

# GENERATIVE ADVERSARIAL NETWORKS WITH DOMAIN-RELATED CONSTRAINTS ARE USED TO MODEL A CALORIMETER DETECTOR FOR HIGH-ENERGY PHYSICS

Dr Anuradha Gupta, Professor, anuradhag@rishiubrcollege.ac.in, Rishi UBR Women's College, Kukatpally, Hyderabad 500085

## ABSTRACT

Generative Adversarial Networks (GANs) have gained notoriety by generating highly realistic images. The present work explores GAN for simulating High Energy Physics detectors, interpreting detector output as three-dimensional images. The demands and requirements of a scientific simulation are quite stringent, as compared to the domain of visual images. Image characteristics such as pixel intensity and sparsity, for example, have very different distributions. Moreover, detector simulation requires conditioning on physics inputs, and domain knowledge becomes essential. We, therefore, adjust the pre-processing and incorporate physics-based constraints in the loss function. We also introduce a multi-step training process based on transfer learning by breaking up the task complexity. Validation of the results primarily consists of a detailed comparison to full Monte Carlo in terms of several physics quantities where a high level of agreement is found (ranging from a few percent up to 10% across a large particle energy range). In addition, we assess the performance by physics unrelated metrics, thereby proving further the variability and pertinence through diverse standpoints. We have demonstrated that an image generation technique from vision can successfully simulate highly complex physics processes while achieving a speedup of more than three orders of magnitude in comparison to the standard Monte Carlo.

## I. INTRODUCTION

Simulation of particle transport through matter is fundamental for interpreting the results of High Energy Physics (HEP) experiments. The particles undergo complex interactions while traversing the detector material with stochastic outcomes. The modeling of these processes is carried out with the help of Monte Carlo (MC) techniques that rely on repeated random sampling. The MC simulation meets the theoretical predictions with a high degree of precision but is both time and resource

intensive. The Worldwide LHC Grid [1] has currently more than 50% of its resources devoted only to simulation [2]. The future High Luminosity LHC [3] will require 100 times more simulated data, thus surpassing the expected resource availability. The HEP community is therefore highly motivated to explore fast alternatives, often trading some accuracy for speed, if only partially, for certain applications. Fast simulation is a set of established techniques that replaces parts of the detailed MC simulation with alternative approaches. Currently parametrized approaches [4]–[6] or lookup

tables [7] can provide between 10 and 100 times speedup, achieving different levels of accuracy. Here we investigate an alternative approach based on Deep Neural Networks.

HEP detectors can be described as 3D cameras, recording pictures of particle collisions. Calorimeters, in particular, detect particles by measuring the energy deposited in interactions with matter. Segmented calorimeters consist of alternate arrays of active sensor material and passive dense layers, to ensure that the incoming (primary) particle will deposit most of its energy inside their volume. The energy depositions in calorimeter cells can be compared to the monochromatic pixel intensities of a 3D image.

The 3DGAN is a convolutional GAN architecture aimed at stimulating the calorimeter's energy response. The calorimeter is a bottleneck in most HEP simulation pipelines due to a large number of complex interactions, taking about 90% of simulation time for some experiments [8]. The model was developed in several steps. The initial proof of concept was a simplified prototype demonstrating successful detector simulation, conditioned by the energy of the primary particle, entering the detector perpendicularly to its surface [9]. This represented a simplification, given that in real-life conditions, particles usually hit the detectors from different directions. The model was revised to learn a joint probability distribution of both the incident particle energy and its direction (quantified by the incident angle). At this stage, we also

increased the image size and added new domain-related features (details in Section IV) in order to improve the accuracy of the results. Some preliminary results for this configuration were presented in [10], where the training was run for a small number of epochs and the result validation was limited. The current work refines the results and describes the architecture, training process, and pre-processing in more detail. The physicsbased analysis of the results is presented in greater detail along with additional performance studies from other viewpoints. These include classification and regression results by a third-party network pre-trained on the same data set, as well as a more detailed investigation using image quality metrics as a function of input conditions.

This paper is organized as follows. Section II will briefly review related work focusing on HEP applications. It will be followed by a description of the training data set and the features used for the evaluation of the results in Section III. The current approach will then be presented in Section IV, together with details on the loss function and the network architecture. The main design choices made during the development process will also be discussed. Validation of the results from a physics perspective is accomplished by a detailed comparison of GAN and Monte Carlo as presented in Section V. This section also includes an investigation of additional validation metrics (e.g. structural similarity index) inspired from the image processing domain. We will conclude by summarizing our main

contributions and suggestions for future work in Section VI.

## II. PREVIOUS WORK

Generative models represent a fundamental part of deep learning. Over the years this field has seen developments in the Generative Stochastic Networks [11], to the Variational AutoEncoders [12], and Generative Adversarial Networks (GAN) [13]. In particular, Generative Adversarial Networks can successfully generate sharp and realistic images with high resolution [13]. Inspired by the Game Theory [14], adversarial training is defined as a competition between two players: a generator and a discriminator. The discriminator distinguishes real from fake images while the generator tries to fool the discriminator by producing an output as realistic as possible. The process eventually results in the generator learning the distribution of the real data if given enough representation capacity and time.

There are many variants of the GAN methodology, such as WGAN [15], StackGAN [16], ProgressiveGAN [17] etc., demonstrating the generation of high quality and high resolution images. The GAN does not rely on the explicit computation of probability densities and is thus suitable for a wider range of applications. Presently this approach has been applied to problems from many domains: ranging from generating musical notes [18], to natural language [19], to medical data [20], [21], to natural scenes [13] and image denoising [22]. The

adversarial approach has also been successfully applied to anomaly detection tasks for industrial [23] and medical [24] applications. Similarly, scientific simulation is another domain where GAN has shown immense success. Astrophysics [25], [26], microrbiology [27], and material composition [28], etc., are some of the avenues explored.

The novelty of our work resides in the high granularity (high spatial resolution) of the detector we simulate and in the use of three-dimensional convolutions that are essential to preserving all spatial correlations between pixels. Our particular pre-processing, loss function, and two-step training results in high accuracy for a more complex scenario in comparison to other GAN applications for HEP calorimeter simulation, involving a wider range of input variables used for conditioning.

## III. DATA SET

The incoming particle creates an avalanche of secondary particles as it traverses the detector thus generating a characteristic energy deposit pattern (called a “shower”).

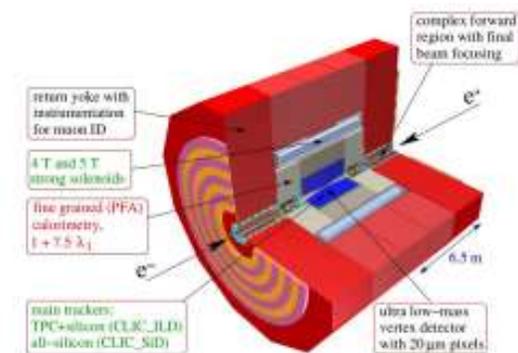


FIGURE 1. Schematic diagram for the CLIC calorimeter [37].

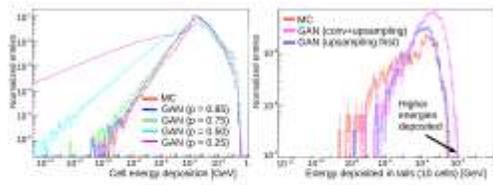
Each entry or event in the data set corresponds to an array of cells centered around the barycenter of the particle shower as a three dimensional  $51 \times 51 \times 25$  pixelized image. The pixel intensities of this image are the energy depositions for the calorimeter cells. The energy (EP) and the incident angle ( $\theta$ ) of the original particle (“primary particle”) are also stored with each event. The data consist of particles with the EP range of 2 GeV to 500 GeV and  $\theta$  range of  $60^\circ$  to  $120^\circ$ . The training uses 137, 342 electron events with EP restricted to the range of 100 GeV to 200 GeV for the first step and 400, 000 events with EP from 2 GeV to 500 GeV for the second training step. The data is divided into a ratio of nine to one for train and test sets. While additional data is used for the validation and detailed analysis presented in Section V (filtered for specific energies and angles from sets of 50, 000 and 100, 000 events respectively).

The physical characteristics of a particle shower are defined by the underlying physics processes, the particle type, its energy, and incident angle. The 2D projections on different planes can be used to convey visual information. The geometry of a shower is essential information for particle type and energy identification. Frequently used geometrical features consist of shower shapes, moments, and ratio of energy deposited in different parts of a shower. The shower shapes are represented by the energies deposited along different axes. Shower moments are another way of defining the shower geometry. The

first moment (M1) corresponds to shower center, the second moment (M2) to shower width, and similarly higher moments to higher-order features. The fraction of energy deposited in different parts of the shower is studied by dividing the shower into three parts and comparing the fraction of total energy deposited in the first eight cells (R1), the middle nine cells (R2), and the last eight cells (R3).

#### IV. THE 3D CONVOLUTIONAL GAN

The 3DGAN represents the first proof of concept for the possibility of using 3D convolutional GANs to simulate high granularity calorimeters. Our network is inspired from the Auxiliary Classifier GANs (ACGAN) [34] and InfoGAN [41] concepts. These architectures are a natural extension of the GAN approach and feature a faster convergence and more stable performance by introducing auxiliary tasks for the discriminator. The GAN application to simulation requires the generation of images conditioned on a set of continuous inputs, thus introducing auxiliary tasks such as regression on the conditioning variables, not only stabilize the training but also provide feedback on the conditioning. With the advent of representation learning through deeper models, domain knowledge is often not required [42]. We present an example of combining deep learning with domain related constraints since scientific simulations must conform to scientific laws.



**FIGURE 3.** The cell energy distribution for Monte Carlo events vs. GAN generated events. (Left) pre-processing by taking the power of pixel intensities for:  $p = 0.85$  (blue);  $p = 0.75$  (green);  $p = 0.5$  (cyan);  $p = 0.25$  (magenta). (Right) the distribution of the highly sparse region near the periphery along the transverse dimensions ( $X$  and  $Y$  axes) for generator with: alternative upsampling and convolutional layers (magenta); upsampling before the convolutions (blue).

## A. PRE-PROCESSING

One of the challenges in applying image generation techniques to the simulation of detector output lies in the large dynamic range of the deposited energies in detector cells as compared to pixel intensities in a typical RGB image. The cell energy deposits can vary over a large range spanning more than ten orders of magnitude as shown in Figure 2 (right). We explored different pre-processing procedures aimed at reducing this dynamic range. The initial tests using the log of the pixel intensity yielded unsatisfactory results (highly distorted images). We also experimented with the power function of pixel intensities using an exponent ( $p$ ) smaller than one. We observed that a smaller exponent means faster convergence but greater distortion. Figure 2 (right panel) compares the pixel intensity distribution for different values of the exponent. An optimum value of 0.85 improve the convergence while retaining accuracy at both ends of the spectrum. The generated images can be transformed back to the original range by simply taking the inverse of the power function.

## B. THE ARCHITECTURE

The design of the 3DGAN architecture required a long and tedious process of trial and error, as well as conventions and suggestions from past efforts. The long training times and practically unlimited choice of architectural hyperparameters do not allow an exhaustive search of the entire design space, although the model development and optimization involved extensive investigation of different model parameters. The 3DGAN initial prototype [9] architecture was inspired by the DCGAN [43] architecture: employing four convolutional layers in both the generator and the discriminator networks. The upsampling layers followed the first two convolutional layers in the generator.

Figure 3 shows the final optimized architecture for the discriminator (D) and the generator (G). The latent space is a vector of 254 random numbers drawn from a Gaussian distribution. The input conditions  $EP$  and  $\theta$  are concatenated to this latent vector to create the generator input. The initial prototype [9] used a Hadamard product of the condition  $EP$  and the latent vector. For the more complex scenario with multiple input conditions, we find that simply concatenating the conditions to the latent vector provides a more compact approach with better overall accuracy. A set of upsampling layers, at the beginning of the generator network, are used to reach the required dimensions before the application of the convolutions [44]. The initial prototype [9] consisted of alternating

upsampling layers and convolutional layers similar to ACGAN [45]. By moving the upsampling layers before convolutions we have improved the generator learning process, obtaining a more realistic output. The improvement is most significant for the simulation of cells with very small energy deposits, occurring mostly in the peripheral regions of the image as presented in Figure 2 left panel. The 3DGAN has a stronger generator (seven convolutional layers) than the discriminator (with four layers) in order to cope with the increase in image complexity. A gradually decreasing kernel is used for the generator (to generate more local features in higher-order layers). The geometry of the energy shower, extending along the length of the longitudinal dimension, while being narrow for the transverse dimensions, influenced the choice of the generator kernels in the Z dimension,

which are kept larger. Experimentation with kernels did not have any significant effect on model performance thus number of filters and kernel sizes are further adjusted in consideration to memory constraints and image dimensions.

Batch normalization [46] (with a smaller fuzz factor of  $10^{-6}$ ) is applied to all except the first convolutional layer in the discriminator and the last two layers in the generator, these exclusions are aimed at encouraging the large dynamic range for the pixel intensities. Relu [47] activation function is used for the generator layers to induce sparsity while the leakyRelu [48] is used for the discriminator hidden layers. The discriminator is regularized by a dropout [49] of 20% and a single average pooling layer after the last convolutional layer (additional pooling layers result in substantial loss of performance).

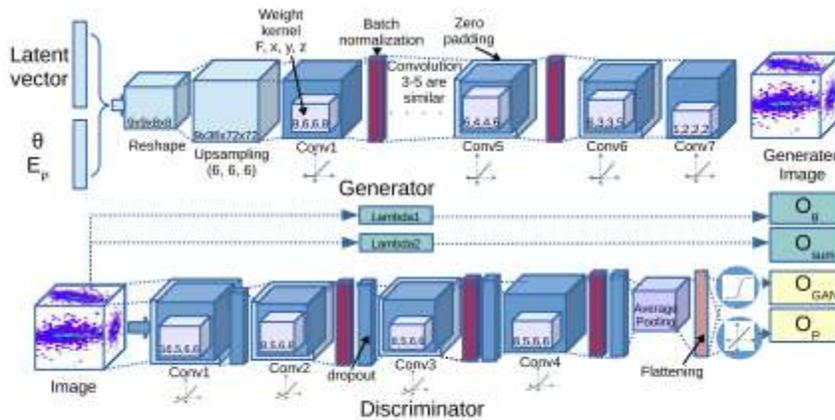


FIGURE 3. The 3DGAN architecture, see the text for details.

### C. LOSS FUNCTION

The 3DGAN loss function is built from the weighted sum of individual terms pertaining to the discriminator outputs and constraints.

The introduction of domain-related constraints is essential to achieve a high level of accuracy. Figure 4 shows the effect of constraining the Esum. Without the

constraint,  $E_{sum}$  for generated images has a roughly uniform distribution, but only after applying the constraint, the sum is correctly mapped to  $E_p$ . It should be noted that by constraining the total deposited energy we make sure that the energy conservation is preserved.

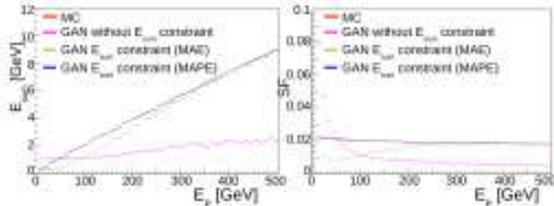


FIGURE 4. The GAN performance without  $E_{sum}$  constraint (magenta) and with the constraint (green and blue), demonstrated by a comparison between Monte Carlo and GAN generated events for  $E_{sum}$  vs.  $E_p$  (left) and SF vs.  $E_p$  (right). (initial prototype [9]).

Equation (1) shows the different components of the discriminator loss: the discriminator real/fake probability as defined in [13] (LG), the primary particle energy regression (LP), the total deposited energy (LE), and incident angle measurement (LA). The losses are balanced by the corresponding weights  $W$ . The LP and  $L_\theta$  provide feedback on the conditioning of the image, while LE and LA help to impose external constraints on the images. The generator loss is implemented as the inverse of LG together with the auxiliary losses and constraints.

$$L_{3DGAN} = W_G L_G + W_P L_P + W_A L_A + W_E L_E \quad (1)$$

The loss components presented in Equation (1) are based on different errors. LG is implemented as binary crossentropy. LP and LE are evaluated as the mean absolute percentage error (MAPE), while LA is evaluated as the mean absolute error (MAE). Figure 4 also demonstrates how the choice

of the error for the loss term affects the quality of results. As Osum and OP are both correlated to pixel intensities, thus a percentage error results in better performance especially at the lower end of the spectrum where smaller energies are deposited. Apart from some degradation in both simulating and predicting these low intensities, a slight energy difference at the lower end would result in a tiny absolute error but a larger percentage error (thus a larger gradient), therefore improving the training for this region.

he loss weights are selected roughly following the guidelines by Chollet [52] according to the type of losses, probable values at convergence, and relative importance. The losses based on mean percentage error are assigned a weight of 0.1. The WA is assigned a value of 25 and the BCE losses are assigned a weight of 3. A high level of accuracy is required for the correct  $E_{sum}$  of the generated images. Increasing the weight  $W_E$  fails to improve the performance on LE as shown in Figure 5 (right), while the performance on LG deteriorates (left). Therefore we add another step to select the final network based on the minimum mean absolute relative error on the sampling fraction (using additional holdout samples), in the last ten epochs so as to retain a high level of accuracy on this quantity.

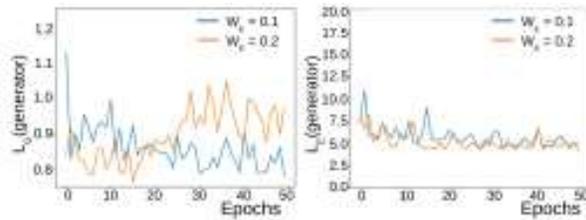


FIGURE 5. The 3DGAN generator test losses for loss weight  $W_L$  equal to 0.1 (cyan) and 0.2 (orange). (Left)  $L_G$  vs. epochs. (Right)  $L_D$  vs. epochs. (initial prototype [9]).

#### D. TRAINING

The GAN training converges when the discriminator cannot discriminate between real and fake samples. The discriminator estimates a real/fake probability close to 50% for both the real and fake images. The convergence for the simulation of the more complex scenario involving both EP and  $\theta$  could not be accomplished directly, only restricting the EP range could finally allow the training to converge. We thus apply a two-step training based on the transfer learning to simulate the full EP range. During the first step, we train 3DGAN on events with EP in the 100 GeV to 200 GeV range. The second training step extends the data to events having EP from the full 2 GeV to 500 GeV range. The final network is further selected based on the minimum mean relative error (see Section V-F2) on SF in the last ten epochs, evaluated on additional holdout samples.

For each iteration, we train the discriminator on a batch of real images and a batch of generated images (applying label switching) [53]. Adopting a balanced approach, we also train the generator twice while freezing the discriminator weights. Another modification applied to the GAN training process

involves generating the fake image batch for the same input conditions as the real image batch instead of randomly sampling the input conditions. Thus not only alleviating the need to find the complex mapping of EP to Esum, as well as having to apply it at run time (as performed in [9]). A speedup of more than 30% could be achieved for the processing time per epoch, as well as removing any chances of error due to inaccurate mapping. A number of optimizers and learning rates were investigated. The RMSProp [54] optimizer with a learning rate of 0.001 was finally selected on the basis of performance, to train the network through Stochastic Gradient Descent in mini-batches of 64 events since larger batch sizes could not be supported due to memory constraints.

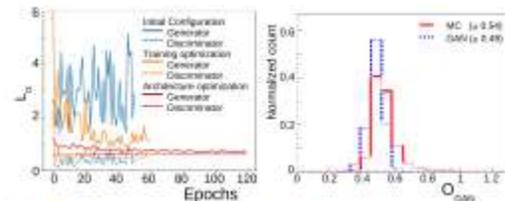


FIGURE 6. The 3DGAN test losses. (Left)  $L_G$  loss for the generator (continuous line) and the discriminator (dashed line): applying the configuration of the initial prototype [9] to the variable angle data does not work (teal); pre-processing for the reduction of the pixel intensity dynamic range and additional loss terms improved convergence (orange); architecture optimization (with upsampling layers at the beginning and more convolutional layers further improved the results (red). See text for additional details. (Right) The real/fake probability ( $O_{GAN}$ ) for both the Monte Carlo (red) and GAN generated events (blue dotted) is close to 0.5.

Figure 6 right panel compares the LG test losses at different stages of the current work for the restricted EP range. Restricting the EP range started showing some improvement in the training losses for the initial configuration [9] but the test losses (blue) remain highly random. Implementing the pre-processing step as explained in section IV-A, adjusting the losses, and generating images for the same inputs as real

images denote the training optimization (orange), that allows the losses to start converging. The architecture optimization involves up sampling before convolutions and additional layers in the generator (red), resulting in decreasing the loss further and improving the convergence (the GAN could still not converge for the full EP range directly thus necessitating the two-step training). The left panel shows how the real/fake probability (OGAN) estimated by the discriminator for the real (red) and the fake (blue) images, has very similar distributions, at the end of training.

## V. RESULTS AND DISCUSSION

As mentioned in Section I, the 3DGAN results cannot be compared directly to previous approaches as the problem is more complex. We thus validate the 3DGAN performance by a detailed comparison to the Monte Carlo simulation. Established fast simulation approaches report an accuracy within 10% of the Monte Carlo, thus we aim to achieve a similar level for the GAN generated events. We test the distributions of several physics quantities as a function of the main inputs to the simulation process (EP and  $\theta$ ). In order to do so, we select events in 5 GeV bins for EP and 0.1 rad ( $5.73^\circ$ ) bins for  $\theta$ , as well as some unbinned events, from the unseen data of 100, 000 random events. The comparisons are performed with GAN events generated for the same EP and  $\theta$  values as the MC events. We have computed the mean absolute relative errors for the histogram bins (BMRE). The detailed physics validation

involves numerous feature distributions compared for different input bins, thus resulting in hundreds of histograms. The GAN generated events have a close agreement with the Monte Carlo, greatly surpassing our aimed accuracy level for most of the physics features. Only a few representative features are presented here in order to simplify the discussion.

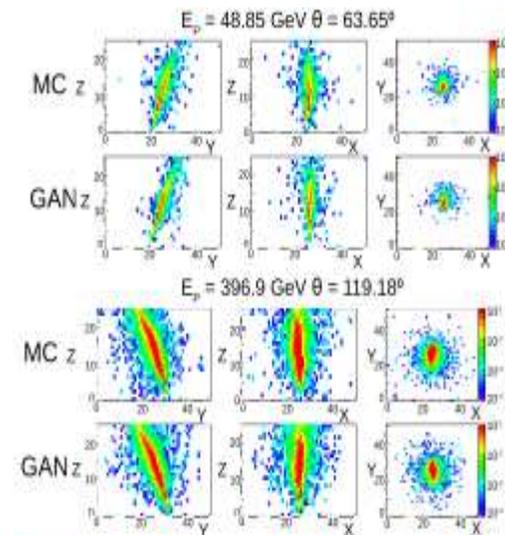


FIGURE 7. Example 2D shower sections [log] on the YZ, XZ, and XY planes with  $E_p$  and  $\theta$  sampled from both ends of their respective spectra. The GAN events are generated for the same  $E_p$  and  $\theta$  values as the MC events.  $E_p = 48.85$  GeV and  $\theta = 63.65^\circ$  (top two rows).  $E_p = 396.9$  GeV and  $\theta = 119.18^\circ$  (bottom two rows).

### A. VISUAL COMPARISON

In order to perform a preliminary visual inspection of the results, we build 2D projections on different planes of three dimensional events. Figure 7 presents an example of 2D sections of the MC and GAN showers corresponding to electrons entering the calorimeter with different energies and angles (selected from the tails of the energy and angle distributions). The images appear visually similar, sharp, and unique.

## B. SAMPLING FRACTION AND CALORIMETER RESPONSE

As explained earlier the Sampling Fraction (SF), is correlated to the amount of energy recorded by the calorimeter and it depends on the particular detector geometry and the primary particle type and energy. Figure 8 demonstrates the close agreement over the entire energy range for both ends of the  $\theta$  spectrum. It should be noted here that although for the  $62^\circ$  bin there is a slightly lower SF for the GAN images, yet the difference is very small with a mean absolute relative error of a few percent for the histogram bins.

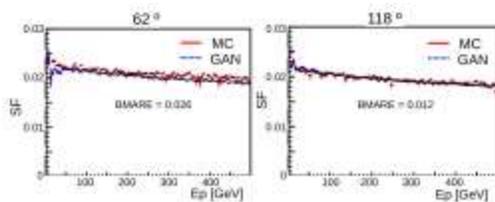


FIGURE 8. The sampling fraction for Monte Carlo vs. GAN events with  $E_p$  from 2 GeV to 500 GeV. The  $\theta$  values are from  $62^\circ$  bin (left) and  $118^\circ$  bin (right).

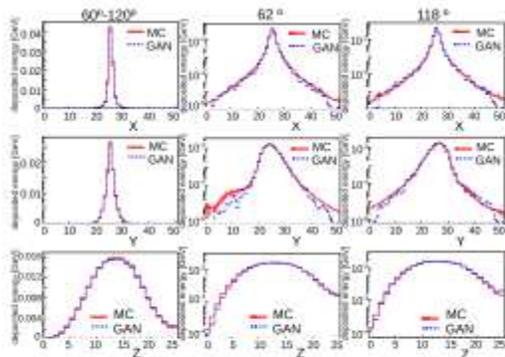


FIGURE 9. Shower Shapes for Monte Carlo vs. GAN events along X, Y, and Z axis with  $E_p$  range of 2 GeV to 500 GeV. The shapes for events with  $\theta$  range of  $60^\circ$  to  $120^\circ$  are plotted on linear y-scale (left column), and bins of  $62^\circ$  and  $118^\circ$  on logarithmic y-scales (two columns on the right).

## C. PARTICLE SHOWER SHAPES

Figure 9 presents the energy profiles for the showers along the X, Y and Z axes, both in

linear (full  $\theta$  range of  $60^\circ$  to  $120^\circ$ ) and log scale ( $62^\circ$  and  $118^\circ$  bins of  $\theta$ ). The network is capable of correctly reproducing the spatial distribution of energy deposits as a function of the incident angle, across a large dynamic range. In the log scale, some discrepancies are observed at the edges of the simulated volumes, where smaller energy depositions occur. The amount of energy expected in this region (well below 10–4 GeV) is below the threshold for this detector and comparable to the pedestal values. Figure 10 presents the longitudinal projection (XY) for the MC and GAN showers for some EP bins proving that the agreement in shapes is maintained along the entire spectrum.

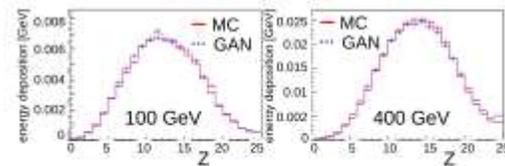


FIGURE 10. Shower shapes along the Z axis for Monte Carlo vs. GAN events in different primary energy bins. (Left)  $E_p = 100$  GeV. (Right)  $E_p = 400$  GeV.

## D. SPARSITY AND CELL ENERGY DISTRIBUTION

The detector cells are mostly empty with energy deposited only in less than 20% of them. Figure 11 (right) shows the fraction of cells where some energy is deposited as a function of the threshold used for the cutoff. It can be seen that the Monte Carlo and GAN images present a similar level of sparsity. The energies deposited in the calorimeter cells are our pixels intensities. Figure 11 (left) shows the MC and the GAN agreement in terms of pixel intensities down

to very low values. It should be noted that at around  $0.2 \times 10^{-3}$  GeV the MC intensities show a sharp, vertical drop. As expected the GAN smooths out this cut.

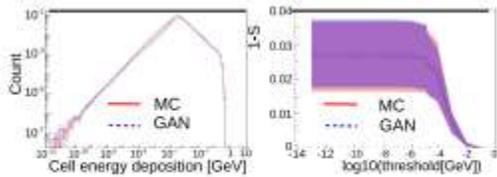


FIGURE 11. Monte Carlo vs. GAN shower features for 2 GeV to 500 GeV electrons show close agreement. (Left) Cell energy distributions. (Right) Sparsity as the fraction of cells above a threshold for different values of threshold.

### E. CORRELATIONS

We study the internal correlation between shower features as well as the correlation between MC and GAN images. The correlations among different shower features and the inputs i.e. energy and angle should be preserved. We evaluate the correlation matrix calculated on different quantities such as the shower shapes,  $E_{sum}$ ,  $EP$ ,  $|90^\circ - \theta|$ , and numbers of hits above a threshold ( $0.2 \times 10^{-3}$  GeV). Figure 12, the left panel, presents the difference between the internal correlations present in the MC data and those of the GAN images. These correlations agree to a considerable extent with a mean error of less than 10%. The incident angle  $\theta$  measured in MC and GAN, shown in Figure 12, right panel, also presents a high level of agreement.

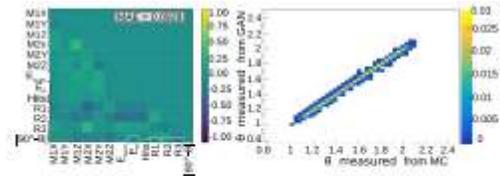


FIGURE 12. (Left) The difference between correlation matrices exhibiting internal correlations among shower inputs and physics features for Monte Carlo and GAN generated events. (Right) Correlation between  $\theta$  measured from Monte Carlo events and that measured from GAN events generated for similar  $\theta$  values.

### F. FURTHER VALIDATION

The main validation of 3DGAN performance involved detailed comparison of physics-based features across the whole range of the primary energies and the incident angles as explained in Section III. We have further explored validation from diverse standpoints including further mathematical formulation of the physics performance, as well as validation by a neural network and image quality metrics.

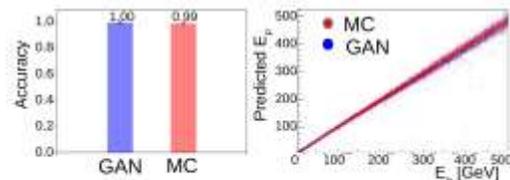


FIGURE 13. Triforce results for GAN vs. Monte Carlo events (9834 electron events for each type). (Left) Classification accuracy. (Right) Scatter plot between true and predicted  $E_p$ .

### 1) VALIDATION BY THIRD PARTY NETWORK

The most popular metrics used in GAN literature involve an external pre-trained network. Similar to what is done in general image generation problems, where the output of the inception network is used to quantify the generated image quality [55], [56], we have used the independently developed pre-trained (trained on the same dataset) classifier and regressor network

from Triforce [57] to validate our results. We stress here that the actual accuracy is not relevant for our test as that depends on the Triforce performance, but obtaining a similar performance level for GAN and MC events is. Figure 13 shows the classification and regression results obtained by running Triforce on 9834 images from both GAN and MC test samples. The GAN images show a similar response from the Triforce network, as MC images.

Thus proving that the GAN images contain most of the features of the MC data, as learnt by the Triforce network.

## 2) QUANTIFYING PHYSICS PERFORMANCE

The physics performance validated in Section V is based on the study of binned distributions, therefore we test several histogram distances. The metrics like Chi2 and Kolmogorov get saturated for small discrepancies in a few bins. These metrics are highly dependent on external factors like binning and ranges and thus are not suitable for combining as averages. On the other hand, mean errors do not describe well the performance over the whole distribution. In order to further quantify the performance, we calculate the relative errors for individual bins of a distribution. The mean of all bin errors is then treated as a metric. We have calculated this metric on physics features like the shower shapes, moments, and sampling fraction. These metrics can also be combined as a single figure for a future hyper-parameter effort.

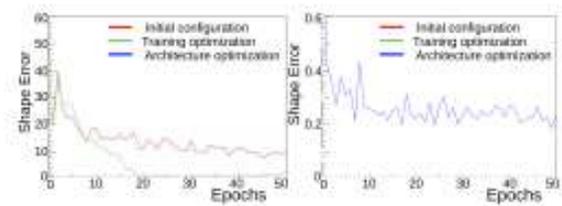
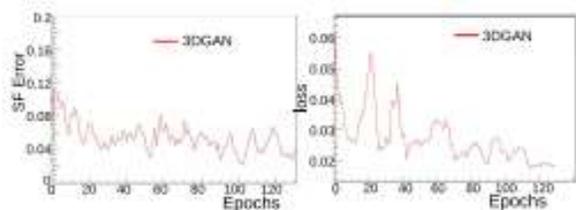


FIGURE 14. Mean relative errors on shower shapes for the first training step with restricted  $E_p$  range. (Left) Initial implementation without optimization (red), with pre-processing and additional losses (green), and after applying architecture modifications (blue). (Right) Reducing the Y axis scale to observe the error for the 3DGAN final version (blue), where it can be seen that the error is now saturated for further analysis.

In the current work we will only explore the metric as a function of training epochs for the different stages of 3DGAN development. The shower shapes provide an example of geometrical features that are learned by the model implicitly. Figure 14 presents the error for shower shapes associated with different versions of 3DGAN as described in Section IV-D for the first step of training (with restricted EP range). The initial configuration has high errors with mean and variance decreasing with epochs (red). The high values of error can be attributed to the tails of the distributions, where the energy deposition is very low. The training optimization greatly reduces the error yet a slight over-fitting can be detected in the later epochs (green). In the final implementation (blue) the shape distributions improve significantly. As observed from Figure 14 right panel, the training has reached a level where the error remains low yet not decreasing much with the epochs and the MRE metric seems to be saturated.

The sampling fraction is an example of analytical functions of the image. The simulation requires a very precise SF and its small variation through epochs is significant

for performance. Figure 15 right panel presents the error on SF averaged over 5 epochs. It can be observed that the error is below 10% for most of the training but in order to further improve the accuracy, we select the final network among the last few epochs (10 for current application) based on the minimum mean relative error, evaluated on holdout samples.



**FIGURE 15.** Validation of the 3DGAN final version for the restricted energy range training step. (Left) Mean relative error for SF vs. epochs (with average smoothing of order 5). (Right) Entropic Gromov Wasserstein Discrepancy calculated on correlation matrix vs. epochs (with average smoothing of order 5).

### 3) IMAGE QUALITY ANALYSIS

The assessment of the similarity/diversity of the generated sample compared to the original sample is an important step for the development of generative models. In the case of simulation, the shower produced by particles with the same EP and  $\theta$  should show similar features while retaining the expected statistical variance. In order to evaluate this aspect, we utilize some metrics for image quality assessment.

The Structural Similarity Index Measure (SSIM) [60] has been used for GAN performance evaluation [34]. The SSIM quantifies the quality of images based on their similarity to a reference. For images  $x$  and  $y$ , the SSIM is computed between windows from both  $x$  and  $y$ . If the mean and standard deviation of both windows are  $\mu_x$ ,

$\sigma_x$  and  $\mu_y$ ,  $\sigma_y$  respectively, then SSIM is given by the following equation (2):

$$SSIM(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \quad (2)$$

where  $C_1 = (k_1 L)^2$  and  $C_2 = (k_2 L)^2$  with  $k_1 = 0.01$  and  $k_2 = 0.03$ . The value of “L” is usually taken to be equal to the dynamic range of the pixel intensities. The parameter L thus determines the fuzz factor added against the weak denominator to avoid zero division. A value of SSIM close to one indicates very similar images while a lower value indicates more diverse images. In the context of GAN, the similarity metric is also treated as a measure of diversity.

The SSIM is sensitive to the pixel dynamic range (by the virtue of parameter L) and should be adjusted accordingly. Figure 11 (left) shows that our “pixel intensities” have a very different dynamic range as compared to the standard RGB case, with a logarithmic difference between the maximum and minimum values. Therefore, the experiment is repeated for different values of L in order to verify the sensitivity of the SSIM to our range of pixel intensities (energy depositions).

We measure SSIM for MC images against MC images, MC images against GAN images, and GAN images against GAN images. The events in each set are random events from the same EP and  $\theta$  bin. The SSIM value from MC vs. MC images can act as a reference to assess the GAN performance.

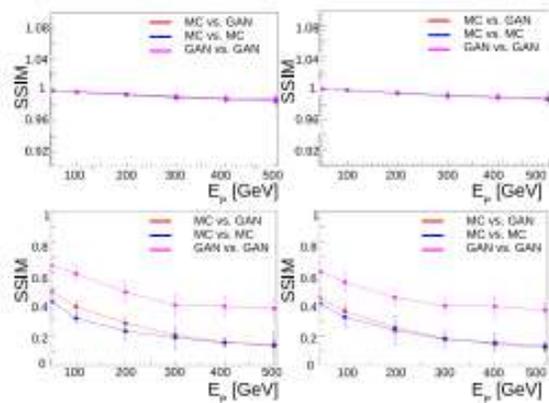


FIGURE 16. SSIM for  $L = 1$  (top row) and  $L = 0.0001$  (bottom row). The SSIM is computed for MC vs. GAN events (red), MC vs. MC events (blue), and GAN vs. GAN events (magenta) as a function of  $E_p$ . (Left column)  $\theta = 62^\circ$  bin. (Right column)  $\theta = 118^\circ$  bin.

The SSIM value for  $L = 1$  (as generally suggested for float intensities) is close to one and identical for all three sets of images as shown in Figure 16 top row, indicating that the index is not sensitive to the difference present between individual samples in a bin at this scale. The SSIM values decrease with  $L$  and a difference can be observed between MC vs. MC and GAN vs. GAN (indicating that the metric is now more sensitive). Decreasing  $L$  below  $10^{-4}$ , has no further effect. Figure 16 bottom row shows the SSIM for  $L$  equal to  $10^{-4}$ . The GAN vs. GAN SSIM is slightly higher than MC vs. MC, indicating that the GAN samples have less diversity in a bin, but is of the same order. The SSIM for MC vs. GAN is similar to MC vs. MC, proving that GAN images come from the same distribution as MC images. The figure also shows that SSIM is constant across the  $\theta$  spectrum.

The Mean Subtracted Contrast Normalized Coefficients (MSCN Coefficients) [61] have been used for blind image quality assessment. The MSCN Coefficients are

computed by taking the mean and the standard deviation of windows in an image. Let  $\mu$  and  $\sigma$  be the mean and standard deviation of an image window then a pixel intensity  $I$  will be converted to coefficient  $I^*$  by the following equation:

$$I^* = \frac{I - \mu}{\sigma} \quad (3)$$

These MSCN Coefficients can then be mapped to an image quality score [61]. The distributions of the MSCN Coefficients from real and fake images have also been used for GAN evaluation [62]. The 3DGAN images are highly sparse with small energy depositions thus the Equation (3) is modified as following :

$$I^* = \frac{(I - \mu_{\text{nonzero}})}{\sigma_{\text{nonzero}} + \text{epsilon}}$$

Here we stabilize by adding a fuzz factor in the denominator (taking epsilon value equal to  $10^{-7}$ ). The mean ( $\mu_{\text{nonzero}}$ ) and the standard deviation ( $\sigma_{\text{nonzero}}$ ) are evaluated considering only the nonzero entries. Figure 17 compares the histograms of the coefficients as computed using Equation (4) for MC and GAN generated images. The nature of our data is very different from natural scenes and thus the distributions are not Gaussian. The MSCN Coefficient distributions can still be treated as statistical signatures and are almost identical for the real and fake images, showing very similar features. The Jensen Shannon Distance (JSD) between the real and fake distributions is very small, thus indicating the histograms to be highly similar.

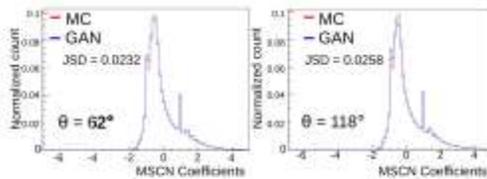


FIGURE 17. MSCN coefficient histograms for MC and GAN images for  $E_p = 2 - 500$  GeV. Left)  $\theta = 62^\circ$ . Right)  $\theta = 118^\circ$ . The JSD denotes the Jensen Shannon Distance between the two distributions.

## VI. CONCLUSION

The Monte Carlo simulation of detector response is a time and resource intensive task. There is a drive in the HEP community to find possible fast alternatives while maintaining physics accuracy. Fastsim is a set of established techniques usually replacing parts of the detailed simulation where some loss in accuracy can be traded for speed. These techniques have been incorporated into existing particle simulation packages frequently employed in practical applications such as GFlash [5], Alftast [4], and FastCaloSim [63] etc. We demonstrate that the fast simulation based on deep learning can surpass these methodologies both in speed and accuracy. The detector response can be generated as images and easily integrated into a detailed simulation, in the same manner as existing fast simulation methodologies.

The 3DGAN model is capable of reproducing single particle (electron) showers for a high granularity calorimeter representative of detectors designed for future particle accelerators. We have validated the 3DGAN performance by comparing it with the classical Monte Carlo in great detail. The agreement is within a few percent over a very large dynamic

range. This impressive level of agreement surpasses similar deep learning-based applications to HEP simulation [30] and fast simulation approaches based on classical methods [5]. The inference time on Intel Xeon 8180 is about 13.4 m sec per particle whereas about 3.5 m sec per particle on GeForce GTX 1080. For comparison, GEANT4 can simulate a similar shower in about 17 seconds per particle on an Intel Xeon 8180 (currently it is not possible to run a full GEANT4-based simulation on GPUs). Therefore, the 3DGAN provides three orders of magnitude speedup.

## ACKNOWLEDGMENT

This work has been conducted with the support of Intel in the framework of the CERN openlab-Intel collaboration agreement. Part of this work was conducted at “Banks,” the AI GPU cluster at Caltech. The authors acknowledge NVIDIA, SuperMicro, and the Kavli Foundation for their support of “iBanks.” They thank Matt Zhang from the University of Illinois at Urbana-Champaign for helping in validation by the Triforce model.

## REFERENCES

- [1] Worldwide LHC Computing Grid. Accessed: Sep. 26, 2021. [Online]. Available: <http://wlcg.web.cern.ch/>
- [2] J. Albrecht et al., “A roadmap for HEP software and computing R&D for the 2020s,” *Comput. Softw. Big Sci.*, vol. 3, no. 1, p. 7, Mar. 2019.

- [3] G. Apollinari, I. A. Béjar, O. Brüning, P. Fessia, M. Lamont, L. Rossi, and L. Tavian, “High-luminosity large hadron collider (HL-LHC): Technical design report V. 0.1,” CERN, Geneva, Switzerland, Tech. Rep. CERN2017-007-M, 2017, pp. 1–516, vol. 4, doi: 10.23731/CYRM-2017-004.
- [4] W. Lucas, “Fast simulation for ATLAS: Atlfast-II and ISF,” in Proc. Int. Conf. Comput. High Energy Nucl. Phys., vol. 396, 2012, pp. 1–10.
- [5] D. Orbaker, “Fast simulation of the CMS detector,” in Proc. Int. Conf. Comput. High Energy Nucl. Phys., vol. 219, 2010, pp. 1–7.
- [6] D. Autiero et al., “Parameterization of e and  $\gamma$  initiated showers in the NOMAD lead-glass calorimeter,” Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip., vol. 425, p. 188, Aug. 1998.
- [7] E. Barberio, J. Boudreau, B. Butler, S. L. Cheung, A. Dell’Acqua, A. D. Simone, E. Ehrenfeld, M. V. Gallas, A. Glazov, Z. Marshall, J. Mueller, R. Plačakyte, A. Rimoldi, P. Savard, V. Tsulaia, A. Waugh, and C. C. Young, “Fast simulation of electromagnetic showers in the ATLAS calorimeter: Frozen showers,” J. Phys., Conf. Ser., vol. 160, Apr. 2009, Art. no. 012082.
- [8] J. Schaarschmidt et al., “The new fast calorimeter simulation in ATLAS,” ATLAS Collaboration, CERN, Geneva, Switzerland, Tech. Rep. ATL-SOFT-PUB-2018-002, Jul. 2018. [Online]. Available: <https://cds.cern.ch/record/2630434>
- [9] G. R. Khattak, S. Vallecorsa, and F. Carminati, “Three dimensional energy parametrized generative adversarial networks for electromagnetic shower simulation,” in Proc. 25th IEEE Int. Conf. Image Process. (ICIP), Oct. 2018, pp. 3913–3917.
- [10] G. R. Khattak, S. Vallecorsa, F. Carminati, and G. M. Khan, “Particle detector simulation using generative adversarial networks with domain related constraints,” in Proc. 18th IEEE Int. Conf. Mach. Learn. Appl. (ICMLA), Dec. 2019, pp. 28–33.
- [11] Y. Bengio, E. Thibodeau-Laufer, and J. Yosinski, “Deep generative stochastic networks trainable by backprop,” CoRR, vol. abs/1306.1091, pp. II-226–II-234, 2013. [Online]. Available: <http://arxiv.org/abs/1306.1091> and <http://dblp.uni-trier.de/rec/bib/journals/corr/BengioT13>
- [12] D. P Kingma and M. Welling, “Auto-encoding variational Bayes,” Dec. 2013, arXiv:1312.6114. [Online]. Available: <http://arxiv.org/abs/1312.6114>