

Lawrence Berkeley National Laboratory

LBL Publications

Title

Score-based Generative Models for Calorimeter Shower Simulation

Permalink

<https://escholarship.org/uc/item/0pm4v4vg>

Authors

Mikuni, Vinicius
Nachman, Benjamin

Publication Date

2022-06-17

Score-based Generative Models for Calorimeter Shower Simulation

Vinicius Mikuni^{1,*} and Benjamin Nachman^{2,3,†}

¹*National Energy Research Scientific Computing Center, Berkeley Lab, Berkeley, CA 94720, USA*

²*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

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

Score-based generative models are a new class of generative algorithms that have been shown to produce realistic images even in high dimensional spaces, currently surpassing other state-of-the-art models for different benchmark categories and applications. In this work we introduce CALOSCORE, a score-based generative model for collider physics applied to calorimeter shower generation. Three different diffusion models are investigated using the Fast Calorimeter Simulation Challenge 2022 dataset. CALOSCORE is the first application of a score-based generative model in collider physics and is able to produce high-fidelity calorimeter images for all datasets, providing an alternative paradigm for calorimeter shower simulation.

I. INTRODUCTION

Detailed detector simulations are an essential component of data analysis in particle and nuclear physics. These simulations are used to fold particle-level predictions for comparisons with data and are used to unfold data for detector-effects in order to compare with predictions and other experiments. Simulations are also a critical input for the design of future experiments.

The most widely used detector simulations are built on the program GEANT [1–3]. Achieving precision requires significant computing time as propagating particles through materials results in a large number of secondary particles each undergoing electromagnetic and/or nuclear interactions. For this reason, the most complex detectors to emulate are calorimeters, whose purpose is to stop particles and measure the deposited energy. An $\mathcal{O}(1)$ fraction of all computing in HEP goes towards simulating particle propagation inside dense materials with GEANT.

The experiments at the Large Hadron Collider (LHC) generate billions of events per run, each of which has hundreds to thousands of individual calorimeter showers. Within the experiments’ computing budgets, it is not possible to run GEANT-based (‘full’) simulation for all events. Therefore, all of the experiments have developed fast simulation methods that replace physics models with simpler parametric models that are tuned to the full simulation. The fast simulation models are constructed with relatively few parameters in order to facilitate efficient optimization and validation. This fundamentally limits their precision, in particular in the ability to model complex correlations in high-dimensions. These correlations may also not be explicitly part of an optimization that uses only a relatively small number of one-dimensional observables.

Deep learning offers a complementary approach to engineered parametric models. Flexible neural networks

are used to transform random numbers (called a *latent space* in machine learning) into structured data. So far, there have been three main strategies for deep generative modeling. Generative Adversarial Networks (GANs) [4] optimize the generator network by means of an auxiliary network (‘discriminator’) that tries to classify generated examples from real examples. Variational Autoencoders (VAEs) [5] learn a stochastic map from the data space to a latent space and back, preserving the statistics of the latent space and data space. Normalizing Flows (NFs) [6] use invertible transformations so that the probability density can be computed and the generator is optimized using the log likelihood. A number of deep generative models have been proposed for emulating calorimeter showers and other particle detectors [7–45].

The first proposals for generating calorimeter showers with deep generative models used GANs (starting with CALOGAN [7–9]). The evaluation of GANs is fast and there are no constraints on the form of the generator function. However, GAN optimization is challenging as it is a minimax problem due to the competition between the generator and discriminator networks. Furthermore, GANs are known to suffer from *mode collapse* where the generator learns to produce only a subset of the possible showers. Despite these challenges, the ATLAS Collaboration become the first experiment to replace part of their fast simulation with a GAN [46] and are already using it to generate billions of events for Run 3. Other experiments are also exploring the use of GANs for detector simulation [12, 18, 20].

Since the first GAN studies, there have been a number of innovations to improve the precision of deep generative calorimeter simulation. This includes variations/modifications to the GAN setup including Wasserstein GANs [20, 47] to help with training stability and mode collapse and refining networks [25, 26] to correct the spectra of an initial generative model. Beyond GANs, recent work with NFs has shown great promise [44, 45]. Normalizing flows tend to be robust to mode collapse and are minimized with a convex loss function. Furthermore, the normalization of the resulting generative model seems to have beneficial regularization properties for the

* vmikuni@lbl.gov

† bpnachman@lbl.gov

training. The authors of the CALOFLOW model [44, 45] showed that a post-hoc classifier struggled to distinguish showers generated from their neural network and showers from GEANT 4, a critical milestone in this field.

While NFs show great promise for fast calorimeter simulation, they have difficulty scaling to higher-dimensional datasets. Such datasets are important for the upgraded CMS forward calorimeter [48] as well as ultra-fine calorimeters proposed for future detectors [49]. Two of the three datasets in the recent *Fast Calorimeter Simulation Challenge 2022* [46, 50–52] have dimensionality at least an order of magnitude beyond what has been studied with NFs. The NF training time for these data is prohibitive and would require significant research and development.

In this paper, we examine the potential of a new class of deep generative algorithms named score-based generative models [53–55]. Like NFs, score-based models minimize a convex loss function with a single generator network that also provides access to the full data likelihood after training (see App. A). However, unlike NFs, the gradient of the data density (the ‘score’) is learned instead of the density. This choice introduces more flexibility to the network architecture, since the Jacobian of the transformation does not need to be computed during training. This additional flexibility, contrary to normalizing flows, also allows the introduction of bottleneck layers (layers with fewer neurons than the previous layer), greatly reducing the number of trainable parameters and improving the scalability of the model. We demonstrate this explicitly on the Fast Calorimeter Simulation Challenge 2022 datasets.

This paper is organized as follows. Sec. II introduces how deep generative models can be constructed using the likelihood gradients instead of likelihoods directly. Different choices of drift and diffusion functions investigated in this work are introduced in Sec. III. Section IV describes how new samples are generated from the trained model. Description of the Fast Calorimeter Simulation Challenge 2022 datasets and network architecture are presented in Secs. V and VI. Finally, numerical results are discussed in Sec. VII. The paper ends with conclusions and outlook in Sec. VIII.

II. GENERATIVE MODELS USING GRADIENTS OF THE DATA

Our approach most closely follows Ref. [55]. Before providing the technical details, we briefly describe the key method components. First, a neural network is learned to approximate the data score, $\nabla_x \log p_{\text{data}}$. In many high energy physics examples, including calorimeter simulation, we do not have access to p_{data} analytically - we can only sample from it. Without access to p_{data} , we cannot use a loss function like the mean squared error to directly approximate (‘match’) the score. NFs circumvent this problem by maximizing the likelihood of

the data. There is no analog of maximum likelihood for score matching and so a different innovation is required.

For data that are smeared, it can be shown that matching the score of the smeared data is equivalent to matching the score of the smearing function [56]. For data purposefully smeared by an analytically tractable smearing function (e.g. a Gaussian), this means that all of components required to compute the loss function are known. The methods explored in this paper make use of purposeful smearing, where the amount of smearing is increased/decreased to estimate the density or generate samples, respectively. A schematic of the idea is shown in Fig. 1.

A stochastic process that continuously corrupts data is described by the following stochastic differential equation (SDE):

$$dx = f(x, t)dt + g(t)dw. \quad (1)$$

The initial data $x(t = 0) := x_0 \in \mathbb{R}^d$ sampled from the distribution p_{data} evolve over time given a set of drift and diffusion coefficients $f(x, t) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $g(t) : \mathbb{R} \rightarrow \mathbb{R}$, respectively. The Wiener process $w(t)$ is indexed by the time parameter $t \in [0, 1]$. The goal of the generative model is to reverse this process, generating new observations starting from a noise distribution and solving the reverse SDE defined as:

$$dx = [f(x, t) - g(t)^2 \nabla_x \log p_t(x)]dt + g(t)d\bar{w}, \quad (2)$$

where \bar{w} is the Wiener process in the reverse time direction [57] and the time-dependent gradient $\nabla_x \log p_t(x)$ is named the *data score*. For fixed choices of the drift and diffusion coefficients, the only term in Eq. 2 that needs to be estimated is the time-dependent score function. However, since the score-function is not tractable, we use neural networks that aim to minimize the difference

$$\frac{1}{2} \mathbb{E}_{p_{\text{data}}(x)} \left[\|s_\theta(x) - \nabla_x \log p_{\text{data}}\|_2^2 \right], \quad (3)$$

with $s_\theta(x)$ the output of a network with trainable parameters θ . Since the true score function is not known, Eq. 3 cannot be readily computed. Instead, a denoising score matching [56] strategy is used. In this strategy, instead of learning the score function of the data, we aim to learn the score function of data that have been perturbed with a known perturbation function since it is sufficient to match the score of the perturbation function. Note that while matching the score of the smeared data requires computing an expectation value over the smeared data, matching the score of the smearing function requires computing an expectation value over the smeared data and the smeared data.

Given a Gaussian perturbation kernel $p_\sigma(\tilde{x}|x) := \mathcal{N}(x, \sigma^2)$ and $p_\sigma(\tilde{x}) := \int p_{\text{data}}(x)p_\sigma(\tilde{x}|x)dx$, the probability density of the perturbed data, the loss function minimized during training is:

$$\frac{1}{2} \mathbb{E}_{p_\sigma(\tilde{x}|x)p_{\text{data}}} \left[\|s_\theta(\tilde{x}) - \nabla_{\tilde{x}} \log p_\sigma(\tilde{x}|x)\|_2^2 \right]. \quad (4)$$




FIG. 1. The score-based generative model is trained using a diffusion process that slowly perturbs the data. Generation of new samples is carried out by reversing the diffusion process using the learned score-function, or the gradient of the data density. For different time-steps, we show the distribution of deposited energies versus generated particle energies (top) and the energy deposition in a single layer of a calorimeter (bottom), generated with our proposed CALOSCORE model.

The advantage of this strategy is that we can directly estimate the last term in Eq. 4, since:

$$\nabla_{\tilde{x}} \log p_{\sigma}(\tilde{x}|x) = \frac{x - \tilde{x}}{\sigma^2} \sim \frac{\mathcal{N}(0, 1)}{\sigma} \quad (5)$$

The time component can be made explicit by rewriting the loss function in Eq. 4 as:

$$\frac{1}{2} \mathbb{E}_t \mathbb{E}_{p(x_t|x_0)p(x_0)} \left[\lambda(t) \|\mathcal{s}_{\theta}(x, t) - \nabla_{x_t} \log p_t(x_t|x_0)\|_2^2 \right]. \quad (6)$$

The weighting function $\lambda(t) : \mathbb{R} \rightarrow \mathbb{R}$ ensures the loss function has the same order of magnitude at all times and is chosen to be inversely proportional to $\mathbb{E} \left[\|\nabla_{x_t} \log p_t(x_t|x_0)\|_2^2 \right]$. When the drift coefficient $f(x, t)$ is chosen to be an affine function of x , the resulting perturbation kernel is always Gaussian [58] and can be chosen such that both mean and variance are known in closed form, making Eq. 6 efficient to compute during training.

III. CHOICE OF DRIFT AND DIFFUSION COEFFICIENTS

In this work we investigate three different choices of drift and diffusion coefficients that result in perturbation kernels that are easy to calculate in closed form. The first SDE, initially proposed in [53], is defined as:

$$dx = \sqrt{\frac{d[\sigma^2(t)]}{dt}} dw. \quad (7)$$

The parameter $\sigma(t) = \sigma_{min} \left(\frac{\sigma_{max}}{\sigma_{min}} \right)^t$ is defined with $\sigma_{min} = 0.01$ and $\sigma_{max} = 50$ to ensure $x(1) \sim \mathcal{N}(0, \sigma_{max}^2)$ is independent from $x(0)$. Since the time-dependent variance of the resulting perturbation explodes when $t \rightarrow \infty$, this SDE is often referred to variance exploding (VE) SDE.

The second SDE is a continuous version of the discrete perturbation introduced in [54], defined as:

$$dx = -\frac{1}{2}\beta(t)xdt + \sqrt{\beta(t)}dw. \quad (8)$$

The parameter $\beta(t) = \beta_{min} + t(\beta_{max} - \beta_{min})$ with $\beta_{min} = 0.1$ and $\beta_{max} = 20$ is used, resulting in $x(1) \sim \mathcal{N}(0, 1)$. The variance of this process is fixed to one when

the initial distribution also has unit variance, hence referred to as variance preserving (VP) SDE.

The last SDE, introduced in [55], is defined as:

$$dx = -\frac{1}{2}\beta(t)xdt + \sqrt{\beta(t)(1 - e^{-2\int_0^t \beta(s)ds})}dw. \quad (9)$$

When the parameter $\beta(t)$ is chosen to be identical as the one used in Eq. 8, the variance of the stochastic process is always smaller than the variance of the VP SDE, hence receiving the name subVP. Similarly to Eq. 8, $x(1)$ also follows a normal distribution.

The resulting perturbation kernels induced by each SDE are listed in Tab. I.

TABLE I. Perturbation kernel induced by different SDE choices.

SDE	Perturbation kernel
VE	$\mathcal{N}(x(0), \sigma^2(t) - \sigma^2(0))$
VP	$\mathcal{N}(x(0)e^{-\frac{1}{2}\int_0^t \beta(s)ds}, 1 - e^{-\int_0^t \beta(s)ds})$
subVP	$\mathcal{N}(x(0)e^{-\frac{1}{2}\int_0^t \beta(s)ds}, (1 - e^{-\int_0^t \beta(s)ds})^2)$

IV. SAMPLE GENERATION

New samples are generated by solving the reverse diffusion process defined in Eq. 2 using a numerical SDE solver. In this work, we use the Euler-Maruyama algorithm [59] followed by an additional corrector step that uses the Langevin MCMC approach [60, 61] to increase the sampling quality. For each time decrement Δt , the updated state of the system is described as:

$$x_{t-\Delta t} = x_t + [f(x_t, t) - g^2(t)s_\theta(x_t, t)]\Delta t + g(t)\sqrt{|\Delta t|}z, \quad (10)$$

where $z \sim \mathcal{N}(0, 1)$ is sampled at each time step. The corrector step takes the updated state from Eq. 10 and applies the correction

$$x'_{t-\Delta t} = x_{t-\Delta t} + \epsilon s_\theta(x_{t-\Delta t}, t)g(t) + \sqrt{2\epsilon}z, \quad (11)$$

where ϵ is a tunable parameter that determines the strength of the correction applied. A dimension-independent expression for ϵ is defined as a function of a signal-to-noise ratio r parameter as:

$$\epsilon = 2r^2 \frac{\|z\|_2^2}{\|s_\theta\|_2^2}, \quad (12)$$

where $\|z\|_2$ and $\|s_\theta\|_2$ are the batch-average norms of the Gaussian noise and trained score function.

High fidelity samples require the time interval Δt to be small, possibly leading to hundreds of iterative steps and consequently hundreds of network evaluations. We decrease the number of function evaluations by reusing the score function evaluated in Eq. 10 during the calculation of the corrector step in Eq. 11. This approach effectively decreases the number of function evaluations by a factor two while no decrease in generation quality is observed.

V. FAST CALORIMETER SIMULATION CHALLENGE 2022

CALOSCORE is trained on the datasets created by the Fast Calorimeter Simulation Challenge 2022 [46, 50–52]. A total of three datasets are provided, representing calorimeter shower simulations with GEANT of different geometries and granularities. Dataset 1 [51] is based on the ATLAS open dataset [46, 62] and is similar to the current ATLAS detector calorimeter geometry. Showers are generated at the calorimeter surface in the pseudo-rapidity range $\eta \in [0.20, 0.25]$. While samples consisting of both photons and pions are provided, we evaluate our model using only the photon dataset. The voxelization procedure is defined such that it minimizes the amount of empty voxels, while maintaining high fidelity compared to the full simulation. This strategy results in different number of voxels per calorimeter layer and a total of 368 voxels to represent the full detector slice. Photon energies are provided in this configuration for 15 incident energies from 256 MeV up to 4 TeV in steps given by powers of two. For each generated energy, 10k samples are provided with this number reduced at higher energies due to long simulation times, resulting in a total of 121k used during training.

Datasets 2 [52] and 3 [50] contain each 100k examples and are simulated using a common detector layout but with different voxelization granularity. The detector simulated has a concentric cylinder geometry with 45 layers, where each layer consists of active (silicon) and passive (tungsten) material, simulated with GEANT4. Electrons are generated at the detector surface with initial energy sampled from a log-uniform distribution ranging from 1 GeV to 1 TeV. In dataset 2, each layer consists of 144 readout cells, with 9 in the radial and 16 in the angular directions. Dataset 3 is more granular, consisting of 900 readout cells in each layer, with 18 in the radial and 50 in the angular directions.

Even though the initial voxelization is provided in Cylindrical coordinates, we found it beneficial to convert the voxels in datasets 2 and 3 to Cartesian coordinates. This change allows the effective usage of 3D convolutional neural networks (CNNs) to build the score model. While CNNs are applicable in polar coordinates, they struggle to learn periodic boundary conditions. Convolutional operations are also less effective, since the majority of the energy depositions are located near $r = 0$, or near the corners of the image. Since the datasets are only available after voxelization, the transformation from Cylindrical to Cartesian coordinates inevitably leads to loss of information¹. Nevertheless, we observe improved generation quality after changing the coordinate system. See

¹ A one-to-one assignment between the two sets of coordinates is possible, but requires the distance interval in Cartesian coordinates to follow a non-linear function since the transformation of coordinates is itself non-linear.

App. B for more details of the transformation. The total number of voxels after the transformation is chosen to be similar to the original number, resulting in $12 \times 12 = 144$ and $32 \times 32 = 1024$ voxels per layer for datasets 2 and 3, respectively.

Energies deposited in each voxel span multiple orders of magnitude, motivating yet another transformation of the inputs before training the generative model. First, each voxel energy E_v is normalized by the value of the generated energy of the particle E_0 times a factor f that ensures the normalized voxel energy $E'_v = \frac{E_v}{fE_0}$ lies between 0 and 1. Naively, the factor f could be taken as 1, since energy conservation should ensure the sum of deposited energies to not exceed the initial particle’s energy. However, the sampling fraction of the calorimeter may lead to a mismatch between these numbers. This effect is particularly important in dataset 1, when particle energies as low as 256 MeV are considered. We take f as the highest deviation between total deposited and generated energies, fixed to $f = 3.1$ for dataset 1 and $f = 2$ for datasets 2 and 3.

The normalized energy depositions are then transformed to log-space, similarly to the strategy used in CALOFLOW. The log-transformed value u_v is defined as:

$$u_v = \log\left(\frac{x}{1-x}\right), x = \alpha + (1-2\alpha)E'_v. \quad (13)$$

The value α is set to 10^{-6} and avoids a discontinuity when $E'_v = 0$. The generated particle energy, used as a conditional input to the model, is also transformed before training. The transformed conditional energy u_0 is defined as:

$$u_0 = \frac{e_0 - e_{\min}}{e_{\max} - e_{\min}}, \quad (14)$$

where e_{\min} and e_{\max} are the minimum and maximum energies available in the dataset.

Finally, dataset 1 is also modified by adding an extra dimension that encodes an overall energy normalization. Before applying the log-transformation in Eq. 13, the total deposited energy is calculated. Instead of normalizing each voxel by the generated energy, the total deposited energy is used, ensuring the sum of all voxels is equal to 1. The additional entry is then defined as the total deposited energy normalized by the initial particle energy times the factor f . This strategy improves the estimation of the total energy deposition during training, now encoded as a single entry rather than the sum of all voxels.

VI. MODEL ARCHITECTURE AND TRAINING DETAILS

The score function is built from a modified version of the U-net model [63] where an encoder-decoder architecture with skip connections is used. 3D convolution operations are used as the basic layers in CALOSCORE for

datasets 2 and 3, leveraging the regular geometry. Dataset 1, on the other hand, is irregular and consists of different number of voxels per detector layer. In this case the score function is built based on 1D convolutional operations. This approach leads to dataset 1 requiring bigger kernel and layer sizes compared to datasets 2 and 3, resulting in a bigger model architecture.

Each convolutional operation uses the swish [64] non-linear activation function, with a kernel size of 5 (dataset 1) and 3 (datasets 2 and 3). The number of dimensions is reduced in the encoder section of the network through average pooling operations, reducing the total number of dimensions by a factor 4 for dataset 1 and a factor $3 \times 2 \times 2 = 12$ for datasets 2 and 3 after each pooling layer. In the opposite direction, upsampling layers are used to increase the dimensionality by repeating entries multiple times.

The different network architectures used for each dataset are shown in Fig. 2.

Conditional inputs used to train the model, namely the time component and generated energy, are first transformed using random Fourier features [65]. The transformed features for each conditional input are then concatenated and passed over 2 fully connected layers of sizes 256 (dataset 1) or 128 (datasets 2 and 3), both followed by a swish activation function. This set of *conditional embeddings* are used during multiple stages of the model architecture. In particular, conditional convolutional operations are created by adding the conditional embeddings as an additional bias at the output of convolutional layers.

The implementation is carried out with TENSORFLOW [66] optimized with ADAM [67] with initial learning rate set to 10^{-4} . The learning rate is reduced by a factor 2 if the loss function does not improve for a period of 100 consecutive epochs, evaluated using an independent dataset. The evaluation dataset is taken as 20% of the total amount of available training events. The models are trained for a total of 2000 epochs. The epoch with lowest evaluation loss is saved for further inspection. For all models, 16 NVIDIA A100 GPUs are used simultaneously interfaced with the Horovod package [68] on the Perlmutter supercomputer. The batch size in each GPU is set to 128 (datasets 1 and 2) and 64 (dataset 3).

Sample generation is performed as described in Sec. IV, with signal-to-noise ratio fixed to 0.2 and total number of time steps set to 100, in dataset 1, and 200, in datasets 2 and 3. See App. C for differences in generation quality for other parameter choices. The total number of trainable weights and the time required to generate 100 calorimeter showers in a single GPU with batch size fixed to 100 are listed in Tab. II.

The lack of a regular geometry resulted in CALOSCORE requiring almost 20 times more trainable parameters for dataset 1 compared to datasets 2. On the other hand, dataset 3 has only $\sim 20\%$ more trainable weights than dataset 2, even though dataset 3 has 7 times more voxels. This observation suggests

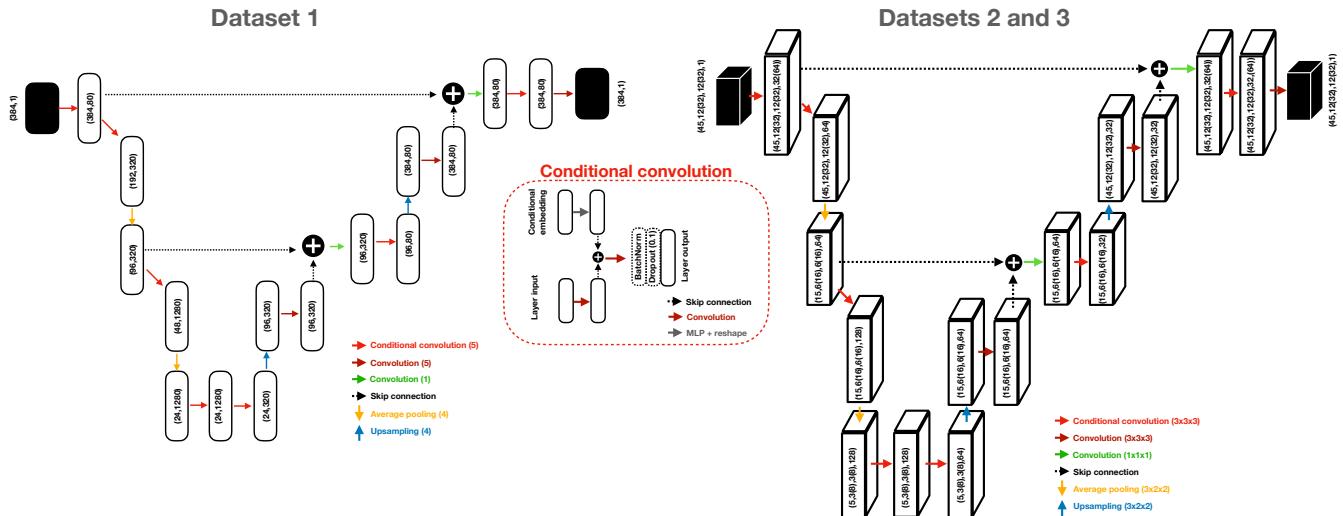


FIG. 2. Network architectures used for datasets 1 (left), 2 and 3 (right) based on the U-Net architecture. Values in parentheses represent the dimensionality of the data at each layer of the network. Parenthesis after convolutional operations represent the kernel size used. The swish non-linear activation function is used after each convolution operation. Conditional inputs (time and energy) are used multiple times in the model to define conditional convolution operations described in the middle.

TABLE II. Number of dimensions, trainable parameter, and time to generate 100 new calorimeter showers for each dataset studied in this work.

Dataset	N. of voxels	N. of weights	Time to 100 showers [s]
dataset 1	384	32M	4.0
dataset 2	6480	1.4M	5.8
dataset 3	46080	1.7M	33.4

that the model complexity is mostly determined by the network architecture rather than the number of voxels in the dataset, contrary to normalizing flows where the complexity from the Jacobian determinant calculation increases at least linearly with the number of voxels.

Although the ultimate goal is for the generation time to be significantly faster than for GEANT, the time to generate new calorimeter showers with CALOSCORE is comparatively slow to other machine learning-based models as a result of the hundreds of function evaluations. In this paper, we have focused on modeling the complex data distributions and we leave explorations of accelerating inference to future studies.

VII. RESULTS

Multiple distributions are used to evaluate the quality of generated samples using different CALOSCORE SDE implementations. The total energy deposited in the detector and the number of calorimeter hits are shown in Fig. 3. A hit is defined as any voxel with energy deposition above a certain energy threshold. The energy thresholds are taken as the minimum energy observed in

each challenge dataset, set to 0.01 keV for dataset 1 and 15.1 keV for datasets 2 and 3.

A good agreement between the GEANT and generated samples is observed for all diffusion models and datasets. At low deposited energies, the difference between CALOSCORE and GEANT increase and is most noticeable for dataset 3. The subVP implementation shows a better agreement overall, followed by VP and VE, indicating that bounding the variance of the diffusion process is beneficial, specially as the number of voxels increase.

A similar conclusion is derived from the average energy deposition as a function of the layer number and Cartesian coordinates x and y shown in Fig. 4.

While the VE implementation agrees with the simulation response in dataset 2, the prediction for dataset 3 is shifted as seen from the layer-dependent distribution.

Similarly, the maximum energy deposited in a single voxel normalized to the total deposited energy is shifted in the VE implementation for dataset 3 as shown in Fig. 5. While the low energy fraction region for dataset 1 is well described by all CALOSCORE implementations, the high energy fraction region starts to show deviations from the GEANT predictions. The maximum energy fraction as a function of the layer number for datasets 2 and 3 shows a good agreement between the different CALOSCORE implementations, with most of the distributions showing deviations within the 10% interval.

The angular distribution of the calorimeter shower is investigated in datasets 2 and 3 in terms of the shower width, shown in Fig. 6. The shower width σ_i with $x_i, i \in [1, 2]$ representing the x- and y- coordinates is calculated

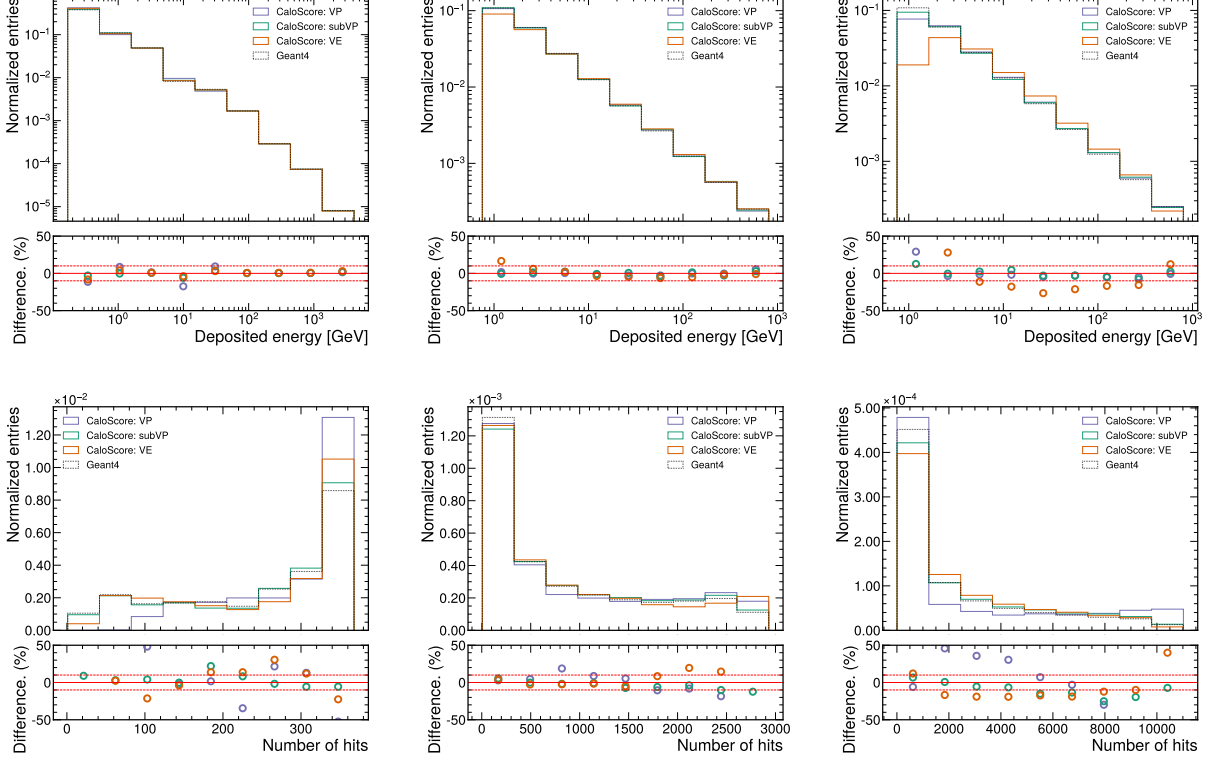


FIG. 3. Comparison of the sum of all voxel energies (top) and number of hits (bottom) for datasets 1 (left), 2 (middle), and 3 (right). Dashed red bands represent the 10% deviation interval of the generated samples when compared to GEANT predictions

as:

$$\sigma_i = \sqrt{\langle x_i^2 \rangle - \langle x_i \rangle^2}, \quad (15)$$

with energy-weighted mean defined as

$$\langle x_i \rangle = \frac{\sum_j x_{i,j} E_j}{\sum_j E_j}. \quad (16)$$

A good agreement between all CALOSCORE implementations and GEANT predictions is observed in dataset 2, with all implementations showing less than 10% deviation in all calorimeter layers. However, for dataset 3, the VP implementations shows a disagreement at the last layers of the detector while the shift observed in Fig. 4 for the VE implementation leads to a similar shift in the shower width. Nevertheless, the subVP implementation maintains the same level of agreement as observed in dataset 2.

A qualitative assessment of the generation is shown in Fig. 7 for datasets 2 and 3. The 2-dimensional distribution of the average energy deposition is shown in the detector layers with highest (layer 10) and lowest (layer 44) mean energy depositions. Empty entries in the GEANT simulation are a result of the initial voxelization combined with the following transformation to Cartesian coordinates. All voxels with an expected energy deposi-

tion above 0 are populated in all CALOSCORE implementations, an indication that CALOSCORE is able to reproduce the shower diversity from the training set. Images at layer 10 are identical for all diffusion models, dominated by the central voxel. Layer 44; however, has more voxels sharing a significant fraction of the layer energy. The subVP implementation shows a visually similar average to GEANT compared to the other diffusion implementations, capturing the high energy depositions along the y-axis in dataset 2 and the isotropic pattern around the center in dataset 3.

Finally, the assessment of generated samples using different conditional energies is investigated in Fig. 8, by comparing the total deposited energy versus the generated particle energy.

All CALOSCORE models show similar mean and spread compared to GEANT, with the exception of the VE implementation that shows a wider spread for dataset 2 and higher mean in dataset 3.

We have also explored the classifier metric introduced in CALOWFLOW whereby a post-hoc classifier is trained to distinguish generated showers from GEANT 4 examples. While the classifier could not exactly identify fake from real showers, it did have an area under the receiver operating characteristic curve (AUC) of about 0.98 for all three models. While this suggests that further (hyperparameter) optimization would be beneficial, it already

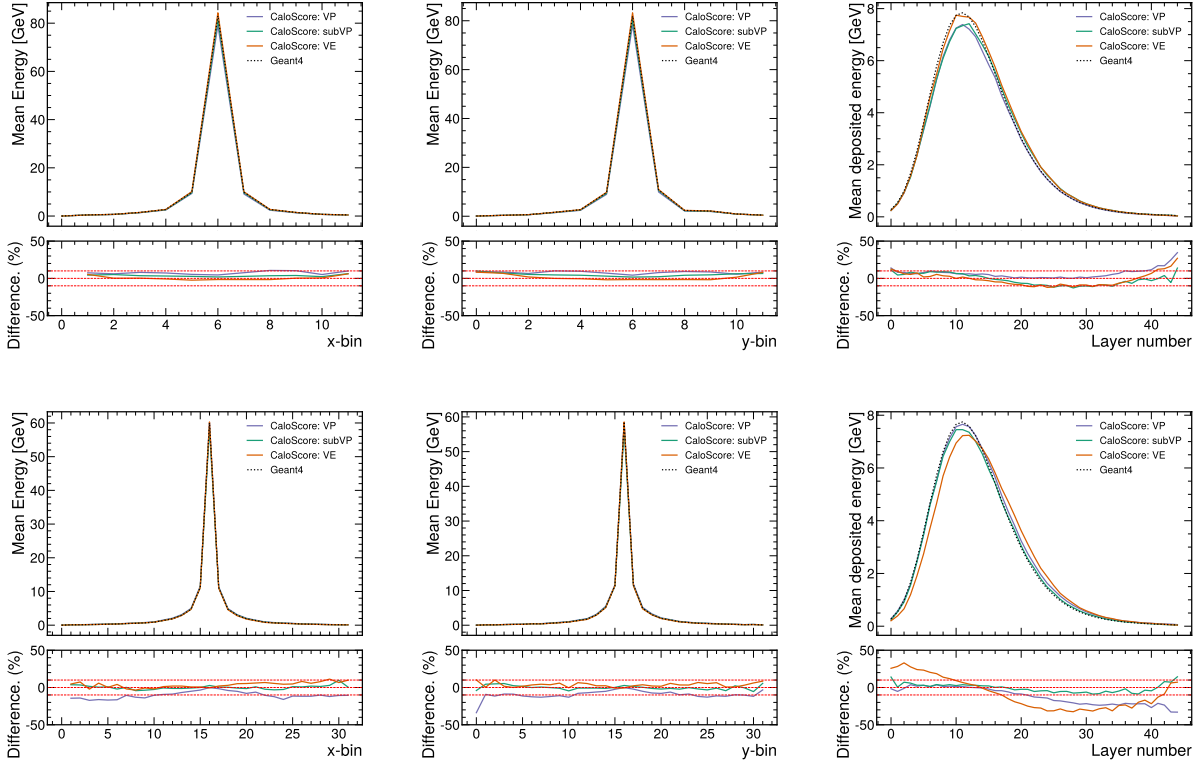


FIG. 4. Comparison of the average deposited energies in the x- (left), y- (middle), and z-coordinates (right) for datasets 2 (top) and 3 (bottom). Dashed red bands represent the 10% deviation interval of the generated samples when compared to GEANT predictions




FIG. 5. Comparison of the maximum deposited energy in a single voxel divided by the sum of deposited energies in datasets 1 (left), 2 (middle), and 3 (right). Dashed red bands represent the 10% deviation interval of the generated samples when compared to GEANT predictions

serves as an important baseline for other methods. It should also be noted that compared to the dataset introduced in the CALOGAN paper, the datasets considered in this work are either more realistic (dataset 1) or higher dimensional (a factor 10 for dataset 2 and 100 for dataset 3) and would be interesting to see the classifier metric obtained from models such as CALOFLOW.

VIII. CONCLUSIONS AND OUTLOOK

In this paper we introduced CALOSCORE, a novel generative model for calorimeter shower simulation based on score-matching, applied for the first time in the context of collider physics.

The performance of CALOSCORE is studied using the

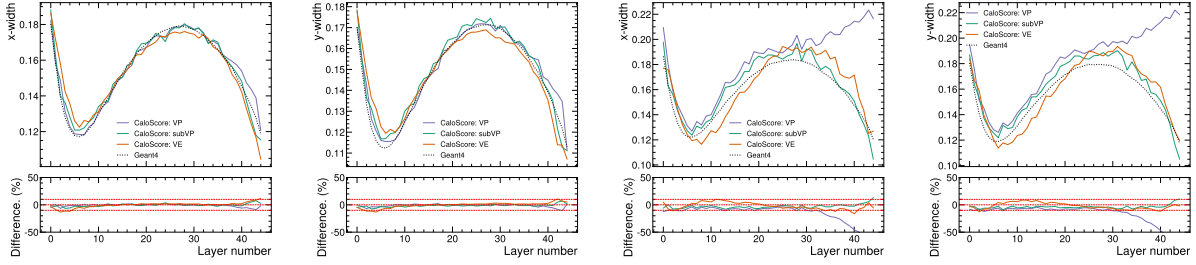


FIG. 6. Comparison of the particle shower width in the x- and y- directions in datasets 2 (first two figures from the left) and 3 (last two figures from the left). Dashed red bands represent the 10% deviation interval of the generated samples when compared to GEANT4 predictions




FIG. 7. Deposited energy versus generated energy in CALOSCORE(blue) and GEANT(orange) for the three different diffusion models: VP (left), subVP(middle), and VE (right). First row samples generated energies from dataset 1, second row from dataset 2 and third row from dataset 3.

Fast Calorimeter Simulation Challenge 2022 datasets and compared to three different model implementations based on different drift and diffusion coefficients. This is the first method to produce results for all three datasets of the challenge and we look forward to comparisons with other models as they become available.

For lower number of voxels, all models are capable of producing realistic calorimeter showers, showing a good agreement with the GEANT simulation in all datasets and for a variety of observables. At the highest dimensional dataset, CALOSCORE with diffusion process described by

the subVP stochastic differential equation is able to produce realistic calorimeter showers while VP and VE SDEs show larger deviations. CALOSCORE is also shown to be scalable, with number of trainable parameters sensitive to the overall model architecture rather than the total dimensionality of the dataset.

While the voxelization strategy used in dataset 1 minimizes the number of empty voxels, the irregular voxelization reduces the geometrical information present in the calorimeter shower. The geometrical information is included as an important inductive bias for datasets 2




FIG. 8. Deposited energy versus generated energy in CALOSCORE (blue) and GEANT (orange) for the three different diffusion models: VP (left), subVP(middle), and VE (right). First row samples are generated using energies from dataset 1, second row from dataset 2 and third row from dataset 3.

and 3 and leads to an order of magnitude fewer trainable parameters in the model compared to the one used in dataset 1. Moreover, a regular voxelization amenable to convolutional operations is obtained in datasets 2 and 3 only after an additional transformation of coordinates. Since the transformation is applied on the voxelized inputs, nonphysical artifacts are introduced, such as empty voxels in regions that are not expected to be empty. All of these issues could be addressed if alternative voxelization schemes were available or if access to the datasets prior to any voxelization was possible.

The addition of inductive biases to the model is also expected to improve the generation capabilities of CALOSCORE, possibly leading to better and even smaller model architectures. Energy conservation in particular is challenging to enforce, since generated samples are not produced during training time, but only at generation time when the reverse stochastic differential equation is solved. We partially address this issue by increasing the dimensionality of dataset 1, introducing an additional entry that stores an overall normalization. Since datasets 2 and 3 rely on the geometrical description of the voxels,

this strategy is not readily applicable and would instead benefit from a two-step approach as used in CALOFLOW, where the overall normalization is determined separately and used as a conditional input to a second model that learns the normalized detector response.

The major challenge to be addressed in CALOSCORE is the generation time, currently requiring hundreds of model evaluations to solve the reverse SDE. While the total generation time of CALOSCORE is still faster compared to the GEANT simulation, we envision future works targeting high fidelity generation with lower number of function evaluations. Indeed, since this limitation is also observed in applications of diffusion models in general, a number of different attempts are currently being proposed to accelerate the generation procedure [69–73], with feasibility for collider physics applications yet to be studied.

Finally, CALOSCORE introduces a new generative paradigm to collider physics with scalable training strategy and able to generate realistic calorimeter showers consisting of tens of thousands of dimensions. While the generation time represents the main challenge to be

overcome, CALOScore is able to incorporate different advantages from other generative models while addressing some of their limitations. These include scalable and stable training schedule, based on the minimization of the convex score-matching loss, and exact likelihood estimation, previously only available with methods such as normalizing flows.

ACKNOWLEDGMENTS

We thank Jean-Roch Vlimant and David Shih for feedback on the manuscript. VM and BN are supported by

the U.S. Department of Energy (DOE), Office of Science under contract DE-AC02-05CH11231. The work of DS was supported by DOE grant DOE-SC0010008. This research used resources of 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.

-
- [1] S. Agostinelli *et al.*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506**, 250 (2003).
- [2] J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce Dubois, M. Asai, G. Barrand, R. Capra, S. Chauvie, R. Chytráček, G. Cirrone, G. Cooperman, G. Cosmo, G. Cuttone, 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. 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. Wellisch, D. Williams, D. Wright, and H. Yoshida, *IEEE Transactions on Nuclear Science* **53**, 270 (2006).
- [3] J. Allison, K. Amako, J. Apostolakis, P. Arce, M. Asai, T. Aso, E. Bagli, A. Bagulya, S. Banerjee, G. Barrand, B. Beck, A. Bogdanov, D. Brandt, J. Brown, H. Burkhardt, P. Canal, D. Cano-Ott, S. Chauvie, K. Cho, G. Cirrone, G. Cooperman, M. Cortés-Giraldo, G. Cosmo, G. Cuttone, G. Depaola, L. Desorgher, X. Dong, A. Dotti, V. Elvira, G. Folger, Z. Francis, A. Galoyan, L. Garnier, M. Gayer, K. Genser, V. Grichine, S. Guatelli, P. Guèye, P. Gumplinger, A. Howard, I. Hrivnáčová, S. Hwang, S. Incerti, A. Ivanchenko, V. Ivanchenko, F. Jones, S. Jun, P. Kaitaniemi, N. Karakatsanis, M. Karamitros, M. Kelsey, A. Kimura, T. Koi, H. Kurashige, A. Lechner, S. 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. Pia, W. Pokorski, J. Quesada, M. Raine, M. Reis, A. Ribon, A. Ristić Fira, F. Romano, G. Russo, G. Santin, T. Sasaki, D. Sawkey, J. Shin, I. Strakovsky, A. Taborda, S. Tanaka, B. Tomé, T. Toshito, H. Tran, P. Truscott, L. Urban, V. Uzhinsky, J. Verbeke, M. Verderi, B. Wendt, H. Wenzel, D. Wright, D. Wright, T. Yamashita, J. Yarba, and H. Yoshida, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **835**, 186 (2016).
- [4] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, (2014), [arXiv:1406.2661 \[stat.ML\]](#).
- [5] D. P. Kingma and M. Welling, (2014), [arXiv:1312.6114 \[stat.ML\]](#).
- [6] D. Rezende and S. Mohamed, *Proceedings of the 32nd International Conference on Machine Learning* **37**, 1530 (2015).
- [7] L. de Oliveira, M. Paganini, and B. Nachman, (2017), [10.1007/s41781-017-0004-6](#), [arXiv:1701.05927 \[stat.ML\]](#).
- [8] M. Paganini, L. de Oliveira, and B. Nachman, *Phys. Rev. Lett.* **120**, 042003 (2018), [arXiv:1705.02355 \[hep-ex\]](#).
- [9] M. Paganini, L. de Oliveira, and B. Nachman, *Phys. Rev.* **D97**, 014021 (2018), [arXiv:1712.10321 \[hep-ex\]](#).
- [10] S. Vallecorsa, F. Carminati, and G. Khattak, *Proceedings, 23rd International Conference on Computing in High Energy and Nuclear Physics (CHEP 2018): Sofia, Bulgaria, July 9-13, 2018* **214**, 02010 (2019).
- [11] C. Ahdida *et al.* (SHiP), (2019), [10.1088/1748-0221/14/11/P11028](#), [arXiv:1909.04451 \[physics.ins-det\]](#).
- [12] V. Chekalina, E. Orlova, F. Ratnikov, D. Ulyanov, A. Ustyuzhanin, and E. Zakharov, *CHEP 2018 (2018)*, [10.1051/epjconf/201921402034](#), [arXiv:1812.01319 \[physics.data-an\]](#).
- [13] *ATL-SOFT-PUB-2018-001 (2018)*.
- [14] F. Carminati, A. Gheata, G. Khattak, P. Mendez Lorenzo, S. Sharan, and S. Vallecorsa, *Proceedings, 18th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2017): Seattle, WA, USA, August 21-25, 2017* **1085**, 032016 (2018).
- [15] S. Vallecorsa, *Proceedings, 18th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2017): Seattle, WA, USA, August 21-25, 2017* **1085**, 022005 (2018).
- [16] P. Musella and F. Pandolfi, *Comput. Softw. Big Sci.* **2**, 8 (2018), [arXiv:1805.00850 \[hep-ex\]](#).
- [17] M. Erdmann, L. Geiger, J. Glombitza, and D. Schmidt, *Comput. Softw. Big Sci.* **2**, 4 (2018), [arXiv:1802.03325 \[astro-ph.IM\]](#).
- [18] K. Deja, T. Trzcinski, and u. Graczykowski, *Proceedings, 23rd International Conference on Computing in High Energy and Nuclear Physics (CHEP 2018): Sofia, Bulgaria, July 9-13, 2018* **214**, 06003 (2019).

- [19] D. Derkach, N. Kazeev, F. Ratnikov, A. Ustyuzhanin, and A. Volokhova, (2019), [10.1016/j.nima.2019.01.031](https://doi.org/10.1016/j.nima.2019.01.031), [arXiv:1903.11788](https://arxiv.org/abs/1903.11788) [hep-ex].
- [20] M. Erdmann, J. Glombitza, and T. Quast, *Comput. Softw. Big Sci.* **3**, 4 (2019), [arXiv:1807.01954](https://arxiv.org/abs/1807.01954) [physics.ins-det].
- [21] L. de Oliveira, M. Paganini, and B. Nachman (2017).
- [22] L. de Oliveira, M. Paganini, and B. Nachman, *J. Phys. Conf. Ser.* **1085**, 042017 (2018), [arXiv:1711.08813](https://arxiv.org/abs/1711.08813) [hep-ex].
- [23] B. Hooberman, A. Farbin, G. Khattak, V. Pacela, M. Pierini, J.-R. Vlimant, M. Spiropulu, W. Wei, M. Zhang, and S. Vallecorsa (2017).
- [24] D. Belayneh *et al.*, (2019), [10.1140/epjc/s10052-020-8251-9](https://doi.org/10.1140/epjc/s10052-020-8251-9), [arXiv:1912.06794](https://arxiv.org/abs/1912.06794) [physics.ins-det].
- [25] E. Buhmann, S. Diefenbacher, E. Eren, F. Gaede, G. Kasieczka, A. Korol, and K. Kruger, (2020), [arXiv:2005.05334](https://arxiv.org/abs/2005.05334) [physics.ins-det].
- [26] S. Diefenbacher, E. Eren, G. Kasieczka, A. Korol, B. Nachman, and D. Shih, *Journal of Instrumentation* **15**, P11004 (2020), [arXiv:2009.03796](https://arxiv.org/abs/2009.03796) [hep-ph].
- [27] R. Kansal, J. Duarte, B. Orzari, T. Tomei, M. Pierini, M. Touranakou, J.-R. Vlimant, and D. Gunopoulos, 34th Conference on Neural Information Processing Systems (2020), [arXiv:2012.00173](https://arxiv.org/abs/2012.00173) [physics.data-an].
- [28] A. Maevskiy, F. Ratnikov, A. Zinchenko, and V. Riabov, (2020), [arXiv:2012.04595](https://arxiv.org/abs/2012.04595) [physics.ins-det].
- [29] F. Rehm, S. Vallecorsa, V. Saletore, H. Pabst, A. Chaibi, V. Codreanu, K. Borrás, and D. Krücker, (2021), [10.5220/0010245002510258](https://doi.org/10.5220/0010245002510258), [arXiv:2103.10142](https://arxiv.org/abs/2103.10142) [physics.data-an].
- [30] F. Rehm, S. Vallecorsa, K. Borrás, and D. Krücker (2021) [arXiv:2103.13698](https://arxiv.org/abs/2103.13698) [hep-ex].
- [31] F. Rehm, S. Vallecorsa, K. Borrás, and D. Krücker (2021) [arXiv:2105.08960](https://arxiv.org/abs/2105.08960) [hep-ex].
- [32] R. Kansal, J. Duarte, H. Su, B. Orzari, T. Tomei, M. Pierini, M. Touranakou, J.-R. Vlimant, and D. Gunopoulos, (2021), [arXiv:2106.11535](https://arxiv.org/abs/2106.11535) [cs.LG].
- [33] G. R. Khattak, S. Vallecorsa, F. Carminati, and G. M. Khan, (2021), [arXiv:2109.07388](https://arxiv.org/abs/2109.07388) [physics.ins-det].
- [34] L. Anderlini (2021) [arXiv:2110.07925](https://arxiv.org/abs/2110.07925) [hep-ex].
- [35] E. Buhmann, S. Diefenbacher, E. Eren, F. Gaede, D. Hundhausen, G. Kasieczka, W. Korcari, K. Krüger, P. McKeown, and L. Rustige, (2021), [arXiv:2112.09709](https://arxiv.org/abs/2112.09709) [physics.ins-det].
- [36] S. Bieringer, A. Butter, S. Diefenbacher, E. Eren, F. Gaede, D. Hundhausen, G. Kasieczka, B. Nachman, T. Plehn, and M. Trabs, (2022), [arXiv:2202.07352](https://arxiv.org/abs/2202.07352) [hep-ph].
- [37] J. W. Monk, (2018), [arXiv:1807.03685](https://arxiv.org/abs/1807.03685) [hep-ph].
- [38] K. Dohi, (2020), [arXiv:2009.04842](https://arxiv.org/abs/2009.04842) [hep-ph].
- [39] E. Buhmann, S. Diefenbacher, E. Eren, F. Gaede, G. Kasieczka, A. Korol, and K. Krüger, (2021), [arXiv:2102.12491](https://arxiv.org/abs/2102.12491) [physics.ins-det].
- [40] A. Hariri, D. Dyachkova, and S. Gleyzer, (2021), [arXiv:2104.01725](https://arxiv.org/abs/2104.01725) [hep-ex].
- [41] C. Fanelli and J. Pomponi, *Mach. Learn. Sci. Tech.* **1**, 015010 (2019), [arXiv:1911.11717](https://arxiv.org/abs/1911.11717) [physics.data-an].
- [42] B. Orzari, T. Tomei, M. Pierini, M. Touranakou, J. Duarte, R. Kansal, J.-R. Vlimant, and D. Gunopoulos, in *38th International Conference on Machine Learning Conference* (2021) [arXiv:2109.15197](https://arxiv.org/abs/2109.15197) [physics.data-an].
- [43] M. Touranakou, N. Chernyavskaya, J. Duarte, D. Gunopoulos, R. Kansal, B. Orzari, M. Pierini, T. Tomei, and J.-R. Vlimant, (2022), [arXiv:2203.00520](https://arxiv.org/abs/2203.00520) [physics.comp-ph].
- [44] C. Krause and D. Shih, (2021), [arXiv:2106.05285](https://arxiv.org/abs/2106.05285) [physics.ins-det].
- [45] C. Krause and D. Shih, (2021), [arXiv:2110.11377](https://arxiv.org/abs/2110.11377) [physics.ins-det].
- [46] G. Aad *et al.* (ATLAS), *Comput. Softw. Big Sci.* **6**, 7 (2022), [arXiv:2109.02551](https://arxiv.org/abs/2109.02551) [hep-ex].
- [47] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein gan,” (2017).
- [48] *The Phase-2 Upgrade of the CMS Endcap Calorimeter*, Tech. Rep. (CERN, Geneva, 2017).
- [49] J. Repond *et al.* (CALICE), *JINST* **13**, P12022 (2018), [arXiv:1809.03909](https://arxiv.org/abs/1809.03909) [physics.ins-det].
- [50] M. Fauci Giannelli, G. Kasieczka, C. Krause, B. Nachman, D. Salamani, D. Shih, and A. Zaborowska, “Fast Calorimeter Simulation Challenge 2022 - Dataset 3,” (2022).
- [51] M. F. Giannelli, G. Kasieczka, C. Krause, B. Nachman, D. Salamani, D. Shih, and A. Zaborowska, “Fast Calorimeter Simulation Challenge 2022 - Dataset 1,” (2022).
- [52] M. Fauci Giannelli, G. Kasieczka, C. Krause, B. Nachman, D. Salamani, D. Shih, and A. Zaborowska, “Fast Calorimeter Simulation Challenge 2022 - Dataset 2,” (2022).
- [53] Y. Song and S. Ermon, “Generative modeling by estimating gradients of the data distribution,” (2019).
- [54] J. Ho, A. Jain, and P. Abbeel, *Advances in Neural Information Processing Systems* **33**, 6840 (2020).
- [55] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, *ArXiv abs/2011.13456* (2021).
- [56] P. Vincent, *Neural Computation* **23**, 1661 (2011).
- [57] B. D. Anderson, *Stochastic Processes and their Applications* **12**, 313 (1982).
- [58] S. Särkkä and A. Solin, *Applied stochastic differential equations*, Vol. 10 (Cambridge University Press, 2019).
- [59] P. E. Kloeden and E. Platen, in *Numerical Solution of Stochastic Differential Equations* (Springer, 1992) pp. 103–160.
- [60] G. Parisi, *Nucl. Phys. B* **180**, 378 (1981).
- [61] U. Grenander and M. I. Miller, *Journal of the royal statistical society series b-methodological* **56**, 549 (1994).
- [62] *Fast simulation of the ATLAS calorimeter system with Generative Adversarial Networks*, Tech. Rep. (CERN, Geneva, 2020).
- [63] O. Ronneberger, P. Fischer, and T. Brox, in *International Conference on Medical image computing and computer-assisted intervention* (Springer, 2015) pp. 234–241.
- [64] P. Ramachandran, B. Zoph, and Q. V. Le, *arXiv preprint arXiv:1710.05941* (2017).
- [65] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng, *Advances in Neural Information Processing Systems* **33**, 7537 (2020).
- [66] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, *et al.*, in *OSDI*, Vol. 16 (2016) pp. 265–283.
- [67] D. Kingma and J. Ba, (2014), [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) [cs].
- [68] A. Sergeev and M. D. Balso, *arXiv preprint arXiv:1802.05799* (2018).

- [69] A. Jolicoeur-Martineau, K. Li, R. Piché-Taillefer, T. Kachman, and I. Mitliagkas, arXiv preprint arXiv:2105.14080 (2021).
- [70] Q. Zhang and Y. Chen, arXiv preprint arXiv:2204.13902 (2022).
- [71] Z. Lyu, X. Xu, C. Yang, D. Lin, and B. Dai, arXiv preprint arXiv:2205.12524 (2022).
- [72] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, arXiv preprint arXiv:2206.00927 (2022).
- [73] Z. Xiao, K. Kreis, and A. Vahdat, in *International Conference on Learning Representations (ICLR)* (2022).
- [74] T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, *CoRR* **abs/1806.07366** (2018), 1806.07366.
- [75] J. Skilling, “The eigenvalues of mega-dimensional matrices,” in *Maximum Entropy and Bayesian Methods: Cambridge, England, 1988*, edited by J. Skilling (Springer Netherlands, Dordrecht, 1989) pp. 455–466.
- [76] M. Hutchinson, *Communications in Statistics - Simulation and Computation* **19**, 433 (1990), <https://doi.org/10.1080/03610919008812866>.

Appendix A: Score function and the connection with continuous normalizing flows

Full data likelihood access is obtained from the trained score-matching model by first identifying the deterministic ordinary differential equation (ODE) associated to the SDE in Eq. 2 that reads:

$$dx = [f(x, t) - \frac{1}{2}g(t)^2 s_\theta(x, t)]dt = \tilde{f}_\theta(x_t, t)dt. \quad (\text{A1})$$

This ODE, named probability flow ODE by [55], has the property of having all trajectories sharing the same marginal probability densities as the SDE in Eq. 2 and is fully determined once the score function is estimated by the generative model. The time evolution of the density is given by the instantaneous change of variables defined in [74]:

$$\log p_0(x_0) = \log p_T(x_T) + \int_0^T \nabla \cdot \tilde{f}_\theta(x_t, t) dt. \quad (\text{A2})$$

Eq. A2 is equivalent to the change of variables often used in continuous normalizing flows. This expression can be estimated efficiently by first noticing that

$$\nabla \cdot \tilde{f}_\theta(x_t, t) = \text{Tr} \left(\nabla \tilde{f}_\theta(x_t, t) \right), \quad (\text{A3})$$

where $\nabla \tilde{f}_\theta(x_t, t)$ represents the Jacobian of \tilde{f}_θ and using algorithms such as the Skilling-Hutchinson trace estimator [75, 76] to approximate the trace calculation.

Appendix B: Cylindrical to Cartesian coordinate transformation

The initial voxelization provided for datasets 2 and 3 are in Cylindrical coordinates (r, α, z') . While this set of coordinates reflect the detector symmetry, we found beneficial to convert the voxelization to Cartesian coordinates (x, y, z) . A voxel initially described in Cylindrical coordinates (r_i, α_i, z_i) is then converted as:

$$x_i = r_i \cos \alpha_i \quad (\text{B1})$$

$$y_i = r_i \sin \alpha_i \quad (\text{B2})$$

$$z_i = z'_i, \quad (\text{B3})$$

where for simplicity we assume $r_i \in [0, 1]$ and $\alpha_i \in [0, 2\pi]$, resulting in $x \in [-1, 1]$ and $y \in [-1, 1]$. Since the overall transformation is not linear, some voxels of the new set of coordinates will always be empty, while others contain the sum of multiple voxels in the initial set of coordinates.

Appendix C: Generation quality for different sampling parameters

Results presented in this work are derived using the value of 0.2 for signal-to-noise-ratio r of the Langevin corrector (Eq. 11) and fixed number of times steps of 100 for dataset 1 and 200 for datasets 2 and 3. This choice of parameters are used to balance the generation quality and the generation time. In Fig. 9 different choices of r are compared while maintaining the same number of steps as before for all datasets and diffusion models, evaluated on the distribution of maximum energy fraction in a single voxel for dataset 1 and average energy deposition per layer for datasets 2 and 3, the distributions that show were observed to be more sensitive to the choice of generation parameters used.

Different choices of r yield similar results for all diffusion models in both datasets 1 and 2. On the other hand, the corrector step has a stronger effect on dataset 3 and is crucial to achieve good generation quality with minimal additional computational complexity.

Contrary to the corrector step, increasing the number of time steps directly affect the generation time, dominated by the number of score function evaluations. Different choices of number of times steps are shown in Fig. 10 with r value fixed to the baseline value.

In all cases using fewer time steps deteriorate the agreement with GEANT while additional steps are able to improve the agreement in both datasets 1 and 3 while dataset 2 shows similar results compared to the baseline. In Tab. III, the time required to generate 100 calorimeter showers using different number of time steps is listed. Even though the additional time steps improve the simulation quality, the time to generate the same amount of new observations increase by more than a factor 2.

TABLE III. Time comparison to generate 100 calorimeter showers using the baseline model and different number of time steps

Dataset	Baseline [s]	N=50 [s]	N=500 [s]
dataset 1	4.0	2.9	14.8
dataset 2	5.8	2.7	13.1
dataset 3	33.4	10.3	80.2

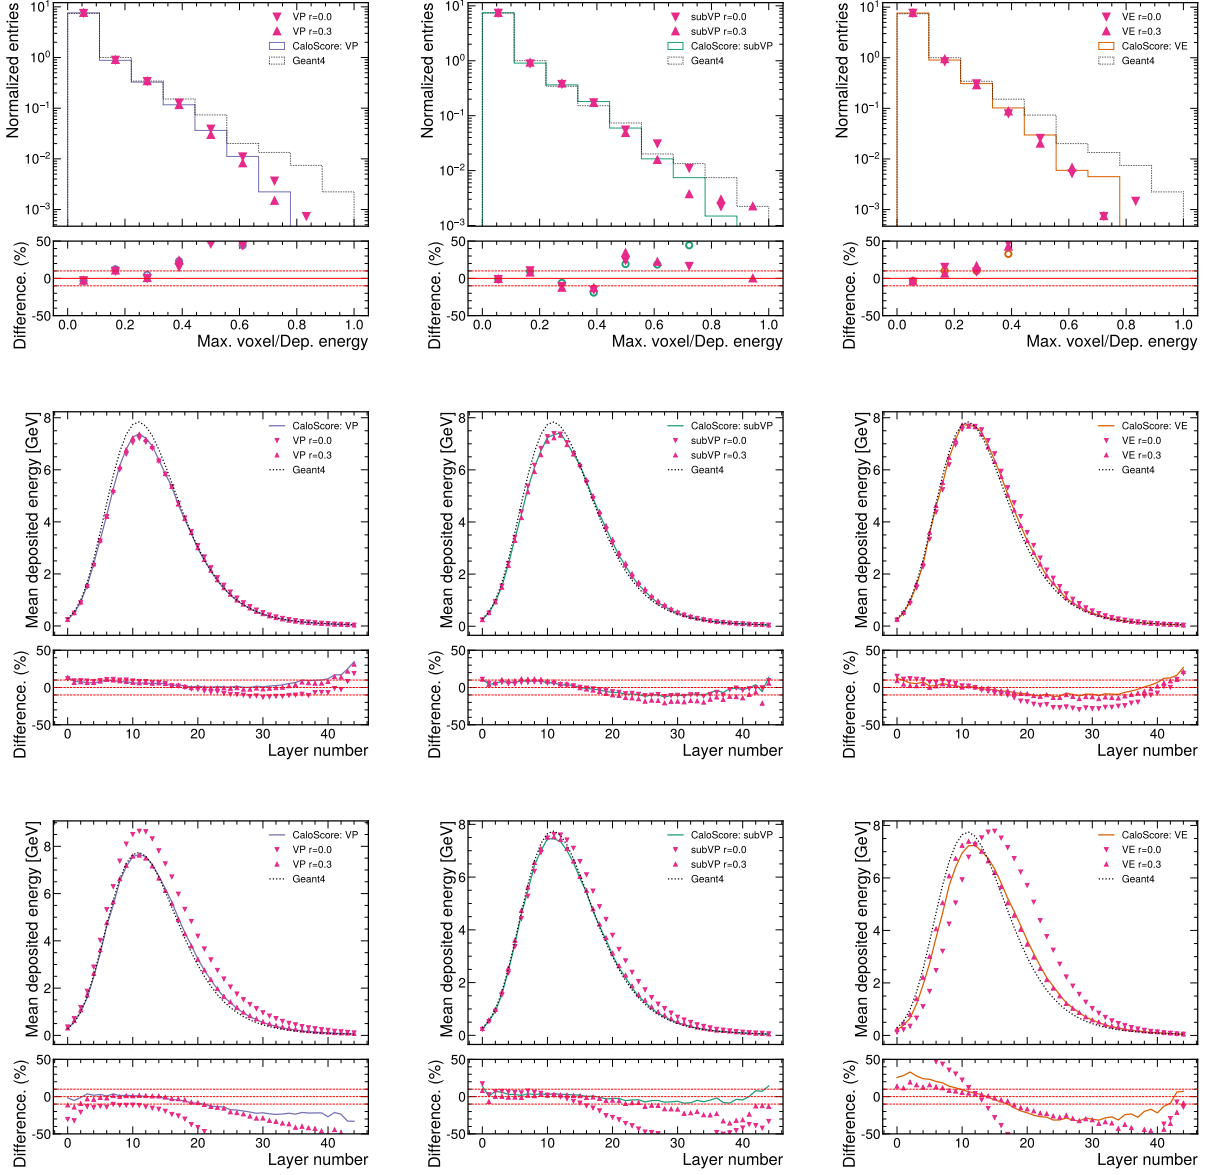


FIG. 9. Comparison of different signal-to-noise choices on the maximum energy fraction deposited in a single voxel for dataset 1 (top) and average energy deposition per layer for datasets 2 (middle) and 3 (bottom) for the three different diffusion models: VP (left), subVP (middle), and VE (right).

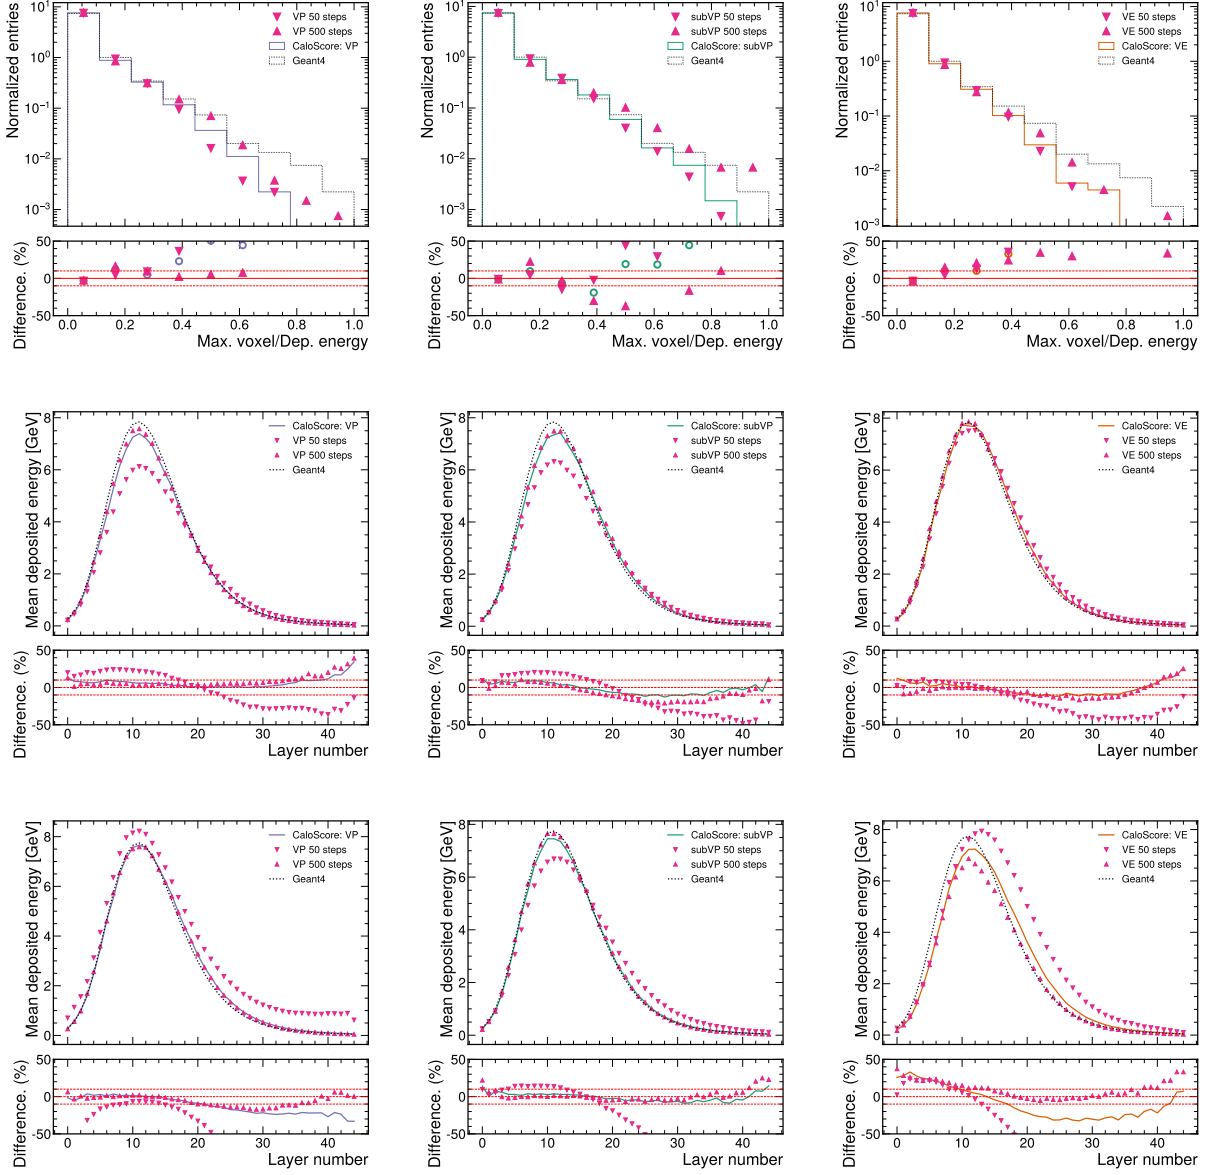


FIG. 10. Comparison of different number of time steps on the maximum energy fraction deposited in a single voxel for dataset 1 (top) and average energy deposition per layer for datasets 2 (middle) and 3 (bottom) for the three different diffusion models: VP (left), subVP(middle), and VE (right).