

Kicking it Off(-shell) with Direct Diffusion

Anja Butter^{1,3}, Tomáš Ježo², Michael Klasen²,
Mathias Kuschick², Sofia Palacios Schweitzer¹, and Tilman Plehn¹

¹ Institut für Theoretische Physik, Universität Heidelberg, Germany

² Institut für Theoretische Physik, Universität Münster, Germany

³ LPNHE, Sorbonne Université, Université Paris Cité, CNRS/IN2P3, Paris, France

Abstract

Off-shell effects in large LHC backgrounds are crucial for precision predictions and, at the same time, challenging to simulate. We present a novel method to transform high-dimensional distributions based on a diffusion neural network and use it to generate a process with off-shell kinematics from the much simpler on-shell one. Applied to a toy example of top pair production at LO we show how our method generates off-shell configurations fast and precisely, while reproducing even challenging on-shell features.

Contents

1	Introduction	2
2	Off-shell vs. on-shell events	3
3	Direct Diffusion	5
4	Classifier control and reweighting	9
5	Outlook	11
	References	14

1 Introduction

Fast and precise theoretical predictions from first principles are required for essentially every experimental LHC analysis. This is especially true for modern inference methods, where the complete phase space coverage ensures an optimal measurement [1,2]. Based on perturbative quantum field theory, precision simulations of the hard scattering process face two challenges, the loop order and the multiplicity of the partonic final state [3]. The latter increases rapidly for example, when we describe the production of decaying heavy particles. Naively, one could expect that describing the decay kinematics close to the mass shell of the heavy particles, for example using a Breit-Wigner propagator, is sufficient. However, given the precision targets of the upcoming LHC runs, on-shell approximations are no longer justified [4,5]. In view of the HL-LHC, simulations need to describe heavy particle production and decay including, both, quantum corrections and off-shell kinematics.

In addition to the theoretical and computational effort behind precise calculations, the question is how they can be made available through multi-purpose event generators. Here we can resort to modern machine learning (ML) [1,6]. Following the modular structure of event generators, the most obvious ML-applications are neural network surrogates for expensive scattering amplitudes [7–13]. Aside from the speed, these surrogates have the advantage that they can always be evaluated in parallel on GPUs [14–17]. Given these amplitudes, we can then improve the phase-space integration and sampling [18–26]. One advantage is that simulations are defined over interpretable physics phase spaces, for example scattering events [27–33], parton showers [34–41], and detector simulations [42–65]. All of these speed improvements maintain the first-principles nature of LHC event generators, which means they take theory predictions and evaluate them using faster ML-methods.

The goal of this study is to develop a precise and fast generative network that maps two LHC phase spaces onto each other [65–67]. If one phase space is a narrow sub-manifold of the other one, like in the case of on-shell and off-shell phase spaces of production of heavy unstable particles, one cannot use regression or classifier reweighting but must resort to generative networks. For LHC applications, we already know that precision-generative networks [32,68–71] are easy to ship and powerful in amplifying their training data [72,73]. Tasks we can solve with their help include event subtraction [74], event unweighting [75,76], or super-resolution enhancement [77,78]. Their conditional versions enable new analysis methods, like probabilistic unfolding [67,79–87], inference [88–90], or anomaly detection [91–96]. Currently, the best-performing network architectures are normalizing flows for invertible tasks and diffusion networks for sampling problems, with help from transformers for combinatorics.

In this paper we present such a mapping and use it to efficiently generate events over an off-shell phase space from given on-shell events. The idea behind this sampling from on-shell events is that the generative network does not have to reproduce the on-shell features and can focus on the additional and relatively smooth off-shell extension. We start by describing our training dataset and the problem of off-shell event generation in Sec. 2. Our generative network setup, based on a Bayesian- Conditional Flow Matching (CFM) architecture [33], is presented in Sec. 3. To control the generative network performance [71] and to improve the precision of the kinematic distributions [32], we apply a classifier reweighting in Sec. 4. We present an Outlook in Sec. 5.

2 Off-shell vs. on-shell events

Our benchmark process is the complete off-shell top pair production followed by leptonic decays,

$$pp \rightarrow be^+ \nu_e \bar{b}\mu^- \nu_\mu. \quad (1)$$

Top pair production with its rich resonance structure is known to challenge phase space sampling, just as generative networks for transition amplitudes [30].

In the factorized approach, in which the production and decay processes decouple, this process is known all the way up to NNLO-QCD [97–100], NLO-EW [101–105] and NNLO-QCD combined with NLO-EW [106] in the production process; NNLO-QCD [107, 108] in the decay; and NNLO-QCD in both production and decay [109–111]. Full off-shell calculations in the dilepton channel are so far only available at NLO-QCD [112–117], but a calculation of full off-shell top pair production in association with an extra jet at NLO is also available in Ref. [118]. In this pilot study we restrict ourselves to leading order in QCD, $\mathcal{O}(\alpha_s^2\alpha^4)$, and reserve the application of our method to higher-order predictions to a future publication.

We generate event samples for 13 TeV proton-proton collisions with NNPDF31_nlo_as_0118 parton distributions [119]. The neutrinos, charged leptons and quarks of the first two generations are treated as massless, and the CKM matrix is assumed to be trivial. All input parameters are given in Tab. 1, and the electromagnetic coupling α and the weak mixing angle are derived from the weak gauge-boson masses and the Fermi constant.

Our study and our results can be extended to higher-order predictions or other processes by combining different jet multiplicities in the final state [90]. Our two benchmark datasets are generated with HVQ [120] and BB4L [121, 122], respectively. The data generated with HVQ only includes approximate off-shell effects using a finite top width and including spin correlations [123] and is referred to as on-shell data. The generator is based on the POWHEG method [124, 125] and is part of the POWHEG BOX V2 [126] package.

The data generated with BB4L takes into account full off-shell effects also including singly-resonant and non-resonant contributions and the corresponding interferences. The generator employs the POWHEGRES method [127], tailored for simulations with unstable particles. In this case, W bosons and b quarks do not always stem from a top decay.

In Fig. 1 we illustrate the size of the off-shell effects for a selection of kinematic distributions of final-state leptons and b -quarks. These distributions are not meant to compare realistic predictions with different treatment of off-shell effects, but rather our two datasets, so no event selection criteria are applied. For the invariant mass of the lepton- b system we ensure, through charge identification, that the two particles come from the same (anti)top decay. Correspondingly, their invariant mass has an upper edge that does not exist for off-shell events. For the reconstructed top mass, or the invariant mass of the three decay products, we clearly see the Breit-Wigner propagator form, with an explicit cutoff. Far below the actual top mass,

m_t	172.5 GeV	Γ_t	1.453 GeV
m_b	4.75 GeV		
m_Z	91.188 GeV	Γ_Z	2.441 GeV
m_W	80.419 GeV	Γ_W	2.048 GeV
m_H	125.0 GeV	Γ_H	0.0403 GeV
$\mathcal{B}(W \rightarrow e\nu/\mu\nu)$		1/9	
G_F		$1.16585 \times 10^{-5} \text{GeV}^{-2}$	

Table 1: Parameters used for the generation of the training datasets.




Figure 1: Example distributions for on-shell and off-shell $t\bar{t}$ event samples, illustrating the different phase space coverage. The right panels show the same distribution as the left panels, but zoomed into the respective bulk region.

the off-shell prediction develops a shoulder at the W -mass. The secondary panels show the ratios of the integrated one-dimensional phase space densities, illustrating that a reweighting strategy between the two samples is unlikely to work.

The strategy behind our surrogate network is to generate an off-shell event dataset and learn its structures relative to a corresponding on-shell event dataset. This strategy can be applied to any process and at any order in perturbation theory. Our two datasets consist of 5M unit-weight events each. The six particles in the final state of Eq.(1) are represented by $\{p_T, \eta, \phi\}$, with fixed external particle masses. We remove three degrees of freedom through a global azimuthal rotation and by enforcing two-dimensional transverse momentum conservation, leaving us with a 15-dimensional phase space. Each on-shell requirement replaces a full phase space dimension by a fixed range, given by the Breit-Wigner shape with a hard-coded cutoff in the reconstructed invariant mass. Moreover, we conceal information relevant for eventual parton showering like the colourflow configuration or resonance history assignment. This is appropriate because the BB4L generator in its default setup does not distinguish between the $t\bar{t}$ and single-top resonance histories and assigns colourflow correspondingly.

For the network input, the kinematic variables are preprocessed: we scale the transverse momenta to $p_T^{1/3}$ and express the azimuthal angles as $\text{arctanh}(\phi/\pi)$. As subsequent classifier input, it turns out that $p_T^{-1/3}$ leads to the best results. Finally, all input dimensions are standardized to zero mean and unit variance.

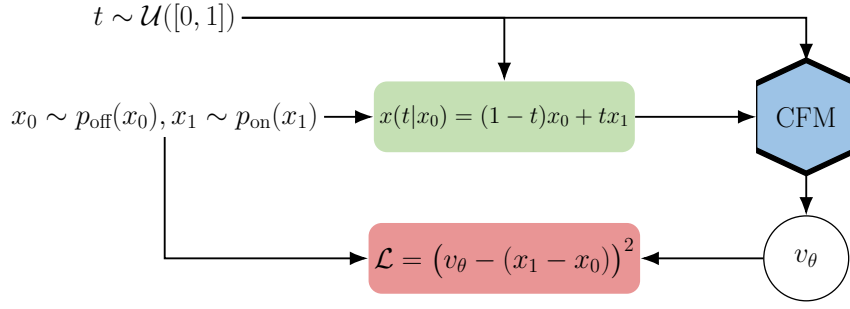


Figure 2: Training procedure for a CFM network mapping between on-shell and off-shell phase space distributions. Diagram adapted from Ref [33].

3 Direct Diffusion

For the network training, we start with the condition that we do not want to train on paired on-shell and off-shell events, because such a pairing does not follow a well-defined algorithm. This does not mean it is impossible to construct such a mapping, but we expect it to lead to artifacts. Instead, we will train a generative Direct Diffusion (DiDi) network on two phase space densities, one from on-shell and one from off-shell events.

Usually, we employ generative networks to learn a mapping between a simple latent space and some kind of phase space. The distribution in the latent space is sampled from, so we typically use a uniform or a Gaussian distribution. In this application, the latent space corresponds to the on-shell phase space, and the mapping is trained to generate off-shell events,

$$x \sim p_{\text{on}}(x) \longleftrightarrow x \sim p_{\text{model}}(x|\theta) \approx p_{\text{off}}(x). \quad (2)$$

Using the conditional flow matching (CFM) setup for LHC events [33, 128], we encode the transformation from on- to off-shell events as a continuous time evolution, which follows an ordinary differential equation (ODE)

$$\frac{dx(t)}{dt} \equiv v(x(t), t). \quad (3)$$

In terms of the related probability density, we can formulate the same task as

$$\frac{\partial p(x, t)}{\partial t} + \nabla_x [p(x, t)v(x, t)] = 0. \quad (4)$$

The ODE and the continuity equation are equivalent, so we can train a network to represent the velocity field and use this velocity field to generate samples using a fast ODE solver for Eq.(3). Defining our time dependent probability density as

$$p(x, t) \rightarrow \begin{cases} p_{\text{off}}(x) & t \rightarrow 0 \\ p_{\text{on}}(x) & t \rightarrow 1, \end{cases} \quad (5)$$

we now need to construct the associated velocity field.

Unlike for the usual generative setup, we define $x(1) = x_1$ as a sample from the on-shell phase space, whereas $x(0) = x_0$ corresponds to the off-shell phase space. Therefore, we adapt the linear trajectory between on-shell and off-shell events to

$$x(t|x_0) = (1-t)x_0 + tx_1 \rightarrow \begin{cases} x_0 & t \rightarrow 0 \\ x_1 \sim p_{\text{on}} & t \rightarrow 1. \end{cases} \quad (6)$$

The true conditional velocity field of our linear trajectory linked to our probability density in Eq. (5) is hence given by

$$v(x(t|x_0), t|x_0) = \frac{d}{dt} [(1-t)x_0 + tx_1] = -x_0 + x_1. \quad (7)$$

The remaining derivation is in complete analogy to Ref. [33], leading to the simple MSE loss

$$\mathcal{L}_{\text{CFM}} = \left\langle [v_\theta((1-t)x_0 + tx_1, t) - (x_1 - x_0)]^2 \right\rangle_{t \sim \mathcal{U}([0,1]), x_0 \sim p_{\text{off}}, x_1 \sim p_{\text{on}}}. \quad (8)$$

Because the CFM setup does not include a likelihood, we can go directly from p_{on} to p_{off} , without extra effort due to detours like in Flows4Flows [66], and without any pairing between $x_0 \sim p_{\text{off}}$ and $x_1 \sim p_{\text{on}}$.

As usual, we use a Bayesian version of the generative network [32, 129], to extract uncertainties on the learned phase space density. The Bayesian CFM loss [33]

$$\mathcal{L}_{\text{B-CFM}} = \left\langle \mathcal{L}_{\text{CFM}} \right\rangle_{\theta \sim q(\theta)} + c \text{KL}[q(\theta), p(\theta)] \quad (9)$$

includes a hyperparameter c to balance the regular CFM loss with the Bayesian regularization term. If the first loss term was a likelihood loss, this factor would be fixed by Bayes' theorem. We have checked that the network performance is stable over many orders of magnitudes for $c = 10^{-10} \dots 10^{-2}$. For larger values we observe that the training becomes unstable, as expected, while for very small values the uncertainty can no longer be captured. In addition, the prior weight distribution $q(\theta)$ is given by a unit Gaussian, where the choice of width hardly affects the network performance. Compared to the deterministic counterpart, this Bayesian network is equally precise.

The network and training setup is visualized in Fig. 2. We use a simple dense network with SiLU activation, where the last layer is initialized at zero. This sets the initial velocity field to zero and induces an identity mapping at the starting point of the training. We do not enforce a prescription for turning a given on-shell event into an off-shell event through the training data or the training procedure. In fact, for each epoch different phase space points will be connected via a linear trajectory. Instead, the network training constructs its own mapping to populate the off-shell phase space. This happens as part of the loss minimization, which means the transport map follows from an implicit measure encoded in the network architecture and loss.

Hyperparameter	
Embedding dimension	64
Layers	8
Intermediate dimensions	768
LR scheduling	OneCycle
Starter LR	10^{-4}
Max LR	10^{-3}
Epochs	1000
Batch size	16384
c	10^{-3}
# Training events	3 M

Table 2: Generative network setup (DiDi) and hyperparameters.




Figure 3: Results from our Direct Diffusion off-shell generator, compared to the on-shell starting point and the off-shell training distributions.

As the CFM training objective is a simple regression, we can train on 17 dimensions, including two redundant degrees of freedom, as this happens to increase the precision. The two additional observables, the transverse momentum and the polar angle of the neutrino, are determined by transverse momentum conservation and are hence multidimensional correlations. Empirically we found that it is easier for the model to learn the behavior of those observables when handed directly. Especially more complicated correlations such as the invariant mass of the reconstructed top benefit greatly from this additional information. While improving the efficiency of the training these dimensions will be ignored for the actual event generation.

The network hyperparameters and the training parameters are given in Tab. 2. We encode t in a higher embedding dimension and following [128] we add batch-wise random noise of scale 10^{-4} to x_0 and x_1 during training. We use the standard *Dopri5* ODE solver to sample from our network. In the interest of precision we use a large batch size. This is a problem for generic optimal-transport networks [130], but can be easily implemented in our architecture. We tested the OT-CFM [130] and found that in our setup the required small batch size indeed led to worse performance.

In Fig. 3 we show a set of one-dimensional kinematic distributions, for the on-shell data we start from, the off-shell training data, and the generated off-shell data. In the first panel we




Figure 4: Migration plot — correlation between generated (off-shell) and starting (on-shell) distributions for two kinematic correlations. It illustrates the mapping found by the network during training.

see how the network learns subtle differences almost perfectly well. The typical precision is around the per-cent level. For complex and sensitive correlations, like the lepton- b invariant mass and the reconstructed top mass we start with a huge deviation between the on-shell distribution and the off-shell target. To generate these distributions correctly, the generative network has to learn correlations in the corresponding 9-dimensional sub phase space.

Our results confirm that in phase space regions covered well by both datasets, the generative network reproduces the target distribution precisely, as one could expect. However, we also see that even in phase space regions not populated by on-shell events the target distribution is reproduced relatively precisely. The typical agreement between the generated and target densities is around 10% in phase space regions with relatively little training data. While this deviation is covered by the uncertainties of the Bayesian network, we propose possible improvements in the next section.

Finally, we can ask how our generative network fills the off-shell phase space from the on-shell events. In Fig. 4 we show the correlations between the generated off-shell distribution and the on-shell starting distributions, i.e. the migration of paired latent and target phase space events from the forward simulation, where we emphasize that the pairing is only defined by the network evaluation, not by the training. In general, the correlation between the kinematic observables should be close to the identity, as confirmed by the $p_{T,\bar{b}}$ distribution in the left panel. This means that the shift from on-shell to off-shell phase space is relatively small and uncorrelated. However, for the reconstructed top mass, some of the events have to be shifted by a larger value, as illustrated in the right panel. While the width of the linear correlation becomes very small around the top mass peak, it rapidly increases away from the peak, demonstrating the large shift required to populate the off-shell phase space. Moreover, our network maps events from each side of the Breit-Wigner peak to the same side of the off-shell distribution.




Figure 5: Histogram of the learned event weights, evaluated on the off-shell training data and the DiDi-generated off-shell events. The two panels show the same weights, on linear and logarithmic scales.

4 Classifier control and reweighting

Because the unsupervised density estimation underlying our generative network is more challenging and less precise than a supervised classifier training, we can use a trained classifier as a function of phase space to systematically check and improve the generative network. A perfectly trained statistical classifier converges to the ratio of likelihoods,

$$C(x) = \frac{p_{\text{data}}(x)}{p_{\text{data}}(x) + p_{\text{model}}(x)}. \quad (10)$$

As a function of phase space we can use this classifier to construct an event reweighting, which improves the precision of the generative networks [32],

$$w(x) = \frac{p_{\text{data}}(x)}{p_{\text{model}}(x)} = \frac{C(x)}{1 - C(x)}. \quad (11)$$

In addition, we can use the same learned event weights to determine the precision of the generative network and systematically search for failure modes by searching for clusters of very small or very large weights in phase space [71].

We train the classifier on 27 observables, our 15 physical dimensions, complemented by the reconstructed top and anti-top masses, the reconstructed W^+ and W^- masses, the reconstructed masses of the $\bar{b}l$ and $b\bar{l}$ systems, and the six corresponding transverse momenta. For

Hyperparameter	
Layers	5
Intermediate dimensions	512
Dropout	0.1
Normalization	BatchNorm1d
LR scheduling	ReduceOnPlateau
Starter LR	1^{-3}
Patience	10
Epochs	100
Batch size	1024
# Training events	2.5 M

Table 3: Classifier network setup and hyperparameters.




Figure 6: Clustering of the classifier trained to distinguish the off-shell training data from DiDi-generated events, for two example distributions.

the input of our classifier we sample events from our Bayesian generator setting each network weight to its mean value. The setup is given in Tab. 3.

We show histograms of learned phase space weights $w(x_i)$ on a linear and a logarithmic axis in Fig. 5. As expected, the distributions peak at unit weights, with a width around 0.3. Following Eq.(11), large weights correspond to phase space regions where the generative network produces a too small density of off-shell points; small weights mark phase space regions where the generative network produces too many off-shell events, compared to the training data. To study both tails of the weight distribution, we evaluate the weights over the combination of 2M training and 2M generated events and confirm that both tails decrease rapidly. We eventually clip the event weights to $w < 15$ to improve the numerical behavior of our generation and avoid sparks in regions of low statistics.

In Fig. 6 we track phase space regions with very small and very large weights. We compare all events of our test sample to the subsets with $w(x) < 0.6$, corresponding to 12.3% of the sample, and $w(x) > 1.6$, corresponding to 6.2% of the generated sample. The two shown distributions illustrate the general feature that the kinematic distributions for both tails are similar. In the $p_{T,\bar{b}}$ -distribution we can identify two limitations of the network training: for small transverse momenta a slight shift of the cliff will lead to large relative weight corrections, while for large transverse momenta the decreasing density of training events will increase the relative size of the noise. For the reconstructed anti-top mass, one of the critical distributions for our generative task, the low-weight and high-weight tails are also comparable with each other and also comparable on both sides of the mass peak. The only distinct feature appears around the peak where events with small weight dominate the slightly over populated sides of the peak. This indicates that all weight tails are generated by noise, there are no missing localized features, and the limiting factor is the statistics of the training data.

Even though the shortcomings of our generative networks, visible in Fig. 3, arise from noisy network training and do not reflect systematic shortcomings, they affect the trained generative network in a systematic, localized manner. As a function of phase space, we can correct them using the event weights from Eq.(11), because the classifier network is more sensitive and more precise than the generator network [71]. In Fig. 7 we show a set of kinematic distributions for reweighted events, where the uncertainty is given by the Bayesian generator. Comparing

the agreement between the reweighted and the target distribution with the unweighted performance from Fig. 3 we see significant improvements. In the secondary panels we show the reweighted predictions from the Bayesian generator, allowing us to compare the statistical uncertainty from the training data with the predictive uncertainty from the generative network. For the entire range of $p_{T,\bar{b}}$ the network agrees with the true distribution within its predictive uncertainties, and almost within the uncertainties of the training data. For $m_{\bar{t}b}$, DiDi has to cover phase space far beyond the on-shell structures, and, again it hardly exceeds the statistical uncertainties of the training data and provides a conservative uncertainty estimate from the Bayesian setup.

The momentum and the mass of the reconstructed W^- match the truth perfectly in the bulk, and roughly within the statistics of the training data in the tails. The former is important, because it confirms the original motivation that the network reproduces the on-shell features with extremely high precision. The same can be seen for the reconstructed anti-top mass, where most of the phase space is filled by extrapolating from the on-shell distribution and the network learns these new phase space regions without degrading the precision in the bulk at all.

5 Outlook

Sufficiently fast and precise event generation for the HL-LHC is a major challenge in the coming decade. A perfect example is the generation of off-shell kinematic configurations for large backgrounds. Even when they can be computed from perturbative QFT, with significant investment in CPU time, they need to be implemented in standard event generators. A promising strategy to provide them in an amortized manner is generative surrogate networks. They are easy to include and ship with event generators and enhance the precision of a limited number of events, just like a fitted function to a statistically limited dataset.

Generating off-shell events is also an exciting problem for modern machine learning, because it cannot be solved by regressing an amplitude at a give phase space point. Instead, it requires a generative surrogate to cover the off-shell phase space. We proposed a new method, namely to generate off-shell configurations relative to given on-shell configurations. Its advantage is that the generative network only learns a controlled deviation from a simple unit transformation. This simplified task allows us to generate a rich resonance structure without the usual challenges in network size and precision training. The approach is enabled by Direct Diffusion (DiDi) networks, which sample from any given distribution to produce another distribution, in our case transforming on-shell events into off-shell events.

For top pair production we have shown that already a relatively small network with limited training effort can reproduce the target off-shell distributions at the 10% level. Most importantly, the DiDi generation does not encounter any issues with learned features. The challenging on-shell peaks are described at the 5% level or better, as seen in Fig. 3.

Using a classifier reweighting we can improve its precision to the few per-cent level, even in challenging kinematic distributions, Fig. 7. Now the bulk precision exceeds 1%, and the precision in the extrapolation region is reliably below or at the 10% level. The same classifier allows us to control the performance of the generative networks and ensures that all local features in phase space are correctly learned. Finally, the Bayesian extension of the DiDi architecture provides us with a conservative uncertainty estimate, tracking the effects of statistically limited training data.

Our approach can be extended to higher-orders as long as the corresponding event samples can be provided as training data. Such an extension is not trivial because with each extra particle in the final state, e.g. due to an extra emission in the NLO real correction, the di-




Figure 7: Results from our re-weighted Direct Diffusion off-shell generator, compared to the on-shell starting point and the off-shell target distributions.

mension of the phase space increases. We do not foresee any conceptual issues however, as increasing the dimension of the phase space in our setup is straightforward and relatively small

event samples were needed at LO, and both diffusion networks as well as Classifiers have been shown to scale well with dimension. The size of the off-shell effect, or more specifically the ratio of the full and approximate off-shell prediction for a given observable, could in principle change dramatically when changing orders.* This is not a problem though, as DiDi is able to reproduce full off-shell predictions also in regions where the approximate off-shell prediction is vanishing. For realistic predictions we will also need to address issues related to matching to the parton shower, like shower starting scale and colourflow configurations not discussed here. For the latter we may have to include single top production in our on-shell data set.

Our fast generative surrogate is ready to be implemented standalone using Les Houches Event (LHE) format for input and output. Generative networks of this kind will also be part of the ML-enhanced ultrafast MADGRAPH event generator [23, 24].

Acknowledgements

We thank Nathan Huetsch for many fruitful discussions. AB and TP are supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant 396021762 – TRR 257 *Particle Physics Phenomenology after the Higgs Discovery*. AB gratefully acknowledges the continuous support from LPNHE, CNRS/IN2P3, Sorbonne Université and Université de Paris Cité. Work at the University of Münster is supported by the BMBF through project *InterKIWWU* and by the DFG through SFB 1225 *Isoquant*, project-id 273811115, and the Research Training Group 2149 *Strong and Weak Interactions - from Hadrons to Dark Matter*. SPS is supported by the BMBF Junior Group *Generative Precision Networks for Particle Physics* (DLR 01IS22079). The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grant no INST 39/963-1 FUGG (bwForCluster NEMO). This work was supported by the DFG under Germany's Excellence Strategy EXC 2181/1 - 390900948 *The Heidelberg STRUCTURES Excellence Cluster*.

*For the example of $t\bar{t}$ production at hand, we compared the size of the off-shell effect at LO vs at NLO, and did not find differences in excess of one order of magnitude.

References

- [1] A. Butter, T. Plehn, S. Schumann, *et al.*, *Machine Learning and LHC Event Generation*, in *2022 Snowmass Summer Study*. 3, 2022. [arXiv:2203.07460 \[hep-ph\]](#).
- [2] K. Cranmer, J. Brehmer, and G. Louppe, *The frontier of simulation-based inference*, *Proc. Nat. Acad. Sci.* **117** (2020) 48, 30055, [arXiv:1911.01429 \[stat.ML\]](#).
- [3] J. M. Campbell *et al.*, *Event Generators for High-Energy Physics Experiments*, in *2022 Snowmass Summer Study*. 3, 2022. [arXiv:2203.11110 \[hep-ph\]](#).
- [4] G. Heinrich, A. Maier, R. Nisius, J. Schlenk, M. Schulze, L. Scyboz, and J. Winter, *NLO and off-shell effects in top quark mass determinations*, *JHEP* **07** (2018) 129, [arXiv:1709.08615 \[hep-ph\]](#).
- [5] S. Ferrario Ravasio, T. Ježo, P. Nason, and C. Oleari, *A theoretical study of top-mass measurements at the LHC using NLO+PS generators of increasing accuracy*, *Eur. Phys. J. C* **78** (2018) 6, 458, [arXiv:1906.09166 \[hep-ph\]](#). [Addendum: *Eur.Phys.J.C* 79, 859 (2019)].
- [6] T. Plehn, A. Butter, B. Dillon, and C. Krause, *Modern Machine Learning for LHC Physicists*, [arXiv:2211.01421 \[hep-ph\]](#).
- [7] F. Bishara and M. Montull, *(Machine) Learning Amplitudes for Faster Event Generation*, [arXiv:1912.11055 \[hep-ph\]](#).
- [8] S. Badger and J. Bullock, *Using neural networks for efficient evaluation of high multiplicity scattering amplitudes*, *JHEP* **06** (2020) 114, [arXiv:2002.07516 \[hep-ph\]](#).
- [9] J. Aylett-Bullock, S. Badger, and R. Moodie, *Optimising simulations for diphoton production at hadron colliders using amplitude neural networks*, *JHEP* **08** (2021) 066, [arXiv:2106.09474 \[hep-ph\]](#).
- [10] D. Maître and H. Truong, *A factorisation-aware Matrix element emulator*, *JHEP* **11** (2021) 066, [arXiv:2107.06625 \[hep-ph\]](#).
- [11] R. Winterhalder, V. Magerya, E. Villa, S. P. Jones, M. Kerner, A. Butter, G. Heinrich, and T. Plehn, *Targeting multi-loop integrals with neural networks*, *SciPost Phys.* **12** (2022) 4, 129, [arXiv:2112.09145 \[hep-ph\]](#).
- [12] S. Badger, A. Butter, M. Luchmann, S. Pitz, and T. Plehn, *Loop Amplitudes from Precision Networks*, [arXiv:2206.14831 \[hep-ph\]](#).
- [13] D. Maître and H. Truong, *One-loop matrix element emulation with factorisation awareness*, [arXiv:2302.04005 \[hep-ph\]](#).
- [14] K. Hagiwara, J. Kanzaki, Q. Li, N. Okamura, and T. Stelzer, *Fast computation of MadGraph amplitudes on graphics processing unit (GPU)*, *Eur. Phys. J. C* **73** (2013) 2608, [arXiv:1305.0708 \[physics.comp-ph\]](#).
- [15] E. Bothmann, W. Giele, S. Hoeche, J. Isaacson, and M. Knobbe, *Many-gluon tree amplitudes on modern GPUs: A case study for novel event generators*, [arXiv:2106.06507 \[hep-ph\]](#).

- [16] S. Carrazza, J. Cruz-Martinez, M. Rossi, and M. Zaro, *MadFlow: automating Monte Carlo simulation on GPU for particle physics processes*, *Eur. Phys. J. C* **81** (2021) 7, 656, [arXiv:2106.10279](https://arxiv.org/abs/2106.10279) [physics.comp-ph].
- [17] A. Valassi, S. Roiser, O. Mattelaer, and S. Hageboeck, *Design and engineering of a simplified workflow execution for the MG5aMC event generator on GPUs and vector CPUs*, *EPJ Web Conf.* **251** (2021) 03045, [arXiv:2106.12631](https://arxiv.org/abs/2106.12631) [physics.comp-ph].
- [18] K. Danziger, T. Janßen, S. Schumann, and F. Siegert, *Accelerating Monte Carlo event generation – rejection sampling using neural network event-weight estimates*, [arXiv:2109.11964](https://arxiv.org/abs/2109.11964) [hep-ph].
- [19] I.-K. Chen, M. D. Klimek, and M. Perelstein, *Improved Neural Network Monte Carlo Simulation*, *SciPost Phys.* **10** (2021) 023, [arXiv:2009.07819](https://arxiv.org/abs/2009.07819) [hep-ph].
- [20] C. Gao, J. Isaacson, and C. Krause, *i-flow: High-dimensional Integration and Sampling with Normalizing Flows*, *Mach. Learn. Sci. Tech.* **1** (2020) 4, 045023, [arXiv:2001.05486](https://arxiv.org/abs/2001.05486) [physics.comp-ph].
- [21] E. Bothmann, T. Janßen, M. Knobbe, T. Schmale, and S. Schumann, *Exploring phase space with Neural Importance Sampling*, *SciPost Phys.* **8** (2020) 4, 069, [arXiv:2001.05478](https://arxiv.org/abs/2001.05478) [hep-ph].
- [22] C. Gao, S. Höche, J. Isaacson, C. Krause, and H. Schulz, *Event Generation with Normalizing Flows*, *Phys. Rev. D* **101** (2020) 7, 076002, [arXiv:2001.10028](https://arxiv.org/abs/2001.10028) [hep-ph].
- [23] T. Heimel, R. Winterhalder, A. Butter, J. Isaacson, C. Krause, F. Maltoni, O. Mattelaer, and T. Plehn, *MadNIS – Neural Multi-Channel Importance Sampling*, [arXiv:2212.06172](https://arxiv.org/abs/2212.06172) [hep-ph].
- [24] T. Heimel, N. Huetsch, F. Maltoni, O. Mattelaer, T. Plehn, and R. Winterhalder, *The MadNIS Reloaded*, [arXiv:2311.01548](https://arxiv.org/abs/2311.01548) [hep-ph].
- [25] T. Janßen, D. Maître, S. Schumann, F. Siegert, and H. Truong, *Unweighting multijet event generation using factorisation-aware neural networks*, *SciPost Phys.* **15** (2023) 107, [arXiv:2301.13562](https://arxiv.org/abs/2301.13562) [hep-ph].
- [26] E. Bothmann, T. Childers, W. Giele, F. Herren, S. Hoeche, J. Isaacson, M. Knobbe, and R. Wang, *Efficient phase-space generation for hadron collider event simulation*, *SciPost Phys.* **15** (2023) 4, 169, [arXiv:2302.10449](https://arxiv.org/abs/2302.10449) [hep-ph].
- [27] S. Otten, S. Caron, W. de Swart, M. van Beekveld, L. Hendriks, C. van Leeuwen, D. Podareanu, R. Ruiz de Austri, and R. Verheyen, *Event Generation and Statistical Sampling for Physics with Deep Generative Models and a Density Information Buffer*, *Nature Commun.* **12** (2021) 1, 2985, [arXiv:1901.00875](https://arxiv.org/abs/1901.00875) [hep-ph].
- [28] B. Hashemi, N. Amin, K. Datta, D. Olivito, and M. Pierini, *LHC analysis-specific datasets with Generative Adversarial Networks*, [arXiv:1901.05282](https://arxiv.org/abs/1901.05282) [hep-ex].
- [29] R. Di Sipio, M. Faucci Giannelli, S. Ketabchi Haghighat, and S. Palazzo, *DijetGAN: A Generative-Adversarial Network Approach for the Simulation of QCD Dijet Events at the LHC*, *JHEP* **08** (2020) 110, [arXiv:1903.02433](https://arxiv.org/abs/1903.02433) [hep-ex].
- [30] A. Butter, T. Plehn, and R. Winterhalder, *How to GAN LHC Events*, *SciPost Phys.* **7** (2019) 6, 075, [arXiv:1907.03764](https://arxiv.org/abs/1907.03764) [hep-ph].

- [31] Y. Alanazi, N. Sato, T. Liu, W. Melnitchouk, M. P. Kuchera, E. Pritchard, M. Robertson, R. Strauss, L. Velasco, and Y. Li, *Simulation of electron-proton scattering events by a Feature-Augmented and Transformed Generative Adversarial Network (FAT-GAN)*, [arXiv:2001.11103 \[hep-ph\]](#).
- [32] A. Butter, T. Heimel, S. Hummerich, T. Krebs, T. Plehn, A. Rousselot, and S. Vent, *Generative Networks for Precision Enthusiasts*, [arXiv:2110.13632 \[hep-ph\]](#).
- [33] A. Butter, N. Huetsch, S. Palacios Schweitzer, T. Plehn, P. Sorrenson, and J. Spinner, *Jet Diffusion versus JetGPT – Modern Networks for the LHC*, [arXiv:2305.10475 \[hep-ph\]](#).
- [34] 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** (2017) 1, 4, [arXiv:1701.05927 \[stat.ML\]](#).
- [35] A. Andreassen, I. Feige, C. Frye, and M. D. Schwartz, *JUNIPR: a Framework for Unsupervised Machine Learning in Particle Physics*, *Eur. Phys. J.* **C79** (2019) 2, 102, [arXiv:1804.09720 \[hep-ph\]](#).
- [36] E. Bothmann and L. Debbio, *Reweighting a parton shower using a neural network: the final-state case*, *JHEP* **01** (2019) 033, [arXiv:1808.07802 \[hep-ph\]](#).
- [37] K. Dohi, *Variational Autoencoders for Jet Simulation*, [arXiv:2009.04842 \[hep-ph\]](#).
- [38] E. Buhmann, G. Kasieczka, and J. Thaler, *EPiC-GAN: Equivariant Point Cloud Generation for Particle Jets*, *SciPost Phys.* **15** (2023) 130, [arXiv:2301.08128 \[hep-ph\]](#).
- [39] M. Leigh, D. Sengupta, G. Quétant, J. A. Raine, K. Zoch, and T. Golling, *PC-JeDi: Diffusion for Particle Cloud Generation in High Energy Physics*, [arXiv:2303.05376 \[hep-ph\]](#).
- [40] V. Mikuni, B. Nachman, and M. Pettee, *Fast point cloud generation with diffusion models in high energy physics*, *Phys. Rev. D* **108** (2023) 3, 036025, [arXiv:2304.01266 \[hep-ph\]](#).
- [41] E. Buhmann, C. Ewen, D. A. Faroughy, T. Golling, G. Kasieczka, M. Leigh, G. Quétant, J. A. Raine, D. Sengupta, and D. Shih, *EPiC-ly Fast Particle Cloud Generation with Flow-Matching and Diffusion*, [arXiv:2310.00049 \[hep-ph\]](#).
- [42] 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** (2018) 4, 042003, [arXiv:1705.02355 \[hep-ex\]](#).
- [43] L. de Oliveira, M. Paganini, and B. Nachman, *Controlling Physical Attributes in GAN-Accelerated Simulation of Electromagnetic Calorimeters*, *J. Phys. Conf. Ser.* **1085** (2018) 4, 042017, [arXiv:1711.08813 \[hep-ex\]](#).
- [44] 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** (2018) 1, 014021, [arXiv:1712.10321 \[hep-ex\]](#).
- [45] 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** (2018) 1, 4, [arXiv:1802.03325 \[astro-ph.IM\]](#).

- [46] M. Erdmann, J. Glombitza, and T. Quast, *Precise simulation of electromagnetic calorimeter showers using a Wasserstein Generative Adversarial Network*, *Comput. Softw. Big Sci.* **3** (2019) 1, 4, [arXiv:1807.01954 \[physics.ins-det\]](#).
- [47] D. Belayneh *et al.*, *Calorimetry with Deep Learning: Particle Simulation and Reconstruction for Collider Physics*, *Eur. Phys. J. C* **80** (2020) 7, 688, [arXiv:1912.06794 \[physics.ins-det\]](#).
- [48] E. Buhmann, S. Diefenbacher, E. Eren, F. Gaede, G. Kasieczka, A. Korol, and K. Krüger, *Getting High: High Fidelity Simulation of High Granularity Calorimeters with High Speed*, *Comput. Softw. Big Sci.* **5** (2021) 1, 13, [arXiv:2005.05334 \[physics.ins-det\]](#).
- [49] E. Buhmann, S. Diefenbacher, E. Eren, F. Gaede, G. Kasieczka, A. Korol, and K. Krüger, *Decoding Photons: Physics in the Latent Space of a BIB-AE Generative Network*, [arXiv:2102.12491 \[physics.ins-det\]](#).
- [50] C. Krause and D. Shih, *CaloFlow: Fast and Accurate Generation of Calorimeter Showers with Normalizing Flows*, [arXiv:2106.05285 \[physics.ins-det\]](#).
- [51] ATLAS Collaboration, *AtlFast3: the next generation of fast simulation in ATLAS*, *Comput. Softw. Big Sci.* **6** (2022) 7, [arXiv:2109.02551 \[hep-ex\]](#).
- [52] C. Krause and D. Shih, *CaloFlow II: Even Faster and Still Accurate Generation of Calorimeter Showers with Normalizing Flows*, [arXiv:2110.11377 \[physics.ins-det\]](#).
- [53] E. Buhmann, S. Diefenbacher, D. Hundhausen, G. Kasieczka, W. Korcari, E. Eren, F. Gaede, K. Krüger, P. McKeown, and L. Rustige, *Hadrons, better, faster, stronger*, *Mach. Learn. Sci. Tech.* **3** (2022) 2, 025014, [arXiv:2112.09709 \[physics.ins-det\]](#).
- [54] C. Chen, O. Cerri, T. Q. Nguyen, J. R. Vlimant, and M. Pierini, *Analysis-Specific Fast Simulation at the LHC with Deep Learning*, *Comput. Softw. Big Sci.* **5** (2021) 1, 15.
- [55] V. Mikuni and B. Nachman, *Score-based generative models for calorimeter shower simulation*, *Phys. Rev. D* **106** (2022) 9, 092009, [arXiv:2206.11898 \[hep-ph\]](#).
- [56] ATLAS Collaboration, *Deep generative models for fast photon shower simulation in ATLAS*, [arXiv:2210.06204 \[hep-ex\]](#).
- [57] C. Krause, I. Pang, and D. Shih, *CaloFlow for CaloChallenge Dataset 1*, [arXiv:2210.14245 \[physics.ins-det\]](#).
- [58] J. C. Cresswell, B. L. Ross, G. Loaiza-Ganem, H. Reyes-Gonzalez, M. Letizia, and A. L. Caterini, *CaloMan: Fast generation of calorimeter showers with density estimation on learned manifolds*, in *36th Conference on Neural Information Processing Systems*. 11, 2022. [arXiv:2211.15380 \[hep-ph\]](#).
- [59] S. Diefenbacher, E. Eren, F. Gaede, G. Kasieczka, C. Krause, I. Shekhzadeh, and D. Shih, *L2LFlows: Generating High-Fidelity 3D Calorimeter Images*, [arXiv:2302.11594 \[physics.ins-det\]](#).
- [60] H. Hashemi, N. Hartmann, S. Sharifzadeh, J. Kahn, and T. Kuhr, *Ultra-High-Resolution Detector Simulation with Intra-Event Aware GAN and Self-Supervised Relational Reasoning*, [arXiv:2303.08046 \[physics.ins-det\]](#).
- [61] A. Xu, S. Han, X. Ju, and H. Wang, *Generative Machine Learning for Detector Response Modeling with a Conditional Normalizing Flow*, [arXiv:2303.10148 \[hep-ex\]](#).

- [62] S. Diefenbacher, E. Eren, F. Gaede, G. Kasieczka, A. Korol, K. Krüger, P. McKeown, and L. Rustige, *New angles on fast calorimeter shower simulation*, *Mach. Learn. Sci. Tech.* **4** (2023) 3, 035044, [arXiv:2303.18150 \[physics.ins-det\]](#).
- [63] E. Buhmann, S. Diefenbacher, E. Eren, F. Gaede, G. Kasieczka, A. Korol, W. Korcari, K. Krüger, and P. McKeown, *CaloClouds: Fast Geometry-Independent Highly-Granular Calorimeter Simulation*, [arXiv:2305.04847 \[physics.ins-det\]](#).
- [64] M. R. Buckley, C. Krause, I. Pang, and D. Shih, *Inductive CaloFlow*, [arXiv:2305.11934 \[physics.ins-det\]](#).
- [65] S. Diefenbacher, V. Mikuni, and B. Nachman, *Refining Fast Calorimeter Simulations with a Schrödinger Bridge*, [arXiv:2308.12339 \[physics.ins-det\]](#).
- [66] T. Golling, S. Klein, R. Mastandrea, B. Nachman, and J. A. Raine, *Flows for Flows: Morphing one Dataset into another with Maximum Likelihood Estimation*, [arXiv:2309.06472 \[hep-ph\]](#).
- [67] S. Diefenbacher, G.-H. Liu, V. Mikuni, B. Nachman, and W. Nie, *Improving Generative Model-based Unfolding with Schrödinger Bridges*, [arXiv:2308.12351 \[hep-ph\]](#).
- [68] R. Winterhalder, M. Bellagente, and B. Nachman, *Latent Space Refinement for Deep Generative Models*, [arXiv:2106.00792 \[stat.ML\]](#).
- [69] B. Nachman and R. Winterhalder, *Elsa: enhanced latent spaces for improved collider simulations*, *Eur. Phys. J. C* **83** (2023) 9, 843, [arXiv:2305.07696 \[hep-ph\]](#).
- [70] M. Leigh, D. Sengupta, J. A. Raine, G. Quétant, and T. Golling, *PC-Droid: Faster diffusion and improved quality for particle cloud generation*, [arXiv:2307.06836 \[hep-ex\]](#).
- [71] R. Das, L. Favaro, T. Heimel, C. Krause, T. Plehn, and D. Shih, *How to Understand Limitations of Generative Networks*, [arXiv:2305.16774 \[hep-ph\]](#).
- [72] A. Butter, S. Diefenbacher, G. Kasieczka, B. Nachman, and T. Plehn, *GANplifying event samples*, *SciPost Phys.* **10** (2021) 6, 139, [arXiv:2008.06545 \[hep-ph\]](#).
- [73] S. Bieringer, A. Butter, S. Diefenbacher, E. Eren, F. Gaede, D. Hundhausen, G. Kasieczka, B. Nachman, T. Plehn, and M. Trabs, *Calomplification — the power of generative calorimeter models*, *JINST* **17** (2022) 09, P09028, [arXiv:2202.07352 \[hep-ph\]](#).
- [74] A. Butter, T. Plehn, and R. Winterhalder, *How to GAN Event Subtraction*, *SciPost Phys. Core* **3** (2020) 009, [arXiv:1912.08824 \[hep-ph\]](#).
- [75] B. Stienen and R. Verheyen, *Phase Space Sampling and Inference from Weighted Events with Autoregressive Flows*, *SciPost Phys.* **10** (2021) 038, [arXiv:2011.13445 \[hep-ph\]](#).
- [76] M. Backes, A. Butter, T. Plehn, and R. Winterhalder, *How to GAN Event Unweighting*, *SciPost Phys.* **10** (2021) 4, 089, [arXiv:2012.07873 \[hep-ph\]](#).
- [77] F. A. Di Bello, S. Ganguly, E. Gross, M. Kado, M. Pitt, L. Santi, and J. Shlomi, *Towards a Computer Vision Particle Flow*, *Eur. Phys. J. C* **81** (2021) 2, 107, [arXiv:2003.08863 \[physics.data-an\]](#).
- [78] P. Baldi, L. Blecher, A. Butter, J. Collado, J. N. Howard, F. Keilbach, T. Plehn, G. Kasieczka, and D. Whiteson, *How to GAN Higher Jet Resolution*, [arXiv:2012.11944 \[hep-ph\]](#).

- [79] K. Datta, D. Kar, and D. Roy, *Unfolding with Generative Adversarial Networks*, [arXiv:1806.00433 \[physics.data-an\]](#).
- [80] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, and R. Winterhalder, *How to GAN away Detector Effects*, [arXiv:1912.00477 \[hep-ph\]](#).
- [81] 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) 18, 182001, [arXiv:1911.09107 \[hep-ph\]](#).
- [82] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, A. Rousselot, R. Winterhalder, L. Ardizzone, and U. Köthe, *Invertible Networks or Partons to Detector and Back Again*, *SciPost Phys.* **9** (2020) 074, [arXiv:2006.06685 \[hep-ph\]](#).
- [83] M. Backes, A. Butter, M. Dunford, and B. Malaescu, *An unfolding method based on conditional Invertible Neural Networks (cINN) using iterative training*, [arXiv:2212.08674 \[hep-ph\]](#).
- [84] M. Leigh, J. A. Raine, K. Zoch, and T. Golling, *ν -Flows: Conditional Neutrino Regression*, *SciPost Phys.* **14** (2023) 159, [arXiv:2207.00664 \[hep-ph\]](#).
- [85] J. A. Raine, M. Leigh, K. Zoch, and T. Golling, *ν^2 -Flows: Fast and improved neutrino reconstruction in multi-neutrino final states with conditional normalizing flows*, [arXiv:2307.02405 \[hep-ph\]](#).
- [86] A. Shmakov, K. Greif, M. Fenton, A. Ghosh, P. Baldi, and D. Whiteson, *End-To-End Latent Variational Diffusion Models for Inverse Problems in High Energy Physics*, [arXiv:2305.10399 \[hep-ex\]](#).
- [87] J. Ackerschott, R. K. Barman, D. Gonçalves, T. Heimel, and T. Plehn, *Returning CP-Observables to The Frames They Belong*, [arXiv:2308.00027 \[hep-ph\]](#).
- [88] S. Bieringer, A. Butter, T. Heimel, S. Höche, U. Köthe, T. Plehn, and S. T. Radev, *Measuring QCD Splittings with Invertible Networks*, *SciPost Phys.* **10** (2021) 6, 126, [arXiv:2012.09873 \[hep-ph\]](#).
- [89] A. Butter, T. Heimel, T. Martini, S. Peitzsch, and T. Plehn, *Two Invertible Networks for the Matrix Element Method*, [arXiv:2210.00019 \[hep-ph\]](#).
- [90] T. Heimel, N. Huetsch, R. Winterhalder, T. Plehn, and A. Butter, *Precision-Machine Learning for the Matrix Element Method*, [arXiv:2310.07752 \[hep-ph\]](#).
- [91] B. Nachman and D. Shih, *Anomaly Detection with Density Estimation*, *Phys. Rev. D* **101** (2020) 075042, [arXiv:2001.04990 \[hep-ph\]](#).
- [92] A. Hallin, J. Isaacson, G. Kasieczka, C. Krause, B. Nachman, T. Quadfasel, M. Schlaffer, D. Shih, and M. Sommerhalder, *Classifying Anomalies THrough Outer Density Estimation (CATHODE)*, [arXiv:2109.00546 \[hep-ph\]](#).
- [93] J. A. Raine, S. Klein, D. Sengupta, and T. Golling, *CURTAINS for your sliding window: Constructing unobserved regions by transforming adjacent intervals*, *Front. Big Data* **6** (2023) 899345, [arXiv:2203.09470 \[hep-ph\]](#).
- [94] A. Hallin, G. Kasieczka, T. Quadfasel, D. Shih, and M. Sommerhalder, *Resonant anomaly detection without background sculpting*, *Phys. Rev. D* **107** (2023) 11, 114012, [arXiv:2210.14924 \[hep-ph\]](#).

- [95] T. Golling, S. Klein, R. Mastandrea, and B. Nachman, *Flow-enhanced transportation for anomaly detection*, *Phys. Rev. D* **107** (2023) 9, 096025, [arXiv:2212.11285 \[hep-ph\]](#).
- [96] D. Sengupta, S. Klein, J. A. Raine, and T. Golling, *CURTAINS Flows For Flows: Constructing Unobserved Regions with Maximum Likelihood Estimation*, [arXiv:2305.04646 \[hep-ph\]](#).
- [97] M. Czakon, P. Fiedler, and A. Mitov, *Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $\mathcal{O}(\alpha_s^4)$* , *Phys. Rev. Lett.* **110** (2013) 252004, [arXiv:1303.6254 \[hep-ph\]](#).
- [98] M. Czakon, D. Heymes, and A. Mitov, *High-precision differential predictions for top-quark pairs at the LHC*, *Phys. Rev. Lett.* **116** (2016) 8, 082003, [arXiv:1511.00549 \[hep-ph\]](#).
- [99] S. Catani, S. Devoto, M. Grazzini, S. Kallweit, J. Mazzitelli, and H. Sargsyan, *Top-quark pair hadroproduction at next-to-next-to-leading order in QCD*, *Phys. Rev.* **D99** (2019) 5, 051501, [arXiv:1901.04005 \[hep-ph\]](#).
- [100] S. Catani, S. Devoto, M. Grazzini, S. Kallweit, and J. Mazzitelli, *Top-quark pair production at the LHC: Fully differential QCD predictions at NNLO*, *JHEP* **07** (2019) 100, [arXiv:1906.06535 \[hep-ph\]](#).
- [101] W. Bernreuther, M. Fucker, and Z.-G. Si, *Weak interaction corrections to hadronic top quark pair production: Contributions from quark-gluon and b anti-b induced reactions*, *Phys. Rev.* **D78** (2008) 017503, [arXiv:0804.1237 \[hep-ph\]](#).
- [102] J. H. Kühn, A. Scharf, and P. Uwer, *Electroweak effects in top-quark pair production at hadron colliders*, *Eur. Phys. J. C* **51** (2007) 37, [arXiv:hep-ph/0610335 \[hep-ph\]](#).
- [103] W. Hollik and D. Pagani, *The electroweak contribution to the top quark forward-backward asymmetry at the Tevatron*, *Phys. Rev.* **D84** (2011) 093003, [arXiv:1107.2606 \[hep-ph\]](#).
- [104] C. Gütschow, J. M. Lindert, and M. Schönherr, *Multi-jet merged top-pair production including electroweak corrections*, *Eur. Phys. J. C* **78** (2018) 4, 317, [arXiv:1803.00950 \[hep-ph\]](#).
- [105] R. Frederix, I. Tsinikos, and T. Vitos, *Probing the spin correlations of $t\bar{t}$ production at NLO QCD+EW*, *Eur. Phys. J. C* **81** (2021) 9, 817, [arXiv:2105.11478 \[hep-ph\]](#).
- [106] M. Czakon, D. Heymes, A. Mitov, D. Pagani, I. Tsinikos, and M. Zaro, *Top-pair production at the LHC through NNLO QCD and NLO EW*, *JHEP* **10** (2017) 186, [arXiv:1705.04105 \[hep-ph\]](#).
- [107] J. Gao, C. S. Li, and H. X. Zhu, *Top Quark Decay at Next-to-Next-to Leading Order in QCD*, *Phys. Rev. Lett.* **110** (2013) 4, 042001, [arXiv:1210.2808 \[hep-ph\]](#).
- [108] M. Brucherseifer, F. Caola, and K. Melnikov, *$\mathcal{O}(\alpha_s^2)$ corrections to fully-differential top quark decays*, *JHEP* **04** (2013) 059, [arXiv:1301.7133 \[hep-ph\]](#).
- [109] J. Gao and A. S. Papanastasiou, *Top-quark pair-production and decay at high precision*, *Phys. Rev.* **D96** (2017) 5, 051501, [arXiv:1705.08903 \[hep-ph\]](#).
- [110] A. Behring, M. Czakon, A. Mitov, A. S. Papanastasiou, and R. Poncelet, *Higher order corrections to spin correlations in top quark pair production at the LHC*, *Phys. Rev. Lett.* **123** (2019) 8, 082001, [arXiv:1901.05407 \[hep-ph\]](#).

- [111] M. Czakon, A. Mitov, and R. Poncelet, *NNLO QCD corrections to leptonic observables in top-quark pair production and decay*, *JHEP* **05** (2021) 212, [arXiv:2008.11133 \[hep-ph\]](#).
- [112] G. Bevilacqua, M. Czakon, A. van Hameren, C. G. Papadopoulos, and M. Worek, *Complete off-shell effects in top quark pair hadroproduction with leptonic decay at next-to-leading order*, *JHEP* **02** (2011) 083, [arXiv:1012.4230 \[hep-ph\]](#).
- [113] A. Denner, S. Dittmaier, S. Kallweit, and S. Pozzorini, *NLO QCD corrections to $WWbb$ production at hadron colliders*, *Phys. Rev. Lett.* **106** (2011) 052001, [arXiv:1012.3975 \[hep-ph\]](#).
- [114] A. Denner, S. Dittmaier, S. Kallweit, and S. Pozzorini, *NLO QCD corrections to off-shell top-antitop production with leptonic decays at hadron colliders*, *JHEP* **10** (2012) 110, [arXiv:1207.5018 \[hep-ph\]](#).
- [115] G. Heinrich, A. Maier, R. Nisius, J. Schlenk, and J. Winter, *NLO QCD corrections to $W^+W^-b\bar{b}$ production with leptonic decays in the light of top quark mass and asymmetry measurements*, *JHEP* **06** (2014) 158, [arXiv:1312.6659 \[hep-ph\]](#).
- [116] R. Frederix, *Top Quark Induced Backgrounds to Higgs Production in the $WW^{(*)} \rightarrow ll\nu\nu$ Decay Channel at Next-to-Leading-Order in QCD*, *Phys. Rev. Lett.* **112** (2014) 8, 082002, [arXiv:1311.4893 \[hep-ph\]](#).
- [117] F. Cascioli, S. Kallweit, P. Maierhöfer, and S. Pozzorini, *A unified NLO description of top-pair and associated Wt production*, *Eur. Phys. J.* **C74** (2014) 3, 2783, [arXiv:1312.0546 \[hep-ph\]](#).
- [118] G. Bevilacqua, H. B. Hartanto, M. Kraus, and M. Worek, *Top Quark Pair Production in Association with a Jet with Next-to-Leading-Order QCD Off-Shell Effects at the Large Hadron Collider*, *Phys. Rev. Lett.* **116** (2016) 5, 052003, [arXiv:1509.09242 \[hep-ph\]](#).
- [119] NNPDF, R. D. Ball *et al.*, *Parton distributions from high-precision collider data*, *Eur. Phys. J. C* **77** (2017) 10, 663, [arXiv:1706.00428 \[hep-ph\]](#).
- [120] S. Frixione, P. Nason, and G. Ridolfi, *A Positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction*, *JHEP* **09** (2007) 126, [arXiv:0707.3088 \[hep-ph\]](#).
- [121] T. Ježo, J. M. Lindert, P. Nason, C. Oleari, and S. Pozzorini, *An NLO+PS generator for $t\bar{t}$ and Wt production and decay including non-resonant and interference effects*, *Eur. Phys. J. C* **76** (2016) 12, 691, [arXiv:1607.04538 \[hep-ph\]](#).
- [122] T. Ježo, J. M. Lindert, and S. Pozzorini, *Resonance-aware NLOPS matching for off-shell $t\bar{t} + tW$ production with semileptonic decays*, *JHEP* **10** (2023) 008, [arXiv:2307.15653 \[hep-ph\]](#).
- [123] S. Frixione, E. Laenen, P. Motylinski, and B. R. Webber, *Angular correlations of lepton pairs from vector boson and top quark decays in Monte Carlo simulations*, *JHEP* **04** (2007) 081, [arXiv:hep-ph/0702198](#).
- [124] P. Nason, *A New method for combining NLO QCD with shower Monte Carlo algorithms*, *JHEP* **11** (2004) 040, [arXiv:hep-ph/0409146](#).
- [125] S. Frixione, P. Nason, and C. Oleari, *Matching NLO QCD computations with Parton Shower simulations: the POWHEG method*, *JHEP* **11** (2007) 070, [arXiv:0709.2092 \[hep-ph\]](#).

- [126] S. Alioli, P. Nason, C. Oleari, and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, *JHEP* **06** (2010) 043, [arXiv:1002.2581 \[hep-ph\]](#).
- [127] T. Ježo and P. Nason, *On the Treatment of Resonances in Next-to-Leading Order Calculations Matched to a Parton Shower*, *JHEP* **12** (2015) 065, [arXiv:1509.09071 \[hep-ph\]](#).
- [128] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, *Flow matching for generative modeling*, [arXiv:2210.02747 \[cs.LG\]](#).
- [129] M. Bellagente, M. Haußmann, M. Luchmann, and T. Plehn, *Understanding Event-Generation Networks via Uncertainties*, *SciPost Phys.* **13** (2022) 003, [arXiv:2104.04543 \[hep-ph\]](#).
- [130] A. Tong, N. Malkin, G. Huguet, Y. Zhang, J. Rector-Brooks, K. Fatras, G. Wolf, and Y. Bengio, *Improving and generalizing flow-based generative models with minibatch optimal transport*, [arXiv:2302.00482 \[cs.LG\]](#).