

Precise Image Generation on Current Noisy Quantum Computing Devices

Florian Rehm^{1,2*}, Sofia Vallecorsa¹, Kerstin Borrás^{2,3}, Michele Grossi¹, Dirk Krücker³ and Valle Varo³

¹*CERN, Geneva, Switzerland.

²RWTH Aachen University, Aachen, Germany.

³DESY, Hamburg, Germany.

*Corresponding author(s). E-mail(s):
florian.matthias.rehm@cern.ch;

Abstract

The Quantum Angle Generator (QAG) is a new full Quantum Machine Learning model designed to generate accurate images on current Noise Intermediate Scale (NISQ) Quantum devices. Variational quantum circuits form the core of the QAG model, and various circuit architectures are evaluated. In combination with the so-called MERA-upsampling architecture, the QAG model achieves excellent results, which are analyzed and evaluated in detail. To our knowledge, this is the first time that a quantum model has achieved such accurate results. To explore the robustness of the model to noise, an extensive quantum noise study is performed. In this paper, it is demonstrated that the model trained on a physical quantum device learns the noise characteristics of the hardware and generates outstanding results. It is verified that even a quantum hardware machine calibration change during training of up to 8% can be well tolerated. For demonstration, the model is employed in indispensable simulations in high energy physics required to measure particle energies and, ultimately, to discover unknown particles at the Large Hadron Collider at CERN.

Keywords: Full Quantum Generative Model, Quantum Image Generation, Detailed Quantum Inference Evaluation, Quantum Noise Study, Quantum Circuit Entanglement Study, Quantum Hardware Training

1 Introduction

Quantum computing has the potential for a new paradigm in future computing to accelerate tasks or even handle classically unsolvable problems [1]. In the current Noise Intermediate Scale Quantum (NISQ) era, quantum devices suffer from non-negligible hardware errors, limited connectivity and a limited number of qubits [2]. While practical quantum advantage is currently extremely difficult to accomplish, finding the best suited algorithms to effectively combat the problems of NISQ-era devices remains a widely researched topic. Quantum Machine Learning (QML) is a domain which achieves acceptable results on NISQ devices due to the observed robustness against noise [3].

High Energy Physics (HEP) experiments, such as those at the Large Hadron Collider (LHC) at CERN, require enormous amounts of simulated data for deriving high precision physics results [4]. To handle this demand, gigantic quantities of computing hardware resources are necessary, which has led to the creation of the world’s largest computing grid operated by CERN [5]. To alleviate this strain on computational resources, Machine Learning (ML) models have been developed that exhibit remarkable speed-ups over current Monte Carlo-based simulations while maintaining the required level of accuracy. In general, QML simulations represent a promising approach to address the only further increasing simulation demands in the future [6, 7]. QML employs quantum circuits which exploit the quantum properties of superposition and entanglement, which possess the potential to outperform neural networks, their classical analogue [8]. In addition, QML might have the advantage of learning more complex distributions with fewer parameters than classical ML due to their wider accessible phase space.

Encoding the classical data into qubit states on quantum computers is a non-trivial task [9]. Currently, many encoding techniques exist, each exhibiting specific advantages and disadvantages, and in practice, identifying “the best” encoding technique remains application dependent [9]. To achieve a potential quantum advantage over classical computing, theoretical studies suggest that at least linear scaling from qubits to features is required [10]. On the other side, models which employ better than linear scaling encoding techniques have drawbacks, making them unsuitable for generating precise images on NISQ devices. For example, amplitude encoding can only generate probability distributions and not absolute pixel entries, i.e. energy values.

At present, there exist several quantum generative models, for example, the Quantum Circuit Born Machine (QCBM) [11], Quantum Variational Autoencoders [12] or variations of quantum Generative Adversarial Networks [13–15]. They all face limitations. Some models either do not scale well in terms of qubits required relative to the number of encoded features, or they do not achieve a satisfying level of fidelity. The Quantum Angle Generator (QAG) presented in this paper and first introduced in reference [16], aims to overcome these problems. Employing angle encoding with linear scaling of qubits to features, it achieves extremely accurate results for a real-world problem on current physical noisy quantum devices.




Fig. 1: An example 3D calorimeter shower image.

The paper content is structured as follows. First, the HEP use case is motivated, and the training data is defined. Next, the QAG model and the employed angle encoding technique are presented. Then, multiple circuit architectures are compared and the advantages of the best ones are highlighted. An in-depth accuracy analysis of the model follows to highlight its excellent precision. The quantum hardware noise behavior is evaluated, including training and inference executed on real quantum devices. Lastly, conclusions are drawn and summarized.

2 High Energy Physics Simulations

Simulations remain a crucial component of HEP analysis to evaluate the results obtained by the processed data of the experiments. Currently, simulations are dominantly performed with Monte Carlo methods such as the Geant4 toolkit [17]. However, Monte Carlo simulations are very hardware resource demanding and occupy half of the worldwide LHC Computing Grid [18]. Future LHC experiments will require more simulations due to more energetic particles, simultaneous collisions and detectors constructed with higher granularity. However, the projected budget for hardware development and computing resources cannot keep pace with these increasing demands [19, 20]. As a result, ML alternatives to Monte Carlo methods are being actively researched. Initial prototypes predict significant reductions in simulation time and hardware resources while retaining acceptable levels of accuracy [21–23]. This research goes one step beyond classical ML. Since HEP data sets are generally created by underlying quantum mechanical effects, performing the computations on quantum devices which likewise make use of quantum effects has the potential to substantially enhance the simulations in accuracy and in terms of sustainable computing.

Electromagnetic calorimeters are constructed as high granularity sensor grids to measure the energy of photons, electrons, and positrons through complex particle shower generation processes in space and time [24]. They constitute a key component of HEP detectors to measure the energy of the




Fig. 2: The structure of the Quantum Angle Generator.

particles produced in the interaction process and occupy most of the simulation time [18]. Calorimeter outputs can be interpreted at lowest order as static and spatial 3D images, which we call “shower images”: the value of each pixel corresponds to the energy measured in a specific calorimeter cell. The initial data from reference [25] consists of $25 \times 25 \times 25$ pixel images. An example of a 3D shower image is visualized in figure 1. To reduce the dimensionality, the images are averaged along two spatial axes (x- and y-direction), resulting in a one-dimensional representation that is further downsampled to eight pixels by averaging three contiguous pixels along the z-direction. Although the initial data set provides many different energies, for simplicity this study focuses on images recorded by particles in the energy range of [225, 275] GeV. The data set is split into a training and test set, each consisting of approximately 1 000 samples. The downsampled data is available in reference [26], and an example image is illustrated later in this paper in figure 6a.

3 The Quantum Angle Generator

The QAG represents a QML model which employs the well established technique of angle encoding [27, 28] to generate extremely precise images. It scales linearly with the number of encoded features. Thus, the generation of n features requires n qubits. In this study, the number of features corresponds to the number of pixels. A comprehensive description of the QAG model, its objective training function, and an evaluation of several quantum circuits are provided below.

3.1 Model Description

The QAG model consists of variational quantum circuits trained by an objective function. The model structure is visualized in figure 2. All qubits are initialized in the basis state $|0\rangle$. The state preparation function implements a Hadamard (H) gate to constitute superposition, followed by a y-rotational (Ry) gate to introduce randomness so that the model can draw new samples at each execution. For this, the Ry gate angles Ω are randomly drawn from a $[-1, 1]$ uniform distribution and pixel-wise multiplied by the pixel standard deviations present in the training data to obtain correct pixel energy variations. To

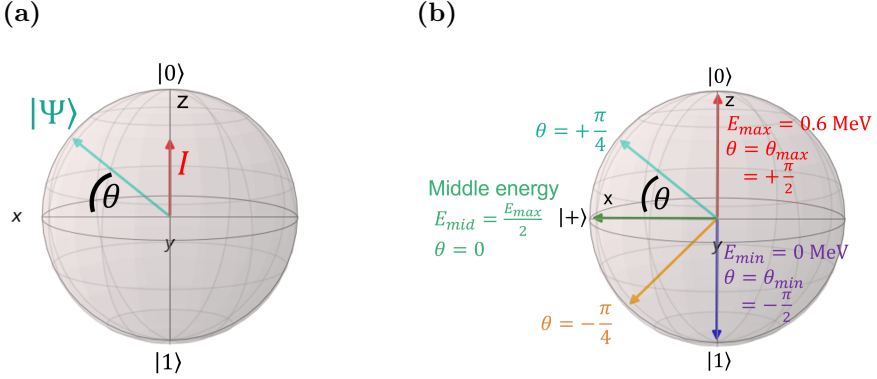


Fig. 3: (a) Decoding of an example state Ψ into an angle θ . (b) Some representative angles θ and its corresponding energy values.

account for all various primary particle energies, all angles Ω are multiplied by a random value between $[-0.25, 0.25]$.

The unitary transformation consists of quantum circuits and constitutes the trainable part of the QAG model. Various circuit architectures were tested as documented in section 3.3.

To convert the quantum states back into classical energy values via angle encoding, the model must be executed multiple times and the quantum states measured. The number of executions is commonly denoted as the number of shots nb_{shots} . It is counted how often state $|0\rangle$ is measured. The scalar intersection I of the vertical axis on the Bloch sphere (z-axis) is calculated with:

$$I = 2 \cdot \frac{\text{counts}(|0\rangle)}{nb_{shots}} - 1. \quad (1)$$

Next, the intersection I is transformed into the rotational angle θ by the trigonometric function $\theta = \arcsin(I)$.

The angle θ operates in the x-z-plane of the Bloch sphere and is defined as zero in the $|+\rangle$ state. Rotating θ clockwise leads to positive angles. The decoding process is visually illustrated in figure 3a with an example state $|\Psi\rangle$, its intersection I and the corresponding angle θ for a single qubit example. The angle θ can then be transformed into a pixel energy E by the linear change of ranges equation:

$$E = \frac{(E_{max} - E_{min})(\theta - \theta_{min})}{\theta_{max} - \theta_{min}}. \quad (2)$$

E_{max} , E_{min} , θ_{max} and θ_{min} are defined in figure 3b: the minimum energy $E_{min} = 0$ MeV is set to $\theta = -\pi/2$ and the maximum energy $E_{max} = 0.6$ MeV is set to $\theta = +\pi/2$. With $E_{min} = 0$ and $\theta_{max} = -\theta_{min}$ equation 2 can be simplified into:

$$E = \left(\frac{E_{max}}{2 \cdot \theta_{max}} \right) \cdot (\theta + \theta_{max}). \quad (3)$$

For multi-qubit quantum circuits: the final quantum states of all qubits are measured, and the outcomes are decoded independently. Although the qubit results are individually decoded, the state of each qubit is entangled with others due to the gates applied within the variational quantum circuit, as lined out later.

It is worth to be noted that the angle θ and, therefore, the corresponding decoded energy E remain in discrete values. Since, θ depends in value and accuracy on the number of shots nb_{shots} : the larger nb_{shots} , the better the achievable energy precision and resolution. Fortunately, with present quantum devices, the nb_{shots} can be easily chosen to be large. Currently, on IBMQ devices the maximum possible number of shots is $nb_{shots} = 100\,000$. For the simplified calorimeter use case, this is more than sufficient. For comparison, in a previous classical reduced precision ML research in reference [29], it is demonstrated that already 256 discrete energy levels are sufficient for correctly reproducing the full-size calorimeter shower image. In this reduced precision research, the parameters of the neural network are quantized from a larger format (floating point 32) down to a smaller number format (integer 8). This study will show that 512 shots provide a sufficient resolution.

3.2 Training Objective Function

The QAG model is trained with two losses employed as objective functions. The first one is the Mean Maximum Discrepancy (MMD) loss, already successfully applied by other quantum models, e.g. the QCBM [11]. Training exclusively with the MMD loss resulted in good average shower distributions. However, when exploring the generated images in more detail, for example in the pixel correlation, the model did not perform satisfactorily. Therefore, a second correlation (Corr) loss is added to help learn the patterns present in the training data (e.g., image pixel correlations). The Corr loss is calculated by the pixel-wise mean squared error (MSE) between the pixel correlation values present in the training data and the ones inside the generated data. The pixel-wise correlations are illustrated in figure 6b.

To train the QAG model, the Simultaneous Perturbation Stochastic Approximation (SPSA) optimizer [30] is employed, which only requires two optimization steps per epoch. The hyperparameters for training were found by extensive hyperparameter searches employing the Optuna [31] library. All tests in this study are executed in Qiskit version 0.26.2. The models are trained for 500 epochs, containing one batch. The dynamic MMD loss weight starts at a value of one and decays linearly with $-0.001 \cdot \text{epoch}$, starting from epoch 100. Opposite, the Corr loss weight increases by the same value starting at zero. The dynamical training batch size is set to generate one image in the first hundred epochs and afterward to 20 images to calculate the Corr loss between multiple images. Each quantum job contains 512 shots for training and inference. The generator SPSA optimizer learning rate is set to $c_0 = 1$ with an exponential learning rate decay of 0.006 starting from epoch 50. All these settings showed the best performance in the tests.

3.3 Quantum Circuit Study

The ideal circuit should contain a certain, optimized to be minimal, number of parameters to achieve a sufficient level of accuracy. Different circuit architectures were compared to each other based on characteristic numbers expressing the power of the circuit. The circuits employ trainable rotational gates and two-qubit entanglement gates. As the angle encoding primarily uses the qubits y-axis component, we predominantly employ Ry gates. For some circuit architectures, we test if additional z-rotational gates (Rz gates) or deeper circuits with depth 2 (denoted as d2) can further improve the results. We use two-qubit controlled-not gates (cx gates) native on IBM Quantum (IBMQ) [32] devices; while other entanglement gates are compositions of multiple native gates of the hardware. Keeping an eye on the goal of executing the training on a real quantum device, the absolute number of decomposed gates should be kept as small as possible.

The characteristic circuit numbers used in this study are the number of trainable parameters N_p , the expressibility X and the entanglement capability E . Larger circuits with more trainable parameters are potentially capable of achieving more accurate results. However, it might be that a plateau is reached at some point where the same task can be solved with similar accuracy by a smaller circuit, which is being investigated here. The definitions for X and E are from reference [33]: X describes how well the circuit can represent the pure states of the representative Hilbert space. For a single qubit the expressibility exhibits how many states of the Bloch sphere can be represented. In this paper, we measure $1 - X$: the closer to 1 the better the expressibility of the model, while the closer to 0 the worse. The entanglement capability E is a measure that expresses the ability of a circuit to generate entangled states between the qubits. Likewise, E ranges from 0 to 1, where 1 represents the best achievable value. The circuit architectures under study are introduced in the appendix A. Their corresponding characteristic numbers and theoretical potential are provided in the appendix B. In the following, the circuits are evaluated for the calorimeter use case.

We start by interpreting the results displayed in figure 4. The MSE accuracy metric is calculated by taking the pixel-wise MSE between the average Geant4 and QAG images. The training is repeated 25 times for each circuit and the mean and standard deviations are plotted. To prevent the influence of outliers, the best and worst two trials are discarded in this analysis. The MSE is given as a function of: N_p on the left, X in the middle, and E on the right. By inspecting the plots, it can be recognized that the MSE does not correlate with any of the characteristic circuit values in the plots, neither do the characteristic values correlate among themselves, as shown in the appendix B.

The MERA-up, MERA-up_d2, and MERA-up_Rz architecture perform best with the lowest MSE. This is consistent with the observation that they maintain a high X and E , as provided in the appendix A. The error bars provide a hint about the training stability. It can be observed that the better the average MSE of a model, the smaller its standard deviation.

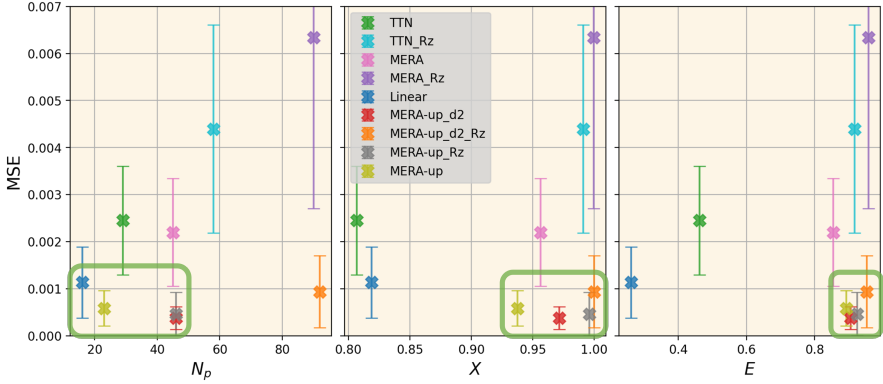


Fig. 4: The MSE uncertainty (the smaller, the better) of the various trained circuit architectures as a function of N_p (left), X (middle) and E (right), see text for explanations. Green rectangles mark optimal areas.

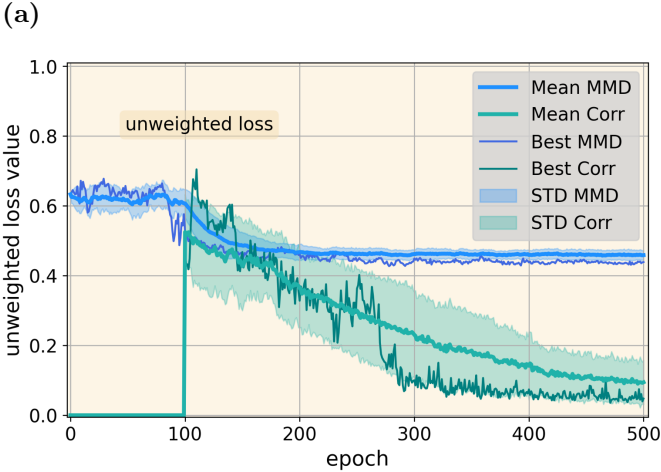


Fig. 5: The unweighted loss values as a function of the training epochs.

All in all, the MERA-up_Rz circuit clearly performs best considering the characteristic circuit values and the achieved accuracy in training. However, with the emphasis on a low number of N_p , the plain MERA-up circuit performs almost as well, while needing only half the number of parameters. Therefore, the following studies employ the MERA-up circuit for training the QAG model.

4 In-depth Accuracy Analysis

In this section, we analyze the results of the QAG model operating the MERA-up circuit architecture. We showcase typical accuracy metrics for the calorimeter simulation in HEP.

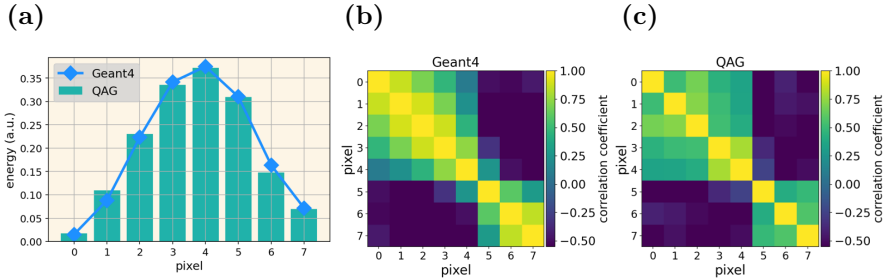


Fig. 6: The average calorimeter shower shapes are visualized in (a). The energy is provided in an arbitrary unit (a.u.) due to image downsampling. (b) The pixel-wise correlation plot of Geant4 and (c) the one of the QAG model.

4.1 Training Evaluation

In a first step, the statistical trends of the objective functions during training are investigated. In figure 5, the unweighted loss functions (excluding loss weights) are plotted as a function of the training epochs. The mean of twenty training repetitions is visualized as a thick solid line and the standard deviation (STD) as a colored band. The Corr loss starts influencing the training only at epoch 100 because its weight is set to zero before.

The MMD loss of all training repetitions converges stable, and the STD band narrows towards the end of the training. Overall, the MMD and Corr loss converge smoothly without strong oscillations, which is a desirable characteristic for stable (Q)ML training. The MMD loss contributes far more than the Corr loss throughout the training. However, the Corr loss plays a significant role in achieving good physics accuracy in the generated shower images.

4.2 Inference Evaluation

In the following, the accuracy in inference is evaluated. The generated images of the best trained model are compared to the Geant4 test data, which consists of 980 images. Likewise, there are 980 images generated by the QAG model to create the following accuracy metrics.

1. Average calorimeter shower shape: The first metric represents the calorimeter shower shape displayed in figure 6a. The shower shape is perfectly reproduced by the QAG model. The MSE corresponds to 0.00059 ± 0.00037 , which is extremely close to zero, indicating a very good accuracy.

2. Pixel-wise correlation: The second metric corresponds to the pixel-wise image correlation. The positive or negative correlation patterns between all the pixels are determined. The baseline represents the correlation from Geant4 in figure 6b. The correlation for the generated data of the QAG model is presented in figure 6c. It can be derived that the overall correlation pattern is accurately reproduced by the QAG model. Like the Geant4 data, it consists of a larger and more compact positively correlated group of pixels.

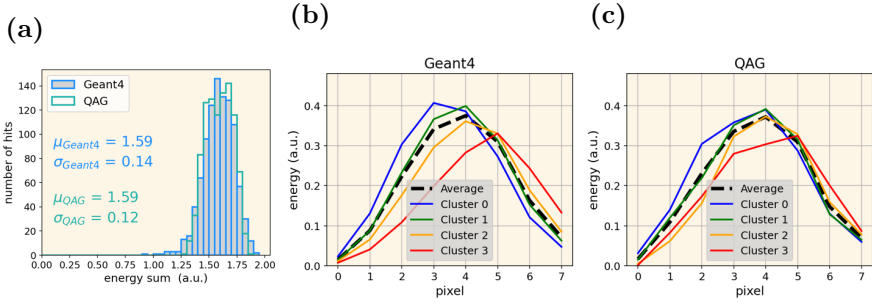


Fig. 7: (a) The energy sum histogram, (b) the four k-means clusters present in the images of Geant4, and (c) for QAG.

The other pixels are negatively correlated. Inspecting the particular details, different color shades indicate some minor deviations. However, the achieved correlation precision by QAG is astonishing. Therefore, it can be concluded, that the quantum circuits are capable to reproduce complex correlation patterns through substantial entanglement strategies, as present in the MERA-up architecture.

3. Energy sum: The energies contained in all pixels calculated for the individual images represent the third accuracy metric. In figure 7a, the energy sum histogram of the Geant4 images reveals a Gaussian shape and is correctly reproduced visually by the QAG model, which is also confirmed by the mean μ and the standard deviation σ .

4. k-means Clusters: This metric evaluates if the QAG model can correctly represent specific image modes. The Geant4 data is clustered with the k-means algorithm [34] to find four clusters or image modes as illustrated in figure 7b. Cluster 0 deposits substantially larger fractions of energy in earlier calorimeter cells than in higher cluster numbers. Further, the particles from clusters 0 and 1 contain a larger energy fraction, estimated by the integral area below the curves. Here we are interested in whether the QAG model can reproduce this behavior. The four clusters of the QAG images are provided in figure 7c. A similar structure can be observed, which indicates a good accuracy in reproducing the energy contents and image modes on average.

5. Pixel-wise energy distribution: The last metric is used to examine the distributions of the energy content of each pixel, as illustrated in figure 8. Overall, the histograms of the QAG model match those of the Geant4 model. Even pixels with non-Gaussian energy distributions in the Geant4 model are correctly reproduced by the QAG model. For example, the longer tail towards smaller energies on the left side of pixel 4 is equally present in the histogram for the QAG model. The large histogram overlaps indicate that not only the averages are reproduced with high accuracy, but also the energy distributions returned for each individual pixel are correct.

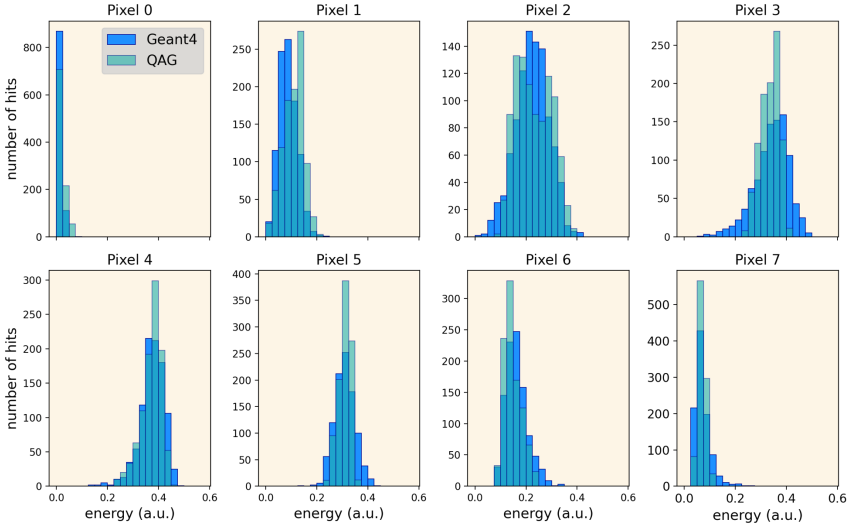


Fig. 8: Energy content histograms for the individual pixels for the Geant4 and the QAG model. The x-axis represents the energy depositions in one pixel and the y-axis the number of counts in each energy bin.

5 Quantum Noise Study

In the current NISQ era, relatively high hardware error levels are one of the primary limitations to effectively employing algorithms on real quantum devices. Similar as in the classical case, QML models appear to be noise resilient to some degree of hardware errors [35–38]. In the following, the robustness of the QAG model to simulated noise is tested in inference and training. Furthermore, training and inference are executed on real quantum devices with measured noise levels and compared to the results with simulated noise.

5.1 Inference

In a first step, quantum noise is applied only to the inference of a model trained without noise. Inference is performed using three different noise configurations: simulated noise at varying levels, simulated noise derived from the later used hardware, and finally, with the real quantum hardware. In the simulated noise configurations, each qubit noise is modeled with the same readout measurement error, and each inter-qubit connection with the same two-gate (CNOT) error. In the combined noise model both are on the same level. In contrast, the hardware noise levels can vary widely for individual qubits as well as for gates.

Multiple noise configurations and error levels from zero up to 15% are tested. The MSE is utilized as the accuracy measure. The results are illustrated in figure 9. For each configuration, the average value of 20 generated images is plotted in dependence of the noise level as a solid line and the standard deviation as a colored band around the mean. The gray horizontal line serves as

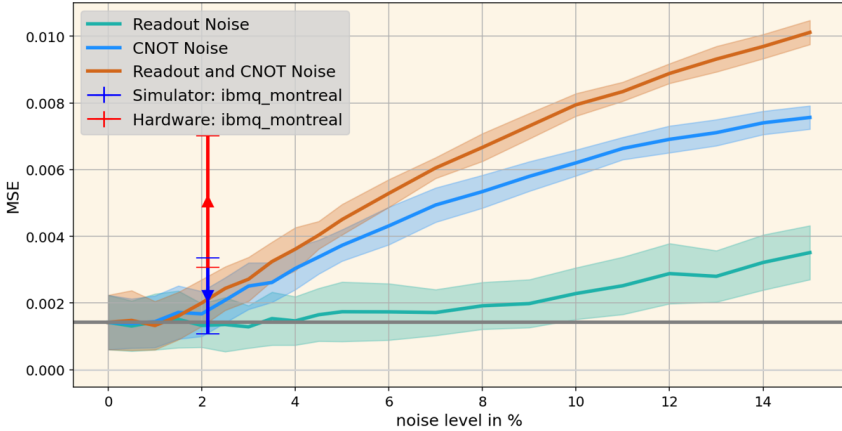


Fig. 9: Inference noise study. On the x-axis the noise levels are provided and on the y-axis the MSE as accuracy metrics (the lower, the better). Inference is executed with various noise configurations and on real quantum hardware.

accuracy reference for the noise-free configuration. All configurations maintain sufficient accuracy up to approximately 1.5% of noise. The configuration with readout noise only (green) is most robust and maintains a stable accuracy of up to 8% of noise. CNOT noise only (blue) and both readout noise and CNOT noise combined (orange) experience a stronger impact. As expected, the combined configuration (orange) performs worst. As a side note, current quantum devices have average noise levels below 5%, which is expected to gradually decrease further. But as discussed in the following, the noise levels are unstable and can sometimes spike up. Therefore, wider noise ranges were investigated.

Next, the inference is run by loading the real hardware noise model from `ibmq_montreal` into the simulator. The `ibmq_montreal` device consists of a Falcon r4 processor with 27 qubits. The average readout noise over the qubits employed at the time of the test is 2.51% and the average gate noise level is 0.97%. The explicit noise model, containing the entries for each qubit, is provided in the appendix in figure 11b. The result is included in figure 9 (blue triangle). The noise level position (x-axis) is determined as the average of readout and CNOT noise. Although the noise levels of the qubits vary strongly, the measured accuracy of the hardware noise simulation agrees well with the simulated noise in mean and standard deviation within the uncertainties. This suggests that a model trained without any noise would theoretically be able to run inference on noisy hardware without a significant drop in accuracy.

Finally, the inference is executed on the real `ibmq_montreal` device. The result is plotted as a red triangle in figure 9. The accuracy on the real hardware is worse than predicted by the simulation, as indicated by a larger MSE value. The decomposition of the circuit to the real hardware includes swap operations, which imply additional two-qubit entanglement gates for the quantum circuit.

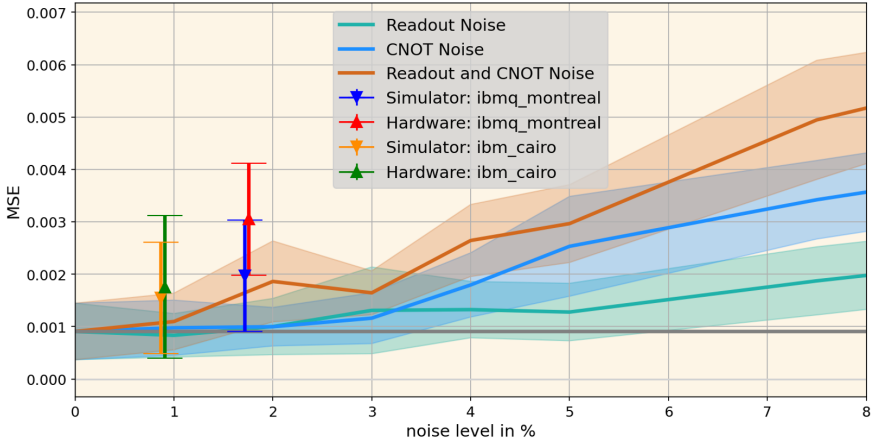


Fig. 10: Training noise study with various noise configurations, inclusive hardware noise simulations and training on the real hardware. The training with fake noise levels are repeated ten times.

It is possible that these are not included in the hardware noise simulation and lead to higher noise influence on real hardware and thus to worse results.

5.2 Training

The noise study is repeated to include noise also during training. We study if the QAG model can learn to compensate for noise in training, especially when running on the real quantum device. Besides, we investigate with which noise values the model can still maintain a reasonable accuracy.

The results are provided in figure 10. The configurations with readout noise (green) and CNOT noise (blue) maintain a similar level of accuracy until approximately 3%. The accuracy of the combination of readout and CNOT noise (orange) decreases marginally from 1% of noise and further with larger noise levels. However, at 3% of noise, its accuracy is still close to the noiseless case, staying within one standard deviation. This indicates that training the model with noise makes the QAG model more robust than only applying noise to a trained model in inference only.

Next, two quantum devices are simulated. `ibmq_montreal` and `ibm_cairo`, which are both 27 qubit devices, but `ibm_cairo` has the more advanced Falcon r5.11 processor. Likewise, the training is repeated ten times and the results are added as blue and orange triangles in figure 10. It can be noted that the average accuracy of the training with hardware noise performs slightly worse than the simulated combined noise model (orange line). The strong overlap between the error from the noise simulation (orange band) and the hardware noise simulation (orange and blue error bar) indicates that the accuracy difference is statistically not significant.

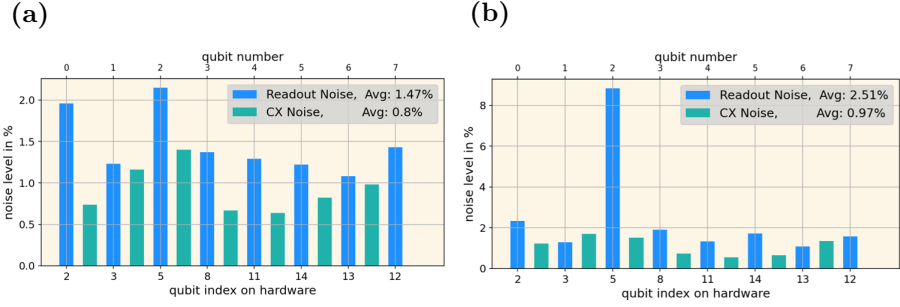


Fig. 11: The noise levels of the `ibmq_montreal` device are shown on the left before and on the right after the calibration change during training. The inference tests are performed with the right noise model. It can be seen, that the readout noise of qubit number five increased from two to over eight percent.

Finally, the entire training was executed on the real quantum device. First, the training was carried out on `ibmq_montreal`. During training, around epoch 280, an unpredicted significant noise change occurred and the readout noise of one qubit increased to 8%, as shown in figure 11. As a result, the MMD loss spikes up, as shown in figure 12a. This negatively influenced the training. However, after the noise change during training, the model recovered and adapted to the new noisy environment. In the still remaining number of epochs the loss decreased to a modest level. The training was repeated on the best performing `ibmq_cairo` device without having a hardware calibration change during training, and the training losses are shown in figure 12b.

The results of both hardware training are included as red and green triangles in figure 10. The error bars correspond to the accuracy deviations within 50 generated validation images. It can be observed that the average accuracy of the `ibmq_montreal` hardware training (red) is visibly worse than the noise simulation (blue). This is most likely due to the large noise increase and the fact that the model did not have enough remaining epochs to fully recover. This is also indicated by the still decreasing losses towards the end of the training. As mentioned above, the training was repeated on the more stable machine `ibmq_cairo`. This time, the training was completed without calibration changes and with about 1% lower hardware noise level (readout noise 0.86% and CNOT noise 0.89%). The accuracy of the hardware training on `ibmq_cairo` (green) is only marginally worse than the one from the simulator (orange). This indicates that the simulated and real hardware results behave similarly for low hardware noise levels. This fulfills the expectations derived from the pure simulation that exhibit only statistically insignificant variations in the very good accuracy at these low noise levels.

Comparing the absolute MSE magnitude of the real hardware training (≈ 0.002 and ≈ 0.003) with the one from inference from the previous section (≈ 0.005), the accuracy improved, suggesting that the QAG model is able to adapt its parameters to the noisy hardware to improve its precision. This is




Fig. 12: The losses during training for `ibmq_montreal` on the left and for `ibmq_cairo` on the right.

also confirmed by the `ibmq_montreal` training, where the accuracy worsened entirely after the calibration change, but then recovered. In the appendix in figure C3a, the average shower image created by the `ibmq_montreal` device is visualized, and in figure C3b by `ibmq_cairo`. The shower image of `ibmq_cairo` agrees well with Geant4, whereas `ibmq_montreal` exhibits some deficiencies because of the missing remaining training epochs after the calibration.

6 Conclusion

The results of this study clearly demonstrate that the newly developed QAG model is capable of generating images with good precision, as measured with a variety of validation metrics. This includes correctly reproducing average values, but most importantly, also complex pixel-wise correlations with the chosen optimal MERA-up quantum architecture. These results reveal that the QAG model with a good entangled circuit is capable of learning intrinsic correlation patterns from the training data.

Our study exhibit the significant impact that quantum hardware noise can have on the accuracy of quantum machine learning models. The results evidence that training the models with noise leads to better performance (stable until 3% noise) because the QAG model adapts to the underlying noise behavior and converges faster in contrast to the situation of applying noise in inference only (stable until 1.5% noise). This was also verified on the real hardware by the `ibmq_cairo` training. Furthermore, our study shows that the QAG model is robust and can produce accurate results even with significant hardware calibration changes with up to 8% noise, as demonstrated by the training on `ibmq_montreal`. Overall, the newly developed QAG model demonstrates that training quantum machine learning models with realistic quantum hardware noise can lead to robust models and accurate results, which is of great importance for the future development of real world quantum machine learning applications.

Acknowledgment

The authors would like to express their sincere gratitude to Simon Schnake and Alexis-Harilaos Verney-Provatas for their help, advice, and proofreading, which improved the quality and outcome of this paper.

This work has been sponsored by the Wolfgang Gentner Programme of the German Federal Ministry of Education and Research and by the CERN Quantum Technology Initiative. Additionally, we would like to acknowledge partial funding from the BMBF Project "NiQ: Noise in Quantum Algorithms" within the BMBF program "Anwendungsnetzwerke für das Quantencomputing" and from the Deutsches Elektronen-Synchrotron DESY, a member of the Helmholtz Association (HGF).

Access to the IBM Quantum Services was obtained through the IBM Quantum Hub at CERN. The views expressed are those of the authors and do not reflect the official policy or position of IBM or the IBM Quantum team.

Appendix A Quantum Circuit Architectures under Study

The quantum circuit architectures investigated in this paper are summarized in figure A1. In general, it was observed that hierarchical architectures perform best while maintaining a reasonable number of quantum gates and parameters. Specifically, we examine the Tree Tensor Network (TTN) architecture and Multi-scale Entanglement Renormalization Ansatz (MERA) introduced in reference [39]. Multiple variations of these circuits are tested: 1) circuits with a depth two (naming scheme "d2"). These circuits contain two layers, with all circuit gates placed twice, to evaluate if deeper circuits perform better. 2) Circuits additional with Rz-gates are employed after each Ry-gate (naming scheme "Rz") to assess if rotations around an additional axis can improve the accuracy. Both architecture variants double the number of parameters of the initial circuit.

For the MERA architecture, further variations are analyzed where only the right circuit half from the original architecture is implemented, denoted as the MERA upsampling (MERA-up) circuit. In the MERA-up circuit, the information is upsampled from the central qubit and spread to all the other qubits, similar to what happens in classical ML generative models, for example, transpose convolutional neural networks. The left half of the MERA circuit, the MERA downsampling (MERA-down) architecture, is not tested because it would rather compress the information as required in classification tasks that are not used in this paper. The last architecture contains a simple linear entanglement strategy.

In reference [33], multiple complex four-qubit architectures are compared. However, these architectures are not exploited in our study. The reason is most of them contain many more gates and parameters, and they do not scale well for more than four qubit circuits. Also, many architectures employ parameterized two-qubit rotational gates which we found, taking the decomposition

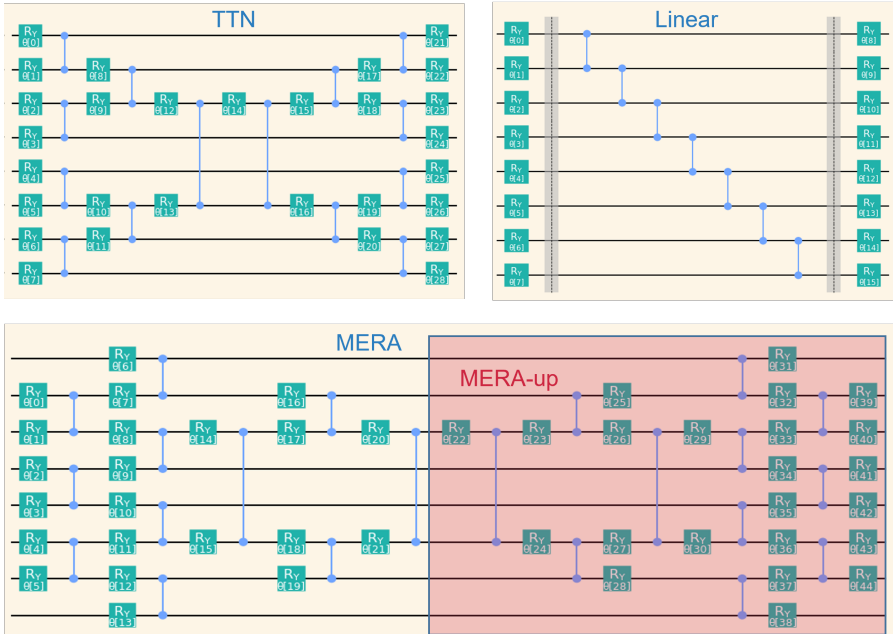


Fig. A1: The basic quantum circuit architectures under study. Top left: TTN architecture, top right: Linear entanglement architecture, bottom: the full circuit corresponds to the MERA architecture, with the red highlighted section corresponding to the MERA-up architecture.

into account, not as effective as the combination of separate rotation and entanglement gates.

Appendix B Characteristic Circuit Numbers

The results of evaluating the characteristic circuit numbers (number of parameters N_p , expressibility X and entanglement capability E) are shown in figure B2. In figure B2a E is plotted as a function of X . The MERA_Rz and MERA-up_d2_Rz architecture perform best and lie almost on top of each other in the top right corner, directly followed by MERA-up_Rz and TTN_Rz. The Linear architecture and the basic TTN perform worst. It can be clearly noted that circuits with more gates (Rz, d2) perform better. N_p is studied in figure B2b and figure B2c. It can be seen that the best two circuits – MERA_Rz and MERA-up_d2_Rz – contain by far the largest number of parameters. However, a more limited number of parameters is desirable to address NISQ hardware limitations. The MERA-up_Rz and TTN_Rz architectures contain approximately half the number of parameters compared to the MERA_Rz and MERA-up_d2_Rz circuits but perform almost similarly accurately. Therefore, they may be preferred in practice. Further reducing the number of parameters, the MERA-up architecture with only $N_p = 23$ parameters maintains adequate

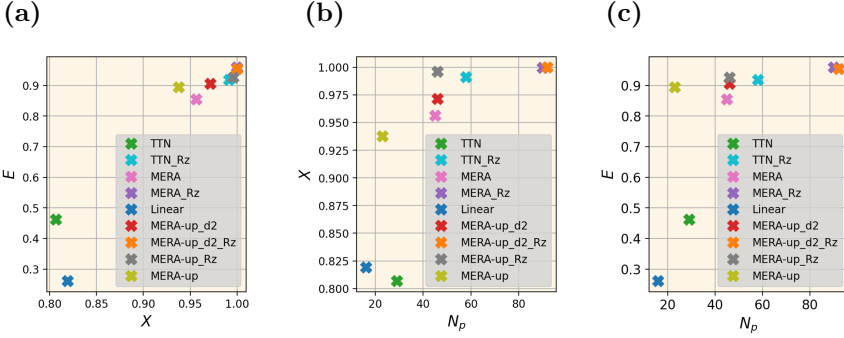


Fig. B2: Study of potential correlations between the three characteristic circuit numbers. On the left: X versus E . In the middle: N_p versus X . On the right: N_p versus E .

Table B1: The characteristic circuit values N_p , X and E are listed. Additionally, the statistics of the MSE accuracy metrics are provided for the individual circuits for repeating the training 25 times on the calorimeter data set.

Circuit Name	N_p	X	E	MSE
Linear	16	0.8191	0.261	0.00113 ± 0.0008
TTN	29	0.8068	0.462	0.00245 ± 0.0012
TTN_Rz	58	0.9912	0.918	0.00440 ± 0.0022
MERA	45	0.9564	0.855	0.00220 ± 0.0011
MERA_Rz	90	0.9997	0.959	0.00634 ± 0.0036
MERA-up	23	0.9377	0.894	0.00059 ± 0.0004
MERA-up_d2	46	0.9715	0.906	0.00038 ± 0.0002
MERA-up_Rz	46	0.9961	0.926	0.00047 ± 0.0005
MERA-up_d2_Rz	92	0.9999	0.954	0.00094 ± 0.0008

values for E and X and is the baseline circuit for the more detailed studies. All characteristic circuit numbers measured in the circuit study are summarized in table B1.

Appendix C Full Training on Quantum Device

The average shower images of the quantum hardware training are correctly reproduced, as provided in figure C3. The correlation plot for the `ibm_cairo` training is provided in figure C3c. The overall correlation pattern is correctly reproduced. All in all, the model trained on `ibm_cairo` shows a good performance.

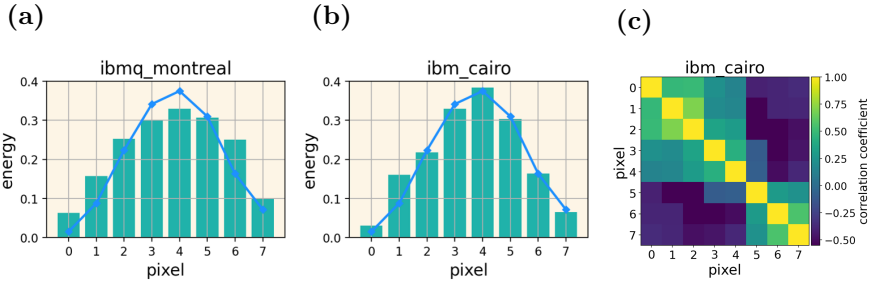


Fig. C3: On the left, the average shower images of the QAG model are trained on `ibmq_montreal` and in the middle on `ibmq_cairo`. The correlation plot on the right is created for the `ibmq_cairo` training.

References

- [1] Zidan, M., Eleuch, H., Abdel-Aty, M.: Non-classical computing problems: Toward novel type of quantum computing problems. *Results in Physics* **21**, 103536 (2021). <https://doi.org/10.1016/j.rinp.2020.103536>
- [2] Preskill, J.: Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018). <https://doi.org/10.22331/q-2018-08-06-79>
- [3] Huang, H.-Y., Broughton, M., Mohseni, M., et al.: Power of data in quantum machine learning. *Nature Communications* **12**(1) (2021). <https://doi.org/10.1038/s41467-021-22539-9>
- [4] Clissa, L.: Survey of Big Data sizes in 2021. *arXiv* (2022). <https://doi.org/10.48550/ARXIV.2202.07659>. <https://arxiv.org/abs/2202.07659>
- [5] CERN: Worldwide LHC Computing Grid. Accessed in June 2022. <https://wlcg-public.web.cern.ch/>
- [6] Zheng, H., Li, Z., Liu, J., Strelchuk, S., Kondor, R.: On the Super-exponential Quantum Speedup of Equivariant Quantum Machine Learning Algorithms with $SU(d)$ Symmetry (2022). <https://doi.org/10.48550/ARXIV.2207.07250>. <https://arxiv.org/abs/2207.07250>
- [7] Wu, S.L., et al.: Application of Quantum Machine Learning to High Energy Physics Analysis at LHC Using Quantum Computer Simulators and Quantum Computer Hardware. *PoS EPS-HEP2021*, 842 (2022). <https://doi.org/10.22323/1.398.0842>
- [8] Khrennikov, A.: Roots of quantum computing supremacy: superposition, entanglement, or complementarity? *The European Physical Journal Special Topics* **230** (2021). <https://doi.org/10.1140/epjs/s11734-021-00061-9>

- [9] Weigold, M., Barzen, J., Leymann, F., Salm, M.: Data encoding patterns for quantum computing. In: Proceedings of the 27th Conference on Pattern Languages of Programs. PLoP '20. The Hillside Group, USA (2020). <https://doi.org/10.5555/3511065.3511068>
- [10] Weigold, M., Barzen, J., Leymann, F., Salm, M.: Expanding Data Encoding Patterns For Quantum Algorithms. In: 2021 IEEE 18th International Conference on Software Architecture Companion (ICSA-C), pp. 95–101 (2021). <https://doi.org/10.1109/ICSA-C52384.2021.00025>
- [11] Liu, J.-G., Wang, L.: Differentiable learning of quantum circuit born machines. *Physical Review A* **98**(6) (2018). <https://doi.org/10.1103/physreva.98.062324>
- [12] Khoshaman, A., Vinci, W., Denis, B., Andriyash, E., Sadeghi, H., Amin, M.H.: Quantum variational autoencoder. *Quantum Science and Technology* **4**(1), 014001 (2018). <https://doi.org/10.1088/2058-9565/aada1f>
- [13] Tong Li, J.X. Shibin Zhang: Quantum generative adversarial network: A survey. *Computers, Materials & Continua* **64**(1), 401–438 (2020). <https://doi.org/10.32604/cmc.2020.010551>
- [14] Rehm, F., Vallecorsa, S., Borrás, K., Krücker, D.: Quantum Machine Learning for HEP Detector Simulations. Conference Proceedings "Distributed Computing and Grid Technologies in Science and Education" (GRID 2021), Dubna, Russia (2021). <http://ceur-ws.org/Vol-3041/363-368-paper-67.pdf>
- [15] Dallaire-Demers, P.-L., Killoran, N.: Quantum generative adversarial networks. *Physical Review A* **98**(1) (2018). <https://doi.org/10.1103/physreva.98.012324>
- [16] Rehm, F., Vallecorsa, S., Grossi, M., Borrás, K., Krücker, D., Schnake, S., Verney-Provatas, A.-H.: Quantum Angle Generator for Image Generation (2022). <https://doi.org/10.1109/SEC54971.2022.00064>. <https://inspirehep.net/literature/2620987>
- [17] Agostinelli, S., *et al.*: GEANT4—a simulation toolkit, vol. 506, pp. 250–303 (2003). [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8)
- [18] Albrecht, A.e.a. Alvesand: A Roadmap for HEP Software and Computing Research and Development for the 2020s, vol. 3 (2019). <https://doi.org/10.1007/s41781-019-0031-6>
- [19] Apostolakis, J.: In: Fabjan, C.W., Schopper, H. (eds.) *Detector Simulation*, pp. 485–531. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-35318-6_11

- [20] Kadochnikov, I., Bird, I., McCance, G., Schovancova, J., Girone, M., Campana, S., Currul, X.E.: WLCG data lake prototype for HL-LHC. Advisory committee, 127 (2018). <https://doi.org/https://ceur-ws.org/Vol-2267/509-512-paper-97.pdf>
- [21] Ju, X., Murnane, D., Calafiura, P., *et al.*: Performance of a geometric deep learning pipeline for hl-lhc particle tracking. The European Physical Journal C **81**(10), 1–14 (2021). <https://doi.org/10.1140/epjc/s10052-021-09675-8>
- [22] Biscarat, C., Caillou, S., Rougier, C., Stark, J., Zahreddine, J.: Towards a realistic track reconstruction algorithm based on graph neural networks for the HL-LHC. In: EPJ Web of Conferences, vol. 251, p. 03047 (2021). EDP Sciences. https://www.epj-conferences.org/articles/epjconf/pdf/2021/05/epjconf_chep2021_03047.pdf
- [23] Sela, E., Huang, S., Horn, D.: Deep learning study of an electromagnetic calorimeter. Algorithms **15**(4), 115 (2022). <https://doi.org/10.3390/a15040115>
- [24] Fabjan, C., Gianotti, F.: Calorimetry for Particle Physics. Reviews of Modern Physics **75** (2003). <https://doi.org/10.1103/RevModPhys.75.1243>
- [25] Pierini, M., Zhang, M.: CLIC Calorimeter 3D images: Electron showers at Fixed Angle. Zenodo (2020). <https://doi.org/10.5281/zenodo.3603122>. <https://doi.org/10.5281/zenodo.3603122>
- [26] Rehm, F.: Downsampled Calorimeter Shower Images to 8 Pixels. Zenodo (2021). <https://doi.org/10.5281/zenodo.7025233>. <https://doi.org/10.5281/zenodo.7025233>
- [27] LaRose, R., Coyle, B.: Robust data encodings for quantum classifiers. Physical Review A **102**(3) (2020). <https://doi.org/10.1103/physreva.102.032420>
- [28] Sierra-Sosa, D., Telahun, M., Elmaghraby, A.: Tensorflow quantum: Impacts of quantum state preparation on quantum machine learning performance. IEEE Access **8**, 215246–215255 (2020). <https://doi.org/10.1109/ACCESS.2020.3040798>
- [29] Rehm, F., Vallecorsa, S., Saletore, V., *et al.*: Reduced Precision Strategies for Deep Learning: A High Energy Physics Generative Adversarial Network Use Case. Proceedings of the 10th International Conference on Pattern Recognition Applications and Methods (2021). <https://doi.org/10.5220/0010245002510258>

- [30] Alhabsi, A.H.: Improved SPSA optimization algorithm requiring a single measurement per iteration. In: 10th International Conference on Information Science, Signal Processing and Their Applications (ISSPA 2010), pp. 263–265 (2010). <https://doi.org/10.1109/ISSPA.2010.5605476>
- [31] Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M.: Optuna: A Next-generation Hyperparameter Optimization Framework (2019). <https://doi.org/https://arxiv.org/abs/1907.10902>
- [32] IBM Quantum Services. IBM (2022). <https://quantum-computing.ibm.com/services?services=systems> Accessed 2022-08-25
- [33] Sim, S., Johnson, P.D., Aspuru-Guzik, A.: Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies* **2**(12), 1900070 (2019). <https://doi.org/10.1002/qute.201900070>
- [34] Hartigan, J.A., Wong, M.A., *et al.*: A k-means clustering algorithm. *Applied statistics* **28**(1), 100–108 (1979). <https://doi.org/https://www.jstor.org/stable/2346830?origin=JSTOR-pdf>
- [35] Wang, H., Ding, Y., *et al.*: QuantumNAS: Noise-Adaptive Search for Robust Quantum Circuits. arXiv (2021). <https://doi.org/10.48550/ARXIV.2107.10845>. <https://arxiv.org/abs/2107.10845>
- [36] Bogdanov, Y.I., Bogdanova, N.A., Fastovets, D.V., Lukichev, V.F.: Study of the effect of quantum noise on the accuracy of the Schrödinger equation simulation on a quantum computer using the Zalka-Wiesner method. arXiv (2022). <https://doi.org/10.48550/ARXIV.2201.03068>. <https://arxiv.org/abs/2201.03068>
- [37] Borrás, K., Chang, S.Y., Funcke, L., Grossi, M., Hartung, T., Jansen, K., Kruecker, D., Kühn, S., Rehm, F., Tüysüz, C., Vallecorsa, S.: Impact of quantum noise on the training of quantum generative adversarial networks. *Journal of Physics: Conference Series* **2438**(1), 012093 (2023). <https://doi.org/10.1088/1742-6596/2438/1/012093>
- [38] Wang, H., Gu, J., Ding, Y., Li, Z., Chong, F., Pan, D.Z., song han: RoQNN: Noise-Aware Training for Robust Quantum Neural Networks (2022). <https://openreview.net/forum?id=wwIBobGFj2V>
- [39] Grant, E., Benedetti, M., Cao, S., *et al.*: Hierarchical quantum classifiers. *npj Quantum Information* **4**(1) (2018). <https://doi.org/10.1038/s41534-018-0116-9>