

Performance Comparison of Full and Hybrid QGANs for Particle Shower Simulation

Bancy Mbura Maranga

Department of Electrical and Information Engineering, University of Nairobi, Nairobi, Kenya.
mburamaranga@uonbi.ac.ke

Article Info

Elaris Computing Nexus
https://elarispublications.com/journals/ecn/ecn_home.html

Received 02 June 2026
Revised from 28 July 2026
Accepted 20 August 2026
Available online 02 September 2026
Published by Elaris Publications.

© The Author(s), 2026.

<https://doi.org/10.65148/ECN/2026014>

Corresponding author(s):

Bancy Mbura Maranga, Department of Electrical and Information Engineering, University of Nairobi, Nairobi, Kenya.
Email: mburamaranga@uonbi.ac.ke

This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Abstract – The use of quantum generative adversarial networks (qGANs) is based on an efficient approach that simulates high-energy particle showers using a reduced model complexity. In this study, we analyze complete and hybrid qGAN architectures for the creation of particle shower images and compare their training behavior, convergence and accuracy. The results show a stable training of the entire qGAN, which can generate the general parameters of the shower well. The hybrid design also exhibits competitive performance with the medium discriminator yielding the highest accuracy of all the tested configurations. However, with a lower or higher discriminator the decrease of performance is due to the imbalance between the capacity of the generator and discriminator. The results show that quantum models can be represented with a much smaller number of trainable parameters, and still capture the representational capabilities of classical models, something that can be studied in greater detail with larger models and more sophisticated physics-based metrics.

Keywords – Quantum Generative Adversarial Networks, Particle Shower Simulation, Hybrid Quantum-Classical Models, Quantum Machine Learning, Generative Models, High-Energy Physics.

I. INTRODUCTION

In the field of particle and nuclear physics, modern experiments have high demands for the development of a complex software program and frameworks, which have to be highly reliable. This enhances the need for large scale, thorough and precise models of the particle detection sensors utilized in this research. The increasing size, complexity and sensitivity of the detectors are driving the need for low-cost, high-capacity computer models that enable the development of larger and more complex models. These considerations are found in many other applications, including the physics of radiation, space research, nuclear medical imaging and, indeed, any application in which particles interact with matter [1].

A particle shower simulation is a computer simulation of the shower of secondary particles that occur when a particle with a high energy hit materials. It is a very important process in the field of high energy physics and radiation science. The technology employs Monte Carlo techniques and applications such as Geant4 to follow electromagnetic and hadronic showers, but often depends on quick simulation techniques such as deep generative AI to reduce the huge processing expenses. It is a significant tool to design the detector, plan the cancer treatment and analyze the information of cosmic ray or satellite radiation [2]. The new simulation framework called Geant4 is the solution to this problem, and is object oriented. The tool-kit provides a variety of fully featured and integrated pieces of software that can be used in many situations. These are pure single literature of basic geometries and processes, detailed detector models for experiments at the LHC (large hadron collider) and other experiments.

The adaptability, versatility and range of Monte Carlo simulation have led to its employment in a variety of disciplines. They include physical sciences, finance, engineering, risk assessment, statistics, etc. Monte Carlo simulation can be used for a variety of reasons, from determining the value of π to examining particular aspects of complicated systems, such as multi-dimensional integrals, computational physics or light beams in a bounded space using MCMC. MCS makes the modeling and simulation easy by applying algorithms that exponentially speed up the process. As the computer industry develops

forward and processing rates grow in the 21st century, MCS will be faster and more adaptable in many areas such that researchers will be able to duplicate practically anything of interest. Furthermore, the use of randomization in the simulation is a proof of the effectiveness of this method for making inferences on a broader population using sample data [3].

To achieve effective and open generative simulation, a promising quantum machine learning concept of qGANs (quantum generative adversarial networks) was proposed as a potential solution, which uses adversarial training with quantum computing to achieve this objective. Qi et al. [4] offer a comprehensive overview of the qGAN models and reveal the tremendous advances from the fundamental ideas to actual applications. The quantum-classical hybrid architectures of existing QGAN designs are classified and their applications are summarized in various areas including image generation, clinical data generation, software bug detection, channel prediction and teaching materials. We are interested in the synergy of QGANs and domain-specific methods including variation circuits, large language models, Wasserstein distance and optimization heuristics, etc.

Boyle and Nikandish [5] consider the problem of efficient transfer of information from classical memory to quantum computers on qGANs. The work highlights the significance of tuning the hyperparameters and the optimizers to address the problems of the longer training time and to increase the accuracy as a whole. The objective of this work is to make qGAN more efficient, by making these components more powerful to transfer the data smoothly and quickly, and also to participate in the discussion about utilizing quantum computing to enhance data loading procedures.

Preda et al. [6] introduced a novel approach for simulating a variety of quantum many-body systems with conditional GANs. With this method, the existence of detailed dynamic excitation spectrum for Hamiltonian is revealed over the entire parameter range of the questions. The conditional GAN technique is astounding in its efficiency, giving in a short time results accurate enough to be compared with those of thorough exact analyzes. This unique approach paves a way to further investigate and understand complex quantum multi-body processes by means of precise and efficient simulations.

Zhou et al. [7] explore how to use hybrid quantum-classical models for GANs to generate images. In this work, they trained the discrete distribution for better image generation and investigated performance of this method. Their research is on the ability to utilize quantum and classical features in GANs, with implications for enhanced image generating skills. The study of hybrid frameworks offers a possibility to use quantum-enabled benefits in discrete probabilistic modeling, thus enriching the emerging discipline of quantum-based machine learning (ML) in image generation.

Based on qGANs, Turkeli and Ozaydin [8] use the quantum gradient calculation to not only efficiently update the circuit parameters but also do not require classical backpropagation. This methodology is motivated by distributions that do not admit a normal representation and thus is potentially quantum advantageous for quantum state learning, cryptography and simulation. A qGAN is composed of two parameterized quantum circuits: generator (G) and discriminator (D), which are competing with each other. The generator employs a label register and noise register and generates a quantum state, which is expected to mimic samples from a real data source. To tell whether the output of the generator is “real” or “fake”, the discriminator is designed to be a unitary quantum circuit and is given as an input the output of the generator or a true state sampled from the source of actual data.

In this paper, we concentrate on complete and hybrid quantum GANs for the simulation of particle showers. We explore their training stability, convergence and generation accuracy, and we show how the complexity of the discriminator affects the performance. Particular emphasis is put on the topic whether fully quantum systems can outperform hybrid quantum-classical models with much less trainable parameters. The remainder of this study can be summarized as follows: Section II describe the dataset, experimental set up, qGAN architectures, training and evaluation methodology. The performance of the full model and the hybrid model is presented in Section III. In Section IV, we cover accuracy, stability, discriminator capacity, parameter efficiency and limits. Finally, the key findings are summarized in Section V and their implications for simulation research are explored.

II. DATA AND METHODS

Dataset and Experimental Framework

We investigate full quantum and hybrid quantum-classical generative adversarial networks (qGANs) for particle-shower simulation in a computational experimental setting. The reference distribution $p_{G4}(x)$ is formed from particle-shower images generated by Geant4. The generator G_{θ_G} receives an input sample $z \sim p_z(z)$ and generates a generated shower image,

$$\hat{x} = G_{\theta_G}(z), \hat{x} \sim p_G(\hat{x}), \quad (1)$$

where θ_G specifies the parameters of the generator. The antagonistic aim is

$$\min_{\theta_G} \max_{\theta_D} V(D, G) = \mathbb{E}_{x \sim p_{G4}} [\log D_{\theta_D}(x)] + \mathbb{E}_{z \sim p_z} \left[\log \left(1 - D_{\theta_D}(G_{\theta_G}(z)) \right) \right] \quad (2)$$

Training and inference run with 1024 shots on all the tests on quantum simulators. Both the full qGAN and each of the hybrid discriminator configurations were trained for 20 independent trials and the results of each configuration for the hybrid discriminator were tested for 10 trials. The repeated trials provide an indication of the variability and stability of training with the model and hyperparameters held constant.




Fig 1. Experimental Process.

The sequence of the experiment is shown in **Fig. 1**. The comparison is done based on the generation accuracy, convergence, discriminator capacity and trainable parameter efficiency. The pixel-wise correlations and the more advanced physics observables are considered as extra metrics, not as the main criteria of evaluation [9]. **Table 1** shows experimental configuration.

Table 1. Experimental Configuration.

Parameter	Full qGAN	Hybrid qGAN
Generator learning rate	0.02	0.01
Discriminator learning rate	0.04	0.006
Generator decay	0.006	N/A
Discriminator decay	0.007	N/A
Joint decay	N/A	0.006
Epochs	1000	1000
Optimization steps per epoch	8	1
Batch size	8	N/A
Shots	1024	1024
Discriminator updates	5 per generator update	1 per generator update
Generator updates	1 per optimization step	1 per optimization step

The hybrid batch size is not specified and is consequently given as “N/A”.

qGAN Architecture and Training

The full qGAN consists of a MERA-up quantum generator and a MERA-down quantum discriminator, both having 20 trainable parameters. The hybrid qGAN is based on the MERA-up generator and the quantum discriminator is replaced with a regular, fully connected network. We explore three capacities of Discriminator: Hybrid-small (153 parameters), Hybrid-medium (433 parameters) and Hybrid-large (1889 parameters). The settings in **Table 2** enable us to decouple the influence of representational capacity of the discriminator on adversarial training balance and accuracy of the generated samples.

The parameterized quantum transition is described by

$$|\psi(\theta, z)\rangle = U(\theta, z) |0\rangle, U(\theta, z) = U_L(\theta_L, z) \cdots U_2(\theta_2)U_1(\theta_1), \quad (3)$$

where $|0\rangle$ is the initial circuit state, U is the parameterized circuit, z is the generator input, and θ includes the executable rotational parameter values. The variation parameters are initialized to 0 and the quantum-circuit assessments are performed using 1024 shots.

Two-sided label smoothing is applied as described in

$$y_{\text{real}} = 0.9, y_{\text{fake}} = 0.1. \quad (4)$$

The discriminator loss for minibatch $\mathcal{B}$ is formulated as

$$\mathcal{L}_D = -\frac{1}{|\mathcal{B}|} \sum_{j \in \mathcal{B}} \left[y_{\text{real}} \log D(x_j) + (1 - y_{\text{real}}) \log (1 - D(x_j)) + y_{\text{fake}} \log D(\hat{x}_j) + (1 - y_{\text{fake}}) \log (1 - D(\hat{x}_j)) \right] \quad (5)$$

where x_j and $\hat{x}_j = G(z_j)$ are the reference and generated images, respectively. The discriminator loss is

$$\mathcal{L}_G = -\frac{1}{|B|} \sum_{j \in B} \log D(\hat{x}_j). \quad (6)$$

The full qGAN utilizes learning rates $\eta_G = 0.02$ and $\eta_D = 0.04$. The discriminator is updated 5 times for each generator update. Training is done for 1000 epochs with 8 optimization phases per epoch and a mini-batch size of 8. Generator decay 0.006 Discriminator decay 0.007 The exponential decay schedule is as follows:

$$\eta_k = \eta_0 \exp(-\lambda k), \quad (7)$$

where η_0 is the initial learning rate, k is the number of the update, and λ is the decay coefficient.

The hybrid qGAN utilizes $\eta_G = 0.01$, $\eta_D = 0.006$ and a joint decay of 0.006. The training is done for 1000 epochs with one optimization step each epoch and both generator and discriminator are updated once per step. For both families of models, hyperparameters are determined via Optuna and empirical testing.

Parameter updates are made using SPSA:

$$\theta_{k+1} = \theta_k - a_k \hat{\nabla}_{\theta} \mathcal{L}(\theta_k), \quad (8)$$

with

$$\hat{\nabla}_{\theta} \mathcal{L} = \frac{\mathcal{L}(\theta_k + c_k \Delta_k) - \mathcal{L}(\theta_k - c_k \Delta_k)}{2c_k} \Delta_k^{-1}, \quad (9)$$

where a_k is the update coefficient, c_k is the perturbation size and Δ_k is the SPSA perturbation vector. The two model architectures are presented in **Fig. 2**. The full architecture is as follows:




Fig 2. Full and Hybrid qGAN Architectures.

whereas the hybrid architecture is shown in **Fig. 3**.

The hybrid discriminator branch is in S , M , and L configurations.

Table 2. Architecture and Parameter Complexity of qGAN

Model	Generator	Discriminator	Trainable Parameters
Full qGAN	MERA-up quantum	MERA-down quantum	20
Hybrid-S	MERA-up quantum	Classical fully connected	153
Hybrid-M	MERA-up quantum	Classical fully connected	433
Hybrid-L	MERA-up quantum	Classical fully connected	1889

Performance Evaluation and Comparative Analysis

Generation quality is mostly evaluated by pixel-wise mean squared error:

$$MSE = \frac{1}{N} \sum_{i=1}^N (x_i - \hat{x}_i)^2, \quad (10)$$




Fig 3. Hybrid Architecture.

where x_i and $\hat{x}_i$ are the reference and created pixel values, respectively, and N is the number of assessed pixel-value pairs. The MSE is calculated for 50 generated photos and the training/reference images. The smaller the numbers the closer the agreement.

For R independent trials the mean and standard deviation are

$$MSE = \frac{1}{R} \sum_{r=1}^R MSE_r, \quad s_{MSE} = \sqrt{\frac{1}{R-1} \sum_{r=1}^R (MSE_r - MSE)^2}. \quad (11)$$

Here $R = 20$ for the full qGAN and $R = 10$ for each hybrid model. The average shower images are also examined using

$$\bar{\hat{x}}_i = \frac{1}{M} \sum_{m=1}^M \hat{x}_{i,m}, \quad (12)$$

where M is the number of created images and $\hat{x}_{i,m}$ is pixel i in image m . This average allows you examine qualitatively the global shape of the shower.

Convergence is measured through the generator objective loss, real discriminator objective loss, fake discriminator loss and cumulative objective discriminator loss. Plots of the different training histories and MSE trajectories are plotted against epochs and independent trials. The hybrid and full models are optimized with separate schedules – curves are interpreted both in terms of number of epochs and number of optimization steps. The comparison is made with respect to mean square error, trial variability, convergence behavior, mean number of shower photos, discriminator capability and number of trainable parameters. It's not meant to be about quantum supremacy or superiority of architecture. Rather, it investigates the effectiveness of the qGAN with the complete set of 20 parameters versus the hybrid methods, and the role of the discriminator capacity on the adversarial balance (Table 3). In Section III, the results of the quantitative analysis and the MSE of four configurations of the model are presented.

Table 3. Quantitative Performance

Model	Number of Trials	MSE	Interpretation
Full qGAN	20	1.77 ± 1.46	Competitive
Hybrid-S	10	6.65 ± 3.11	Lower performance
Hybrid-M	10	1.50 ± 2.15	Best MSE
Hybrid-L	10	3.80 ± 4.32	Lower performance

III. RESULTS

Here, we provide results from the qGAN experiments: training performance, convergence, inference accuracy and comparison of the full and hybrid quantum-classical architectures. The results explore the influence of some hyperparameters and learning rates, discriminator configurations and model complexity. The accuracy is mainly evaluated with the MSE over the generated shower images.

Full qGAN Training and Performance

Distinct training rates are used for the discriminator and the generator based on experimental evidence where the training rate for the generative network is calibrated to 0.021 and for the discriminative network it is configured to 0.04. Learning congregates much easier if the discriminative network is learned more often compared to the generative network. Tsang et al. [10] reported a similar result for a hybrid qGAN system. The discriminative network is learned 5 epochs for every optimization iteration in the full qGAN framework, followed by a single learning phase of the generator. The system undergoes training for 1000 epochs with 8 optimization phases (mini-batch size of 8) every epoch.

The discriminator and generator are trained at different rates (0.006 and 0.007 respectively) and they are also exponentially decayed. Training and Inference is done on 1024 samples. The tuning parameters are decided via hyperparameter optimization searches with the Optuna optimization framework [11] and experimental experience. In this study, all the trials mentioned are performed on the quantum simulation platforms. For training we use two-sided label smoothing where a true label of 1 is set to 0.9 and a false label of 0 is set to 0.1. This results in improved precision train/final convergence. All rotation measurement will be set to 0 in order to have a stable start of training. The SPSA optimizer [11] is used to determine the parameter modification.




Fig 4. The Accuracy Metric MSE Across Training Epochs for 20 Experiments.

The results of 20 training trials (with the same hyperparameters) with respect to the MSE criterion are displayed in **Fig. 4**. We computed MSE per pixel of 50 engendered images and learning epochs. As a result of tuning parameter optimization, the average estimate (green) rises up quickly and remains constant. The light purple frequency band reflects a limited standard deviation while the green solid line shows an ideal trial with little variations. All these metrics highlight an integrated, coherent and effective learning.




Fig 5. Unique Error Values of The Ideal Trial.

The learning losses of the optimal experimentations are plotted separately in **Fig. 5**. The graph shows the generator objective function loss (red), the true loss of the discriminator objective function (purple), the false loss of the discriminator (blue) and the total losses of the discriminator (light green). As a result of minimal (0) configuration of the hardware constraints, the hybrid discriminator objective function error is smaller compared to the error in the beginning of the initial iteration. The errors are aligned across epochs to converge towards similar levels of losses, which leads to effective training convergence.

Moreover, when training is complete, the losses will stabilize due to the use of a reduced learning rate. The losses are within a wider range between single epochs, probably as a result of a low batch number. Increasing the number of batches each iteration may be able to reduce the loss. The number of generated images is roughly equivalent to the number of images in the Geant4 test data collection to test the accuracy of the inference and the mean shower image (MSI) of the optimal trial is shown in **Fig. 6**. The complete qGAN system is able to regenerate the MSI of Geant4 accurately, resulting in a total MSE of: 1.76 ± 1.45 .

Hybrid qGAN Performance and Model Assessment

A thorough study of a quantum-classical GAN hybrid has been conducted, which consists of a MERA-up quantum generative network and conventional discriminative network. This analysis may allow to see possible benefits of quantum behavior in comparison to classical behavior or the other way round. The conventional discriminator is a fully connected network. We look at three different sizes of discriminators: large with 1889 model parameters to train, medium with 433 parameters to train and a small with 153 parameters to train. The tuning parameters used for training convergence in the hybrid qGAN model are: generator learning rate of 0.01, discriminator training rate of 0.006 and combined learning rate decay of 0.006. The training is performed over 1000 “epochs”, where one “epoch” is one optimization phase. Both the discriminator and the generator are augmented together for each iteration of training.




Fig 6. The Typical Shower Image Formed at Inference.

We performed 10 training experiments for each of the 3 models (see **Fig. 7**). For all models, the mean squared error (MSE) levels off. Training the hybrid large model results in the lowest MSE (1.51 ± 2.14) at the end of training, which means it is the most accurate model. The worse performance of the small (6.64 ± 3.10) and large (3.81 ± 4.31) model architectures is explained by the GAN principle, the best accuracy in adversarial training is achieved when there is a balance between the discriminator and generator model architectures. The accuracy decreases due to the fact that the large (small) discriminator is more apt than the generative network. The medium model shows a representational capacity most similar to that of the generator.




Fig 7. MSE as a Performance Metric W.R.T. The Learning Iteration for the 3 Composite Qgan Architectures.

The mean shower image of the best hybrid medium model is displayed in **Fig. 8**. It demonstrates a good agreement with a Geant4 forecasting. The precision of the entire qGAN architecture (1.76 ± 1.45) is the same as that of the optimal integrated hybrid architecture medium (1.51 ± 2.14), showing that the representative capabilities of the full quantum discriminator and the hybrid discriminator are almost the same according to the data. This also means that the MERA-based quantum discriminator architecture consisting of 20 modeled parameters is working as well as a traditional neural system with more than 400 learnable model parameters. This raises an important question: Can parametric quantum circuits with much lower learned parameters achieve a comparable representational energy to artificial neural networks?

Other similar works, including [12] and [13] also experiment with quantum circuits, which are said to be more efficient in computing, have better storing and modeling strength, using quantum features. Importantly, the trends of correlation at the pixel level cannot be reproduced by the complete qGAN nor the mixed qGAN architectures, which are more complex physics precision measurements that are performed beyond the MSI.




Fig 8. The Mean Shower Configuration Plot from the Best Trial of the Hybrid Qgan Model.

As a comparison, an integrated qGAN was optimized with 3 different discriminator dimensions according to the parameter values. The prototype with a medium-sized discriminative network was the only model to achieve MSE numerical values similar to the complete qGAN. The results were poor when both a large discriminator and a small one was trained due to the energy imbalance between the discriminator and generator architectures. It should be noted that the complete qGAN architecture with only 20 model coefficients, has achieved a statistical equivalent precision to the integrated model (medium) with 433 model coefficients.

IV. DISCUSSION

Accuracy and Training Stability

Our results demonstrate that the overall qGAN can learn the global structure of the Geant4 particle-shower distribution and still be stable from the adversarial training. The mean MSE tends to converge, and the variability between runs is not so high, indicating that the training choice used here results in a stable training regime. This result is in line with the recent full-qGAN study for the average calorimeter shower image using a MERA-up and a MERA-down quantum-based generator [14] which also showed stable training.

This is relevant correspondence to the previous work because the same general full-qGAN architecture and training technique is used in the current study. Liu and Hsieh [15] also observed that discriminator needed to be optimized more than the generator and that different learning rates were needed, which is indicative of the training strategy used here. They also noted that they saw oscillations as a result of the small number of batches, which would indicate that the changes they saw during training were not necessarily as a result of poor optimization.

However, copying of the average shower appearance does not mean that they are copying the whole particle shower distribution. To validate the classical GAN based calorimeter simulation physically, observables other than image-level similarity such as energy deposition, fluctuations and correlations should be analyzed in shower events. Kotwal and Hays [16] have, for example, investigated the physical observables of the showers built in the calorimeter, and compared their distributions with Geant4 for certain aspects of the comparison. The stable MSE, therefore, in the present study indicates good image level reconstruction, but not necessarily physical integrity per se.

Capacity and Parameter Efficiency of The Discriminator

The hybrid tests provide an interesting insight into the discriminator ability and generative performance. The medium classical discriminator has the lowest MSE of all the hybrid configurations tested, with the smaller and greater discriminators

doing badly. This shows that the performance of qGANs does not necessarily improve with the increase in the number of parameters in the discriminator. Instead, the discriminator should have at least as powerful representational capacity as the generator.

This is consistent with the full-qGAN study that is closely related to this work, Pathak, Yilmazer, and Patel [17] which compared a quantum discriminator with classical fully connected discriminators of varying sizes. Similarly, in that investigation, the medium hybrid setup yielded the highest accuracy, while the smaller and larger discriminators performed worse as a result of the difference in capacities of the discriminator and the generator. This renders the agreement more robust than a comparison to other GAN studies which are not directly related: the relation between the capacity of the discriminator and the adversarial performance has been identified in both studies.

In this context, the comparison has the more basic sense of the outcome of parameter-efficiency. It was found that the entire quantum discriminator has an order of magnitude lower number of trainable parameters, and achieves statistically equivalent MSE performance as the medium classical discriminator based on the provided experimental data. This finding consistent with the interest in learning quantum generative models as a small model choice for classical neural architecture. In related HEP research, quantum generative techniques have also been reported to be able to do competitive work, with one example being the ability of a qGAN-based approach to obtain comparable accuracy in detecting anomalies, but using much less training data in that specific HEP application [18].

However, the result achieved here can only be considered as evidence for the efficiency of parameters, but not as a proof of quantum computational advantage. Reduced number of trainable parameters does not imply reduced computation cost and increased scalability. Other factors that are important are the circuit depth, the number of circuit evaluations, shots of measurement, complexity of optimization and the cost of simulators and the noise of the final hardware. Moreover, the experiments of this paper are performed on quantum simulators. Competitive precision with a small quantum discriminator should not be interpreted as quantum advantage, but instead as a motivation to continue investigating this.

Limitations

However, we do extend the case to full quantum structure to produce particle showers while also hinting to a balanced architecture as opposed to the previous work on qGAN in high-energy physics. Rehm et al. demonstrated that one can produce full qGANs for single calorimeter shower images, and directly compare the quantum architecture with a hybrid one with classical discriminators. This representational efficiency of compact quantum discriminators, they found, is not a one-off discovery, but something that can be reproduced by researchers.

At the same time, current work is a critical qualification to the interpretation of successful qGAN picture production. Singh, Shankar and Prakash [19] have shown that quantum QGANs suffer from the deficiencies of generalization over data sets and conveyance to the average/dominating data representation of the training set. They learn that full quantum generative architectures today still have generalization as a big problem. This conclusion is directly relevant to the present work as up to now the best performance evidence is from MSE and average shower-image agreement. A model can be able to reproduce the overall shape of the shower, but it cannot capture higher order components of the shower distribution.

This limitation is also in line with original work on full-qGAN for HEP, which stated that average shower images and MSE are not a complete picture of physics correctness and that pixel-wise correlation patterns are an interesting avenue for future research. It is thus not possible to state the whole qGAN is the whole physical distribution of particle showers in the present study. The results rather show competitive image-level performance, and recommend a wider physics validation.

Another limitation is the number of samples, which is statistical in nature. The repeated training trials provide information regarding the variability in optimization but increase the number of training trials and/or the amount of evaluation data would provide better estimates of model uncertainty. This is especially significant in light of the relatively small difference between the full qGAN and the medium hybrid model, when compared to the spread between trials. Therefore, similarity in performance should be recognized in the light of the statistical accuracy of the present experiments.

V. CONCLUSION

Our results show the potential of building complete and hybrid qGANs for particle-shower simulation, and highlight the importance of properly tuning the discriminator capacity to the generator capacity. The most important consequence is that a much smaller quantum discriminator is able to attain an accuracy similar to that of a larger classical discriminator for the measured task. The evidence, however, is consistent with a competitive representational efficiency claim not necessarily with a proven quantum advantage claim. Future studies will need to be done not only at MSE, but also at the pixel level, and on energy dependent observables, shower shape distributions and other physics quantities. Here, it would be possible to see the parameter-efficiency on a scale on which the difference between the representational capacity of the quantum and the classical is more pronounced.

CRedit Author Statement

The author reviewed the results and approved the final version of the manuscript.

Data Availability

The datasets generated during the current study are available from the corresponding author upon reasonable request.

Conflicts of Interests

The authors declare that they have no conflicts of interest regarding the publication of this paper.

Funding

No funding was received for conducting this research.

Competing Interests

The authors declare no competing interests.

References

- [1]. Patera and I. Mattei, "Nuclear interactions and medicine," *Eur. Phys. J. Plus*, vol. 134, no. 1, Jan. 2019, doi: 10.1140/epjp/i2019-12484-6.
- [2]. G. C. Pereira, M. Traughber, and R. F. Muzic, "The role of imaging in radiation therapy planning: Past, present, and future," *Biomed. Res. Int.*, vol. 2014, pp. 1–9, Jan. 2014, doi: 10.1155/2014/231090.
- [3]. Bajard, S. Chabaud, D. Pérol, J.-P. Boissel, and P. Nony, "Revisiting the level of evidence in randomized controlled clinical trials: A simulation approach," *Contemp. Clin. Trials*, vol. 30, no. 5, pp. 400–410, Jul. 2009, doi: 10.1016/j.cct.2009.06.005.
- [4]. H. Qi, Y. Xu, H. Wang, A. Gani, and L. Y. Por, "Quantum generative adversarial networks: A comprehensive survey of theories, applications, and challenges in the NISQ era," *Quantum Inf. Process.*, vol. 25, no. 6, May 2026, doi: 10.1007/s11128-026-05186-1.
- [5]. O. Boyle and R. Nikandish, "A hybrid quantum-classical generative adversarial network for near-term quantum processors," *IEEE Access*, vol. 12, pp. 102688–102701, Jan. 2024, doi: 10.1109/ACCESS.2024.3433383.
- [6]. T. Preda et al., "Design of nanoscale quantum interconnects aided by conditional generative adversarial networks," *Appl. Sci.*, vol. 14, no. 3, Art. no. 1111, Jan. 2024, doi: 10.3390/app14031111.
- [7]. N.-R. Zhou, T.-F. Zhang, X.-W. Xie, and J.-Y. Wu, "Hybrid quantum-classical generative adversarial networks for image generation via learning discrete distribution," *Signal Process.: Image Commun.*, vol. 110, Art. no. 116891, Nov. 2022, doi: 10.1016/j.image.2022.116891.
- [8]. M. Islam, S. Turkeli, and F. Ozaydin, "Quantum generative adversarial networks: Architectures, use cases, and real-world implementations," *Quantum Inf. Process.*, vol. 25, no. 1, Jan. 2026, doi: 10.1007/s11128-025-05044-6.
- [9]. V. S. Ngairangbam, A. Bhardwaj, P. Konar, and A. K. Nayak, "Invisible Higgs search through vector boson fusion: A deep learning approach," *Eur. Phys. J. C*, vol. 80, no. 11, Nov. 2020, doi: 10.1140/epjc/s10052-020-08629-w.
- [10]. S. L. Tsang, M. T. West, S. M. Erfani, and M. Usman, "Hybrid quantum-classical generative adversarial network for high-resolution image generation," *IEEE Trans. Quantum Eng.*, vol. 4, pp. 1–19, Jan. 2023, doi: 10.1109/TQE.2023.3319319.
- [11]. V. S. Kumar, M. P. Chinnasamy, and S. R. Kumar, "Deep learning framework with OPTUNA-based hyperparameter tuning for predicting dry turning process performance of 42CrMo4 steel," *IEEE Access*, vol. 14, pp. 2119–2133, Dec. 2025, doi: 10.1109/ACCESS.2025.3649243.
- [12]. M. Yousif, B. Al-Khateeb, and B. Garcia-Zapirain, "A new quantum circuits of quantum convolutional neural network for X-ray images classification," *IEEE Access*, vol. 12, pp. 65660–65671, Jan. 2024, doi: 10.1109/ACCESS.2024.3396411.
- [13]. S. Chakraborty, T. Das, S. Sutradhar, M. Das, and S. Deb, "An analytical review of quantum neural network models and relevant research," in *Proc. 5th Int. Conf. Commun. Electron. Syst. (ICCES)*, 2020, pp. 1395–1400, doi: 10.1109/ICCES48766.2020.9137960.
- [14]. Zoufal, A. Lucchi, and S. Woerner, "Quantum generative adversarial networks for learning and loading random distributions," *npj Quantum Inf.*, vol. 5, no. 1, Nov. 2019, doi: 10.1038/s41534-019-0223-2.
- [15]. X. Liu and C.-J. Hsieh, "Rob-GAN: Generator, discriminator, and adversarial attacker," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 11226–11235, doi: 10.1109/CVPR.2019.01149.
- [16]. V. Kotwal and C. Hays, "Electromagnetic shower properties in a lead-scintillator sampling calorimeter," *Nucl. Instrum. Methods Phys. Res. A*, vol. 729, pp. 25–35, Jun. 2013, doi: 10.1016/j.nima.2013.05.185.
- [17]. P. Pathak, N. Yilmazer, and D. P. Patel, "Quantum GAN-based adversarial attack generation for network intrusion detection system," in *Proc. IEEE Int. Carnahan Conf. Secur. Technol. (ICCST)*, 2025, pp. 1–6, doi: 10.1109/ICCST63435.2025.11295591.
- [18]. R.-Z. He, J.-J. Su, S.-J. Qin, Z.-P. Jin, and F. Gao, "QGAN-based data augmentation for hybrid quantum-classical neural networks," *Chin. J. Phys.*, vol. 97, pp. 1453–1463, Jul. 2025, doi: 10.1016/j.cjph.2025.07.017.
- [19]. Singh, G. Shankar, and V. Prakash, "Quantum generative AI hybrid framework for securing general-purpose cyber-physical systems," in *Proc. Int. Conf. Electron. Renewable Syst. (ICEARS)*, 2026, pp. 354–361, doi: 10.1109/ICEARS67481.2026.11416624.

Publisher's note: The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations. The content is solely the responsibility of the authors and does not necessarily reflect the views of the publisher.

ISSN (Online): 3105-9082