

Impact of quantum noise on the training of quantum Generative Adversarial Networks

Kerstin Borrás^{1,2}, Su Yeon Chang^{3,4}, Lena Funcke⁵, Michele Grossi³, Tobias Hartung^{6,7}, Karl Jansen¹, Dirk Kruecker¹, Stefan Kühn⁶, Florian Rehm^{3,2}, Cenk Tüysüz^{1,8}, and Sofia Vallecorsa³

¹ Deutsches Elektronen-Synchrotron DESY, Platanenallee 6, 15738 Zeuthen, Germany

² RWTH Aachen University, Templergraben 55, 52062 Aachen, Germany

³ CERN, 1211 Geneva 23, Switzerland

⁴ Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland

⁵ Center for Theoretical Physics, Co-Design Center for Quantum Advantage, and NSF AI Institute for Artificial Intelligence and Fundamental Interactions, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

⁶ Computation-Based Science and Technology Research Center, The Cyprus Institute, 20 Kavafi Street, 2121 Nicosia, Cyprus

⁷ Department of Mathematical Sciences, 4 West, University of Bath, Claverton Down, Bath BA2 7AY, UK

⁸ Institut für Physik, Humboldt-Universität zu Berlin, Newtonstr. 15, 12489 Berlin, Germany

E-mail: su.yeon.chang@cern.ch

Preprint number: MIT-CTP/5400

Abstract. Current noisy intermediate-scale quantum devices suffer from various sources of intrinsic quantum noise. Overcoming the effects of noise is a major challenge, for which different error mitigation and error correction techniques have been proposed. In this paper, we conduct a first study of the performance of quantum Generative Adversarial Networks (qGANs) in the presence of different types of quantum noise, focusing on a simplified use case in high-energy physics. In particular, we explore the effects of readout and two-qubit gate errors on the qGAN training process. Simulating a noisy quantum device classically with IBM's Qiskit framework, we examine the threshold of error rates up to which a reliable training is possible. In addition, we investigate the importance of various hyperparameters for the training process in the presence of different error rates, and we explore the impact of readout error mitigation on the results.

1. Introduction

In classical deep learning, an extensive amount of studies have proven that noise plays a crucial role for the training of neural networks. Artificial noise injection is an efficient regularization method to speed up convergence and to improve the stability of the training process [1, 2].

Meanwhile, quantum computing, which is a completely new paradigm of computation, is characterized by statistical uncertainty from its probabilistic nature. Furthermore, on current and near-term quantum hardware one encounters the challenge to overcome the noise due to gate errors, readout errors, and interactions with the environment [3]. The presence of this



intrinsic quantum noise suggests the possibility of replacing the artificial noise in the context of classical machine learning with the noise of the quantum hardware.

In this paper, we study the impact of quantum noise on QML, more specifically on quantum Generative Adversarial Networks (qGANs), which are the quantum analog of classical GANs. Focusing on a simplified high-energy physics (HEP) use case, we first explore the impact of readout noise on the qGAN model by classically simulating a noisy quantum device with IBM's Qiskit software development kit [4]. In particular, we study the importance of several training hyperparameters at various levels of noise. We also apply different readout mitigation techniques and study their impact on the performance of qGAN training. Finally, we assess the performance of qGAN training in the presence of the full noise model, including two-qubit gate errors. This work provides broad insights into the impact of noise on QML, which is particularly relevant for applying current noisy intermediate-scale quantum (NISQ) devices [3] to QML.

2. Quantum Generative Adversarial Networks

First, we briefly review the qGAN model introduced in Ref. [5], which is a hybrid quantum-classical algorithm illustrated in Fig. 1a. This model consists of a classical discriminator and a quantum generator applied to n qubits. The generator is a variational quantum circuit, which consists of alternating layers of single-qubit rotation gates and two-qubit entangling gates. In addition, an initial layer of Hadamard gates provides an equal-weight superposition of all computational basis states. By comparing the qGAN performance for different gate choices, we found optimal performance for $R_Y(\theta_i)$ rotation gates and controlled- Z entangling gates.

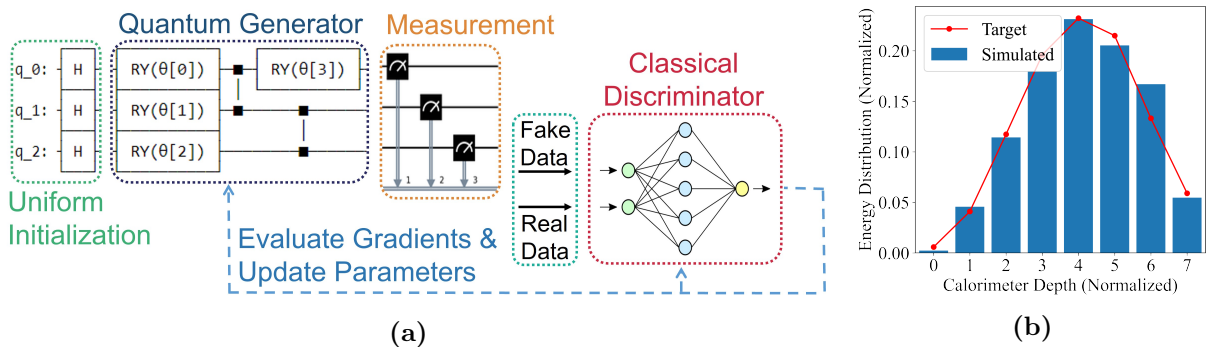


Figure 1: (a) Schematic diagram of the qGAN model, (b) exemplary results of the trained qGAN model applied to a simplified HEP problem using a statevector simulator.

The quantum generator is parameterized by the angles θ_i of the $R_Y(\theta_i)$ gates, while the classical discriminator is parameterized by the neural network parameters ϕ_i . During training, the generator and discriminator are optimized alternately. After training, the generator reproduces a probability distribution $Q(x)$ in the computational basis $|x\rangle \in \{|0\rangle, \dots, |2^n - 1\rangle\}$, which approximates the target distribution P_{target} . The input for the discriminator contains both real data, which are continuous scalars following the distribution P_{target} , and fake data, which are discrete integers between 0 and $2^n - 1$, chosen by an affine mapping between the computational basis states produced by the generator and an equidistant grid over the real data.

In this paper, we use the qGAN model to reproduce electrons energy profiles in HEP calorimeters. To simplify the problem, we sum up the energy distribution along the longitudinal profile, which contains the most essential information to classify the incoming particle. We average over the training set, binning the resulting probability distribution into $2^n = N$ pixels. Figure 1b displays the result of an exemplary qGAN training with $n = 3$ qubits using a statevector simulator, in the absence of statistical fluctuations and quantum noise.

3. Quantum Noise Study

In order to study the impact of quantum noise on the training process, we use a qGAN model with a quantum generator consisting of $n = 3$ qubits and 2 learning layers. Readout errors dominate over two-qubit gate errors for shallow circuits; thus, we choose to focus on the former. We assume that readout errors, i.e., misidentifying a measurement outcome as 1 given it was 0 and vice versa, are uncorrelated between different qubits and that the bit-flip probability p for misidentifying the measurement outcome is the same for each qubit. For the following simulations, we evaluate the performance of the qGAN model using the relative entropy

$$D_{\text{KL}}(P||Q) = \sum_{x=0}^{N-1} P(x) \log \frac{P(x)}{Q(x)}, \quad (1)$$

also called the Kullback-Leibler (KL) divergence. Here, $Q(x)$ is the output distribution of the quantum generator and $P(x)$ is the discretized version of the target distribution P_{target} .

3.1. Hyperparameter Scan

In this section, we study different values of the bit-flip probability, $p = \{0.01, 0.05, 0.1\}$, as well as different subsets of hyperparameters: the generator learning rate lr_g , the discriminator learning rate lr_d , and the exponential decay rate γ for the learning layers. In Fig. 2, we plot our results for the relative entropy using the best (orange) and worst (blue) choices of hyperparameters, as well as an average over all runs (green). For an appropriate choice of hyperparameters, the qGAN model converges for all choices of p . However, the relative entropy increases with a higher noise level, which cannot be fully mitigated even when choosing optimal hyperparameters.

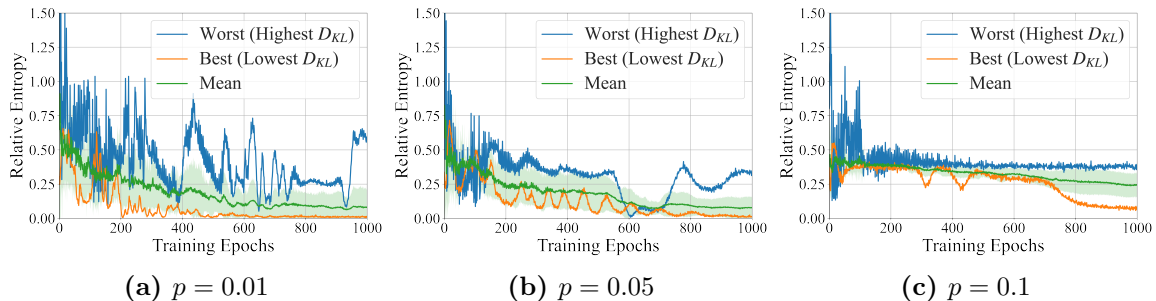


Figure 2: Relative entropy defined in Eq. (1) as a function of training epochs, shown for different hyperparameter choices and for different bit-flip probabilities p .




Figure 3: Importance of hyperparameters for different strengths of the readout error.

To examine which of the three hyperparameters has the strongest impact on the training for different noise levels, we display the so-called importance value of the hyperparameters for the

objective value in Fig. 3, which we calculated based on the fANOVA hyperparameter importance evaluation algorithm [6]. We find that the generator learning rate lr_g has the highest impact on the training, which demonstrates the difficulty of training the quantum generator. Thus, as expected, the impact of lr_g grows as the bit-flip probability increases.

3.2. Instability of qGAN Training

Instabilities, which manifest as large fluctuations of the relative entropy when choosing different initial training parameters, are one of the major challenges that have to be overcome in GAN training, both in the classical and the quantum case. In this section, we investigate the instability of qGAN training in presence of readout noise, using the optimal hyperparameters found in Section 3.1. The initial training parameters θ_i of the quantum generator circuit are sampled randomly from a normal distribution with a mean of 0.1 and a standard deviation of 0.2.




Figure 4: Relative entropy defined in Eq. (1) as a function of training epochs, averaged over different numbers of repetitions n_{rep} , shown for different bit-flip probabilities p .

Figure 4 summarizes the results of the qGAN training for different numbers of repetitions, $n_{\text{rep}} = 20$ (blue) and $n_{\text{rep}} = 100$ (orange), and for different readout-noise levels. Interestingly, for all values of p , both the average relative entropy and the standard deviation do not substantially change as we increase the statistics by increasing n_{rep} . Thus, although the model converges in a stable manner on the ensemble of simulations, individual unstable runs occur with a constant probability. These results demonstrate that training instabilities occur independently from the readout-error level and cannot be alleviated by quantum noise or increasing statistics.

3.3. Error Mitigation

In the following, we study two different methods to mitigate readout errors: conventional bit-flip (BF) mitigation [7] and independent bit-flip (IBF) mitigation [8]. These methods rely on a calibration matrix (M), which maps the measured probability distribution to the expected probability distribution ($P_{\text{noisy}} = MP_{\text{ideal}}$). In this way, the inverse of M can be used to obtain the expected outcome from measured probabilities. The matrix M for a single qubit reads

$$M = \begin{bmatrix} 1 - p_{01} & p_{10} \\ p_{01} & 1 - p_{10} \end{bmatrix}, \quad (2)$$

where p_{01} (p_{10}) is the bit-flip probability of misidentifying 0 as 1 (1 as 0) during readout.

For the conventional BF method [7], the computational resources required to obtain the inverse of M grow exponentially with respect to the number of qubits (M is a $2^n \times 2^n$ matrix). In contrast, the IBF method [8] assumes that the single-qubit matrix M in Eq. (2) can be extended to multiple qubits with a tensor product over all qubits ($M = M_1 \otimes M_2 \otimes \dots \otimes M_N$). Thus, this method assumes uncorrelated multi-qubit readout errors, which is an assumption

that has recently been verified experimentally [9]. This way, computing the inverse of M can be done by computing the inverse of n many 2×2 matrices using polynomial resources.

Correct estimation of the bit-flip probabilities is essential for successful error mitigation. Therefore, one needs to execute calibration circuits on the quantum device to obtain these values. The precise estimation of these values might be limited by their time dependence and by the finite sampling size (shot noise), e.g., $\mathcal{O}(1/\epsilon^2)$ shots are required for ϵ precision.

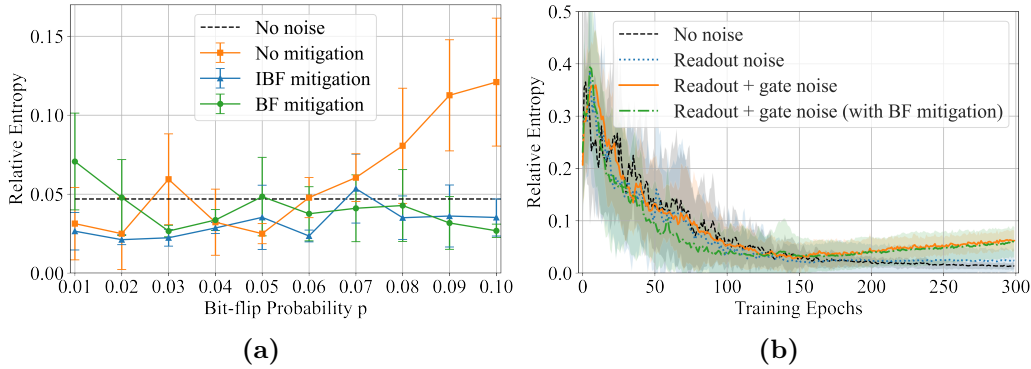


Figure 5: (a) Final relative entropy after qGAN training in the presence of readout noise, as a function of the bit-flip probability $p = p_{01} = p_{10}$, averaged over $n_{\text{rep}} = 20$ repetitions, shown with and without error mitigation. (b) Relative entropy as a function of training epochs, shown with different types of noise, with and without readout-error mitigation.

In Fig. 5a, we examine the impact of readout errors and error mitigation on the final relative entropy after qGAN training, comparing the cases of no noise (black), readout noise (orange), IBF mitigation (blue), and BF mitigation (green). When focusing on the average value of the relative entropy, the results for training with a small readout error ($p < 0.06$) seem to outperform the results without noise, which aligns with our initial suggestion of enhancing the qGAN training with quantum noise. However, within the statistical uncertainty, the noisy results for $p < 0.06$ are compatible with both the noise-free and the error-mitigated results. For larger levels of readout error ($p \geq 0.06$), we observe that the noisy results become worse compared to the noise-free and error-mitigated cases. This implies that error mitigation plays a crucial role for qGAN training in the presence of the large readout errors on current NISQ devices.

When comparing the two different error mitigation methods in Fig. 5a, we find that the resulting values for the relative entropy agree within the statistical uncertainty, as expected. For small values of p , we find a small deviation between the mean values for the relative entropy, which is most likely caused by using exact probabilities to estimate the calibration matrix M in the case of the IBF method and using a finite number of samples (3000 shots) in the case of the BF method. In the latter case, the finite number of samples results in a precision error of $\epsilon \sim \mathcal{O}(0.02)$ for estimating $p = p_{01} = p_{10}$, in agreement with our results in Fig. 5a.

Finally, we study the impact of both readout and two-qubit gate errors on the qGAN training. In Fig. 5b, we compare the results for the relative entropy (orange) to the results for the no-noise case (black), the readout-error-only case (blue), and the case of BF-mitigated readout errors and unmitigated gate errors (green). For this comparison, we consider error rates of 2.5% for the readout error and 1.5% for the gate error. As shown in Fig. 5b, new optimized hyperparameters are found to reduce the number of training epochs to only 300 epochs until reaching convergence. Furthermore, a new training schema is adapted, by updating the quantum generator once and the classical discriminator 10 times at each step. Similar to our findings in Fig. 5a, the inclusion of quantum noise does not improve the results for the relative entropy, as the results for the noise-free, noisy, and error-mitigated cases agree within the statistical uncertainty. Interestingly,

the average values of the relative entropy quickly converge in the noise-free and readout-error-only cases, but do not seem to converge when including two-qubit gate errors. These findings reveal the critical effect of two-qubit gate errors on qGAN training with current NISQ devices.

4. Conclusion

In this paper, we studied the impact of quantum noise on qGAN training, focusing on a simplified use case in high-energy physics. We first investigated different levels of readout noise and demonstrated that the qGAN model can be successfully trained even in the presence of this noise, although the common intrinsic instability of qGAN training cannot be overcome. We then examined the impact of different hyperparameters on the training and found that the learning rate of the quantum generator has the strongest impact, which becomes more pronounced with increasing readout-error rates. We also demonstrated that readout-error probabilities smaller than $p \sim 6\%$ do not significantly impact the qGAN training, while slightly larger readout errors still yield training convergence but substantially worse training results. To enhance these results, we studied the impact of two different readout-error mitigation methods, which both improved the qGAN training to a performance similar to the noise-free case. For two-qubit gate errors, we found that their effect on qGAN training is even more pronounced, preventing the convergence of training already for small error rates of $p \sim 1.5\%$. These results demonstrate that the mitigation of readout and gate errors plays a crucial role for qGAN training on current NISQ devices.

In future work, we will extend our current study to qGAN training on real quantum hardware. Moreover, we will investigate mitigation techniques to tackle the common issue of training instabilities for qGAN models. Finally, we are planning to study the impact of quantum noise and error mitigation methods on other QML models beyond qGANs.

Acknowledgments

This project is supported by CERN's Quantum Technology Initiative. L.F. is supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C²QA) under contract number DE-SC0012704, by the DOE QuantiSED Consortium under subcontract number 675352, by the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, <http://iaifi.org/>), and by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under grant contract numbers DE-SC0011090 and DE-SC0021006. S.K. acknowledges financial support from the Cyprus Research and Innovation Foundation under project "Future-proofing Scientific Applications for the Supercomputers of Tomorrow (FAST)", contract no. COMPLEMENTARY/0916/0048.

References

- [1] Bishop C M 1995 *Neural Comput.* **7** 108
- [2] Karras T, Laine S, Aittala M, Hellsten J, Lehtinen J and Aila T 2020 *CVPR 2020*
- [3] Preskill J 2018 *Quantum* **2** 79 ISSN 2521-327X URL <http://dx.doi.org/10.22331/q-2018-08-06-79>
- [4] Abraham H *et al.* 2019 *Zenodo*
- [5] Zoufal C, Lucchi A and Woerner S 2019 *npj Quantum Inf.* **5** 103 ISSN 2056-6387 URL <https://doi.org/10.1038/s41534-019-0223-2>
- [6] Hutter F, Hoos H and Leyton-Brown K 2014 An Efficient Approach for Assessing Hyperparameter Importance *Proceedings of the 31st International Conference on Machine Learning* (Beijing, China) URL <https://proceedings.mlr.press/v32/hutter14.html>
- [7] Qiskit documentation and source code accessed: 09-02-2022 URL <https://qiskit.org/textbook/ch-quantum-hardware/measurement-error-mitigation.html>
- [8] Funcke L, Hartung T, Jansen K, Kühn S, Stornati P and Wang X 2020 ArXiv:2007.03663 URL <http://arxiv.org/abs/2007.03663>
- [9] Alexandrou C, Funcke L, Hartung T, Jansen K, Kühn S, Polykratis G, Stornati P and Wang X 2021 ArXiv:2111.08551