

Particle Generative Adversarial Networks for full-event simulation at the LHC and their application to pileup description

Jesús Arjona Martínez^{1,2,3}, Thong Q Nguyen², Maurizio Pierini³,
Maria Spiropulu² and Jean-Roch Vlimant²

¹ University of Cambridge, Trinity Ln, Cambridge CB2 1TN, UK

² California Institute of Technology, 1200 E. California Blvd, Pasadena, CA 91125

³ CERN, CH-1211 Geneva, Switzerland

E-mail: ja618@cam.ac.uk, {thong, smaria, jvlimant}@caltech.edu,
maurizio.pierini@cern.ch

Abstract. We investigate how a Generative Adversarial Network could be used to generate a list of particle four-momenta from LHC proton collisions, allowing one to define a generative model that could abstract from the irregularities of typical detector geometries. As an example of application, we show how such an architecture could be used as a generator of LHC parasitic collisions (pileup). We present two approaches to generate the events: unconditional generator and generator conditioned on missing transverse energy. We assess generation performances in a realistic LHC data-analysis environment, with a pileup mitigation algorithm applied.

1. Introduction

The simulation of subatomic particle collisions, their subsequent detector interaction, and their reconstruction is a computationally demanding task for the computing infrastructure of the experiments operating at the CERN Large Hadron Collider (LHC). The high accuracy of state-of-the-art Monte Carlo (MC) simulation software, typically based on the GEANT4 [1], has a high cost: MC simulation amounts to about one half of the experiments' computing budget and to a large fraction of the available storage resources [2], the other half being largely used to process simulated and real data (event reconstruction).

Following their invention in 2014, GANs gained traction as generative models, often superior to Variational Autoencoders [3] and with very impressive results in image production [4, 5]. Due to their high inference speed, GANs can be used as fast-simulation libraries. This approach has been successfully investigated with proof-of-principle studies related to particle showers in multilayer calorimeters [6, 7, 8] and particle jets [9], as well as in similar application to different HEP domains [10, 11, 12, 13, 14, 15]. All these studies formalized the simulation task in terms of either image generation or analysis-specific high-level features.

In this work, we present pGAN, a full-event particle-based generative model that can be used to emulate pileup simulation at the LHC.

2. Pileup simulation

The majority of LHC proton-proton collisions result in so-called *minimum-bias* (MB) events, i.e., in low-energy (*soft*) interactions between proton constituents. These events are characterized



by low- p_T particles, as opposed to the head-on collision processes typically studied at the LHC (so-called *hard* or high- p_T interactions). Any hard interaction happens simultaneously to many parasitic MB events, generically referred to as *pileup*. Pileup simulation is a fundamental aspect of a realistic LHC simulation software. The current implementation of pileup simulation consists in overlapping a set of MB events to the main high- p_T collision. Events could be generated on demand or be sampled from a pre-generated library. The former is computationally expensive, while the latter is inflexible and suffers from I/O issue.

GAN emerges as a possible solution to speed up the on-demand generation of MB events and remove the need for a pre-generated library. To our knowledge, the only application of machine learning to pileup simulation is the work presented in Ref. [16], where pileup images are generated using a Deep Convolutional GAN model (DCGAN) [17]. However, an image-based event representation can't be used as input for downstream reconstruction algorithms. In addition, our proposed pGAN, which uses particle-based event representation, can abstract from the details of the detector geometry (e.g., its irregularities) and better scales with the foreseen increase of detector complexity.

3. Dataset

Synthetic MB events from proton-proton collisions are produced using the PYTHIA8 [18] event generator. The center-of-mass energy of the collision is set to 13 TeV, corresponding to the LHC Run II (2015-2018). All soft QCD processes are activated, allowing for both initial- and final-state radiation as well as multiple parton interactions. The produced events are passed to DELPHES [19], to emulate the detector response. We take as a reference the upgraded design of the CMS detector, foreseen for the High-Luminosity LHC phase. The DELPHES particle-flow reconstruction algorithm is applied, returning the list of charged particles, photons, and neutral hadrons in the event. Minimum bias events are then combined to simulate a per-event pileup contribution, with mean number of MB events $\bar{n}_{\text{PU}} = 20$ following a Poisson distribution. We randomly sample n_{PU} events from the MB dataset and mix them by merging the list of charged particles, neutral hadrons, and photons across the events.

Due to the complexity of training on long sequences, we restrict a maximum of 150 particles per event: the 50 charged particles, 50 photons and 50 neutral hadrons with the highest p_T value. This choice is mainly due to technical limitations that a more powerful training setup might help to overcome. On the other hand, cutting the sequence after p_T ordering is a well motivated simplification of the problem: a typical physics analysis would be based on a pileup mitigation algorithm, which usually removes the majority of the soft pileup contamination.

4. Network architectures

A GAN consists of two neural networks, a generator \mathcal{G} and a discriminator \mathcal{D} . Given a set of samples x , the aim of a GAN training is to learn the function $p_{\text{data}}(x) \in A$ under which the x samples are distributed. We define an n -dimensional prior of input noise $z \sim p_z(z) \in \mathbb{R}^n$. The generator \mathcal{G} is a differentiable function with trainable parameters $\theta_{\mathcal{G}}$, mapping \mathbb{R}^n to A . The discriminator \mathcal{D} , with trainable parameters $\theta_{\mathcal{D}}$, is a map between A and $[0, 1]$, returning the probability that a given sample belongs to the set of real samples rather than originating from \mathcal{G} . \mathcal{D} is trained to assign the correct probability to both real and generated ("fake") data; \mathcal{G} is trained to produce samples such that they maximize the probability of them being real $\mathcal{D}(\mathcal{G}(z))$.

We develop 2 models for pGAN: an unconditional pGAN where the generator starts from purely random noise z , and a conditional pGAN [20] where a label acts as an extension of z to allow for generation of events based on an initial condition. In our use case, the missing transverse energy E_T^{miss} is chosen as the label due to its importance in most physics analyses.

The loss function to train \mathcal{G} and \mathcal{D} in an unconditional pGAN is described as:

$$\mathcal{L}^{\text{uncond}} = \mathcal{L}_{\text{adv}} = \mathbb{E}_{z \sim p_z(z)}[\log(\mathbb{P}(\mathcal{D}(\mathcal{G}(z)) = 0))] + \mathbb{E}_{I \sim A}[\log(\mathbb{P}(\mathcal{D}(I) = 1))] , \quad (1)$$

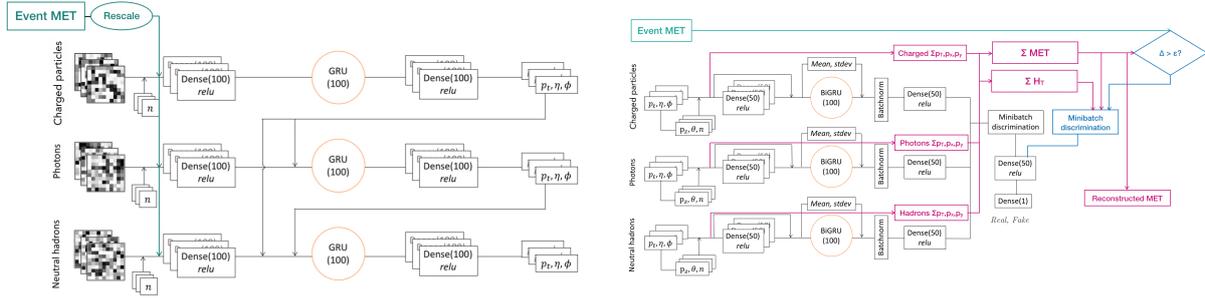


Figure 1: The architecture of conditional pGAN: generator $\mathcal{G}^{\text{cond}}$ (left) and discriminator $\mathcal{D}^{\text{cond}}$ (right). Arrows signify concatenation. Details are described in the text.

while for conditional pGAN, the loss function becomes:

$$\mathcal{L}^{\text{cond}} = \mathcal{L}_{\text{adv}} + \alpha \mathcal{L}_{\text{aux}}, \text{ where} \quad (2)$$

$$\mathcal{L}_{\text{aux}} = \mathbb{E}_{(z, E_T^{\text{miss}}) \sim (p(z) \times f(A))} \left[\frac{|E_T^{\text{miss}} - \hat{E}_T^{\text{miss}}(G(z|E_T^{\text{miss}}))|}{E_T^{\text{miss}}} \right] + \mathbb{E}_{x \sim A} \left[\frac{|E_T^{\text{miss}}(x) - \hat{E}_T^{\text{miss}}(x)|}{E_T^{\text{miss}}(x)} \right] \quad (3)$$

\hat{E}_T^{miss} is the missing transverse energy value computed from list of particles input to the discriminator, $f(A)$ is the empirical E_T^{miss} distribution of the dataset. In the conditional GAN setting, the dataset is transformed such that $\phi(\vec{E}_T^{\text{miss}}) = 0$ and the ϕ value of each particle is computed as azimuth angle between its momenta and \vec{E}_T^{miss} . This way a single scalar value E_T^{miss} can describe the full vector \vec{E}_T^{miss} since the direction is chosen as the coordinate axis.

4.1. Unconditional pGAN

4.1.1. Generator The \mathcal{G} model takes as input a set of time steps, corresponding to the number of particles to generate, sampled from a uniform distribution from 0 to 50 for each particle class. Along with the noise, the current particle number n is given to the network. We represent each particle as a (p_T, ϕ, η) tuple. We enforce the ϕ -symmetry constraint through a mod- 2π activation function for ϕ . The input in each branch from the noise input is processed via a block of Dense layers, a GRU layer, and another block of Dense layers as the final outputs. Each block of Dense consists 3 concatenated Dense layers that represent the particle's p_T , ϕ , and η respectively.

Due to the peculiar shapes of the η distributions, we pre-train a small dense network that fed with a Gaussian-distributed input mimics the η distribution as output. This network is used to process the η output of the generator. This can be viewed as an activation function parametrizing the η distribution in an unbiased way and decouples learning the distribution from the already complex particle generation task.

4.1.2. Discriminator The discriminator \mathcal{D} is of a binary classifier whose first step consists in a Physics layer, which takes each particle's defining features (p_T, ϕ, η, n) and concatenates to them other redundant features (θ, p_z, p) . This Physics layer is introduced to maximize the information given to the discriminator without adding redundant information to the particle representation returned by the generator. This prevents the generator for having to learn dependencies between different features of the particle representation, while allowing the discriminator to exploit them. The discriminator makes use of a Bidirectional recurrent layer, which reduces the dependence on the long term memory of the GRU cell. A layer that calculates the mean and standard deviation of the features in a dense layer is included, in a similar spirit to feature matching [21]. Mini-batch discrimination is used to prevent mode collapse.

4.2. Conditional pGAN

4.2.1. Generator The generator $\mathcal{G}^{\text{cond}}$ for the conditional pGAN, as shown in Fig. 1, is built on top of the unconditional generator \mathcal{G} . An initial value of the E_T^{miss} , sampled from real data, is injected as the input of the generator along with the noise after being rescaled to same range of the noise. This initial E_T^{miss} is also used as the input to the discriminator. The rest of the generator architecture is similar to the unconditional pGAN's \mathcal{G} .

4.2.2. Discriminator The discriminator $\mathcal{D}^{\text{cond}}$, as shown in Fig. 1, takes as inputs the 3 lists of particles along with the event E_T^{miss} , which is either the initial condition E_T^{miss} (in the case of the generated lists of particles) or the actual E_T^{miss} (in the case of sampling the lists of particles from the training data). In addition to usual computation flow in the unconditional \mathcal{D} described in Sec. 4.1, $\mathcal{D}^{\text{cond}}$ also computes a few high-level features out of the input particle lists, in particular the reconstructed E_T^{miss} and H_T (the scalar sum of all input particles' momenta), which are used as inputs to the final prediction. A binary flag Δ , which returns 1 if the absolute difference between reconstructed E_T^{miss} and the initial E_T^{miss} is greater than ϵ and returns 0 otherwise, is also concatenated to the input of the final prediction. This can be viewed as an attempt to let the discriminator to learn some global kinematic features of the inputs. Additionally, the reconstructed E_T^{miss} are compared with the initial input E_T^{miss} to construct the additional term \mathcal{L}_{aux} in the loss function, as described in Eq. (3).

4.3. Implementation details

Adam [22] is used for optimization with a batch size of 32, a learning rate of 1×10^{-4} for the generator and of 2×10^{-4} for the discriminator. For the conditional pGAN, we choose $\epsilon = 10$ GeV and $\alpha = 0.05$. We train on cropped sequences of variable length, forcing the discriminator to learn to distinguish even very short arrays, and make use of dropout in the discriminator. Batch normalization [23] is included in the discriminator but not in the generator. Hyperparameter optimization was performed on the learning rate of both networks. We find Least-Square GAN [24] offers the most stable training, outperforming both Wasserstein GAN [25] and the original GAN implementation.

5. Experimental evaluation

GAN performance evaluation is a challenging task. Due to their unsupervised nature, domain and task-specific metrics are often required. Our evaluation technique is three-fold: (i) we plot distributions of relevant physical features and quantify the matching between the ground-truth (GT), PYTHIA8 + DELPHES, and our network; (ii) we make use of high-level global event-shape variables such as the transverse thrust and the E_T^{miss} ; (iii) we evaluate the effect of using our proposed generation technique in a real analysis environment. To this purpose, we apply a state-of-the-art pileup removal algorithm (*SoftKiller* [26]) and cluster the remaining particles in the event, characterizing the agreement between real simulation and pGAN on jet kinematic properties (e.g., the jet p_T). These three different performance assessments are discussed in the rest of this section.

5.1. Histogram matching

Fig. 2 shows how the pGAN generators succeed in learning the main aspects of the particles' p_T , η and ϕ distributions. The comparison is limited to the first 50 high- p_T particles of each class, in order for the representation of the generated event to be consistent with the GAN generator output. We observe a remarkable agreement in p_T and ϕ : the long tail of the transverse momentum distribution is well described across four orders of magnitude and the ϕ rotation invariance of the physical process is reproduced. On the other hand, while the qualitative features of the pseudorapidity distribution are learned, the agreement is not completely satisfactory. A better match should be the goal of future work. Fig. 3 shows comparisons for some of the event-related high level features used in model evaluation: transverse thrust and E_T^{miss} . We observe

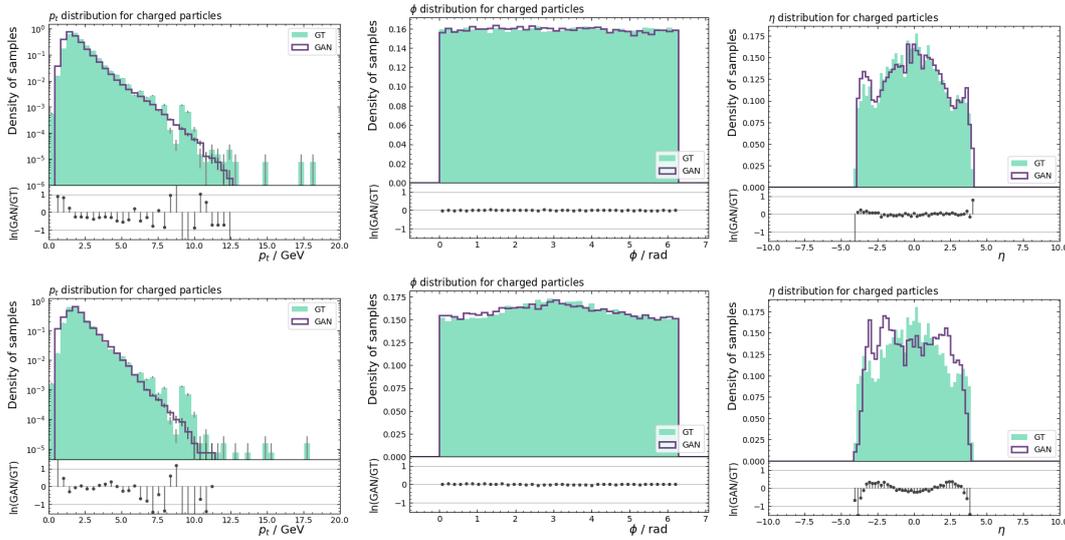


Figure 2: Comparison of the transverse momentum p_T (left), azimuth angle ϕ (center) and pseudorapidity η (right) for charged particles between the test data and the events generated by unconditional pGAN (top) and conditional pGAN (bottom). For the conditional pGAN, ϕ is transformed to be the azimuth angle between the particles' momenta and \vec{E}_T^{miss} .

some discrepancy being associated to overall scale shifts, related to the different truncation criteria applied to the two sets of events.

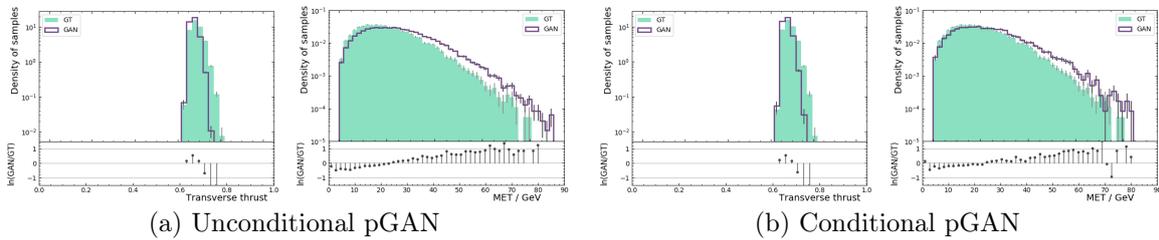


Figure 3: Comparison of the transverse thrust and E_T^{miss} distributions between the test data and the generated events by pGANs.

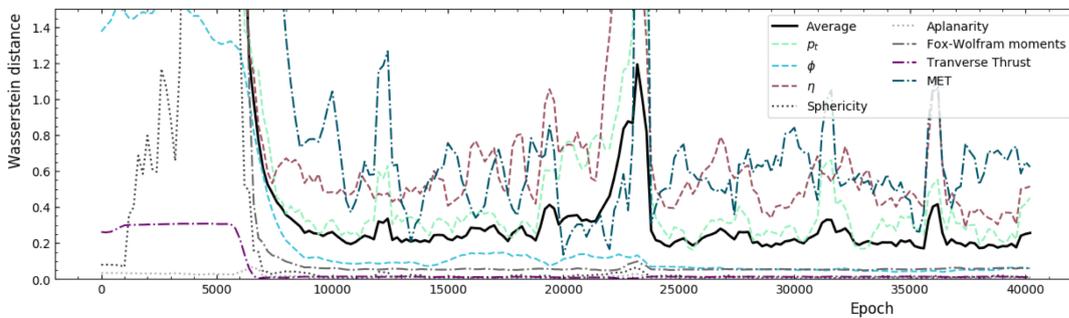


Figure 4: Evolution of our performance metric (solid black) as a function of training. EM distances for some of the individual quantities are superposed.

5.2. Wasserstein metric

We use the previously described histograms to quantify the difference between the target and generated distributions through the Wasserstein or Earth Mover's (EM) distance. We define our performance metric as the weighted average of the EM distances over the feature distributions: p_T , η , ϕ for all three particle types, sphericity, transverse thrust, aplanarity, E_T^{miss} , and the first and second transverse Fox-Wolfram moments [27]. The use of this metric tackles the problem of lack of interpretability of the loss function: we observe that the metric decreases steadily as the training progresses, as shown in Fig. 4, providing a way of monitoring progress, performing early-stopping and tracking training failure. Based on this metric, we perform model comparison, hyperparameter tuning, and the final best-model choice.

5.3. Pileup subtraction

Typical LHC analyses are performed after applying a pileup removal algorithm, which aims to subtract soft radiation from QCD. It is then important to demonstrate that the pGAN is good in modeling the residual pileup contribution, after such a subtraction algorithm is applied. Since this residual pileup contribution is the only relevant effect for physics analyses, it is acceptable for a pileup emulation software to have a non-accurate pileup simulation as long as the pileup effect is well model after applying a pileup mitigation technique. For this purpose, we consider a sample of $Z \rightarrow \nu\bar{\nu}$ events, generated using *Pythia8*. Events are processed with *Delphes*, using the same setup adopted to generate the pileup reference sample, both with and without activating the pileup emulation at $\bar{n}_{\text{PU}} = 20$. The generated no-pileup events are mixed with the pileup emulation returned by the generator. The *SoftKiller* [26] algorithm with a grid size of $a \approx 0.5$ is then applied to both these events and those with a full pileup simulation.

Table 1: Mean leading-jet p_T for events with no pileup and pileup generated by *Pythia8* and by the network (pGAN), before and after running *SoftKiller*.

	$\langle p_T \rangle / \text{GeV}$		$\langle p_T \rangle / \text{GeV}$	
No PU	136.8			
Pileup - GT	146.6	Pileup - GT - subtracted	135.0	
Pileup - pGAN	141.1	Pileup - pGAN - subtracted	135.7	

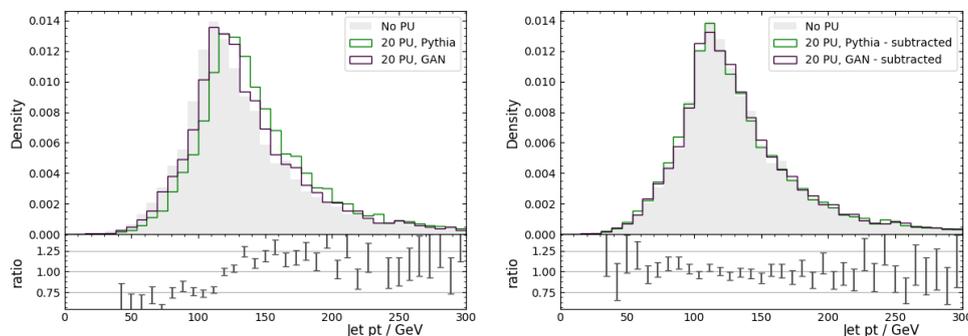


Figure 5: Comparison between leading jet p_T distributions for events with no pileup (solid black) and pileup generated by *Pythia8* (green) and by the network (magenta). Distributions are shown both before (left) and after (right) running the *SoftKiller* pileup mitigation algorithm. The bottom plot shows the *Pythia/GAN* ratio.

Fig. 5 shows the p_T distribution of the highest- p_T jet in the event in various configurations. The main effect of pileup contamination is a shift in the p_T distribution towards larger values. The shift is underestimated when the pileup is described through the pGAN generator, which is due to the fact that pGAN only returns the first 150 particles per event, instead of the usual ~ 900 . After processing the event with the *SoftKiller* algorithm, the leading jet p_T distribution for Pythia8 and our GAN match within 0.7 GeV. The agreement could be further improved by increasing the number of neutral hadrons and charged particles returned by the pGAN.

6. Conclusions

We presented GAN models based on a recurrent unit, capable of generating lists of particles with variable lengths and also with an initial condition on E_T^{miss} . Such a model could be used for particle-based simulation software, such as those of experiments using particle-flow reconstruction algorithms. This model could be used to replace ordinary rule-based algorithms in specific aspects of jet generation. In this paper, we show its application to pileup emulation in LHC collisions. While technical limitations forced us to reduce the length of the returned particle chain, the network is capable of emulating the effect of pileup on a realistic data analysis, after applying a pileup mitigation algorithm.

Acknowledgments

We are grateful to Caltech and the Kavli Foundation for their support of undergraduate student research in cross-cutting areas of machine learning and domain sciences. This work was conducted at "*iBanks*", the AI GPU cluster at Caltech. We acknowledge NVIDIA, SuperMicro and the Kavli Foundation for their support of "*iBanks*". This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement n^o 772369). This project is partially supported by the United States Department of Energy, Office of High Energy Physics Research under Caltech Contract No. DE-SC0011925.

References

- [1] Agostinelli S *et al.* (GEANT4) 2003 *Nucl. Instrum. Meth.* **A506** 250–303
- [2] Allton C *et al.* 2017 Tech. Rep. CERN-RRB-2017-125 CERN Geneva
- [3] Kingma D P and Welling M 2013 *CoRR (Preprint 1312.6114)*
- [4] Karras T *et al.* 2018 *ICLR (Preprint 1710.10196)*
- [5] Zhu J *et al.* 2017 *CoRR (Preprint 1703.10593)*
- [6] Paganini M *et al.* 2018 *Phys. Rev.* **D97** 014021 (*Preprint 1712.10321*)
- [7] Carminati F *et al.* 2017 *Deep Learning for Physical Sciences NIPS workshop*
- [8] Erdmann M, Glombitza J and Quast T 2018 (*Preprint 1807.01954*)
- [9] Musella P and Pandolfi F 2018 (*Preprint 1805.00850*)
- [10] Erdmann M *et al.* 2018 *Comput. Softw. Big Sci.* **2** 4 (*Preprint 1802.03325*)
- [11] de Oliveira L, Paganini M and Nachman B 2017 (*Preprint 1701.05927*)
- [12] Chekalina V *et al.* 2018 (*Preprint 1812.01319*)
- [13] Lin J, Bhimji W and Nachman B 2019 (*Preprint 1903.02556*)
- [14] Sipio R D *et al.* 2019 (*Preprint 1903.02433*)
- [15] Hashemi B *et al.* 2019 (*Preprint 1901.05282*)
- [16] Farrell S *et al.* 2018 *CHEP 2018 Conference, Sofia, Bulgaria*
- [17] Radford A, Metz L and Chintala S 2015 *CoRR (Preprint 1511.06434)*
- [18] Sjöstrand T *et al.* 2015 *Comput. Phys. Commun.* **191** 159–177 (*Preprint 1410.3012*)
- [19] de Favereau J *et al.* 2014 *JHEP* **02** 057 (*Preprint 1307.6346*)
- [20] Mirza M and Osinder S 2014 (*Preprint 1411.1784*)
- [21] Salimans T *et al.* 2016 *NeurIPS 29* 2234–2242 (*Preprint 1606.03498*)
- [22] Kingma D P and Ba J 2014 *CoRR (Preprint 1412.6980)*
- [23] Ioffe S and Szegedy C 2015 (*Preprint 1502.03167*)
- [24] Mao X *et al.* 2016 *CoRR (Preprint 1611.04076)*
- [25] Arjovsky M, Chintala S and Bottou L 2017 *ICML 34* vol 70 (PMLR) pp 214–223 (*Preprint 1701.07875*)
- [26] Cacciari M, Salam G P and Soyez G 2015 *Eur. Phys. J.* **C75** 59 (*Preprint 1407.0408*)
- [27] Fox G C and Wolfram S 1979 *Nuclear Physics B* **149** 413–496