the determinant of the Jacobian easy to compute. This constrains the level of complexity that can be modeled by the normalizing flow. The complexity can be increased using so-called coupling layers, which were first introduced in [28,29]. An alternative technique is based on autoregressive models [36,37]. In this study we focus on [26,27] which generalizes the design of the coupling layers proposed by [29].

A. Coupling layers

A coupling layer is a special design of a bijector, first proposed in [28,29]. Figure 1 shows its basic structure. For each bijective mapping, the input variable $\vec{x} = \{x_1, \dots, x_D\}$ is partitioned into two subsets, $\vec{x}_A = \{x_1, \dots, x_d\}$ and $\vec{x}_B = \{x_{d+1}, \dots, x_D\}$. Under the bijective map, g , the resulting variable transforms as

$$x'_A = x_A, \quad A \in [1, d], \quad (5)$$

$$x'_B = C(x_B; m(\vec{x}_A)), \quad B \in [d+1, D].$$

In this context, m represents a neural network that takes x_A as an input and then outputs the parameters of the invertible “coupling transform,” C , that will be applied to x_B . The inverse map G^{-1} is given by

$$x_A = x'_A, \quad (6)$$

$$x_B = C^{-1}(x'_B; m(\vec{x}_A)) = C^{-1}(x'_B; m(\vec{x}'_A)),$$

which leads to the simple Jacobian

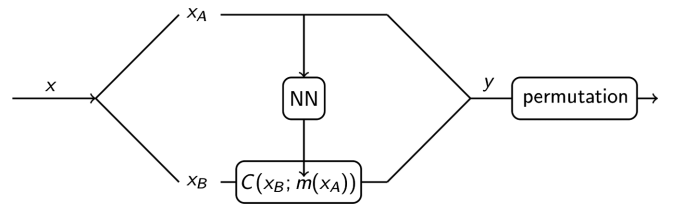


FIG. 1. Structure of a coupling layer. m is the output of a neural network and defines the coupling transform, C , that will be applied to x_B .

$$\left| \frac{\partial G(\vec{x})}{\partial \vec{x}} \right|^{-1} = \left| \begin{pmatrix} \vec{1} & 0 \\ \frac{\partial C}{\partial m} \frac{\partial m}{\partial \vec{x}_A} & \frac{\partial C}{\partial \vec{x}_B} \end{pmatrix} \right|^{-1} = \left| \frac{\partial C(\vec{x}_B; m(\vec{x}_A))}{\partial \vec{x}_B} \right|^{-1}. \quad (7)$$

Note that Eq. (7) does not require the computation of the gradient of $m(\vec{x}_A)$, which would scale as $\mathcal{O}(D^3)$ with D the number of dimensions. In addition, since C is diagonal, the computation of the determinant of $\partial C / \partial \vec{x}_B$ scales linearly with the number of dimensions, and is therefore tractable even for high dimensional problems. The normalizing flow method is thus evidently superior to existing integration techniques based on neural networks. To construct a complete normalizing flow, one simply compounds a series of coupling layers with the freedom of choosing any of the input dimensions to be transformed in each layer. We show in Ref. [38] that at most $2\lceil \log_2 D \rceil$ coupling layers are required in order to express arbitrarily complicated, nonseparable structures of the integrand.

In order to implement a normalizing flow integrator in practice, the user must provide a neural network, represented by $m(\vec{x}_A)$, a function f to integrate, and the definition of a loss function. This is discussed in more detail in Ref. [38].

B. Piecewise Polynomial and rational quadratic spline coupling transforms

So far we have not yet specified the invertible coupling transforms C . The design of this function impacts the flexibility of coupling layer based normalizing flow algorithms and is an active field of research. A very powerful definition of C was introduced in Ref. [26]. Both the domain and codomain of each coupling layer are defined to be the unit hypercube. If the random variable \vec{x}' is

uniformly distributed, it follows from Eq. (7) that the initial variable \vec{x} follows the distribution $|\partial C / \partial \vec{x}_B|$. Thus the coupling transform can be interpreted as the cumulative distribution function (CDF) of \vec{x} . Each dimension is then divided into K bins and models this CDF with a monotonically increasing polynomial function per bin. In particular, Ref. [26] experiments with piecewise linear and piecewise quadratic coupling transforms. In the implementation of a piecewise quadratic coupling transform, the bin width is allowed to vary in order to increase the flexibility of the coupling layer.

It may seem natural to generalize the piecewise polynomial coupling transform to include even higher order terms in order to increase the expressivity of the coupling layer. This has been proposed in Ref. [27], which generalized the piecewise quadratic coupling transform to allow a monotonically increasing rational-quadratic function in each bin of the coupling transform. To implement this, the bin heights, bin widths and also the derivatives in between each bin are allowed to vary and are predicted by the neural network.

III. PHASE-SPACE INTEGRATION

In this section we briefly summarize the diagram-based [39] recursive [31] multichannel [30] integration used in our numerical routines. The latter are designed to cope with especially large numbers of outgoing particles and exhibit exponential scaling, reduced from factorial scaling by means of dynamic programming.

Consider a $2 \rightarrow n$ scattering process, where we denote the incoming particles by a and b and the outgoing particles by $1 \dots n$. The corresponding n -particle differential phase space element reads

$$d\Phi_n(a, b; 1, \dots, n) = \left[\prod_{i=1}^n \frac{d^4 p_i}{(2\pi)^3} \delta(p_i^2 - m_i^2) \Theta(p_{i0}) \right] (2\pi)^4 \delta^{(4)} \left(p_a + p_b - \sum_{i=1}^n p_i \right), \quad (8)$$

where m_i are the on-shell masses of outgoing particles. The full phase space can be factorized as [35]

$$d\Phi_n(a, b; 1, \dots, n) = d\Phi_{n-m+1}(a, b; \pi, m+1, \dots, n) \frac{ds_\pi}{2\pi} d\Phi_m(\pi; 1, \dots, m), \quad (9)$$

where $\pi = \{1, \dots, m\}$ corresponds to a set of particle indices. Denoting the missing subset as $\bar{\alpha} = \{a, b, 1, \dots, n\} \setminus \alpha$ for all $\alpha \subset \{a, b, 1, \dots, n\}$, Eq. (9) allows one to decompose the complete phase space into building blocks corresponding to the t - and s -channel decay processes $T_{\alpha, b}^{\pi, ab\pi} = d\Phi_2(\alpha, b; \pi, \overline{ab\pi})$ and $S_{\pi}^{\rho, \pi \setminus \rho} = d\Phi_2(\pi; \rho, \pi \setminus \rho)$ and an s -channel production process $D_{\alpha, b}$, which corresponds to overall momentum conservation and the associated overall weight factor. These objects have been introduced as phase space vertices in Ref. [31], while

the integral $P_\pi = ds_\pi / 2\pi$, in Eq. (9), was called a phase space propagator. In this notation, there is a one-to-one correspondence between the computation of hard matrix elements in the Berends-Giele recursion [40–43] and the computation of phase-space weights. Thus, the computation of phase-space weights can be carried out in a recursive fashion, yielding the same (exponential) scaling as the computation of the hard matrix element. The basic building blocks of phase space integration can be summarized as

$$\begin{aligned}
P_\pi &= \frac{ds_\pi}{2\pi}, \\
S_\pi^{\rho,\pi\setminus\rho} &= \frac{\lambda(s_\pi, s_\rho, s_{\pi\setminus\rho})}{16\pi^2 2s_\pi} d\cos\theta_\rho d\phi_\rho, \\
T_{\alpha,b}^{\pi,\overline{ab}\pi} &= \frac{\lambda(s_{ab}, s_\pi, s_{\overline{ab}\pi})}{16\pi^2 2s_{ab}} d\cos\theta_\pi d\phi_\pi, \\
D_{\alpha,b} &= (2\pi)^4 d^4 p_{\overline{ab}} \delta^{(4)}(p_\alpha + p_b - p_{\overline{ab}}). \quad (10)
\end{aligned}$$

Here, λ is given by the Källén function $\lambda(a, b, c) = \sqrt{(a-b-c)^2 - 4bc}$. Note that, in the context of Monte Carlo integration, each basic integral ds , $d\cos\theta$ and $d\phi$ in Eq. (10) corresponds to a random variable,

$$\begin{aligned}
d\Phi_S(\pi) &= \sum_{(\pi_1, \pi_2) \in \mathcal{OP}(\pi)} \omega_\pi^{\pi_1, \pi_2} S_\pi^{\pi_1, \pi_2} P_{\pi_1} d\Phi_S(\pi_1) P_{\pi_2} d\Phi_S(\pi_2), \\
d\Phi_T^{(b)}(\alpha) &= \sum_{(\pi_1, \pi_2) \in \mathcal{OP}(\overline{ab})} \omega_\alpha^{\pi_1, \alpha\pi_1} T_{\alpha,b}^{\pi_1, \pi_2} P_{\pi_1} d\Phi_S(\pi_1) P_{\pi_2} d\Phi_T^{(b)}(\alpha\pi_1) + \omega_{\alpha,b} D_{\alpha,b} d\Phi_S(\overline{ab}). \quad (11)
\end{aligned}$$

In this context we defined the one- and no-particle phase space as $d\Phi(i) = 1$ and $d\Phi(\emptyset) = 0$. The numbers ω correspond to vertex-specific weights, which are normalized as $\sum_{\pi_1} \omega_\pi^{\pi_1, \pi\setminus\pi_1} = 1$ and $\sum_{\pi_1} \omega_\alpha^{\pi_1, \alpha\pi_1} + \omega_{\alpha,b} = 1$ and can be adapted to optimize the integrator. The sums run over all possible S - and T -type vertices which have a correspondence in the matrix element. The full differential phase space element is given by

$$d\Phi_n(a, b; 1, \dots, n) = d\Phi_T^{(b)}(a). \quad (12)$$

The recursive integrator is typically improved by combining it with the VEGAS algorithm [6,7]. In this configuration, VEGAS generates the input random numbers, \vec{x} , that are used to perform the basic integrals ds , $d\cos\theta$ and $d\phi$ in Eq. (10). We adopt the same strategy and simply replace VEGAS by the neural network + normalizing flow algorithm proposed in Refs. [28,29] and extended in Refs. [26,27], which is implemented in the I-FLOW package [38]. We include the multichannel weights, ω , in the integration, which allows us to work with a single network for all channels. All other components of the integrator remain the same. This makes it possible to directly compare the performance to existing algorithms. As the normalizing flow method is capable of capturing correlations in the multidimensional phase space, while still exhibiting polynomial scaling with the problem size, one would expect a performance improvement particularly when the number of dimensions in the problem is large.

chosen in the appropriate range. Details on the construction of the phase-space recursion are given in Ref. [31]. Here we simply recall the result in a schematic form:

- (1) Select an ordered partition of the multi-index $\pi \rightarrow (\pi_1, \pi_2)$ (s -channel) or $\overline{ab} \rightarrow (\alpha\pi_1 b, \pi_2)$ (t -channel).
- (2) If the partition is an s -channel, insert an s -channel vertex $S_\pi^{\pi_1, \pi_2}$ else insert a t -channel vertex $T_{\alpha,b}^{\pi_1, \pi_2}$.
- (3) If π_1 (π_2) is a multi-index, insert a propagator P_{π_1} (P_{π_2}).
- (4) Proceed until there is no multi-index left.

We can improve this procedure by forming an average over all possible ordered partitions (\mathcal{OP}) of each multi-index. Assigning each splitting an adjustable weight defines the recursive multichannel. This is formalized as follows:

IV. NUMERICAL RESULTS

This section presents a numerical study of the novel phase-space integrator. All computations are performed using the event generation framework SHERPA [44–46], with matrix elements provided by COMIX [31]. This use case differs slightly from Ref. [47], where matrix elements are computed by AMEGIC [48]. AMEGIC performs an explicit sum over color degrees of freedom, while COMIX uses Monte Carlo sampling. The neural network + normalizing flow integrator could in principle be trained on a combined color-momentum space when used with COMIX. We have tested this approach but have not found an advantage for processes beyond $W/Z + 1j$. Hence we refrain from using this technique.

A. Definition of efficiency

In order to assess the performance of the new integrator we investigate its unweighting efficiency. The most basic definition of unweighting efficiency would be the average weight during event generation [i.e., the integral, Eq. (1)], divided by the maximum weight,

$$\eta_{\text{basic}} = \frac{\langle f/g \rangle_G}{\max\{f/g\}}. \quad (13)$$

Its inverse corresponds to the number of events to be drawn on average, before an event is accepted with unit weight. If the distribution of weights does not exhibit a sharp upper edge, the denominator in Eq. (13) will depend on the sample size, and the efficiency will decrease with increasing number of events. This is particularly worrisome when events are generated in a distributed computing approach

(e.g., on the LHC computing grid). Each individual computing job will have its own, individual maximum, say $\max\{f/g\}_i$. In order to realize the accuracy of the combined event sample, each subsample must then be weighted by $\max\{f/g\}_i/\max_j\{\max\{f/g\}_j\}$.

Here we follow a different approach. It has recently been pointed out that a weighted combination of event samples is prone to outliers in the weight distribution, unless adaptive algorithms continue to be optimized during event generation [49]. If event generation is performed in a distributed fashion, this cannot be achieved, as the individual compute jobs do not communicate. The efficiency should therefore be computed based on the number of points during the last iteration of the optimization. This definition is still prone to possible outliers, which are removed by performing a bootstrap:

- (1) Assuming the number of events during optimization was N_{opt} , draw nN_{opt} events.
- (2) From these events, select m replicas of N_{opt} events each and compute their maximum weight.
- (3) Compute the total maximum, w_{max} , as the median of the individual maxima.

The efficiency is then given by

$$\eta_{\text{bs}} = \frac{\langle f/g \rangle_G}{w_{\text{max}}}. \quad (14)$$

There will be a number of event weights that exceed w_{max} . We can account for the mismatch on an event-by-event basis by recording their relative weight. Formally, if an event is generated with weight w_i , we keep it with weight \tilde{w}_i , where

$$\tilde{w}_i = w_{\text{max}} \Theta\left(\frac{w_i}{w_{\text{max}}} - R\right) \left[\frac{w_i}{w_{\text{max}}} \Theta\left(\frac{w_i}{w_{\text{max}}} - 1\right) + \Theta\left(1 - \frac{w_i}{w_{\text{max}}}\right) \right], \quad (15)$$

with $R \in [0, 1]$ a uniformly distributed random number. The event sample will then be partially weighted, unweighted against a maximum that corresponds—on a statistical basis—to the largest weight probed by the adaptive integrator during optimization. This allows one to reach in principle arbitrary precision by applying the bootstrap and jackknife techniques of Ref. [50]. Possible practical implementations of this method are discussed in Ref. [49].

B. Hyperparameter optimization

We use a quasirandom search strategy to optimize the hyperparameters of the neural networks. The basic idea of random search has been proposed in Ref. [51], and follows the known strategy used in tuning Monte Carlo event generators [52]. It was pointed out that a random search

strategy typically covers the space of hyperparameters better than a grid-based search, since the influence of hyperparameters is often uncorrelated. Since computing resources are limited and the computational cost of the training is fairly large, we resort to Sobol sequences to adequately populate the hyperparameter space, \mathcal{H} .

The figure of merit for a configuration $c \in \mathcal{H}$ is its unweighting efficiency. Due to the limited computing resources we redefine for the purpose of tuning w_{max} in Eq. (14) to be the 99.9th percentile of the distribution of nonzero absolute values of weights. We have tested that this makes the definition of w_{max} stable against variations in the random seed and gives results that are qualitatively compatible with the technique described in Sec. IV A. The main advantage of the alternative definition is that it allows one to determine the target number of points for the computation of w_{max} as $1/(\eta_{\text{cut}}\sigma^2\varepsilon)$, where η_{cut} is the cut efficiency, σ is the desired Monte Carlo accuracy, and ε is the desired percentile defining w_{max} .

Table I shows the sampling boundaries for \mathcal{H} , and the optimal choice for $W + 1\text{jet}$. Figure 2 shows the distribution of sampling points in the $N_{\text{epochs}}\text{-LR}$ and $N_{\text{bins}}\text{-}N_{\text{samples}}$ planes. The first suggests a high learning rate, coupled with a large number of epochs to be beneficial. The second suggests a strong preference for a small number of bins. We also performed a parameter scan for $W + 2\text{jet}$ and found the best configuration to be comparable to the above and to yield similar unweighting efficiency as the best configuration in the $W + 1\text{jet}$ setup. Due to limited computing resources, we did not perform a separate hyperparameter scan for $W + \geq 3$.

C. Comparison to existing approaches

In this section we compare the performance of our integrator based on the normalizing flow technique to the best alternative method available in the public event generation program SHERPA [44–46]. The basic computational setup is analogous to Ref. [4]. We consider W^\pm - and Z -boson production in proton-proton collisions at the high-luminosity LHC at $\sqrt{s} = 14$ TeV. We use the

TABLE I. Parameter ranges and prior distributions in \mathcal{H} and values of the optimal point in the scan. We use LR as abbreviation for the learning rate.

Parameter	c_{min}	c_{max}	prior	c_{best}
LR	10^{-5}	10^{-2}	log	4.55×10^{-4}
LR decay rate	0	1	lin	0.534
LR decay e -folds	2	10	lin	5
$N_{\text{nodes}}^{\text{max}}$	2^5	2^9	lin	2^9
N_{samples}	1000	10000	log	4148
N_{epochs}	500	5000	log	3824
N_{bins}	10	100	log	16
N_{layers}	6	10	lin	8

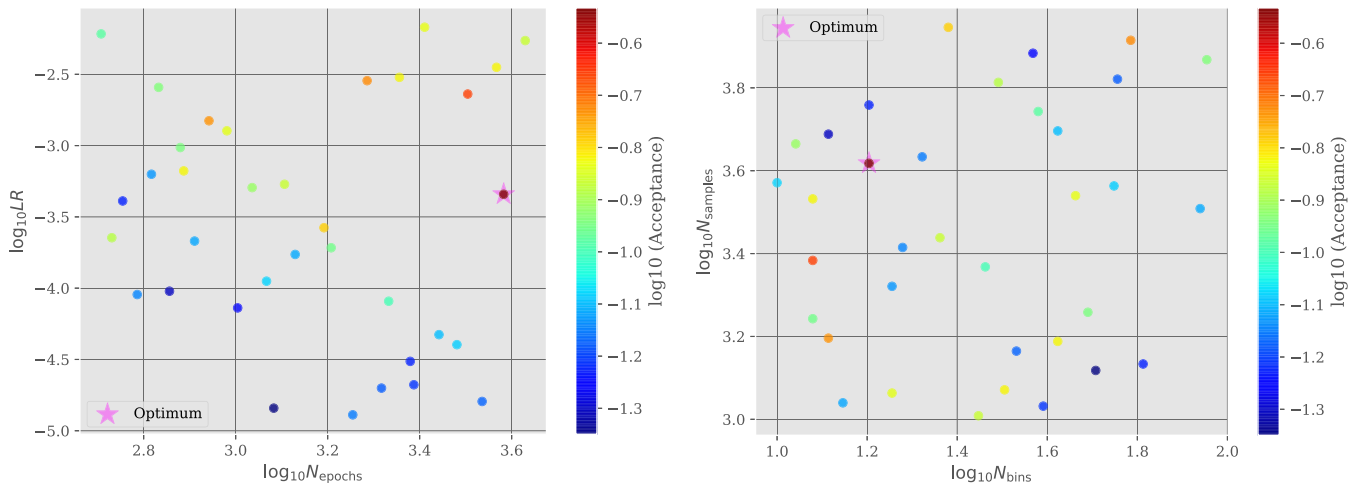


FIG. 2. Projections of the sampled parameters and color coded acceptances. The plot on the left suggest a high learning rate, coupled with a large number of epochs to be beneficial. The plot on the right suggests a strong preference for a small number of bins. The best performing configuration is indicated with a star.

NNPDF3.0 NNLO PDF set [53] and evaluate the strong coupling accordingly. Jets are defined using the k_T clustering algorithm with $R = 0.4$, $p_{T,j} > 20$ GeV and $|\eta_j| < 6$. Following the good agreement between parton-level and particle-level results established in Ref. [54], and the good agreement between fixed-order and MINLO [55] results established in Ref. [56], the renormalization and factorization scales are set to $\hat{H}_T/2$ [57].

Table II shows a comparison of unweighting efficiencies defined according to Eq. (14), where $N_{\text{opt}} = 2 \times 10^4$, $n = 50$, and $m = 10^3$. We give results for both leading-order cross sections, and for the subtracted real-emission corrections to next-to-leading order cross sections, using the dipole method of Catani and Seymour [58]. The subtracted real-emission corrections typically present the biggest challenge in cross-section calculations at next-to-leading order in the perturbative expansion, and therefore drive the computing demands of precision simulations for

LHC experiments. The new integrator based on neural networks and normalizing flows gives a much larger unweighting efficiency than SHERPA in processes with few jets, both at LO and at NLO precision. In processes with more final-state jets it performs similarly to the existing integration techniques in SHERPA. The neural network technique generally performs better when we do not combine it with a multichannel approach for initial-state integration. Only the major features of the final state should be mapped out, for example the Breit-Wigner resonance in W^\pm production. This indicates, unsurprisingly, that the normalizing flow approach is more efficient in approximating smooth structures of the integrand than in differentiating between effectively independent integration domains. It leads to the strikingly lower efficiencies in Z -boson production processes, where we have combined a $1/\hat{s}$ integrator and a Breit-Wigner mapping. In general, the new method is best applied to low-multiplicity problems,

TABLE II. Unweighting efficiencies at the LHC at $\sqrt{s} = 14$ TeV using the NNPDF 3.0 NNLO PDF set and a correspondingly defined strong coupling. Jets are identified using the k_T clustering algorithm with $R = 0.4$, $p_{T,j} > 20$ GeV and $|\eta_j| < 6$. In the case of Z/γ^* production, we also apply the invariant mass cut $66 < m_{ll} < 116$ GeV.

Unweighting efficiency		LO QCD					NLO QCD (RS)	
		$n = 0$	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 0$	$n = 1$
$W^+ + n$ jets	SHERPA	2.8×10^{-1}	3.8×10^{-2}	7.5×10^{-3}	1.5×10^{-3}	8.3×10^{-4}	9.5×10^{-2}	4.5×10^{-3}
	NN + NF	6.1×10^{-1}	1.2×10^{-1}	1.0×10^{-2}	1.8×10^{-3}	8.9×10^{-4}	1.6×10^{-1}	4.1×10^{-3}
	Gain	2.2	3.3	1.4	1.2	1.1	1.6	0.91
$W^- + n$ jets	SHERPA	2.9×10^{-1}	4.0×10^{-2}	7.7×10^{-3}	2.0×10^{-3}	9.7×10^{-4}	1.0×10^{-1}	4.5×10^{-3}
	NN + NF	7.0×10^{-1}	1.5×10^{-1}	1.1×10^{-2}	2.2×10^{-3}	7.9×10^{-4}	1.5×10^{-1}	4.2×10^{-3}
	Gain	2.4	3.3	1.4	1.1	0.82	1.5	0.91
$Z + n$ jets	SHERPA	3.1×10^{-1}	3.6×10^{-2}	1.5×10^{-2}	4.7×10^{-3}		1.2×10^{-1}	5.3×10^{-3}
	NN + NF	3.8×10^{-1}	1.0×10^{-1}	1.4×10^{-2}	2.4×10^{-3}		1.8×10^{-3}	5.7×10^{-3}
	Gain	1.2	2.9	0.91	0.51		1.5	1.1

where the training of the neural networks can be performed at reasonable speed with relatively few samples per epoch. We expect that in the high multiplicity cases the neural network + normalizing flow technique will also outperform SHERPA, if it can be trained over sufficiently many epochs with sufficiently many sample points. However, due to restricted computing resources, we were not able to verify this claim for the case of $W^\pm/Z + 4j$ production. The picture might be altered by future implementations of matrix element generators on accelerators. For exploratory work on this topic, see Refs. [59,60].

V. CONCLUSIONS

We have presented a novel approach to phase-space integration for collider physics simulations, which is based on neural networks and normalizing flows. The integrator is implemented as an add-on to the existing event generator SHERPA and can be used with both internal matrix-element generators, COMIX and AMEGIC. Neural network hyperparameters were tuned using a quasirandom search strategy. For the optimal set of parameters, the unweighting efficiency of the integrator exceeds that of conventional methods by a factor 2–3 in simple processes. In high-multiplicity processes, traditional techniques tend to perform similarly well, while also requiring fewer computing resources. We expect this picture to change as implementations of matrix element generators on accelerators such as GPUs and TPUs become available. These findings are corroborated by the results presented in Refs. [38,47].

ACKNOWLEDGMENTS

We are grateful to John Campbell and Tobias Neumann for many interesting discussions on the definition of unweighting efficiency. We thank Enrico Bothmann, Marek Schönherr, Steffen Schumann and Frank Siegert for comments and for discussions on unweighted event generation and adaptive integration using neural networks. We thank Tilman Plehn for discussions on neural networks and Joao M. Goncalves Caldeira, Felix Kling, Luisa Lucie-Smith, Nhan Tran, and the participants of the Aspen workshop “The Energy Frontier Beyond the LHC Run 2” for their comments. This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. The work of H.S. and S.H. was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) program, grants “HEP Data Analytics on HPC,” No. 1013935 and “HPC framework for event generation at colliders.” C.K. acknowledges the support of the Alexander von Humboldt Foundation. This work was performed in part at Aspen Center for Physics, which is supported by National Science Foundation Grant No. PHY-1607611. It used computing resources of SLAC National Accelerator Laboratory, and of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

-
- [1] B. Webber, *Annu. Rev. Nucl. Part. Sci.* **36**, 253 (1986).
 - [2] A. Buckley *et al.*, *Phys. Rep.* **504**, 145 (2011).
 - [3] The ATLAS Collaboration, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ComputingandSoftwarePublicResults> (2019).
 - [4] S. Höche, S. Prestel, and H. Schulz, *Phys. Rev. D* **100**, 014024 (2019).
 - [5] A. Buckley, [arXiv:1908.00167](https://arxiv.org/abs/1908.00167).
 - [6] G. P. Lepage, *J. Comput. Phys.* **27**, 192 (1978).
 - [7] G. P. Lepage, Report No. CLNS-80/447, 1980, <http://cds.cern.ch/record/123074>.
 - [8] J. H. Friedman and M. H. Wright, *J. Assoc. Comput. Mach.* **7**, 76 (1981).
 - [9] W. H. Press and G. R. Farrar, *Comput. Phys.* **4**, 190 (1990).
 - [10] T. Ohl, *Comput. Phys. Commun.* **120**, 13 (1999).
 - [11] S. Jadach, *Comput. Phys. Commun.* **130**, 244 (2000).
 - [12] T. Hahn, *Comput. Phys. Commun.* **168**, 78 (2005).
 - [13] K. Kröniger, S. Schumann, and B. Willenberg, *Comput. Phys. Commun.* **186**, 1 (2015).
 - [14] J. Bendavid, [arXiv:1707.00028](https://arxiv.org/abs/1707.00028).
 - [15] M. D. Klimek and M. Perelstein, [arXiv:1810.11509](https://arxiv.org/abs/1810.11509).
 - [16] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, *J. High Energy Phys.* **07** (2014) 079.
 - [17] F. Bishara and M. Montull, [arXiv:1912.11055](https://arxiv.org/abs/1912.11055).
 - [18] S. Otten, S. Caron, W. de Swart, M. van Beekveld, L. Hendriks, C. van Leeuwen, D. Podareanu, R. R. de Austri, and R. Verheyen, [arXiv:1901.00875](https://arxiv.org/abs/1901.00875).
 - [19] B. Hashemi, N. Amin, K. Datta, D. Olivito, and M. Pierini, [arXiv:1901.05282](https://arxiv.org/abs/1901.05282).
 - [20] R. Di Sipio, M. F. Giannelli, S. K. Haghighat, and S. Palazzo, *J. High Energy Phys.* **08** (2019) 110.
 - [21] A. Butter, T. Plehn, and R. Winterhalder, *SciPost Phys.* **7**, 075 (2019).
 - [22] S. Carrazza and F. A. Dreyer, *Eur. Phys. J. C* **79**, 979 (2019).
 - [23] C. Ahdida *et al.* (SHiP Collaboration), *J. Instrum.* **14**, P11028 (2019).

- [24] A. Butter, T. Plehn, and R. Winterhalder, [arXiv:1912.08824](https://arxiv.org/abs/1912.08824).
- [25] K. T. Matchev and P. Shyamsundar, [arXiv:2002.06307](https://arxiv.org/abs/2002.06307).
- [26] T. Müller, B. McWilliams, F. Rousselle, M. Gross, and J. Novák, *ACM Trans. Graph.* **38**, 145:1 (2019).
- [27] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, [arXiv:1906.04032](https://arxiv.org/abs/1906.04032).
- [28] L. Dinh, D. Krueger, and Y. Bengio, in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, 2015, Workshop Track Proceedings* (2015), <http://arxiv.org/abs/1410.8516>.
- [29] L. Dinh, J. Sohl-Dickstein, and S. Bengio, [arXiv:1605.08803](https://arxiv.org/abs/1605.08803).
- [30] R. Kleiss and R. Pittau, *Comput. Phys. Commun.* **83**, 141 (1994).
- [31] T. Gleisberg and S. Höche, *J. High Energy Phys.* **12** (2008) 039.
- [32] J. Song, S. Zhao, and S. Ermon, [arXiv:1706.07561](https://arxiv.org/abs/1706.07561).
- [33] D. Levy, M. D. Hoffman, and J. Sohl-Dickstein, [arXiv:1711.09268](https://arxiv.org/abs/1711.09268).
- [34] M. Hoffman, P. Sountsov, J. V. Dillon, I. Langmore, D. Tran, and S. Vasudevan, [arXiv:1903.03704](https://arxiv.org/abs/1903.03704).
- [35] F. James, CERN Report No. CERN-68-15, 1968, <http://cds.cern.ch/record/275743>.
- [36] D. P. Kingma, T. Salimans, and M. Welling, [arXiv:1606.04934](https://arxiv.org/abs/1606.04934).
- [37] G. Papamakarios, T. Pavlakou, and I. Murray, [arXiv:1705.07057](https://arxiv.org/abs/1705.07057).
- [38] C. Gao, J. Isaacson, and C. Krause, [arXiv:2001.05486](https://arxiv.org/abs/2001.05486).
- [39] E. Byckling and K. Kajantie, *Nucl. Phys.* **B9**, 568 (1969).
- [40] F. A. Berends and W. T. Giele, *Nucl. Phys.* **B306**, 759 (1988).
- [41] F. Caravaglios and M. Moretti, *Phys. Lett. B* **358**, 332 (1995).
- [42] A. Kanaki and C. G. Papadopoulos, *Comput. Phys. Commun.* **132**, 306 (2000).
- [43] C. Duhr, S. Höche, and F. Maltoni, *J. High Energy Phys.* **08** (2006) 062.
- [44] T. Gleisberg, S. Höche, F. Krauss, A. Schälicke, S. Schumann, and J. Winter, *J. High Energy Phys.* **02** (2004) 056.
- [45] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, F. Siegert, and J. Winter, *J. High Energy Phys.* **02** (2009) 007.
- [46] E. Bothmann *et al.*, *SciPost Phys.* **7**, 034 (2019).
- [47] E. Bothmann, T. Janßen, M. Knobbe, T. Schmale, and S. Schumann, [arXiv:2001.05478](https://arxiv.org/abs/2001.05478).
- [48] F. Krauss, R. Kuhn, and G. Soff, *J. High Energy Phys.* **02** (2002) 044.
- [49] J. Campbell and T. Neumann, *J. High Energy Phys.* **12** (2019) 034.
- [50] B. Efron, *J. R. Stat. Soc. Ser. B* **54**, 83 (1992).
- [51] J. Bergstra and Y. Bengio, *J. Mach. Learn. Res.* **13**, 281 (2012), <http://www.jmlr.org/papers/v13/bergstra12a.html>.
- [52] A. Buckley, H. Hoeth, H. Lacker, H. Schulz, and J. E. von Seggern, *Eur. Phys. J. C* **65**, 331 (2010).
- [53] R. D. Ball *et al.* (NNPDF), *Eur. Phys. J. C* **77**, 663 (2017).
- [54] J. Bellm *et al.*, *Eur. Phys. J. C* **80**, 93 (2020).
- [55] K. Hamilton, P. Nason, and G. Zanderighi, *J. High Energy Phys.* **10** (2012) 155.
- [56] F. R. Anger, F. Febres Cordero, S. Höche, and D. Maître, *Phys. Rev. D* **97**, 096010 (2018).
- [57] C. F. Berger, Z. Bern, L. J. Dixon, F. Febres-Cordero, D. Forde, T. Gleisberg, H. Ita, D. A. Kosower, and D. Maître, *Phys. Rev. Lett.* **106**, 092001 (2011).
- [58] S. Catani and M. H. Seymour, *Nucl. Phys.* **B485**, 291 (1997).
- [59] W. Giele, G. Stavenga, and J.-C. Winter, *Eur. Phys. J. C* **71**, 1703 (2011).
- [60] K. Hagiwara, J. Kanzaki, Q. Li, N. Okamura, and T. Stelzer, *Eur. Phys. J. C* **73**, 2608 (2013).