

Lawrence Berkeley National Laboratory

LBL Publications

Title

New directions for surrogate models and differentiable programming for High Energy Physics detector simulation

Permalink

<https://escholarship.org/uc/item/6tx6q2n7>

Authors

Adelmann, Andreas
Hopkins, Walter
Kourlitis, Evangelos
et al.

Publication Date

2022-03-15

New directions for surrogate models and differentiable programming for High Energy Physics detector simulation

ANDREAS ADELMANN

Paul Scherrer Institute, 5232 Villigen PSI, Switzerland

WALTER HOPKINS, EVANGELOS KOURLITIS

Argonne National Laboratory, Lemont, IL 60439, USA

MICHAEL KAGAN

SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

GREGOR KASIECZKA

Institut für Experimentalphysik, Universität Hamburg, Germany

CLAUDIUS KRAUSE, DAVID SHIH

NHETC, Department of Physics & Astronomy, Rutgers University, Piscataway, NJ 08854, USA

VINICIUS MIKUNI, BENJAMIN NACHMAN

Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

KEVIN PEDRO

Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

DANIEL WINKLEHNER

Massachusetts Institute of Technology, Cambridge, MA 02139, USA

ABSTRACT

The computational cost for high energy physics detector simulation in future experimental facilities is going to exceed the current available resources. To overcome this challenge, new ideas on surrogate models using machine learning methods are being explored to replace computationally expensive components. Additionally, differentiable programming has been proposed as a complementary approach, providing controllable and scalable simulation routines. In this document, new and ongoing efforts for surrogate models and differential programming applied to detector simulation are discussed in the context of the 2021 Particle Physics Community Planning Exercise ('Snowmass').

1 Introduction

Experiments in high energy physics (HEP) rely heavily on simulation for a wide array of tasks, including data selection, statistical inference, and design optimization for new experiments. On the other hand, the computational demands for simulation of current and next generation HEP experiments have inspired investigation of surrogates, or approximations of the detector simulation, using deep generative models to decrease simulation time while maintaining fidelity. Usually, the most computationally intensive step of the simulation is the modeling of the detector response. Interactions between particles and the detector material are simulated in large experimental collaborations such as ATLAS [1] and CMS [2] using the GEANT4 [3–5] software package. While full simulation ensures high fidelity samples, the computational cost becomes prohibitive as many billions of simulated events are required to describe different Standard Model and Beyond the Standard Model processes. For comparison, detector simulation in the ATLAS and CMS experiments consumed 40% of the grid central processing unit (CPU) during Run 2 of the LHC experiment [6, 7], and the expected CPU time needed to simulate an event increasing by a factor of three [8] or more after the HL-LHC upgrade in the upcoming years.

Generative models leveraging recent advancements in machine learning (ML) are able to build surrogate models capable of generating high fidelity samples with reduced computational cost. Common software frameworks for ML research, like TensorFlow [9], JAX [10], or PyTorch [11] benefit from strong community support and highly efficient implementations on hardware accelerators, such as Graphics Processing Units (GPUs). This flexibility is easily ported to experimental facilities and lowers the barrier of entry for software development, support, and maintenance. These improvements have the possibility to accelerate the comparison between measurements and theoretical predictions while decreasing the need for methods such as unfolding [12–14] once an efficient detector simulation is available.

Traditional simulation routines can be improved in multiple ways, leveraging modern software frameworks and hardware accelerators, such as GPUs [15–17]. Alternatively, differentiable programming (DP) software can also enable GPU support to traditional simulation routines. Differentiable programs track gradients with respect to simulation parameters or input variables at each step of the simulation program. While DP is not required for an algorithm to benefit from modern hardware acceleration (and not all DP frameworks are inherently GPU-compatible), differentiable programs provide additional advantages. For example, optimization of simulation inputs can be directly inferred from experimental data by finding the simulation parameters that jointly minimize the difference between synthetic

and experimental data. The optimization step is performed by propagating the gradients back through the simulation chain, thus reducing the required scale of simulated datasets used for alternative setups [18]. DP also opens new directions for simulation modeling that incorporate physics knowledge, which is critical for developing more robust, interpretable, and generalizable domain-aware scientific ML [19].

Additional usage of surrogate modeling and differentiable programming for HEP in the context of Snowmass are covered in detail in the Snowmass’21 LOIs [20, 21], the upcoming Snowmass’21 whitepaper by the Beam and Accelerator Modeling Interest Group (BAMIG) [22], two recent ICFA newsletters [23, 24], and the MODE collaboration [25].

This document is divided as follows. Sections 2 and 3 present brief introductions to surrogate models based on ML techniques and differentiable programming. In Secs. 4, 5, and 7, several ongoing projects are described and discussed. These examples are not meant to be comprehensive, but instead illustrative of the scope of research in this area. Finally, in Sec. 9, future directions and synergies in the short and long term future are explored.

2 Surrogate Models

The existing landscape of detector simulation consists of two primary approaches. The first is the accurate, but computationally intensive “full simulation” using GEANT4. The second is typically called “fast simulation” and may be considered a classical version of a surrogate model. Decreased simulation time is achieved by replacing computationally intensive parts of the simulation with simplified detector assumptions, resulting in speed improvements of more than 100 times compared to the full simulation. The resulting simulation, however, is less realistic and may be unsuitable for physics measurements that rely on detailed detector effects. Within the category of classical fast simulation, there are experiment-specific solutions [7, 26, 27] and the ultra-fast generic simulation DELPHES [28].

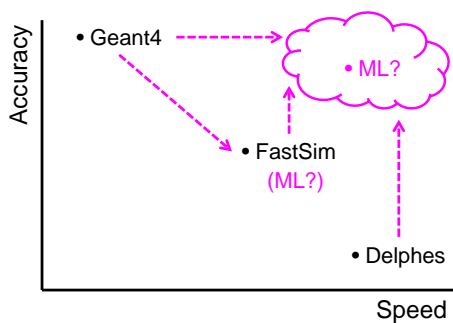


Figure 1: Depiction of different ways to incorporate ML in detector simulation workflows.

Figure 1 shows the different ways to introduce ML to this landscape: by replacing or augmenting part or all of GEANT4, or part or all of a classical fast simulation. Each option has a different goal: increasing speed while preserving accuracy, or preserving speed while increasing accuracy, respectively. ML could also be used to create a faster but less accurate simulation, similar to existing classical fast simulations. Alternatively, different ML surrogate models approaches may be classified based on what input data they require to produce simulated events. This leads to two categories:

1. fully generative models that entirely replace classical simulation engines, taking generated particle data or random noise as input; and
2. refinement techniques that are applied

during or after the event simulation step, taking lower-quality simulated events as input. Popular deep learning architectures for fully generative models are divided into three main categories including generative adversarial networks (GANs) [7, 29–53], variational autoencoders (VAEs) [40, 54–57], and normalizing flows [58–61]. Refinement techniques may be based on classification [45, 62] or regression [63, 64]. The generative models apply a stochastic approach, while the refinement techniques are usually deterministic.

Initially proposed in [65], GANs are trained following a minimax game:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_x(x)}[\log D(x)] + \mathbb{E}_{z \sim p_z(z)}[\log(1 - D(G(z)))], \quad (1)$$

where a generator network G is tasked to generate new samples from a noise distribution $p_z(z)$ while the discriminator network D judges the quality of the generated samples by comparing with target events sampled from $p_x(x)$. The adversarial loss function can lead to unstable training, often requiring additional fine tuning to achieve realistic results. An alternative to the loss function was proposed in [66] named Wasserstein GAN (WGAN). In the WGAN framework, the discriminator network is replaced with a critic network that uses the Wasserstein distance between generated and data samples as a metric to be minimized during training. There are many GAN variations that go beyond the vanilla and WGAN approaches.

Autoencoders are composed of two components: an encoder that compresses a set of input features into a smaller latent space, and a decoder that uses the information in that latent space to attempt to reconstruct the input features. VAEs combine autoencoders with a tractable latent space to generate new and realistic samples. Even though the probability density of the data is not tractable, VAEs minimize the evidence lower bound loss:

$$\mathcal{L}_{\text{VAE}} = -\mathbb{E}_{z \sim q(z|x)}[\log p_x(x|z)] + \mathcal{D}_{KL}(q(z|x)||p_z(z)). \quad (2)$$

The approximate posterior probability density $q(z|x)$ is enforced to be a tractable distribution through the Kullback–Leibler divergence term \mathcal{D}_{KL} . The reconstruction loss $\log p_x(x|z)$ is often defined as the mean squared error loss, in case of continuous distributions, or the cross-entropy loss, in case of discrete distributions.

As an alternative approach to handling data with an intractable probability distribution, normalizing flows [67–69] define a bijective transformation between a tractable base distribution, such as a normal or uniform distribution, to the data using the transformation of variables:

$$\log p_x(x) = \log p_z(z) - \log \det \left| \frac{\partial f(z)}{\partial z} \right|, \quad (3)$$

with terms $\frac{\partial f(z)}{\partial z}$ representing the Jacobian matrix of the transformation f . The loss function to be minimized is then defined as $-\log p_x(x)$, which is equivalent to minimizing the KL divergence between the transformed tractable base distribution and the data distribution.

The general calculation of the determinant in Eq. 3 has $\mathcal{O}(D^3)$ computational cost. This limitation is mitigated by restricting the bijective transformation f to the family of functions with triangular Jacobian matrix, bringing the computational complexity down to

$\mathcal{O}(D)$. There exist two main architectures to ensure a triangular Jacobian: bipartite [70] flows based on so-called coupling layers or autoregressive [71, 72] flows based on masked neural networks (NNs) [73]. Practically speaking, their main difference is the speed in which both directions of the bijector can be evaluated. In bipartite flows accessing the log-likelihood of data and sampling is equally fast. Autoregressive flows have a fast and a slow direction. Masked Autoregressive Flows (MAFs) [71] are fast in density estimation, but a factor D , given by the dimension of data space, slower in sampling. Inverse Autoregressive Flows (IAFs) [72] are fast in sampling, and a factor D slower in estimating the density of data points. See [68, 69] for more details on normalizing flows. An application of density estimation for calorimeter simulation is described in Sec. 5. A simplified depiction of the different generative strategies is shown in Fig. 2.

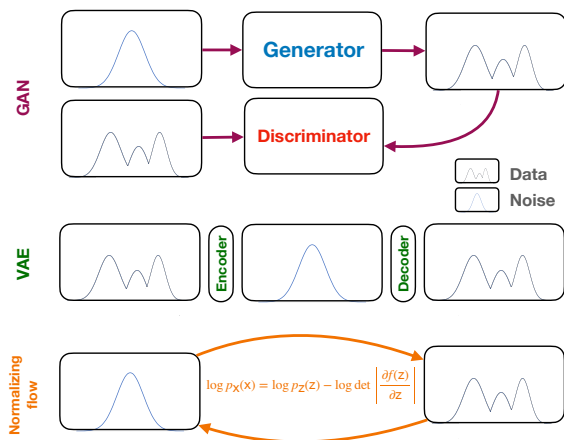


Figure 2: Summary of different machine learning methods used for generative models.

A natural question about surrogate models is to what extent the generated samples increase the statistical power with respect to the training data. At its core, the benefit from deep generative models comes from their ability to interpolate in high dimensions. One source of statistical amplification from the training dataset is from combinatorics - there are combinatorially many ways to attach showers to N particles in an event and deep generative models can naturally interpolate from the training dataset to have the correct kinematic properties. Interpolation also can result in improved statistical precision from the smoothness properties of neural networks (a form of ‘inductive bias’) [74, 75].

Further refinements to generated samples can be derived to improve generation quality. Those corrections can be coupled either to ML-based surrogate models or to classical fast simulation routines. Advantages of refinement in the former case include that the training is often more stable than for the original generative model, and correspondingly, the generative model may not need to match the precision of GEANT4. Alternatively, replacing the generative model with a classical simulator and relying on ML only for refinement may decrease the probability of unphysical output and provide better extrapolation beyond the training data. Some example applications of this type are discussed in Sec. 4.

3 Differentiable Programming

An alternative but complementary direction for surrogates lies in recent advancements in differentiable programming (DP). In DP, software is written in, or transformed into, dif-

ferentiable code via the use of automatic differentiation (AD) [76], an algorithmic way to efficiently evaluate derivatives of computer programs. When software is written in DP frameworks, access to the dependence of predictions on inputs is enabled through gradients.

These gradients are a significant addition to the information typically provided by simulators and crucially can be used in downstream modeling and inference tasks. This approach is flexible and optimizable; differentiable HEP software and ML tools can be mixed, for instance to use ML surrogates of non-differentiable computations, and can be jointly optimized to improve speed and prediction accuracy. When developed with DP, HEP simulation tools, and the physics knowledge they encode, can be used as physics prediction engines directly within ML pipelines for developing physics-informed ML tools. An illustration of differentiable programming is shown in Fig. 3. We note that deep generative models are a type of differentiable detector simulation, since gradients are readily available for neural networks. More details on differentiable programming for detector simulation are given in Sec. 7.




Figure 3: An illustration of differentiable programming for detector simulation.

4 ML-based Correction to Accelerate Geant4 Calorimeter Simulations

In full simulation routines, particles can be fully tracked using the complete underlying physics knowledge (FullSim) or approximate parametrizations can be used to simplify and accelerate the process (FastSim). Although future experiments plan to be heavily based on FastSim methods, the usage of FullSim is still imperative [77] (including FastSim tuning).

Focusing on the FullSim, among the most computationally demanding apparatuses to simulate are dense highly segmented particle physics detectors (e.g. calorimeters). This is because highly energetic particles produce cascades of secondary particles, resulting in an exponential number of particles with respect to the particle energy. The actual limit on the lowest energy particle simulated is controlled by *range cuts*. Increased range cuts correspond to increased production energy thresholds, thus reducing the number of produced secondary particles. As an immediate effect, the computational demands of the simulation are reduced. A side effect can be the reduction of the accuracy of the simulation. The extent of the inaccuracy increases as the range cut grows relative to the scale of the sensitive elements of the detector. While other parameters in GEANT4 may also be varied with effects on the simulation computing time and accuracy, range cuts have been found to be the most impactful.

This section outlines an approach to accelerate the FullSim execution time. One possibility is to use a deep generative model as a base that is then refined [45]. Another approach

is applying aggressive range cuts and then correcting the reduced accuracy simulation (referred to as *modified*) using ML methods. There are several complementary techniques to derive these corrections.

One method relies on event-level weights. A neural network is trained to classify the nominal versus the modified simulation and the classification score is used to calculate a multi-dimensional density ratio. This ratio is finally used to reweight the modified observables back to the nominal ones. The classification score is then used to approximate the density ratio $r(x)$. There are many ways to do this (see Ref [78]), but the most common approach is to use the binary cross entropy loss function and then derive r as:

$$r(\bar{x}) = \frac{\rho(y = 1|x)}{1 - \rho(y = 1|x)} \approx \frac{p_{\text{full}}(x)}{p_{\text{fast}}(x)}, \quad (4)$$

where ρ is the classifier and x is a set of observables used in the reweighting. The approximation in the above equation is a well-known result from statistics (see e.g., [79,80]). Additional post-processing can improve the approximation [81].

There is no unique way to pick x . One possibility is to refrain from choosing any specific high-level observables and instead learn directly from the lowest level inputs (e.g. energy deposits per calorimeter cell). An advantage of using high-level features is that it provides some regularization so that if the original model has phase space gaps in high-dimensions, there will not be infinities in the likelihood ratio [45].

Serving as a proof of concept, an example reweighting application using the lowest level inputs has been developed for the International Large Detector (ILD) electromagnetic barrel calorimeter [82]. The multilayer calorimeter consist of 30 layers, each one segmented in 30×30 cells. The data are projected into 3D images, where the color of each of the 27,000 voxels represent the energy deposit in the cell. A convolutional neural network (CNN) utilizing 3D convolution operations is trained to discriminate nominal from modified simulation events in order to approximate the ratio of Equation 4. Preliminary results shown in Figure 4 showcase the improvement of the reweighted modified simulation, resembling the high accuracy nominal GEANT4 simulation. The tradeoff of the correction via reweighting is the statistical dilution of the simulation sample [45].

An alternative method directly modifies the simulated event contents. A first proof of concept for this method is described in Ref. [64], based on an approach used in industry to accelerate MC ray-tracing [83], has recently been published. GEANT4 with an increased range cut provides the modified simulation, and a CNN is used to regress the energy value of each pixel, with the detector represented as a digitized grid. Figure 5 shows the promising results for photon showers in the CMS electromagnetic calorimeter. Regression approaches can also be applied to improve high-level variables, which may complement the low-level approach.

These correction approaches can reduce the computational complexity for two reasons: 1) the absolute time reduction of the simulation from a faster surrogate (either a deep generative model or higher range cuts) can reduce the calculation time by orders of magnitude and 2) the correction may be applied in parallel to many events utilizing parallel computations in accelerator hardware, such as GPUs. The overall speedup from this approach may

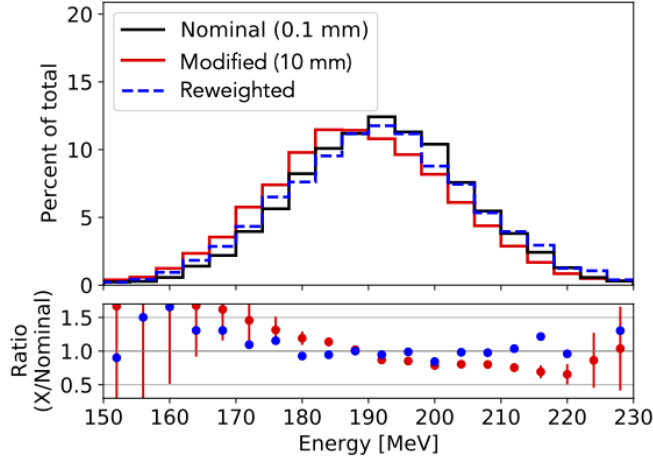


Figure 4: Comparison of the nominal, modified, and reweighted event energy deposit at the ILD barrel calorimeter induced by 10 GeV electron showers. The nominal distribution uses a 0.1 mm range cut, while the modified uses 10 mm, leading to $\sim 15\%$ simulation CPU speedup.

be limited by the throughput of the classical fast simulation engine, in this case the modified GEANT4. However, the potential for greater accuracy and reliability may make such a trade-off worthwhile, if the overall speedup is enough to meet the computing challenges of the HL-LHC and future colliders.

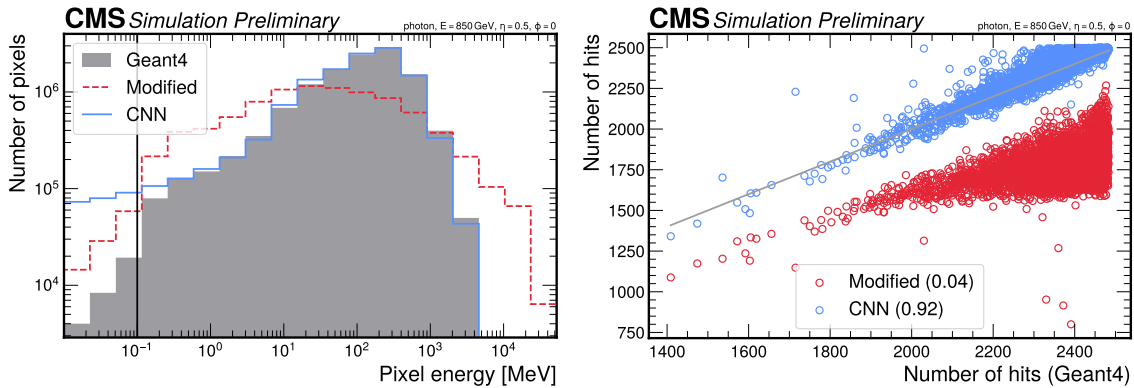


Figure 5: Left: A comparison of the per-pixel energy distribution for the modified simulation, the CNN output, and GEANT4. Right: Per-event comparisons of the number of hits, with the concordance correlation [84] between GEANT4 and the other simulations listed in parentheses and the gray line indicating exact agreement. These figures are reproduced from Ref. [64].

5 Detector Simulation with Normalizing Flows

Ideal surrogate models are fast and at the same time indistinguishable from the full simulation based on GEANT4. The latter can be tested by training a neural network classifier on “real” (based on GEANT4) and “fake” (generated from the surrogate) samples [85]. Previous surrogate models, based simple on GANs or VAEs, have failed such a test and yielded samples that were separable to nearly 100%. Normalizing Flows (NFs) provide an alternative approach to generative modeling, since they learn the likelihood of the data explicitly, in contrast to GANs and VAEs that only implicitly learn the data distribution. Maximizing the log-likelihood of the training data directly is more stable and not prone to mode-collapse. In addition, picking the model with the lowest validation loss seems to be an effective model selection strategy, which is a challenge for deep generative adversarial models.

The proof of concept of this approach is given in [59, 60] (called CALOFLOW), based on the same detector geometry that was studied in [29, 30]. This geometry is a simplified version of the ATLAS electromagnetic (ECAL) calorimeter, consisting of 3 layers with 288, 144, and 72 voxel, respectively. A new instance of CALOFLOW was trained for each particle type (e^+ , γ , π^+). In CALOFLOW, the data likelihood is learned in two steps, with two separate NFs. The first step only learns how the total deposited energy is distributed across the three calorimeter layers, conditioned on the incident energy, $p_1(E_i|E_{\text{inc}})$, with $E_{\text{tot}} = \sum_i E_i$. CALOFLOW uses a MAF for flow 1. The second step learns the normalized shower shape, i.e. how the energy deposited in each layer is distributed into the voxels, conditioned on the energy deposition of each layer and the incident energy, $p_2(\mathcal{I}|E_i, E_{\text{inc}})$. Both autoregressive architectures, MAFs and IAFs, have been applied to this step in [59] and [60] respectively. However, the high dimensionality of the voxel space made a training based on the log-likelihood prohibitive for the IAF. Instead, the flow of [60] was trained using probability density distillation, a method originally developed for speech synthesis in [86]. In generation, one first samples E_i from flow 1. These energies are then given to flow 2 to generate the showers. After shower generation, the resulting showers are renormalized to have the energies according to the E_i of flow 1.

Table 1 shows the main results of that approach, given by the training of a binary neural classifier. While this GAN-based model yields samples that are distinguishable from the GEANT4 samples, the NF-based model has a much higher fidelity and can fool the classifier much more often. * The generation of samples, especially with CALOFLOW v2, is as fast as the GAN. Differences in training time become irrelevant once more than 10^9 showers are generated, see Fig. 6. Figure 7 shows some example distributions for π^+ showers, comparing GEANT4 to CaloGAN [29, 30] to CaloFlow [59, 60].

In order for normalizing flow-based models to be used by the experimental collaborations, they have to prove their performance in more realistic setups as well. These will either have a higher number of voxels (like the ILD or CMS high granularity [87] calorimeters) and/or a conditioning on incident angle and/or position. Additional studies on the specific

*It is also possible that more recent, state-of-the-art GANs or VAEs perform better on this dataset, which is an interesting topic for future studies.

Table 1: AUC and JSD metrics for the classification of GEANT4 vs CALOGAN, CALOFlow v1, and CALOFlow v2 showers. Classifiers were trained on each particle type (e^+ , γ , π^+) separately. All entries show mean and standard deviation of 10 runs and are rounded to 3 digits (lower numbers are better). Taken from [60].

AUC / JSD		DNN-based classifier		
		GEANT4 vs.		
		CALOGAN	CALOFlow v1	CALOFlow v2
e^+	unnormalized	1.000(0) / 0.995(1)	0.859(10) / 0.365(14)	0.786(7) / 0.201(11)
	normalized	1.000(0) / 0.997(0)	0.870(2) / 0.378(5)	0.824(4) / 0.257(8)
γ	unnormalized	1.000(0) / 0.998(0)	0.756(48) / 0.174(68)	0.758(14) / 0.162(18)
	normalized	1.000(0) / 0.994(1)	0.796(2) / 0.216(4)	0.760(3) / 0.158(4)
π^+	unnormalized	1.000(0) / 0.993(0)	0.649(3) / 0.060(2)	0.729(2) / 0.144(3)
	normalized	1.000(0) / 0.997(1)	0.755(3) / 0.153(3)	0.807(1) / 0.230(3)




Figure 6: Comparison of shower generation times of GEANT4, CaloGAN [29, 30], CaloFlow v1 [59], and CaloFlow v2 [60].

NF architecture, such as autoregressive vs. bipartite flows, convolutional and other types of networks to give parameters of the transformation, or new bijective transformations might reveal more efficient (in terms of memory usage and/or sampling time) setups. Novel setups might also circumvent the 2-step approach of CALOFlow.

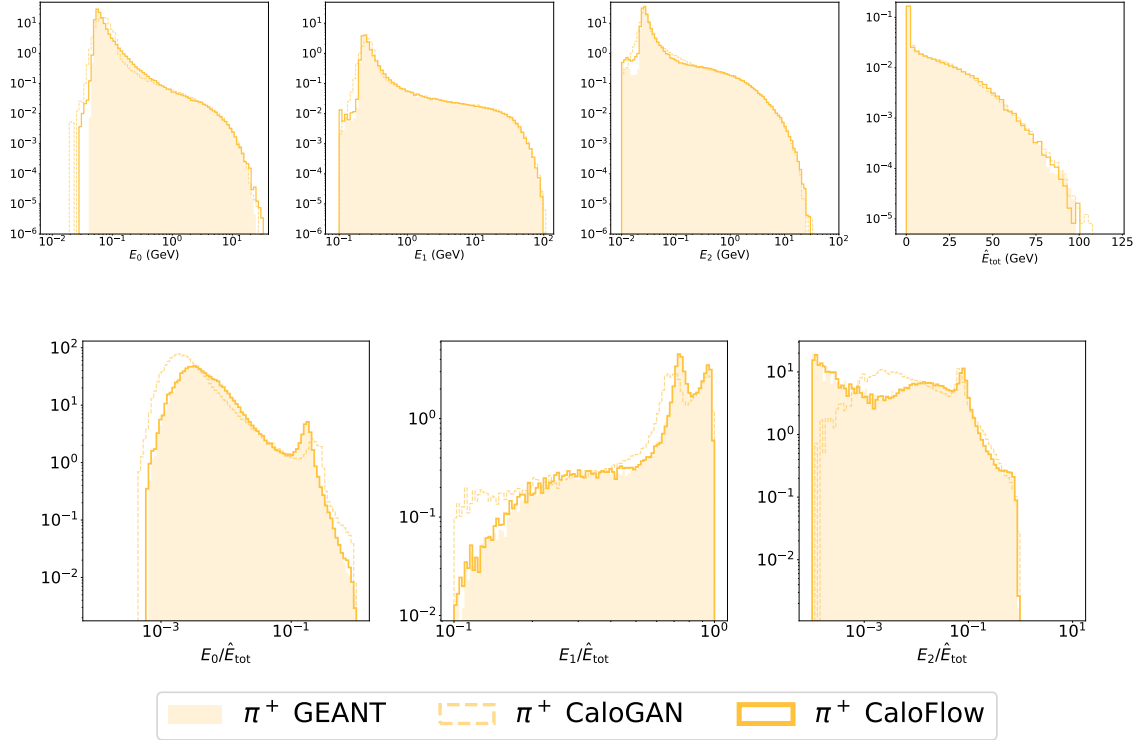


Figure 7: Distributions of energies in the 3 Calorimeter layers and total deposited energy (top) and ratio of layer energies to total deposited energy (bottom) for incident π^+ particles, comparing GEANT4 to CaloGAN [29,30] to CaloFlow [59,60].

6 Simulation of increasingly complex detectors

A number of challenges is encountered when moving from simplified detectors towards realistic simulations of energy deposits in modern calorimeters.

A primary issue lies in the substantial number of hits that need to be simulated. For example, the planned CMS High-Granularity Calorimeter (HGCal) [88] will have ≈ 6 million individual read-out channels with a similar order of magnitude for calorimeters in the future International Large Detector (ILD). Similarly, due to other design constraints, cells in a realistic calorimeter are not arranged in a regular grid but in more complex geometric patterns.

Simulating calorimeters with more than 10k cells using generative models was first attempted in [89] (65k channels) and [90] (27k channels). While these numbers are still much smaller than the entire calorimeter, they allow simulating a slice in η - ϕ large enough to fully contain a shower with realistic granularity. Based on such slices, entire calorimeters could e.g. be simulated by conditioning on impact position and angle.

The most accurate generative architecture tested by [90] was the co-called bounded

information-bottleneck autoencoder (BIB-AE) [91]. It essentially is a VAE with additional GAN-like critic networks. A key result of this contribution was the correct description of the single hit energy spectrum around the energy deposited by a minimally ionizing particle (MIP). To this end, an additional post-processing network was trained to fine-tune the output of the generative model. Further improvement of the fidelity of the generated data was possible by including a secondary density estimation step in the latent space [54], following the Buffer-VAE approach from [92].

Another challenge lies in simulating the more complex (compared to purely electromagnetic showers) showers initiated by hadrons. Here, [93] considered WGAN and BIB-AE architectures for the simulation of positively charged pions in the highly-granular Analogue Hadron Calorimeter (AHCAL). A potentially important observation was that while the BIB-AE yielded a more accurate initial description of showers, this difference — at least for energy response and resolution — largely vanished after processing with standard particle flow reconstruction. This implies that, depending on the intended downstream use, also simpler generative models might be able to capture relevant characteristics of a shower.

Looking towards the future, a number of challenges remains to be solved:

- Simultaneous simulation of different detector geometries and materials for the full depth of a highly granular calorimeter.
- Use of non-grid-based architectures (sets, graphs) to capture the geometry of realistic detectors (see e.g. [51])
- Multi-dimensional conditioning on energy, impact position, impact angle, and particle type.
- Integration in generation workflows of large experimental collaborations (see e.g. [40] for ATLAS).
- Solid treatment of the statistical properties and uncertainties of generated calorimeter data [75, 94].

Nevertheless, the large possible speed-up over alternative methods afforded by generative models makes them a crucial tool in understanding collider data at the highest precision.

7 Differentiable Programming for Detector Simulation, Design, and Inference

Differentiable programming for simulation relies on building AD-aware HEP simulation tools. AD uses the chain rule to evaluate derivatives of a function that is represented as a computer program. AD takes as input program code, whose derivative can be defined, and produces new code for evaluating the program and derivatives. AD typically builds a computational graph, or a directed acyclic graph of mathematical operations applied to an

input. Gradients are defined for each operation of the graph, and the total gradient can be evaluated from input to output, called forward mode, and from output to input, called reverse mode or backpropagation in ML.

For HEP simulation tools, ideally one would not rewrite the software but instead use AD tools which can merge easily with the existing software. For instance, recent work on `madjax` [95], a differentiable matrix element generator, augments python matrix element code generated by `MadGraph` [96] and merges it with `JAX`.

Surrogate models trained to mimic the behavior of high fidelity detector simulators can also be used within DP pipeline. For instance, iteratively trained surrogate models of a `GEANT4` magnet simulation were used for estimating gradients in a gradient descent optimization of the magnet system for the SHiP experiment [97]. This optimization found more performant and lighter weight designs for a magnetic shield. Similarly, detector surrogates were used to model the smearing induced by detectors on jets, and subsequently to provide gradients for gradient-based unfolding of jet distributions in [98]. Additional applications for fast surrogate models include optimization and design of particle accelerators [99–101], real-time feedback during commissioning and tuning of an accelerator facility [100,101], and uncertainty quantification of simulated parameters [102,103]. These examples show the large potential for such surrogate systems to be used in differentiable inference pipelines for tasks beyond only data generation, see also [25].

8 Synergies and a Joint Framework for Detector Simulation

Different ideas surveyed in this document have shown promising results on individual challenges in detector simulation for HEP. One of the challenges in the future is to identify how different ideas can be combined in a way that benefits the overall scientific community. One of the biggest advantages of `GEANT4` is the flexibility the software provides, resulting in widespread usage.

Providing a joint framework for detector simulation supports the testing and benchmarking of new methods as an effective way to promote collaboration between researchers and an ideal environment to keep track of new developments. This direction also streamlines the combination of multiple methods, such as individual detector surrogates, that combined create a full detector simulation.

Data challenges are also an effective method to build collaborations between different scientific communities. Researchers of different backgrounds have the opportunity to discuss and cooperate, promoting new developments. Data challenges are also a good opportunity for transparent comparison of new algorithms. This goal is currently being pursued in “Fast Calorimeter Simulation Challenge 2022” [104]. However, challenges by themselves are not sufficient - resources are required for integrating tools into simulation frameworks (experiment-independent or experiment-specific).

A joint software framework also opens the possibility for shared development between experiments. Machine learning based models often require large amounts of data for training,

restricting the number of users with access to computing centers with available resources for large scale development.[†] However, in a shared software environment, large ML models can be pre-trained in dedicated computing facilities using generic detector geometries. These models can then be later fine-tuned to include experiment-specific information, decreasing the computational burden required to achieve state-of-the-art results.

One of the biggest challenges of having a unified framework covering multiple experimental facilities is to cope with the differences in computational resources available and experiment-specific software. A possible solution is support for containerized images [106] from experimental collaborations to reproduce their detector simulation routines. This option reduces the need for experiment-specific knowledge while improving software portability.

Maintenance of the software also becomes crucial. New job positions for HEP software development should also be promoted to ensure future usability and continuity, in order to accelerate future generations of experiments and to ensure that legacy data and results are still accessible.

9 Future Directions

The computational complexity required for full detector simulation in high energy physics far exceeds the predicted resources available in future experimental facilities, requiring innovative strategies to accelerate the simulation process while preserving generation quality. Surrogate models are proposed as fast alternatives to replace part of full simulation routines, leveraging advancements in machine learning implemented in heterogeneous computing architectures.

While realistic simulations were used in some projects, primarily studies with simplified calorimeters were used to demonstrate the feasibility of new models. These include calorimeter geometries with a very regular structure or with a reduced amount of readout channels, such as the ILD example described in Section 4 or CaloGAN dataset described in Section 5. However, in a realistic detector this is not usually the case. The number of cells can be large and the geometry irregular. For example, the ATLAS detector calorimeter consists of 173,952 channels of variable size and shape [107], and the CMS High Granularity calorimeter will be constructed using hexagonal wafers [87]. Additionally, during a typical shower evolution into the calorimeter only a small portion of the cells ($\mathcal{O}(0.1\%)$) register a signal, leading to a very sparse dataset. Novel data structures and neural network architectures are required to account for the properties of the data. An example is to represent the calorimeter data in the form of a graph and use a Graph Neural Network to operate on it [46, 51]. This approach also detaches the method from a particular geometry; data from any type/shape of calorimeter can be converted into a universal graph data structure.

Differential programming can provide powerful new directions in simulator modeling.

[†]An alternative solution to this problem in the context of method development and prototyping was recently discussed in [105].

Building a fully differentiable HEP simulation chain would open a realm of new schemes for optimizing simulations, improving simulation speed, inference and design optimization tasks, and for building physics-informed HEP-ML system that utilize the physics knowledge within HEP simulation software. Dedicated automatic differentiation tools capable of augmenting existing software, rather than requiring complete software rewrites, are needed. New compiler-based source-translation based AD tools, such as `enzyme` [108] and `CLAD` [109], are promising for such tasks.

Applications to realistic scenarios for all ideas will be crucial to identify current limitations and future research directions. While examples in this document have shown promising results, one needs to consider the software environment required to maintain, support, and develop new algorithms. Maintenance of the software is imperative to ensure that algorithms used within experimental collaborations are up to date with the ones available to the wider scientific community.

ACKNOWLEDGEMENTS

VM and BN are supported by the U.S. Department of Energy (DOE), Office of Science under contract DE-AC02-05CH11231. MK is supported by the US Department of Energy (DOE) under grant DE-AC02-76SF00515. CK and DS are supported by DOE grant DOE-SC0010008. KP is supported by the Fermi National Accelerator Laboratory, managed and operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy. DW is supported by NSF grant PHY-1912764 and funding from the Heising-Simons Foundation and the Bose Foundation. Gregor Kasieczka is supported by the Deutsche Forschungsgemeinschaft under Germany’s Excellence Strategy – EXC 2121 Quantum Universe – 390833306.

10 References and bibliography

- [1] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003, 2008. doi:10.1088/1748-0221/3/08/S08003.
- [2] CMS Collaboration. The CMS Experiment at the CERN LHC. *JINST*, 3:S08004, 2008. doi:10.1088/1748-0221/3/08/S08004.
- [3] S. Agostinelli et al. Geant4 - a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250 – 303, 2003. URL: <http://www.sciencedirect.com/science/article/pii/S0168900203013688>, doi:[https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).
- [4] J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce Dubois, M. Asai, G. Barraud, R. Capra, S. Chauvie, R. Chytracek, G.A.P. Cirrone, G. Cooperman, G. Cosmo,

- G. Cuttone, G.G. Daquino, M. Donszelmann, M. Dressel, G. Folger, F. Foppiano, J. Generowicz, V. Grichine, S. Guatelli, P. Gumplinger, A. Heikkinen, I. Hrivnacova, A. Howard, S. Incerti, V. Ivanchenko, T. Johnson, F. Jones, T. Koi, R. Kokoulin, M. Kossov, H. Kurashige, V. Lara, S. Larsson, F. Lei, O. Link, F. Longo, M. Maire, A. Mantero, B. Mascialino, I. McLaren, P. Mendez Lorenzo, K. Minamimoto, K. Murakami, P. Nieminen, L. Pandola, S. Parlati, L. Peralta, J. Perl, A. Pfeiffer, M.G. Pia, A. Ribon, P. Rodrigues, G. Russo, S. Sadilov, G. Santin, T. Sasaki, D. Smith, N. Starkov, S. Tanaka, E. Tcherniaev, B. Tome, A. Trindade, P. Truscott, L. Urban, M. Verderi, A. Walkden, J.P. Wellisch, D.C. Williams, D. Wright, and H. Yoshida. Geant4 developments and applications. *IEEE Transactions on Nuclear Science*, 53(1):270–278, 2006. doi:[10.1109/TNS.2006.869826](https://doi.org/10.1109/TNS.2006.869826).
- [5] J. Allison, K. Amako, J. Apostolakis, P. Arce, M. Asai, T. Aso, E. Bagli, A. Bagulya, S. Banerjee, G. Barrand, B.R. Beck, A.G. Bogdanov, D. Brandt, J.M.C. Brown, H. Burkhardt, Ph. Canal, D. Cano-Ott, S. Chauvie, K. Cho, G.A.P. Cirrone, G. Cooperman, M.A. Cortés-Giraldo, G. Cosmo, G. Cuttone, G. Depaola, L. Desorgher, X. Dong, A. Dotti, V.D. Elvira, G. Folger, Z. Francis, A. Galoyan, L. Garnier, M. Gayer, K.L. Genser, V.M. Grichine, S. Guatelli, P. Guèye, P. Gumplinger, A.S. Howard, I. Hřivnáčová, S. Hwang, S. Incerti, A. Ivanchenko, V.N. Ivanchenko, F.W. Jones, S.Y. Jun, P. Kaitaniemi, N. Karakatsanis, M. Karamitros, M. Kelsey, A. Kimura, T. Koi, H. Kurashige, A. Lechner, S.B. Lee, F. Longo, M. Maire, D. Mancusi, A. Mantero, E. Mendoza, B. Morgan, K. Murakami, T. Nikitina, L. Pandola, P. Paprocki, J. Perl, I. Petrović, M.G. Pia, W. Pokorski, J.M. Quesada, M. Raine, M.A. Reis, A. Ribon, A. Ristić Fira, F. Romano, G. Russo, G. Santin, T. Sasaki, D. Sawkey, J.I. Shin, I.I. Strakovsky, A. Taborda, S. Tanaka, B. Tomé, T. Toshito, H.N. Tran, P.R. Truscott, L. Urban, V. Uzhinsky, J.M. Verbeke, M. Verderi, B.L. Wendt, H. Wenzel, D.H. Wright, D.M. Wright, T. Yamashita, J. Yarba, and H. Yoshida. Recent developments in geant4. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 835:186–225, 2016. URL: <https://www.sciencedirect.com/science/article/pii/S0168900216306957>, doi:<https://doi.org/10.1016/j.nima.2016.06.125>.
- [6] J Apostolakis et al. HEP Software Foundation Community White Paper Working Group - Detector Simulation. 3 2018. arXiv:[1803.04165](https://arxiv.org/abs/1803.04165).
- [7] Georges Aad et al. AtlFast3: the next generation of fast simulation in ATLAS. 9 2021. arXiv:[2109.02551](https://arxiv.org/abs/2109.02551).
- [8] Pedro, Kevin. Current and future performance of the CMS simulation. *EPJ Web Conf.*, 214:02036, 2019. doi:[10.1051/epjconf/201921402036](https://doi.org/10.1051/epjconf/201921402036).
- [9] Martín Abadi et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. URL: <http://tensorflow.org/>.
- [10] James Bradbury et al. JAX: composable transformations of Python+NumPy programs, 2018. URL: <http://github.com/google/jax>.

- [11] Adam Paszke et al. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [12] A. N. Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Soviet Math. Dokl.*, 4:1035–1038, 1963.
- [13] G. D'Agostini. A Multidimensional unfolding method based on Bayes' theorem. *Nucl. Instrum. Meth. A*, 362:487–498, 1995. doi:10.1016/0168-9002(95)00274-X.
- [14] Andreas Hocker and Vakhtang Kartvelishvili. SVD approach to data unfolding. *Nucl. Instrum. Meth. A*, 372:469–481, 1996. arXiv:hep-ph/9509307, doi:10.1016/0168-9002(95)01478-0.
- [15] Stefano Carrazza and Juan M. Cruz-Martinez. VegasFlow: accelerating Monte Carlo simulation across multiple hardware platforms. *Comput. Phys. Commun.*, 254:107376, 2020. arXiv:2002.12921, doi:10.1016/j.cpc.2020.107376.
- [16] Stefano Carrazza, Juan Cruz-Martinez, Marco Rossi, and Marco Zaro. MadFlow: automating Monte Carlo simulation on GPU for particle physics processes. *Eur. Phys. J. C*, 81(7):656, 2021. arXiv:2106.10279, doi:10.1140/epjc/s10052-021-09443-8.
- [17] Zhihua Dong, Heather Gray, Charles Leggett, Meifeng Lin, Vincent R. Pascuzzi, and Kwangmin Yu. Porting hep parameterized calorimeter simulation code to gpus. *Frontiers in Big Data*, 4, 2021. URL: <https://www.frontiersin.org/article/10.3389/fdata.2021.665783>, doi:10.3389/fdata.2021.665783.
- [18] Johann Brehmer, Gilles Louppe, Juan Pavez, and Kyle Cranmer. Mining gold from implicit models to improve likelihood-free inference. *Proc. Nat. Acad. Sci.*, 117(10):5242–5249, 2020. arXiv:1805.12244, doi:10.1073/pnas.1915980117.
- [19] Nathan Baker, Frank Alexander, Timo Bremer, Aric Hagberg, Yannis Kevrekidis, Habib Najm, Manish Parashar, Abani Patra, James Sethian, Stefan Wild, Karen Willcox, and Steven Lee. Workshop report on basic research needs for scientific machine learning: Core technologies for artificial intelligence. 1(1), 2 2019. URL: <https://www.osti.gov/biblio/1478744>, doi:10.2172/1478744.
- [20] R. Lehe et al. Machine learning and surrogates models for simulation-based optimization of accelerator design. *Snowmass21 LOI*, 2020. URL: https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF2_CompF3-AF1_AF6_Lehe-075.pdf.
- [21] D. Winklehner and A. Adelman. Application of Machine Learning to Particle Accelerator Simulations. *Snowmass21 LOI*, 2020. URL: https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF3_CompF0-AF1_AF0_Winklehner-108.pdf.

- [22] J-L Vay et al. Snowmass21 Accelerator Modeling Community White Paper. *Snowmass'21 Community Exercise*, 2022. White Paper Submitted to Snowmass 2021, Computational & Accelerator Frontiers.
- [23] J.-L. Vay, A. Huebl, R. Lehe, N. M. Cook, R. J. England, U. Niedermayer, P. Piot, F. Tsung, and D. Winklehner. Modeling of advanced accelerator concepts. *Journal of Instrumentation*, 16(10):T10003, 2021. Publisher: IOP Publishing. URL: <https://iopscience.iop.org/article/10.1088/1748-0221/16/10/T10003/meta>, doi:10.1088/1748-0221/16/10/T10003.
- [24] D. Sagan, M. Berz, N. M. Cook, Y. Hao, G. Hoffstaetter, A. Huebl, C.-K. Huang, M. H. Langston, C. E. Mayes, C. E. Mitchell, C.-K. Ng, J. Qiang, R. D. Ryne, A. Scheinker, E. Stern, J.-L. Vay, D. Winklehner, and H. Zhang. Simulations of future particle accelerators: issues and mitigations. *Journal of Instrumentation*, 16(10):T10002, 2021. Publisher: IOP Publishing. URL: <https://doi.org/10.1088/1748-0221/16/10/t10002>, doi:10.1088/1748-0221/16/10/T10002.
- [25] The MODE Collaboration. Toward the End-to-End Optimization of Particle Physics Instruments with Differentiable Programming. 2022.
- [26] G. Aad et al. The ATLAS Simulation Infrastructure. *Eur. Phys. J. C*, 70:823–874, 2010. [arXiv:1005.4568](https://arxiv.org/abs/1005.4568), doi:10.1140/epjc/s10052-010-1429-9.
- [27] Sezen Sekmen. Recent Developments in CMS Fast Simulation. *PoS, ICHEP2016*:181, 2016. [arXiv:1701.03850](https://arxiv.org/abs/1701.03850), doi:10.22323/1.282.0181.
- [28] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi. DELPHES 3, a modular framework for fast simulation of a generic collider experiment. *JHEP*, 02:057, 2014. [arXiv:1307.6346](https://arxiv.org/abs/1307.6346), doi:10.1007/JHEP02(2014)057.
- [29] Michela Paganini, Luke de Oliveira, and Benjamin Nachman. Accelerating Science with Generative Adversarial Networks: An Application to 3D Particle Showers in Multilayer Calorimeters. *Phys. Rev. Lett.*, 120(4):042003, 2018. [arXiv:1705.02355](https://arxiv.org/abs/1705.02355), doi:10.1103/PhysRevLett.120.042003.
- [30] Michela Paganini, Luke de Oliveira, and Benjamin Nachman. CaloGAN : Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks. *Phys. Rev.*, D97(1):014021, 2018. [arXiv:1712.10321](https://arxiv.org/abs/1712.10321), doi:10.1103/PhysRevD.97.014021.
- [31] Luke de Oliveira, Michela Paganini, and Benjamin Nachman. Controlling Physical Attributes in GAN-Accelerated Simulation of Electromagnetic Calorimeters. *J. Phys. Conf. Ser.*, 1085(4):042017, 2018. [arXiv:1711.08813](https://arxiv.org/abs/1711.08813), doi:10.1088/1742-6596/1085/4/042017.
- [32] Riccardo Di Sipio, Michele Faucci Giannelli, Sana Ketabchi Haghghat, and Serena Palazzo. DijetGAN: A Generative-Adversarial Network Approach for the Simulation of QCD Dijet Events at the LHC. 2019. [arXiv:1903.02433](https://arxiv.org/abs/1903.02433), doi:10.1007/JHEP08(2019)110.

- [33] Steven Farrell, Wahid Bhimji, Thorsten Kurth, Mustafa Mustafa, Deborah Bard, Zarija Lukic, Benjamin Nachman, and Harley Patton. Next Generation Generative Neural Networks for HEP. *EPJ Web Conf.*, 214:09005, 2019. doi:[10.1051/epjconf/201921409005](https://doi.org/10.1051/epjconf/201921409005).
- [34] Martin Erdmann, Lukas Geiger, Jonas Glombitza, and David Schmidt. Generating and refining particle detector simulations using the Wasserstein distance in adversarial networks. *Comput. Softw. Big Sci.*, 2(1):4, 2018. arXiv:[1802.03325](https://arxiv.org/abs/1802.03325), doi:[10.1007/s41781-018-0008-x](https://doi.org/10.1007/s41781-018-0008-x).
- [35] Martin Erdmann, Jonas Glombitza, and Thorben Quast. Precise simulation of electromagnetic calorimeter showers using a Wasserstein Generative Adversarial Network. *Comput. Softw. Big Sci.*, 3(1):4, 2019. arXiv:[1807.01954](https://arxiv.org/abs/1807.01954), doi:[10.1007/s41781-018-0019-7](https://doi.org/10.1007/s41781-018-0019-7).
- [36] Kamil Deja, Tomasz Trzcinski, and Lukasz Graczykowski. Generative models for fast cluster simulations in the TPC for the ALICE experiment. *EPJ Web Conf.*, 214:06003, 2019. doi:[10.1051/epjconf/201921406003](https://doi.org/10.1051/epjconf/201921406003).
- [37] Pasquale Musella and Francesco Pandolfi. Fast and Accurate Simulation of Particle Detectors Using Generative Adversarial Networks. *Comput. Softw. Big Sci.*, 2(1):8, 2018. arXiv:[1805.00850](https://arxiv.org/abs/1805.00850), doi:[10.1007/s41781-018-0015-y](https://doi.org/10.1007/s41781-018-0015-y).
- [38] S. Vallecorsa. Generative models for fast simulation. *J. Phys. Conf. Ser.*, 1085(2):022005, 2018. doi:[10.1088/1742-6596/1085/2/022005](https://doi.org/10.1088/1742-6596/1085/2/022005).
- [39] F. Carminati, A. Gheata, G. Khattak, P. Mendez Lorenzo, S. Sharan, and S. Vallecorsa. Three dimensional Generative Adversarial Networks for fast simulation. *J. Phys. Conf. Ser.*, 1085(3):032016, 2018. doi:[10.1088/1742-6596/1085/3/032016](https://doi.org/10.1088/1742-6596/1085/3/032016).
- [40] Deep generative models for fast shower simulation in ATLAS. *ATL-SOFT-PUB-2018-001*, Jul 2018. URL: <http://cds.cern.ch/record/2630433>.
- [41] Viktoria Chekalina, Elena Orlova, Fedor Ratnikov, Dmitry Ulyanov, Andrey Ustyuzhanin, and Egor Zakharov. Generative Models for Fast Calorimeter Simulation.LHCb case. *CHEP 2018*, 2018. arXiv:[1812.01319](https://arxiv.org/abs/1812.01319), doi:[10.1051/epjconf/201921402034](https://doi.org/10.1051/epjconf/201921402034).
- [42] C. Ahdida et al. Fast simulation of muons produced at the SHiP experiment using Generative Adversarial Networks. 2019. arXiv:[1909.04451](https://arxiv.org/abs/1909.04451), doi:[10.1088/1748-0221/14/11/P11028](https://doi.org/10.1088/1748-0221/14/11/P11028).
- [43] Jesus Arjona Martinez, Thong Q. Nguyen, Maurizio Pierini, Maria Spiropulu, and Jean-Roch Vlimant. Particle Generative Adversarial Networks for full-event simulation at the LHC and their application to pileup description. *ACAT 2019*, 2019. arXiv:[1912.02748](https://arxiv.org/abs/1912.02748), doi:[10.1088/1742-6596/1525/1/012081](https://doi.org/10.1088/1742-6596/1525/1/012081).
- [44] Anja Butter and Tilman Plehn. Generative Networks for LHC events. 8 2020. arXiv:[2008.08558](https://arxiv.org/abs/2008.08558).

- [45] S. Diefenbacher, E. Eren, G. Kasieczka, A. Korol, B. Nachman, and D. Shih. DC-TRGAN: Improving the Precision of Generative Models with Reweighting. *Journal of Instrumentation*, 15:P11004, 2020. [arXiv:2009.03796](#), [doi:10.1088/1748-0221/15/11/p11004](#).
- [46] Raghav Kansal, Javier Duarte, Breno Orzari, Thiago Tomei, Maurizio Pierini, Mary Touranakou, Jean-Roch Vlimant, and Dimitrios Gunopoulos. Graph Generative Adversarial Networks for Sparse Data Generation in High Energy Physics. *34th Conference on Neural Information Processing Systems*, 11 2020. [arXiv:2012.00173](#).
- [47] A. Maevskiy, F. Ratnikov, A. Zinchenko, and V. Riabov. Simulating the Time Projection Chamber responses at the MPD detector using Generative Adversarial Networks. 12 2020. [arXiv:2012.04595](#).
- [48] Suyong Choi and Jae Hoon Lim. A Data-driven Event Generator for Hadron Colliders using Wasserstein Generative Adversarial Network. 2 2021. [arXiv:2102.11524](#), [doi:10.1007/s40042-021-00095-1](#).
- [49] Florian Rehm, Sofia Vallecorsa, Vikram Saletore, Hans Pabst, Adel Chaibi, Valeriu Codreanu, Kerstin Borrás, and Dirk Krücker. Reduced Precision Strategies for Deep Learning: A High Energy Physics Generative Adversarial Network Use Case. 3 2021. [arXiv:2103.10142](#), [doi:10.5220/0010245002510258](#).
- [50] Florian Rehm, Sofia Vallecorsa, Kerstin Borrás, and Dirk Krücker. Validation of Deep Convolutional Generative Adversarial Networks for High Energy Physics Calorimeter Simulations. 3 2021. [arXiv:2103.13698](#).
- [51] Raghav Kansal, Javier Duarte, Hao Su, Breno Orzari, Thiago Tomei, Maurizio Pierini, Mary Touranakou, Jean-Roch Vlimant, and Dimitrios Gunopoulos. Particle Cloud Generation with Message Passing Generative Adversarial Networks. 6 2021. [arXiv:2106.11535](#).
- [52] Gul Rukh Khattak, Sofia Vallecorsa, Federico Carminati, and Gul Muhammad Khan. Fast Simulation of a High Granularity Calorimeter by Generative Adversarial Networks. 9 2021. [arXiv:2109.07388](#).
- [53] Carlos Bravo-Prieto, Julien Baglio, Marco Cè, Anthony Francis, Dorota M. Grabowska, and Stefano Carrazza. Style-based quantum generative adversarial networks for Monte Carlo events. 10 2021. [arXiv:2110.06933](#).
- [54] Erik Buhmann, Sascha Diefenbacher, Engin Eren, Frank Gaede, Gregor Kasieczka, Anatolii Korol, and Katja Krüger. Decoding Photons: Physics in the Latent Space of a BIB-AE Generative Network. 2 2021. [arXiv:2102.12491](#).
- [55] Ali Hariri, Darya Dyachkova, and Sergei Gleyzer. Graph Generative Models for Fast Detector Simulations in High Energy Physics. 4 2021. [arXiv:2104.01725](#).
- [56] Wei Mu, Alexander I. Himmel, and Bryan Ramson. Photon detection probability prediction using one-dimensional generative neural network. 9 2021. [arXiv:2109.07277](#).

- [57] Breno Orzari, Thiago Tomei, Maurizio Pierini, Mary Touranakou, Javier Duarte, Raghav Kansal, Jean-Roch Vlimant, and Dimitrios Gunopulos. Sparse Data Generation for Particle-Based Simulation of Hadronic Jets in the LHC. In *38th International Conference on Machine Learning Conference*, 9 2021. [arXiv:2109.15197](#).
- [58] Yadong Lu, Julian Collado, Daniel Whiteson, and Pierre Baldi. Sparse autoregressive models for scalable generation of sparse images in particle physics. *Phys. Rev. D*, 103:036012, 2021. [arXiv:2009.14017](#), [doi:10.1103/PhysRevD.103.036012](#).
- [59] Claudius Krause and David Shih. CaloFlow: Fast and Accurate Generation of Calorimeter Showers with Normalizing Flows, 6 2021. [arXiv:2106.05285](#).
- [60] Claudius Krause and David Shih. CaloFlow II: Even Faster and Still Accurate Generation of Calorimeter Showers with Normalizing Flows, 10 2021. [arXiv:2110.11377](#).
- [61] Anja Butter, Theo Heimel, Sander Hummerich, Tobias Krebs, Tilman Plehn, Armand Rousset, and Sophia Vent. Generative Networks for Precision Enthusiasts. 10 2021. [arXiv:2110.13632](#).
- [62] Ramon Winterhalder, Marco Bellagente, and Benjamin Nachman. Latent Space Refinement for Deep Generative Models. 6 2021. [arXiv:2106.00792](#).
- [63] Cheng Chen, Olmo Cerri, Thong Q. Nguyen, Jean-Roch Vlimant, and Maurizio Pierini. Data augmentation at the LHC through analysis-specific fast simulation with deep learning. 10 2020. [arXiv:2010.01835](#).
- [64] Sunanda Banerjee, Brian Cruz Rodriguez, Lena Franklin, Harold Guerrero De La Cruz, Tara Leininger, Scarlet Norberg, Kevin Pedro, Angel Rosado Trinidad, and Yiheng Ye. Denoising convolutional networks to accelerate detector simulation, 2022. [arXiv:2202.05320](#).
- [65] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL: <https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf>.
- [66] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017.
- [67] Danilo Jimenez Rezende and Shakir Mohamed. Variational Inference with Normalizing Flows. *arXiv e-prints*, page arXiv:1505.05770, May 2015. [arXiv:1505.05770](#).
- [68] Ivan Kobyzev, Simon J. D. Prince, and Marcus A. Brubaker. Normalizing Flows: An Introduction and Review of Current Methods. *arXiv e-prints*, page arXiv:1908.09257, August 2019. [arXiv:1908.09257](#).

- [69] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing Flows for Probabilistic Modeling and Inference. *arXiv e-prints*, page arXiv:1912.02762, December 2019. [arXiv:1912.02762](#).
- [70] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. *arXiv e-prints*, page arXiv:1605.08803, May 2016. [arXiv:1605.08803](#).
- [71] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked Autoregressive Flow for Density Estimation. *arXiv e-prints*, page arXiv:1705.07057, May 2017. [arXiv:1705.07057](#).
- [72] Diederik P. Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improving Variational Inference with Inverse Autoregressive Flow. *arXiv e-prints*, page arXiv:1606.04934, June 2016. [arXiv:1606.04934](#).
- [73] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. MADE: Masked Autoencoder for Distribution Estimation. *arXiv e-prints*, page arXiv:1502.03509, February 2015. [arXiv:1502.03509](#).
- [74] G. Kasieczka A. Butter, S. Diefenbacher, B. Nachman, and T. Plehn. GANplifying Event Samples. *SciPost Physics*, 10:139, 2021. [arXiv:2008.06545](#).
- [75] Sebastian Bieringer, Anja Butter, Sascha Diefenbacher, Engin Eren, Frank Gaede, Daniel Hundhausen, Gregor Kasieczka, Benjamin Nachman, Tilman Plehn, and Mathias Trabs. Calomplification - The Power of Generative Calorimeter Models. 2 2022. [arXiv:2202.07352](#).
- [76] Atilim Gunes Baydin, Barak A. Pearlmutter, and Alexey Andreyevich Radul. Automatic differentiation in machine learning: a survey. *CoRR*, abs/1502.05767, 2015. URL: <http://arxiv.org/abs/1502.05767>, [arXiv:1502.05767](#).
- [77] P Calafiura, J Catmore, D Costanzo, and A Di Girolamo. ATLAS HL-LHC Computing Conceptual Design Report. Technical report, CERN, Geneva, Sep 2020. URL: <https://cds.cern.ch/record/2729668>.
- [78] Benjamin Nachman and Jesse Thaler. Learning from many collider events at once. *Phys. Rev. D*, 103(11):116013, 2021. [arXiv:2101.07263](#), [doi:10.1103/PhysRevD.103.116013](https://doi.org/10.1103/PhysRevD.103.116013).
- [79] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York Inc., New York, NY, USA, 2001.
- [80] Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. *Density Ratio Estimation in Machine Learning*. Cambridge University Press, 2012. [doi:10.1017/CBO9781139035613](https://doi.org/10.1017/CBO9781139035613).
- [81] Kyle Cranmer, Juan Pavez, and Gilles Louppe. Approximating Likelihood Ratios with Calibrated Discriminative Classifiers. 6 2015. [arXiv:1506.02169](#).

- [82] Halina Abramowicz et al. International Large Detector: Interim Design Report. 3 2020. [arXiv:2003.01116](https://arxiv.org/abs/2003.01116).
- [83] Steve Bako, Thijs Vogels, Brian McWilliams, Mark Meyer, Jan Novák, Alex Harvill, Pradeep Sen, Tony Derose, and Fabrice Rousselle. Kernel-predicting convolutional networks for denoising monte carlo renderings. *ACM Trans. Graph.*, 36(4), jul 2017. [doi:10.1145/3072959.3073708](https://doi.org/10.1145/3072959.3073708).
- [84] Lawrence I-Kuei Lin. A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, 45(1):255–268, 1989. URL: <http://www.jstor.org/stable/2532051>.
- [85] David Lopez-Paz and Maxime Oquab. Revisiting Classifier Two-Sample Tests. *arXiv e-prints*, page arXiv:1610.06545, October 2016. [arXiv:1610.06545](https://arxiv.org/abs/1610.06545).
- [86] Aaron van den Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George van den Driessche, Edward Lockhart, Luis C. Cobo, Florian Stimberg, Norman Casagrande, Dominik Grewe, Seb Noury, Sander Dieleman, Erich Elsen, Nal Kalchbrenner, Heiga Zen, Alex Graves, Helen King, Tom Walters, Dan Belov, and Demis Hassabis. Parallel WaveNet: Fast High-Fidelity Speech Synthesis. *arXiv e-prints*, page arXiv:1711.10433, November 2017. [arXiv:1711.10433](https://arxiv.org/abs/1711.10433).
- [87] The Phase-2 Upgrade of the CMS Endcap Calorimeter. 2017.
- [88] The Phase-2 Upgrade of the CMS Endcap Calorimeter. Technical report, CERN, Geneva, Nov 2017. URL: <https://cds.cern.ch/record/2293646>.
- [89] Dawit Belayneh et al. Calorimetry with deep learning: particle simulation and reconstruction for collider physics. *Eur. Phys. J. C*, 80(7):688, 2020. [arXiv:1912.06794](https://arxiv.org/abs/1912.06794), [doi:10.1140/epjc/s10052-020-8251-9](https://doi.org/10.1140/epjc/s10052-020-8251-9).
- [90] Erik Buhmann, Sascha Diefenbacher, Engin Eren, Frank Gaede, Gregor Kasieczka, Anatolii Korol, and Katja Krüger. Getting High: High Fidelity Simulation of High Granularity Calorimeters with High Speed. *Comput. Softw. Big Sci.*, 5(1):13, 2021. [arXiv:2005.05334](https://arxiv.org/abs/2005.05334), [doi:10.1007/s41781-021-00056-0](https://doi.org/10.1007/s41781-021-00056-0).
- [91] Slava Voloshynovskiy, Mouad Kondah, Shideh Rezaeifar, Olga Taran, Taras Holotyak, and Danilo Jimenez Rezende. Information bottleneck through variational glasses. *arXiv preprint arXiv:1912.00830*, 2019.
- [92] Sydney Otten, Sascha Caron, Wieske de Swart, Melissa van Beekveld, Luc Hendriks, Caspar van Leeuwen, Damian Podareanu, Roberto Ruiz de Austri, and Rob Verheyen. Event Generation and Statistical Sampling for Physics with Deep Generative Models and a Density Information Buffer. *Nature Commun.*, 12(1):2985, 2021. [arXiv:1901.00875](https://arxiv.org/abs/1901.00875), [doi:10.1038/s41467-021-22616-z](https://doi.org/10.1038/s41467-021-22616-z).
- [93] Erik Buhmann, Sascha Diefenbacher, Engin Eren, Frank Gaede, Daniel Hundhausen, Gregor Kasieczka, William Korcari, Katja Krüger, Peter McKeown, and Lennart Rustige. Hadrons, Better, Faster, Stronger. 12 2021. [arXiv:2112.09709](https://arxiv.org/abs/2112.09709).

- [94] Anja Butter, Sascha Diefenbacher, Gregor Kasieczka, Benjamin Nachman, and Tilman Plehn. GANplifying event samples. *SciPost Phys.*, 10(6):139, 2021. [arXiv:2008.06545](#), [doi:10.21468/SciPostPhys.10.6.139](#).
- [95] Lukas Heinrich and Michael Kagan. Differentiable Matrix Elements with MadJax. In *20th International Workshop on Advanced Computing and Analysis Techniques in Physics Research: AI Decoded - Towards Sustainable, Diverse, Performant and Effective Scientific Computing*, 2 2022. [arXiv:2203.00057](#).
- [96] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP*, 07:079, 2014. [arXiv:1405.0301](#), [doi:10.1007/JHEP07\(2014\)079](#).
- [97] Sergey Shirobokov, Vladislav Belavin, Michael Kagan, Andrey Ustyuzhanin, and Atılım Güneş Baydin. Black-Box Optimization with Local Generative Surrogates. 2 2020. [arXiv:2002.04632](#).
- [98] Maxime Vandegar, Michael Kagan, Antoine Wehenkel, and Gilles Louppe. Neural Empirical Bayes: Source Distribution Estimation and its Applications to Simulation-Based Inference. 11 2020. [arXiv:2011.05836](#).
- [99] Auralee Edelen, Nicole Neveu, Matthias Frey, Yannick Huber, Christopher Mayes, and Andreas Adelman. Machine learning for orders of magnitude speedup in multiobjective optimization of particle accelerator systems. *Physical Review Accelerators and Beams*, 23(4):044601, April 2020. Publisher: American Physical Society. URL: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.23.044601>, [doi:10.1103/PhysRevAccelBeams.23.044601](#).
- [100] Daniel Koser, Loyd Waites, Daniel Winklehner, Matthias Frey, Andreas Adelman, and Janet Conrad. Input beam matching and beam dynamics design optimization of the IsoDAR RFQ using statistical and machine learning techniques. *arXiv:2112.02579 [physics]*, 2021. (Submitted to *Frontiers in Physics*). URL: <http://arxiv.org/abs/2112.02579>, [arXiv:2112.02579](#).
- [101] Frederik Van Der Veken, Gabriella Azzopardi, Fred Blanc, Loic Coyle, Elena Fol, Massimo Giovannozzi, Tatiana Pieloni, Stefano Redaelli, Belen Maria Salvachua Ferrando, Michael Schenk, Rogelio Tomas Garcia, and Gianluca Valentino. Machine learning in accelerator physics: applications at the CERN Large Hadron Collider. In *Proceedings of Artificial Intelligence for Science, Industry and Society PoS(AISIS2019)*, volume 372, page 044. SISSA Medialab, July 2020. URL: <https://pos.sissa.it/372/044/>.
- [102] Andreas Adelman. On Nonintrusive Uncertainty Quantification and Surrogate Model Construction in Particle Accelerator Modeling. *SIAM/ASA Journal on Uncertainty Quantification*, 7(2):383–416, January 2019. URL: <https://epubs.siam.org/doi/10.1137/16M1061928>, [doi:10.1137/16M1061928](#).

- [103] Daniel Winklehner, Janet M. Conrad, Devin Schoen, Maria Yampolskaya, Andreas Adelmann, Sonali Mayani, and Sriramkrishnan Muralikrishnan. Order-of-magnitude beam current improvement in compact cyclotrons. *New Journal of Physics*, 24(2):023038, 2022. Publisher: IOP Publishing. doi:10.1088/1367-2630/ac5001.
- [104] Michele Fauci Gianelli, Gregor Kasieczka, Claudius Krause, Ben Nachman, Dalila Salamani, David Shih, and Anna Zaborowska. Fast calorimeter simulation challenge 2022. 2022. URL: <https://calochallenge.github.io/homepage/>.
- [105] Atul Kumar Sinha, Daniele Paliotta, Bálint Máté, Sebastian Pina-Otey, John A. Raine, Tobias Golling, and François Fleuret. SUPA: A Lightweight Diagnostic Simulator for Machine Learning in Particle Physics. 2 2022. arXiv:2202.05012.
- [106] Dirk Merkel. Docker: lightweight linux containers for consistent development and deployment. *Linux journal*, 2014(239):2, 2014.
- [107] *ATLAS liquid-argon calorimeter: Technical Design Report*. Technical design report. ATLAS. CERN, Geneva, 1996. URL: <https://cds.cern.ch/record/331061>.
- [108] William Moses and Valentin Churavy. Instead of rewriting foreign code for machine learning, automatically synthesize fast gradients. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 12472–12485. Curran Associates, Inc., 2020. URL: <https://proceedings.neurips.cc/paper/2020/file/9332c513ef44b682e9347822c2e457ac-Paper.pdf>.
- [109] V. Vassilev, M. Vassilev, A. Penev, L. Moneta, and V. Ilieva. Clad — Automatic Differentiation Using Clang and LLVM. In *Journal of Physics Conference Series*, volume 608 of *Journal of Physics Conference Series*, page 012055, May 2015. doi:10.1088/1742-6596/608/1/012055.