

Lawrence Berkeley National Laboratory

LBL Publications

Title

Comparison of point cloud and image-based models for calorimeter fast simulation

Permalink

<https://escholarship.org/uc/item/3bb77150>

Journal

Journal of Instrumentation, 19(05)

ISSN

1748-0221

Authors

Acosta, Fernando Torales

Mikuni, Vinicius

Nachman, Benjamin

et al.

Publication Date

2024-05-01

DOI

10.1088/1748-0221/19/05/p05003

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Comparison of Point Cloud and Image-based Models for Calorimeter Fast Simulation

Fernando Torales Acosta,^{1,*} Vinicius Mikuni,² Benjamin Nachman,^{1,3} Miguel Arratia,^{4,5} Bishnu Karki,⁴ Ryan Milton,⁴ Piyush Karande,⁶ and Aaron Angerami⁷

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

²*National Energy Research Scientific Computing Center, Berkeley Lab, Berkeley, CA 94720, USA*

³*Berkeley Institute for Data Science, University of California, Berkeley, CA 94720, USA*

⁴*Department of Physics and Astronomy, University of California, Riverside, CA 92521, USA*

⁵*Thomas Jefferson National Accelerator Facility, Newport News, Virginia 23606, USA*

⁶*Computational Engineering Division, Lawrence Livermore National Laboratory, Livermore CA 94550*

⁷*Nuclear and Chemical Science Division, Lawrence Livermore National Laboratory, Livermore, CA 94550*

Score based generative models are a new class of generative models that have been shown to accurately generate high dimensional calorimeter datasets. Recent advances in generative models have used images with 3D voxels to represent and model complex calorimeter showers. Point clouds, however, are likely a more natural representation of calorimeter showers, particularly in calorimeters with high granularity. Point clouds preserve all of the information of the original simulation, more naturally deal with sparse datasets, and can be implemented with more compact models and data files. In this work, two state-of-the-art score based models are trained on the same set of calorimeter simulation and directly compared.

CONTENTS

I. Introduction	1
II. Deep Learning Models	2
III. Detector and Data Descriptions	3
A. Calorimeter Simulation	3
B. Datasets	3
IV. Results	4
V. Conclusion and Outlook	7
Code Availability	9
Acknowledgments	9
References	9

I. INTRODUCTION

Detector simulations are essential tools for data analysis by connecting particle and nuclear physics predictions to measurable quantities. The most precise detector simulations are computationally expensive. This is especially true for calorimeters, which are designed to stop most particles and thus require modeling interactions from the highest accessible energies down to the lowest ones. Well-established experiments typically have bespoke fast simulations that capture the salient aspects of the precise simulations (usually based on GEANT [1–3]) at a fraction of the computational cost. Traditionally,

fast simulations are constructed to reproduce a series of low-dimensional observables. Furthermore, assembling an effective fast simulation is time intensive. If there was a way to build a fast simulation automatically and using the full detector dimensionality, then data analysis at existing and developing experiments could be greatly enhanced.

Deep learning (DL) has been used to build automated and high-dimensional fast simulations (‘surrogate models’) for calorimeters. Starting from Generative Adversarial Networks (GANs) [4] [5–20] and now including Variational Autoencoders [21] [18, 19, 22–24], Normalizing Flows [25] [26–32], and Diffusion Models [33] [34, 35], deep learning based calorimeter simulations have rapidly improved over the last years. They are even starting to be used in actual experimental workflows, such as the ATLAS Collaboration fast simulation [20]. The recent CaloChallenge [36] community comparison showcased the state-of-the-art methods deployed to increasingly granular current and future detectors. As segmented detectors, calorimeters are naturally represented as (possibly irregular) images. Nearly all proposed methods for DL-based calorimeter simulations are based on an image format (fixed grid of pixels). However, these data are unlike natural images in a number of ways, most notably in their sparsity. As such, image-based approaches pioneered in industry may not be the most effective for particle interactions.

Since most cells in a calorimeter image are empty, a more natural representation of these data may be a point cloud. Point clouds are a set of attributes assigned to locations in space. In the calorimeter case, the attribute is energy and the location is the cell coordinates. A calorimeter point cloud would require far fewer numbers to specify than an image representation, since only cells with non-zero energy would be recorded. The main challenges for point cloud models in contrast to image-based approaches is that they must cope with variable-length

* ftoralesacosta@lbl.gov

outputs that respect permutation invariance. With a lag compared to image-based approaches, point cloud generative models for particle/nuclear physics applications have seen a rapid development in recent years [37–42]. However, until recently, these models have never been applied to calorimeter simulations.

The first (and until now, only) publication describing point cloud generative models applied to calorimeters is Ref. [35], which proposed generating GEANT ‘hits’ (deposits of energy) prior to their discretization into cells. This innovative idea enables the separation of material interactions from readout geometry. However, the number of hits vastly exceeds the number of non-zero cells which makes this task difficult. In this paper, we explore point cloud generative models applied directly to cell-level information. In other words, we take calorimeter images and compare state-of-the-art generative models that represent the same inputs as either images or (zero-suppressed) point clouds. As a case study, the two representations are compared using simulations of a high-granularity hadronic calorimeter, similar to the design planned for the ePIC detector at the future Electron-Ion Collider [43–45].

This paper is organized as follows. Section II describes the DL models used for the comparison. Both the image-based and point-cloud representations are generated with diffusion models in order to make the comparison as direct as possible. The simulation of the calorimeter dataset is found in Sec. III. Discussion of the advantages and disadvantages of both representation, as well as numerical results are presented in Sec. IV. The paper ends with conclusions and outlook in Sec. V.

II. DEEP LEARNING MODELS

Generative models for detector simulation aim to precisely emulate physics-based models, like those based on GEANT, but using far less time than the full simulation. With $\mathcal{O}(100)$ detector components, neural network architectures solely based on fully connected layers can efficiently produce high fidelity samples, resulting in surrogate models thousands of times faster than the standard simulation routines [18–20, 27]. For higher detector granularity ($\mathcal{O}(1k)$ - $\mathcal{O}(10k)$), the use of data symmetries becomes crucial to achieve precision. These can be directly included in the model design through dedicated neural network architectures or included in the data pre-processing [26]. For generative models such as normalizing flows, introducing flexible network architectures is often not trivial as the model invertibility and tractable Jacobian of the transformation places a strong restriction on the model design. A second difficulty is to achieve a stable training routine of the surrogate model. At finer granularities, neural network models tend to become larger to accommodate the data complexity, often resulting in unstable training schedules. This issue becomes more prominent in generative models such as variational

autoencoders, where the latent space can vary rapidly, leading to an unstable response of the decoder network, or GANs, where the adversarial training requires careful tuning of the model hyperparameters to achieve a stable training.

Diffusion models are a class of generative neural networks that allow for stable training paired with high flexibility in the model design. Data is slowly perturbed over time using a time parameter $t \in \mathbb{R}$ that determines the perturbation level. The task of the neural network is to approximate the gradients of the log probability of the data, or the score function $\nabla_{\mathbf{x}} p(\mathbf{x}) \in \mathbb{R}^D$, based on data observations $\mathbf{x} \in \mathbb{R}^D$ in the D -dimensional space. This can be approximated by a denoising score-matching strategy [46]. In the implementation used in this paper, data observations $\mathbf{x} \sim p_{\text{data}}(\mathbf{x})$ are perturbed using the kernel $\mathbf{x}_t \sim q(\mathbf{x}_t|\mathbf{x}) = \mathcal{N}(\mathbf{x}_t; \alpha_t \mathbf{x}, \sigma_t^2 \mathbf{I})$, with time-dependent parameters α and σ determining the strength of the perturbation to be applied. In the variance-preserving setting of diffusion processes, $\sigma_t^2 = 1 - \alpha_t^2$. For the time-dependence, a cosine schedule is used such that $\alpha_t = \cos(0.5\pi t)$. The loss function to be minimized is implemented using a *velocity* parameterization:

$$\mathcal{L}_\theta = \mathbb{E}_{\epsilon, t} \|\mathbf{v}_t - \hat{\mathbf{v}}_{t, \theta}\|^2, \quad (1)$$

where the time-dependent network output with trainable parameters θ , $\hat{\mathbf{v}}_{t, \theta}$, is compared with the velocity of the perturbed data at time t , $\mathbf{v}_t \equiv \alpha_t \epsilon - \sigma_t \mathbf{x}$, with $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The score function is then identified as

$$\nabla_{\mathbf{x}} \log \hat{p}_\theta(\mathbf{x}_t) = -\mathbf{x}_t - \frac{\alpha_t}{\sigma_t} \hat{\mathbf{v}}_{t, \theta}(\mathbf{x}_t). \quad (2)$$

The data generation from the trained diffusion models is implemented using the DDIM sampler proposed in Ref. [47] that can be interpreted as an integration rule [48] with update rule specified by:

$$\mathbf{x}_s = \alpha_s \hat{\mathbf{x}}_\theta(\mathbf{x}_t) + \sigma_s \frac{\mathbf{x}_t - \alpha_t \hat{\mathbf{x}}_\theta(\mathbf{x}_t)}{\sigma_t}. \quad (3)$$

For a fair comparison, all diffusion models are trained using the same score-matching strategy and fixed number of 512 time steps during sampling.

The fast point cloud diffusion model (FPCD) follows [41], where a permutation equivariant estimation of the score function is obtained by the combination of a DEEPSSETS [49] architecture with attention layers [50]. During the point cloud simulation, two models are also defined: one that learns the number of non-empty cells, conditioned on the initial energy of the incoming particle, and one model that learns the score function of the normalized point cloud, also conditioned on the energy of the particle to be simulated and the number of hits to be generated. This model is trained on Dataset 1, described in Sec. III B.

The model trained on the image dataset (CALOSCORE) is adapted from [34] with a few modifications. Compared

to the original implementation, the calorimeter simulation task is now broken down in two diffusion models: one that learns only the energy deposits in each layer of the calorimeter, conditioned on the initial energy of the particle to be simulated, and one model that learns to generate normalized voxels per layer, conditioned on the energy deposition in each layer and the initial energy of the particle to be simulated. Additionally, the original U-NET [51] model is combined with attention layers. These changes increase the model expressiveness and the generation fidelity. This model is trained on Dataset 2, described in Sec. III B

III. DETECTOR AND DATA DESCRIPTIONS

A. Calorimeter Simulation

The DD4HEP framework [52] is used to run GEANT simulations of a high-granularity iron-scintillator calorimeter (based on the CALICE-style design [53]), which has dimensions similar to those of the forward hadronic calorimeter in the future ePIC detector (LFH-CAL [44]) at the EIC. Specifically, the sampling structure comprises 0.3 cm scintillator tiles sandwiched between 2.0 cm thick steel plates. It consists of a total of 55 layers. The transverse area of the scintillator is set to 10 cm \times 10 cm, somewhat larger than in Ref. [44]. It adopts a non-projective geometry with tower elements arranged in parallel to the z axis and has its front face at $z=3.8$ m.

1.7 million events of single π^+ particles incident on the center of the calorimeter are simulated. The incident momentum, $P_{\text{Gen.}}$, was generated uniformly in \log_{10} space in the range $1.0 < P_{\text{Gen.}} < 125$ GeV/ c . In order to hit the center of the calorimeter, the pions were generated with a polar angle of $\theta_{\text{Gen.}} = 17^\circ$. Because the detector is symmetric about ϕ , the particles are generated in the range $0^\circ < \phi_{\text{Gen.}} < 360^\circ$. An energy threshold corresponding to 0.3 MeV are used to select hits for further analysis.

B. Datasets

Dataset 1 is the point cloud representation of the GEANT showers, while Dataset 2 represents the same showers using the image representation. Both Dataset 1 and Dataset 2 used in training share the same parent GEANT simulation, such that the fast point cloud diffusion model and the image model are trained on different representations of the same set of calorimeter showers.

Dataset 1 is created by taking the GEANT simulation and converting it to a format based on JetNet data [54], that stores information on jets and their constituents in a zero-suppressed point cloud representation. The GEANT data is stored in files containing two datasets, *clusters*

and *cells*. The cluster dataset contains the $P_{\text{Gen.}}$ of the incident pion, as well as the number of hits in the calorimeter. The cell data is comprised of a constant number of 200 cells per event. Empty cells, or cells with deposited energy below the threshold are masked, with all values set to 0.0, and ignored during training.

The x , y , and z distributions of the GEANT simulation are initially discrete, resulting from the digitization step of the simulation, with values equal to the centers of the cells in each dimension. The point cloud model struggles to learn extremely sharp features, as the score function is not well-defined for discrete inputs. To circumvent this, a uniform smearing within a cell-width is applied to the cells along each dimension to obtain continuous distributions for the final point cloud dataset. This maintains the same distributions at histogram-level when binning according to the cell-width, but yields a point cloud dataset with smooth x , y , and z distributions. Without this smearing, the distributions in x , y , and z resemble a series of delta functions that the point cloud model struggles. The point cloud model is trained on this smeared point cloud representation of the GEANT simulation.

Dataset 2 is created by converting the point cloud dataset into an image format. Images at the original granularity would be too large for the generative model. The calorimeter cells were therefore clustered into groups of 5 along each axis of the detector to create voxels, where $5 \times 5 \times 5$ cells = 1 voxel. Energy in each of the cells making up the voxel were summed and assigned to the final voxel's total energy. The final image format consists of 11×11 voxels. A hit in the voxelized dataset, and referenced in Section IV, is defined as any voxel with energy deposition above threshold.

For the final comparison, generated samples from the point cloud model are voxelized using the same method for Dataset 2. All comparisons are in this image format, at the same resolution of $11 \times 11 \times 11$ voxels per image.

Images representing the full resolution of the calorimeter with $55 \times 55 \times 55$ voxels were not used, as this would result in unmanageably large datasets (see Table I), and would represent the largest calorimeter image training ever done. The point cloud model was trained on the full resolution because point clouds naturally represent the calorimeter at full granularity. Training the point cloud model on this more natural representation is in line with the goal of this work to investigate advantages/disadvantages of two representations of the calorimeter data. It is also for this reason that the generated point cloud distributions are shown separately, while the direct comparisons between models are done in the image representation. Investigating possible advantages of a point-cloud model trained directly on the voxelized dataset is left to future work.

IV. RESULTS

All generated samples along with GEANT are converted to the same image format at the same resolution of $11 \times 11 \times 11$ voxels per event for fair comparison. A variety of distributions are used to evaluate the quality of the generated images. After comparing calorimeter images generated by both models, the point cloud representation of GEANT is compared to the generated samples of the point-cloud model to provide additional insight to the previous image-based comparison. For all comparisons, the Earth mover’s distance (EMD) [55], also known as the 1-Wasserstein distance [56], between generated distributions and GEANT) distributions is calculated. The EMD score a distance-like measure of the dissimilarity between two distributions. It roughly represents the minimum amount of work needed to transform one distribution to another. While this is not the only possible metric, it is a standard and widely-used statistic that was also the main distance deployed in [34], where an image based model was compared to a Wasserstein-GAN. All EMD scores in Figures 2, 3 and 4 are calculated on final voxelized distributions

Figure 1 shows a qualitative assessment of the generative models using the 2-dimensional distribution of the average energy deposition in three layers. All voxels with an expected energy deposition above 0 are populated in both the image and point cloud based models, with very few additional hits. The calorimeter shower will have diverse shapes, as well as different overall distribution of voxels due to the variation of $\phi_{\text{Gen.}}$. The qualitative similarities in each image in Fig 1 indicate that models reproduce the various showers from the training dataset well. Each image contains a ring due to $\theta_{\text{Gen.}}$ being fixed while varying $\phi_{\text{Gen.}}$.

Table I shows the model size, size of each dataset, and time to generate 100k calorimeter showers. The disk size and sample time under the point cloud model are for showers in the point cloud representation. The AUC is obtained from a classifier trained to distinguish the samples of both models only in the voxelized image format. Both models have very good AUC, reasonably close to 0.5, with the image model having the lower AUC. The point cloud model is smaller by a factor of 4 compared to the image based model, and samples events 3 times faster. Lastly, the point cloud dataset requires over 100 times less disk space than the image format at full granularity.

Figure 2 compares the total energy deposited in the calorimeter and total number of calorimeter hits, where a hit is defined as any voxel with energy above threshold. The EMD is also calculated between GEANT and the different generative models.

Both the image-based diffusion model and the point-cloud based diffusion model are in good agreement with GEANT at small deposited energies, deviating no more than 10%. At the highest deposited energies, however, both diffusion models begin to fall away from GEANT,

with the point-cloud model generating less energy, and the image based model generating slightly more energy than GEANT. These trends begin at about 10 GeV, with the point-cloud model deviating slightly earlier. The point-cloud model also shows a slightly higher EMD score than the image based model. The region where the deviations are largest, past 20 GeV of deposited energy are rare, and statistical fluctuations begin to dominate the GEANT distributions.

The number of hits shows a similar trend, though with larger deviations. At a small number of hits, both show good agreement with GEANT, with deviations slightly above 10%. At 15 or more hits, both models begin to deviate well past 10%, with the point cloud model over-sampling the number of hits, and the image based model generating less hits than GEANT.

Figure 3 and 4 shows the average deposited energy x , y , and z -coordinates. Both models struggle in the first and last layers in x and y coordinates, but show good agreement in the middle layers. While the image-based model shows larger deviations in the first and last layers of the calorimeter compared to the point-cloud model, it has an overall lower EMD in both distributions. The two-pronged feature of these distributions is a result of generating the pions at a fix polar angle and varying ϕ . It should be noted that there are little to no hits in the first and last x and y layers of the calorimeter, so even a very small deviation from GEANT will result in a large deviation percentage (bottom panels of Fig. 3 and 4). Similarly, as there are fewer hits towards the back of the detector, deviations increase slightly for the very last layers. However, The z -distributions show both models in very good agreement with the original GEANT predictions, a possible effect of the z -distribution of hits being less dependant on the generated θ and ϕ ranges.

All three distributions show the point cloud samples are systematically lower than the original GEANT distributions. This indicates the point cloud model would benefit from learning the energy per layer directly, as is done in the image model described Sec. II. This difference likely explains why this small bias is observed in the point cloud model, but not in the image model, and is an avenue for improving the point cloud.

Following [26], a classifier was trained to distinguish between generated showers and GEANT showers. The classifier is comprised of two fully connected layers of size 256 using the RELU activation function. The classifier is trained only on vectors of voxelized images of each dataset. The area under the receiver-operator curve (AUC) for the image model was 0.673. The AUC for the point-cloud model was 0.726. Generally, being closer to 0.5, where the classifier is maximally confused, is the target. However the AUC obtained by both models is very promising, as having an AUC even slightly below 1.0 is non-trivial.

A key advantage of the point cloud model is that the distributions at the sub-voxel level can be shown. The point cloud model already simulates the data at the orig-

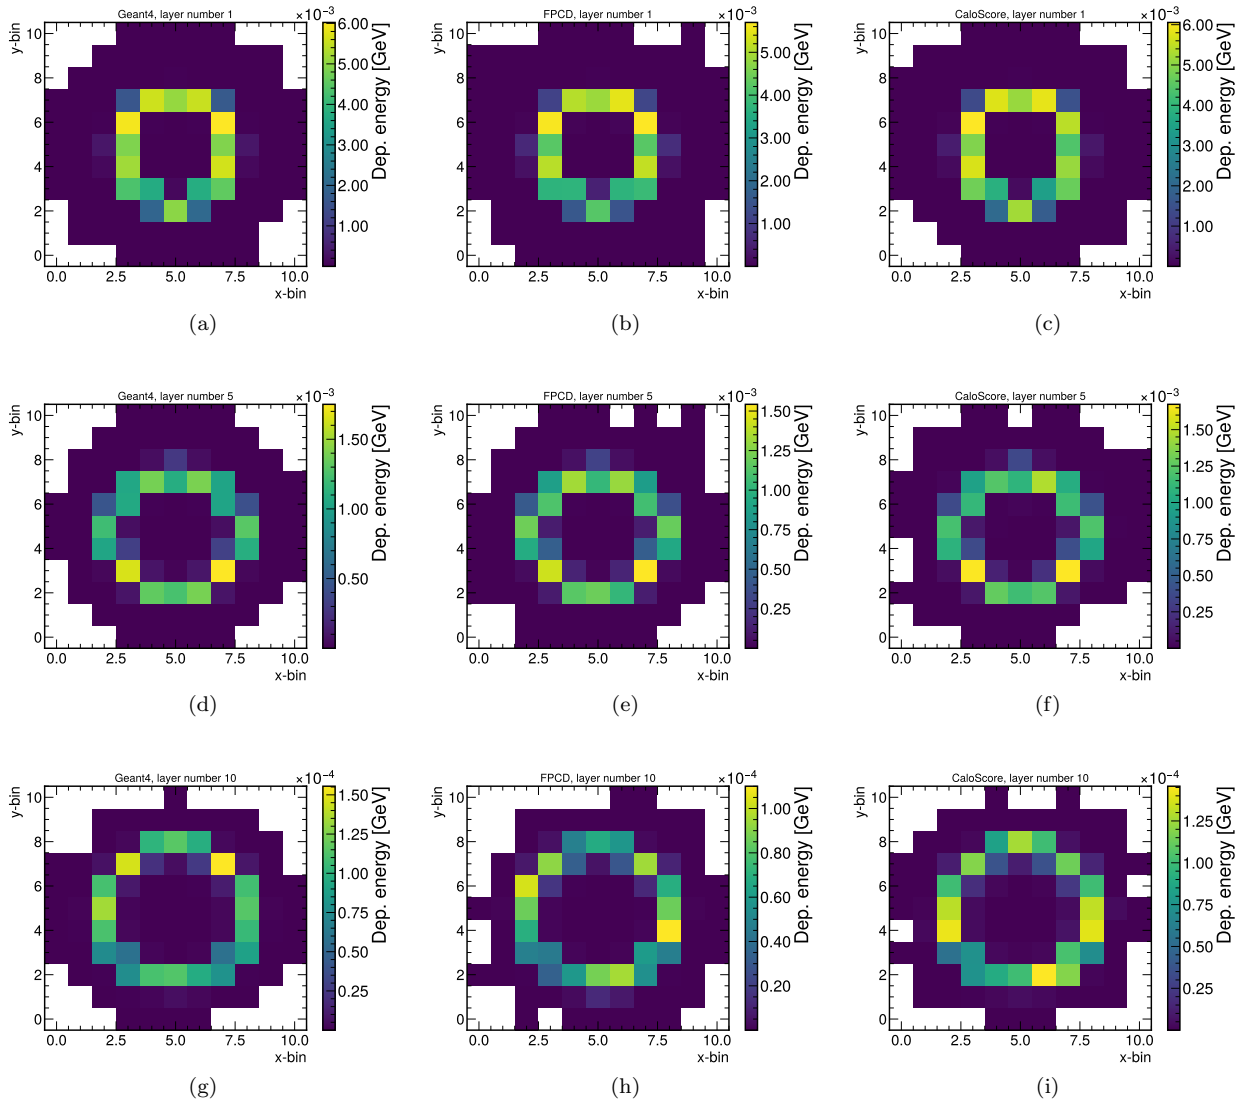


FIG. 1. The 2-dimensional distribution of the mean deposited energy in the 1st, 5th, and 10th voxelized layer of the calorimeter. The first column is the original Geant simulation. The second column is the fast point-cloud based diffusion model (FPCD), and the 3rd column is the image-based model (CALOSCORE).

Model	# Parameters	Disk Size (Full)	Sample Time	AUC
Image	2,572,161	1016MB (62GB)	8036.19s	0.673
Point Cloud	620,678	509 MB	2631.41s	0.726

TABLE I. Comparison of model size, size of data representation on disk, generation time, and AUC of the same classifier trained to distinguish between the model and the original GEANT showers. All comparisons are done for 100k calorimeter showers. The all results in the image row were obtained with the scaled down, $11 \times 11 \times 11$ voxel images, however the disk size of the image dataset at full granularity is shown in parenthesis.

inal granularity of the calorimeter, and voxelization is only necessary for the image representation. The original output of the point cloud model is compared to the continuous (or smeared) GEANT distributions. Figure 5 shows the number of hits in the point cloud representation of the calorimeter showers. In the point-cloud rep-

resentation, a hit is defined as any *cell* that has a energy deposited above threshold.

The point-cloud model reproduces the total number of cell hits well, much better than the voxel hit distribution, shown in Fig. 2. This may indicate that while the point cloud model is overall similar to GEANT in both represen-

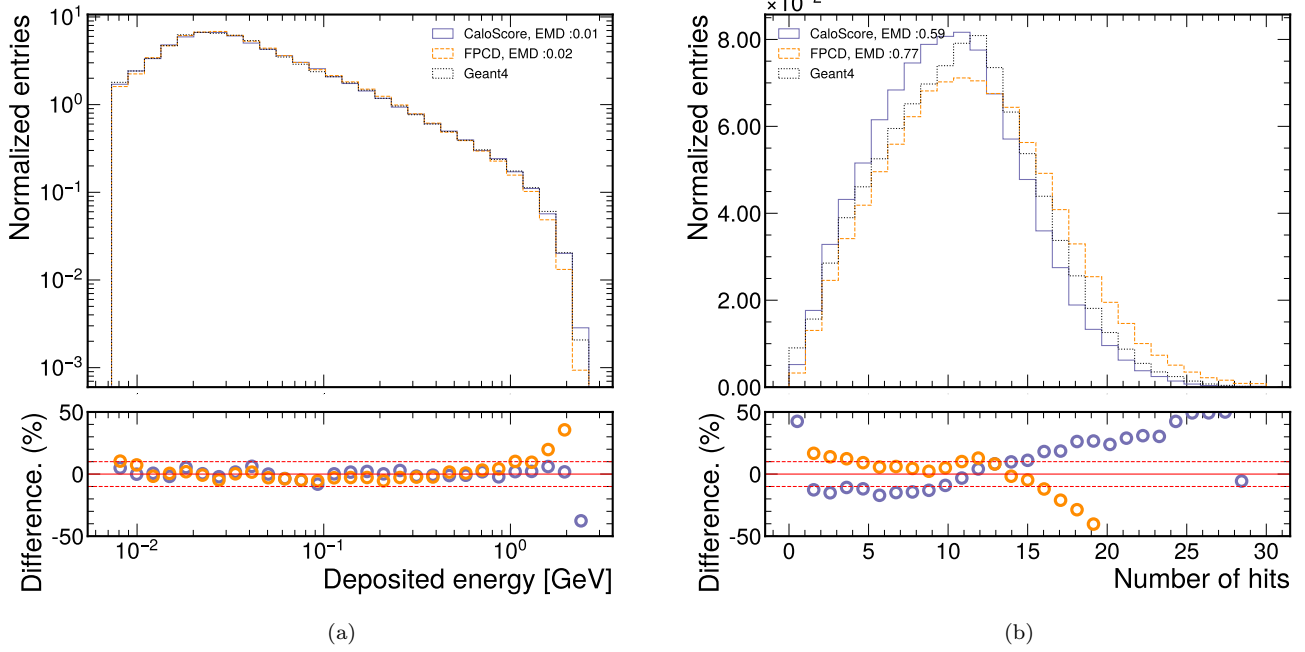


FIG. 2. Sum of all voxel energies of the image representation of FPCD model, shown in orange, and the image based model, shown in grey-blue. The parent GEANT distributions are shown as a dotted black line in the top panels. The dashed red lines in the bottom panel of each figure represent the 10% deviation interval of the generated samples from the original GEANT simulation. The earth mover's distance (EMD) between each distribution and the GEANT distribution is also shown.

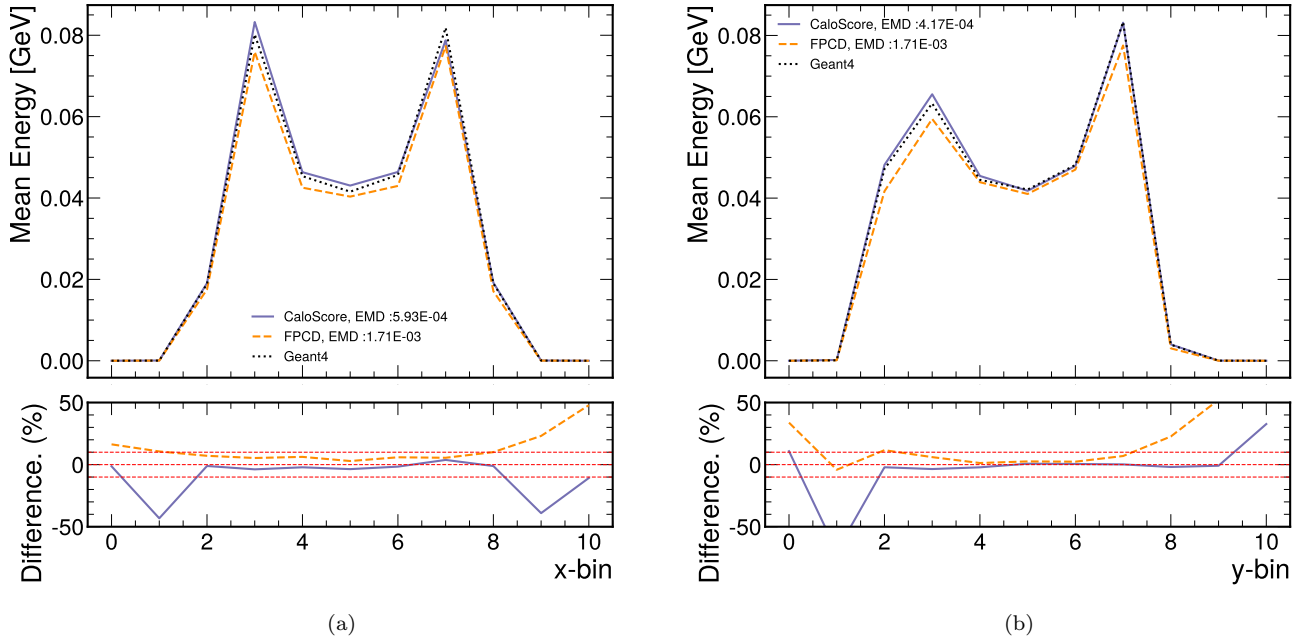


FIG. 3. Comparison of the average deposited energies in the x (left), y (center), and z (right) coordinates. The dashed red lines in the bottom panel of each figure represent the 10% deviation interval of the generated samples from the original GEANT simulation.

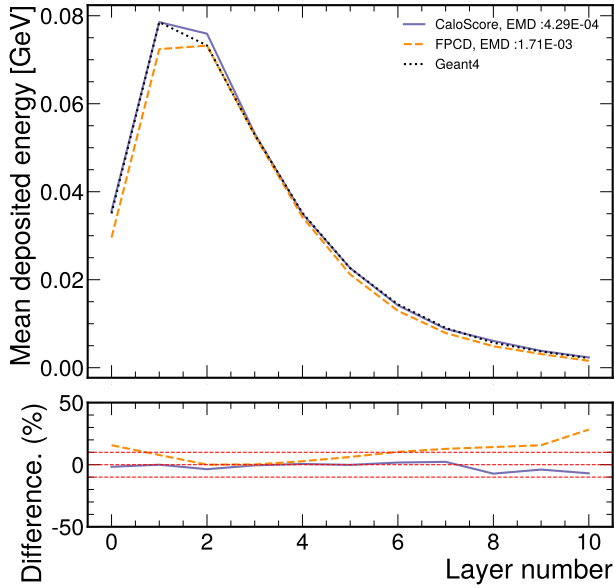


FIG. 4. Comparison of the average and z coordinate. The dashed red lines in the bottom panel of each figure represent the 10% deviation interval of the generated samples from the original GEANT simulation.

tations, small deviations in point cloud distributions can be summed into larger deviations during the voxelization process, where 125 individual cells are combined into a single voxel. However, there is a large symmetry group under which mismodelings in the bigger space may not affect the modeling in the coarser space, so further investigation is needed. However, the very good agreement with GEANT in the number of cell hits and degrading agreement in the number of voxel hits indicates that the first diffusion model of the point cloud model architecture is performing well, while the second model, responsible for sampling the cell distributions, would likely benefit from additional tuning.

Similar conclusions can be derived from Fig. 6, show the generated point samples at the full detector granularity and in good agreement with GEANT. Fig. 6 shows the average x , y , and z coordinate distributions, as well as the cell $\log_{10}E$ distribution in the point representation. Again, there are larger relative deviations in the first and last layers in x , y , and z , coordinates where there are very few hits, just as in the image representation. However, there is very good agreement with the GEANT simulation in layers containing a reasonable number of hits.

V. CONCLUSION AND OUTLOOK

In this paper, we make the first direct comparison between two score based generative models using either

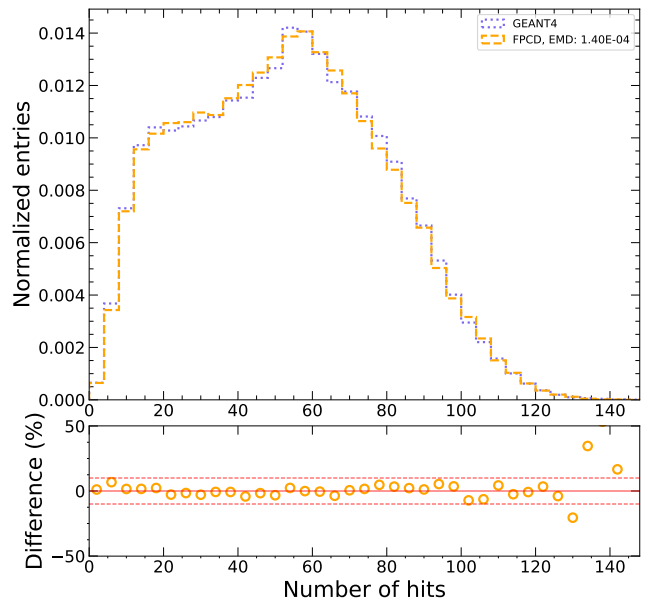


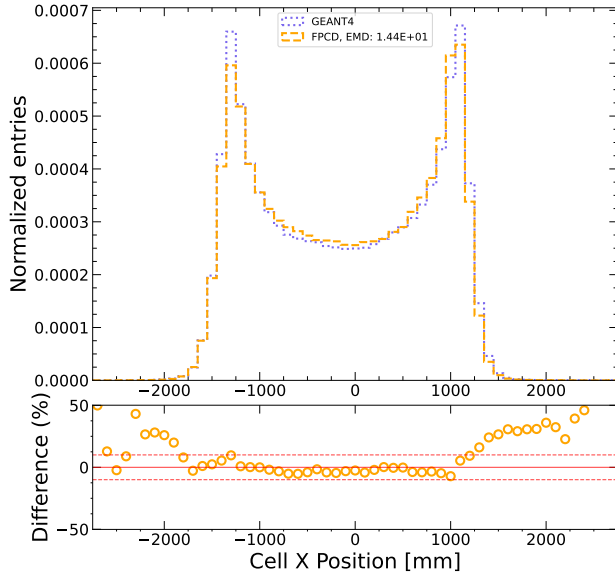
FIG. 5. The total number of hits in the point cloud representation of calorimeter showers, at full granularity. The dashed red lines in the bottom panel of each figure represent the 10% deviation interval of the generated samples from the original Geant simulation.

images or point clouds as representations of the same training data. We use GEANT calorimeter simulations of a high-granularity hadronic calorimeter. Both models perform well for most distributions, with very similar AUCs, but the image-based diffusion model invariably has a lower EMD in each comparison to GEANT.

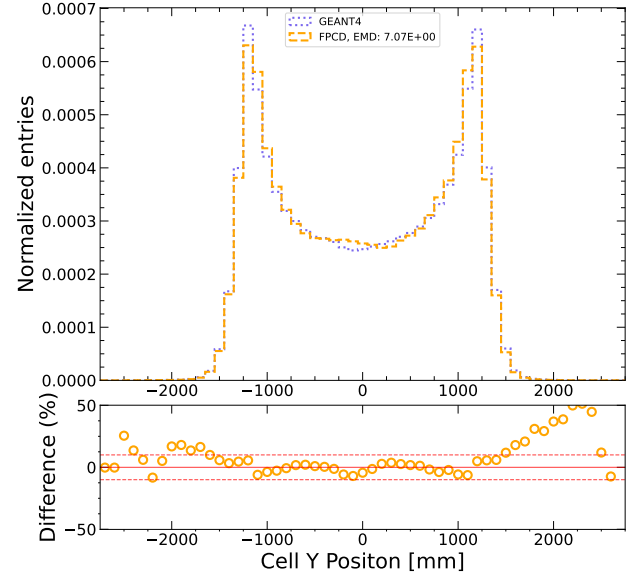
Overall, the performance of the point-cloud diffusion model is very close to the image model. This is despite the point cloud model being disadvantaged in this work in a few important ways.

First, the calorimeter showers from the FPCD model are closest to GEANT in the point cloud representation at the full calorimeter granularity, as shown in Fig. 5 and 6. But it is later voxelized for comparison. This may compound mismodeling during the voxelization, however further investigation is needed.

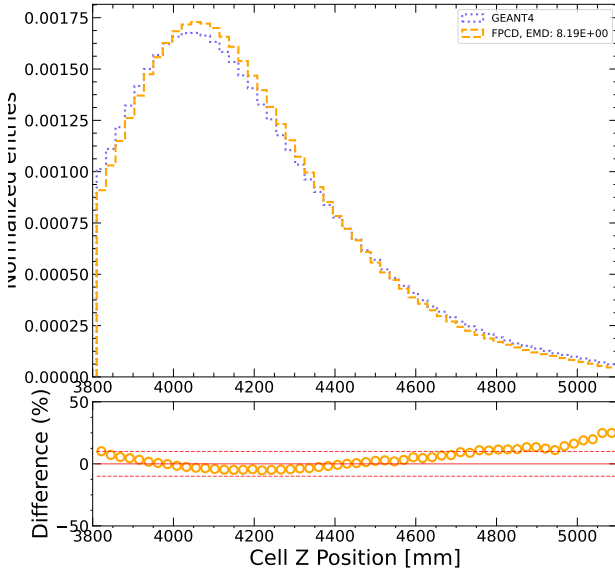
Second, the point cloud model is adapted from a model architecture initially designed for jet data from the JetNet datasets. While the high-level structure of the datasets are very similar, the data itself are quite different. For example, the first diffusion model making up the point cloud model was initially much larger, as predicting the jet multiplicity is in general a more difficult problem than the number of non-empty cells in a calorimeter shower. Reducing the size of the first diffusion model of the point cloud model architecture had no impact on performance while speeding up training. The



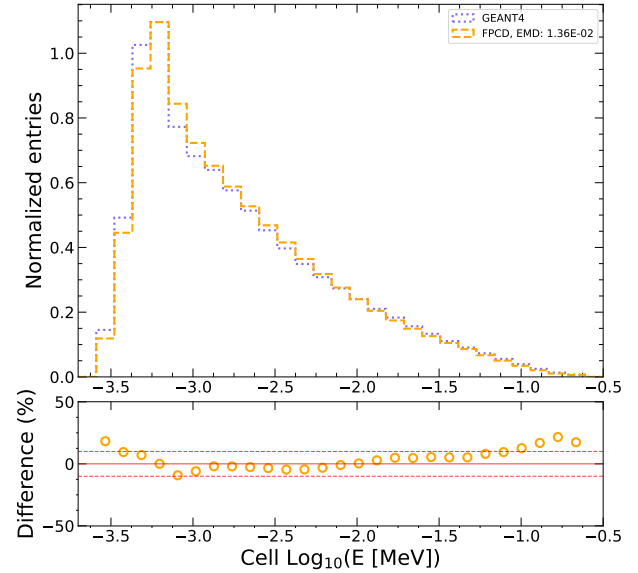
(a)



(b)



(c)



(d)

FIG. 6. Comparison of the average cell x (top left), y (top right), z (bottom left) and $\log_{10}E$ (bottom right) distributions in the point cloud datasets. Each distribution is binned according to the cell-width to show the full granularity of the detector. The dashed red lines in the bottom panel of each figure represent the 10% deviation interval of the generated samples from the original GEANT simulation.

second diffusion model making up the point cloud model architecture that is responsible for sampling the cell x , y , z , and E was directly adapted from [41]. Further tuning of the point cloud model, particularly the cell-model can likely close the small remaining gap in performance. The image model, in contrast, is based on CALOSCORE, which was tuned specifically for calorimeter showers.

Lastly, the image-based model uses the energy deposition in each layer in addition to the generated particle momentum to condition the second diffusion model making up its architecture. The second diffusion model making up the point cloud model is solely conditioned on the generated particle momentum. This might explain why the point cloud model has systematically lower mean energy distributions (see Fig. 3 and 4) compared to both GEANT and the image based model.

These potential sources of improvement in the point cloud model should not detract from its already very reasonable performance, deviating from GEANT more 10% only in the sparsest of layers, where the image based model also struggles. At the same time, the point cloud model offers several advantages over the image model.

First, the sheer size of the data. The point cloud data saved to HDF5 files is a factor of 100 times smaller using the same zlib compression as the image based dataset at full granularity, with no voxelization. As calorimeters continue to increase in granularity, this difference will only increase.

Second, information is lost during voxelization process; cell hits with the same x , y , z coordinates, but different energies are summed over in the image representation. This is true even if images are produced at the full granularity of the calorimeter, where hits within the single cells are summed over. This means that voxelized datasets cannot naturally be reverted back to a point cloud representation.

Additionally, as was showed in this work, the generated point clouds can be voxelized afterwards, or converted into other representations that better fit specific use cases.

This work establishes a benchmark for future research

on generative models, offering valuable insights into the challenges of modeling hadronic showers in highly granular calorimeters using image-based techniques, while also exploring the potential of point-cloud methods. The current advantages of point clouds, in combination with improvements to close the remaining performance gap described earlier, will likely make point cloud based models a clear choice for highly granular calorimeters. This work should serve as a reference for studies utilizing future calorimeters based on the CALICE design, including those intended for use in CMS at the LHC and ePIC at the EIC.

CODE AVAILABILITY

The code used to produce the point cloud results shown in this document are available at https://github.com/ftoralesacosta/GSGM_for_EIC_Calo. The code for the image based model and comparisons of images is available at <https://github.com/ViniciusMikuni/Calo4EIC>. Example GEANT4 datasets and generated samples are available at <https://zenodo.org/record/8128598>.

ACKNOWLEDGMENTS

We acknowledge support from DOE grant award number DE-SC0022355. This research used resources from the LLNL institutional Computing Grand Challenge program and the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC award HEP-ERCAP0021099. M.A acknowledges support through DOE Contract No. DE-AC05-06OR23177 under which Jefferson Science Associates, LLC operates the Thomas Jefferson National Accelerator Facility. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract No. DE-AC52-07NA27344.

-
- [1] 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.
 - [2] J. Allison et al. Geant4 developments and applications. *IEEE Transactions on Nuclear Science*, 53(1):270–278, 2006.
 - [3] J. Allison et al. Recent developments in geant4. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 835:186–225, 2016.
 - [4] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks. *arXiv e-prints*, page arXiv:1406.2661, June 2014.
 - [5] 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.
 - [6] 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. D*, 97(1):014021, 2018.
 - [7] 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.
- [8] 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.
- [9] 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.
- [10] Dawit Belayneh et al. Calorimetry with Deep Learning: Particle Simulation and Reconstruction for Collider Physics. 12 2019.
- [11] Sofia Vallecorsa, Federico Carminati, and Gulrukh Khatkhat. 3D convolutional GAN for fast simulation. *Proceedings, 23rd International Conference on Computing in High Energy and Nuclear Physics (CHEP 2018): Sofia, Bulgaria, July 9-13, 2018*, 214:02010, 2019.
- [12] C. Ahdida et al. Fast simulation of muons produced at the SHiP experiment using Generative Adversarial Networks. 2019.
- [13] Viktoria Chekalina, Elena Orlova, Fedor Ratnikov, Dmitry Ulyanov, Andrey Ustyuzhanin, and Egor Zakharov. Generative Models for Fast Calorimeter Simulation.LHCb case. *CHEP 2018*, 2018.
- [14] F. Carminati, A. Gheata, G. Khatkhat, P. Mendez Lorenzo, S. Sharan, and S. Vallecorsa. Three dimensional Generative Adversarial Networks for fast simulation. *Proceedings, 18th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2017): Seattle, WA, USA, August 21-25, 2017*, 1085(3):032016, 2018.
- [15] S. Vallecorsa. Generative models for fast simulation. *Proceedings, 18th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2017): Seattle, WA, USA, August 21-25, 2017*, 1085(2):022005, 2018.
- [16] Pasquale Musella and Francesco Pandolfi. Fast and Accurate Simulation of Particle Detectors Using Generative Adversarial Networks. *Comput. Softw. Big Sci.*, 2(1):8, 2018.
- [17] Kamil Deja, Tomasz Trzcinski, and Lukasz Graczykowski. Generative models for fast cluster simulations in the TPC for the ALICE experiment. *Proceedings, 23rd International Conference on Computing in High Energy and Nuclear Physics (CHEP 2018): Sofia, Bulgaria, July 9-13, 2018*, 214:06003, 2019.
- [18] Deep generative models for fast photon shower simulation in ATLAS. 10 2022.
- [19] Deep generative models for fast shower simulation in ATLAS. *ATL-SOFT-PUB-2018-001*, Jul 2018.
- [20] Georges Aad et al. AtlFast3: The Next Generation of Fast Simulation in ATLAS. *Comput. Softw. Big Sci.*, 6(1):7, 2022.
- [21] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv e-prints*, page arXiv:1312.6114, December 2013.
- [22] 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.
- [23] 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.
- [24] Sascha Diefenbacher, Engin Eren, Frank Gaede, Gregor Kasieczka, Anatolii Korol, Katja Krüger, Peter McKeown, and Lennart Rustige. New Angles on Fast Calorimeter Shower Simulation. 3 2023.
- [25] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1530–1538, Lille, France, 07–09 Jul 2015. PMLR.
- [26] Claudius Krause and David Shih. CaloFlow: Fast and Accurate Generation of Calorimeter Showers with Normalizing Flows. 6 2021.
- [27] Claudius Krause and David Shih. CaloFlow II: Even Faster and Still Accurate Generation of Calorimeter Showers with Normalizing Flows. 10 2021.
- [28] Matthew R. Buckley, Claudius Krause, Ian Pang, and David Shih. Inductive CaloFlow. 5 2023.
- [29] Claudius Krause, Ian Pang, and David Shih. CaloFlow for CaloChallenge Dataset 1. 10 2022.
- [30] Sascha Diefenbacher, Engin Eren, Frank Gaede, Gregor Kasieczka, Claudius Krause, Imahn Shekhzadeh, and David Shih. L2LFlows: Generating High-Fidelity 3D Calorimeter Images. 2 2023.
- [31] Jesse C. Cresswell, Brendan Leigh Ross, Gabriel Loaizaganem, Humberto Reyes-Gonzalez, Marco Letizia, and Anthony L. Caterini. CaloMan: Fast generation of calorimeter showers with density estimation on learned manifolds. In *36th Conference on Neural Information Processing Systems*, 11 2022.
- [32] Junze Liu, Aishik Ghosh, Dylan Smith, Pierre Baldi, and Daniel Whiteson. Generalizing to new geometries with Geometry-Aware Autoregressive Models (GAAMs) for fast calorimeter simulation. 5 2023.
- [33] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. *arXiv e-prints*, page arXiv:2011.13456, November 2020.
- [34] Vinicius Mikuni and Benjamin Nachman. Score-based generative models for calorimeter shower simulation. *Phys. Rev. D*, 106(9):092009, 2022.
- [35] Erik Buhmann, Sascha Diefenbacher, Engin Eren, Frank Gaede, Gregor Kasieczka, Anatolii Korol, William Korcari, Katja Krüger, and Peter McKeown. CaloClouds: Fast Geometry-Independent Highly-Granular Calorimeter Simulation. 5 2023.
- [36] Fast calorimeter simulation challenge 2022.
- [37] 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.
- [38] Erik Buhmann, Gregor Kasieczka, and Jesse Thaler. EPiC-GAN: Equivariant Point Cloud Generation for Particle Jets. 1 2023.
- [39] Benno Käch, Dirk Krücker, Isabell Melzer-Pellmann, Moritz Scham, Simon Schnake, and Alexi Verney-Provatas. JetFlow: Generating Jets with Conditioned and Mass Constrained Normalising Flows. 11 2022.
- [40] Rob Verheyen. Event Generation and Density Estimation with Surjective Normalizing Flows. 5 2022.

- [41] Vinicius Mikuni, Benjamin Nachman, and Mariel Pettee. Fast Point Cloud Generation with Diffusion Models in High Energy Physics. 4 2023.
- [42] Matthew Leigh, Debajyoti Sengupta, Guillaume Quétant, John Andrew Raine, Knut Zoch, and Tobias Golling. PC-JeDi: Diffusion for Particle Cloud Generation in High Energy Physics. 3 2023.
- [43] R. Abdul Khalek et al. Science Requirements and Detector Concepts for the Electron-Ion Collider: EIC Yellow Report. *Nucl. Phys. A*, 1026:122447, 2022.
- [44] F. Bock et al. Design and Simulated Performance of Calorimetry Systems for the ECCE Detector at the Electron Ion Collider. 7 2022.
- [45] Miguel Arratia et al. A high-granularity calorimeter insert based on SiPM-on-tile technology at the future Electron-Ion Collider. *Nucl. Instrum. Meth. A*, 1047:167866, 2023.
- [46] Pascal Vincent. A Connection Between Score Matching and Denoising Autoencoders. *Neural Computation*, 23(7):1661–1674, 07 2011.
- [47] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *CoRR*, abs/2010.02502, 2020.
- [48] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022.
- [49] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, and Alexander Smola. Deep Sets. *arXiv e-prints*, page arXiv:1703.06114, March 2017.
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017.
- [51] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing.
- [52] Markus Frank, F. Gaede, C. Grefe, and P. Mato. DD4hep: A Detector Description Toolkit for High Energy Physics Experiments. *J. Phys. Conf. Ser.*, 513:022010, 2014.
- [53] Design, Construction and Commissioning of a Technological Prototype of a Highly Granular SiPM-on-tile Scintillator-Steel Hadronic Calorimeter. 9 2022.
- [54] Anonymous. Jetnet, May 2021.
- [55] Y. Rubner, C. Tomasi, and L.J. Guibas. A metric for distributions with applications to image databases. In *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*, pages 59–66, 1998.
- [56] E. Levina and P. Bickel. The earth mover’s distance is the mallows distance: some insights from statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 251–256 vol.2, 2001.