PROCEEDINGS OF SCIENCE

Quantum Generative Modeling for Calorimeter Simulations on Noisy Quantum Devices

Saverio Monaco,[a,b,*] **Jamal Slim**,[a] **Florian Rehm**,[c] **Moritz Scham**, **Dirk Krücker**[a] **and Kerstin Borras**[a,b]

[a] *DESY,*
Notkestraße 85, Hamburg, Germany
[b] *RWTH University,*
Templergraben 55, Aachen, Germany
[c] *CERN,*
Esplanade des Particules 1, Meyrin, Switzerland

E-mail: saverio.monaco@desy.de

Quantum-based generative models offer a promising approach for simulating complex phenomena in high-energy physics, such as calorimeter showers used for particle identification and energy reconstruction at the LHC. We propose the Quantum Angle Generator (QAG), a variational quantum model trained with a Maximum Mean Discrepancy loss to generate images from the probabilistic outputs of quantum circuits. We study model and training hyperparameters, assess the impact of quantum noise during training and inference, and evaluate the QAG in both noiseless simulations and simulated Noisy Intermediate-Scale Quantum hardware. Results show that the QAG can adapt through learning to hardware-induced noise, yielding stable, high-quality outputs even under significant noise and calibration variability.

The European Physical Society Conference on High Energy Physics (EPS-HEP2025)
7-11 July 2025
Marseille, France

*Speaker

 ISSN 1824-8039 . Published by SISSA Medialab.

PoS(EPS-HEP2025)639

1. Quantum Computing in HEP

Simulating particle showers in electromagnetic calorimeters is a central task in High Energy Physics (HEP). Classical computing approaches, however, can be computationally demanding, especially in view of the upcoming High-Luminosity LHC and other future colliders. Quantum computing offers alternative computational paradigms that may help address these challenges. In particular, superposition and entanglement allow quantum systems to explore an exponential amount of configurations in parallel, a feature that is difficult to replicate classically.

The datasets used to explore the model's capabilities consist of condensed representations of energy deposits from single-particle showers in the ECAL+HCAL calorimeters of the CLIC detector within the $[225, 275]$ GeV energy range [3].

2. Quantum Angle Generator

The Quantum Angle Generator (QAG) is a variational quantum model designed to generate continuous-valued outputs suitable for representing energy hits from a detector. Unlike a Quantum Circuit Born Machine (QCBM), which produce single-shot measurement outcomes, hence bitstrings, the QAG outputs expectation values computed over multiple measurements. This distinction is crucial in the HEP context, as calorimeter signals are inherently continuous and benefit from smooth representations.

2.1 Model

The QAG architecture can be divided into two components, as illustrated in Figure 1: the *randomizing unitary* and the *parametrized unitary*. The randomizing unitary ensures that the circuit starts from a different state at each run, so that the expectation values and consequently the outputs vary with each generation. The parametrized unitary is trained to evolve the initial random state into one that produces the desired expectation values.

In the model output, the expectation value of each qubit is transformed into the energy of the corresponding pixel through a process called *decoding*, since expectation values range from -1 to 1, whereas energies are positive. This decoding step is inspired by classical generative deep learning [7] but adapted to the constraints of quantum measurement.

2.2 Training

The QAG learns to generate images through an iterative procedure. At each iteration, we sample a set of k random vectors $\{\vec{x}_k\} \sim \mathcal{N}(\mu, \sigma^2)$, with $\mu = 0$ and σ^2 as a hyperparameter. From these k random inputs, the model generates a corresponding set of encoded fake samples $\{\tilde{A}_k\}$, and the Maximum Mean Discrepancy (MMD) [4] is used to compute the distance between the generated set and a subset of the same size drawn from the true encoded images $\{A_k\}$, as illustrated in Figure 2. Based on the MMD loss value, a new set of trained parameters $\vec{\theta}$ is obtained. Since gradient estimation on a quantum architecture is extremely slow [8], the optimizer employed is COBYLA [5], which does not require gradients and only needs a single evaluation of the loss function at each iteration.

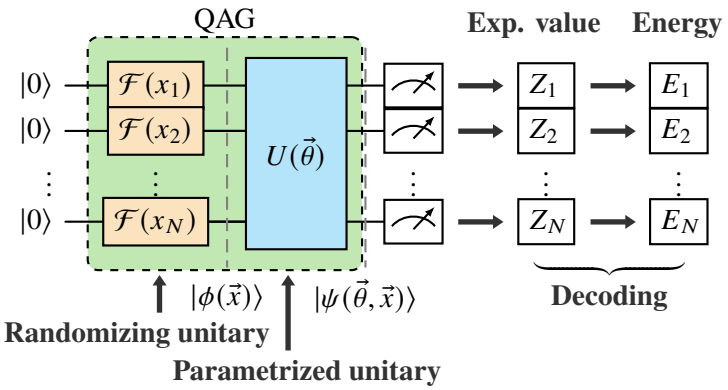

Figure 1: Schematic representation of the QAG. Classical pixel energies are generated by applying random qubit rotations, evolving them through a parameterized unitary, and then decoding the outcomes via measurements in the Z basis, averaged over many shots.

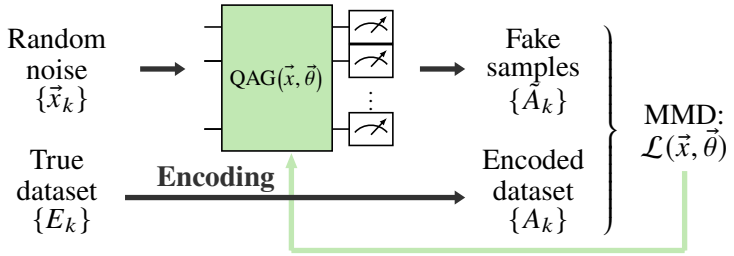

Figure 2: Training loop of the QAG. Encoded real datasets are compared with generated samples using the MMD loss. The optimizer updates circuit parameters to minimize the distance between real and generated distributions.

An example of the training procedure is shown in Figure 3, illustrating the minimization of the MMD loss alongside other monitored quantities used to assess the training quality. The training appears slightly unstable, this is due to the limited number of images generated per epoch (50) and the use of the COBYLA optimizer. After the initial phase of instability, the model is able to effectively minimize the MMD, which also indirectly reduces the other monitored quantities.

2.3 Results

The training procedure described before results from a curated hyperparameter optimization, summarized in Table 1. In particular, the most critical hyperparameter is the *Ansatz*, which defines the structure of the trainable unitary, as illustrated in Figure 4. In Quantum Machine Learning (QML), the Ansatz is analogous to the architecture of a deep neural network, specifying the number of layers, the arrangement of gates, and the pattern of qubit connections, all of which determine the expressivity and performance of the model. Our Ansatz was chosen to efficiently capture correlations between pixels while remaining shallow enough to be executable on all modern quantum architectures.

The final MMD value of the QAG-generated dataset over the test set is $2.41 \pm 0.14\,(\times 10^{-3})$. The quality of the generated dataset can be further assessed using the energy sum distribution,

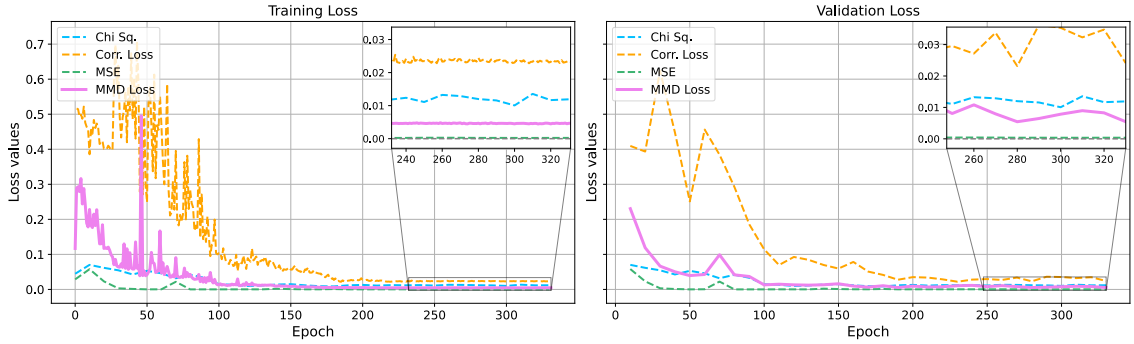

Figure 3: Loss curves during training of the QAG. The plots show the minimization of the MMD loss together with other monitored quantities: Correlation Loss [2], Mean Squared Error (MSE), and χ^2, computed over the encoded training and validation sets (left and right, respectively).

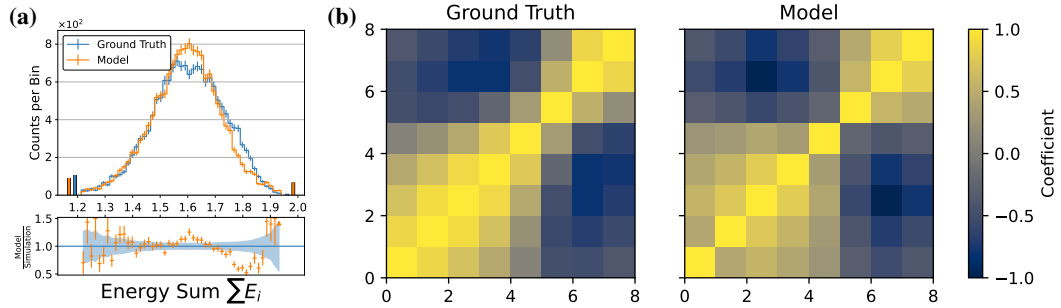

Figure 5: Comparison between ground-truth samples and the QAG-generated dataset: (a) distribution of total shower energy (orange: generated, blue: true), (b) pixel–pixel correlation matrices of the ground-truth (left) and generated dataset (right).

shown in Figure 5a, and the correlation matrix, shown in Figure 5b. Together, these metrics confirm that the QAG is capable of capturing both global and local features of the calorimeter data.

	Hyperparameter	Value
Training	Number of images	50
	Number of epochs	600
Model	Circuit design	clic8
	Number of shots	2048
	σ^2	0.5
COBYLA	tol	0.004
	rhobeg	5.593

Table 1: Hyperparameters used for training.

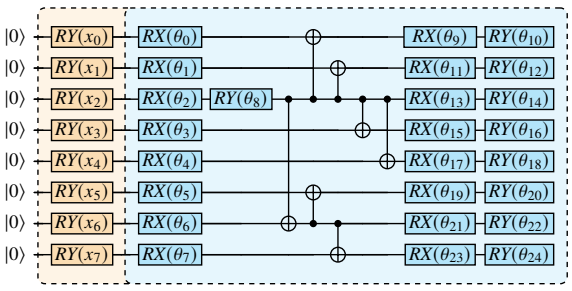

Figure 4: The clic8 Ansatz that is used here opposite to [2].

2.4 Noise Study

In addition, a noise study was conducted to assess the impact of realistic Noisy Intermediate-Scale Quantum (NISQ) noise levels on the training and evaluation of the architecture. The results

are shown in Figure 6. The noise profile of the *ibm_fez* device was used as reference, and its noise levels were rescaled on a simulator to enable training and evaluation of the model under different noise conditions. For each noise level, two scenarios were investigated and repeated 30 times to account for statistical variance: in the first case, the model was both trained and then evaluated under noise, with performance measured via the MMD between the generated and test datasets (violet); in the second, the model was trained on a noiseless simulator and subsequently evaluated on the noisy instance (blue).

The study shows two key outcomes. First, during training the model is able to adapt to the underlying noise profile, leading to noticeable improvements in performance at higher noise levels. Second, the performance gap between noiseless execution ($x = 0$) and realistic noise levels ($x = 1$) is minimal and falls within statistical fluctuations, suggesting that current NISQ architectures are already capable of handling such tasks. This resilience is particularly encouraging for the deployment of quantum generative models in near-term experiments, where hardware noise cannot be entirely avoided.

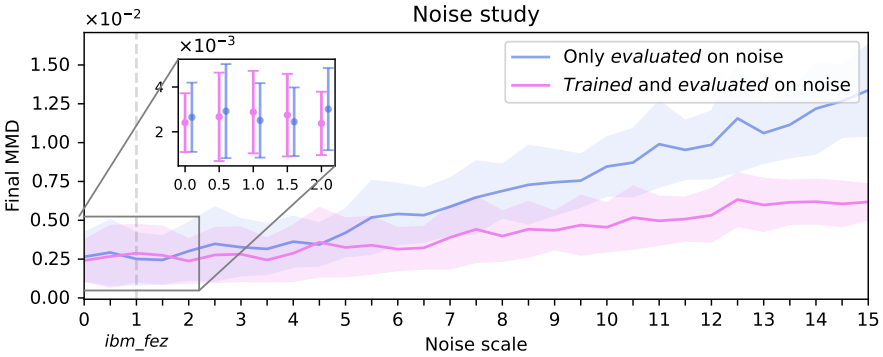

Figure 6: Impact of quantum noise on the QAG performance. Final MMD values for models trained on noiseless simulators (blue), and for models trained directly under noise (violet).

3. Conclusions and Outlook

In this work we successfully applied the QAG, a fully quantum generative model capable of producing realistic downscaled calorimeter shower images. The architecture demonstrated strong agreement with classical datasets and, importantly, robustness to realistic quantum noise. These findings strengthen the case for near-term quantum computing in HEP simulations, particularly in tasks where stochasticity and continuous distributions play a central role. However, scalability remains a crucial challenge: even if errors on current NISQ devices are manageable, practical applications of such models at the scale required for full detector simulations are still a long way off.

Looking ahead, several directions appear promising. Scaling to higher-dimensional datasets, such as full 2D or 3D calorimeter images, would provide a natural next step to test the expressive power of the approach. Moreover, hybrid training strategies that combine quantum circuits with classical neural networks could further enhance stability and scalability. Finally, benchmarking the QAG against state-of-the-art classical generative models will allow us to quantify any true quantum advantage in simulation quality and computational efficiency.

PoS(EPS-HEP2025)639

References

[1] F. Rehm, K. Borras and O. Pooth, *Deep learning and quantum generative models for high energy physics calorimeter simulations*, Tech. Rep., Lehrstuhl für Experimentelle Elementarteilchenphysik (DESY), 2023.

[2] F. Rehm, S. Vallecorsa, K. Borras, D. Krücker, M. Grossi and V. Varo, *Precise image generation on current noisy quantum computing devices*, *Quantum Science and Technology* **9** (2023) 015009.

[3] M. Pierini and M. Zhang, *CLIC Calorimeter 3D images: electron showers at fixed angle*, Jan. 2020.

[4] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf and A. Smola, *A kernel two-sample test*, *Journal of Machine Learning Research* **13** (2012), 723–773.

[5] M. J. D. Powell, *A direct search optimization method that models the objective and constraint functions by linear interpolation*, in *Advances in Optimization and Numerical Analysis*, 51–67, Springer, 1994.

[6] J.-G. Liu and L. Wang, *Differentiable learning of quantum circuit Born machines*, *Physical Review A* **98** (2018), 062324.

[7] C. Krause, et al *CaloChallenge 2022: A Community Challenge for Fast Calorimeter Simulation*, arXiv:2410.21611 [physics.ins-det], 2024.

[8] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac and N. Killoran, *Evaluating analytic gradients on quantum hardware*, *Physical Review A* **99** (2019), 032331.

PoS(EPS-HEP2025)639