Exploring Quantum Generative Models: A Benchmark Analysis

Natalia Muñoz Moruno (1), Artur Garcia-Saez (2) and Miguel A. González Ballester (1,2,3)

(1) BCN Medtech, Department of Engineering, Universitat Pompeu Fabra, Barcelona, Spain;

(2) Quantic, Barcelona Supercomputing Center, Barcelona, Spain; (3) ICREA, Barcelona, Spain

In recent years, numerous contributions have been made to the field of quantum machine learning, driven by its promising potential. A significant discipline within this field is quantum generative models. However, in terms of image generation, most of the proposed projects have been focused on the creation of images with simple distributions, few pixels or substantial downscaling, making their evaluation challenging. Hence, the primary objective of this project is to perform a benchmark analysis to assess how these models behave with different image resolutions. To achieve this goal, we have developed an open-source codebase that contains several quantum generative models, including Quantum Generative Adversarial Networks (GANs) and Quantum Diffusion models. This repository enables the training of the different architectures on two different biomedical datasets: the Blood MedMNIST [1], characterized by images of low dimensionality and the Brain Tumor Segmentation (BraTS) challenge dataset [2], which contains images of a considerable dimensionality. Regarding the architectures used, in the case of Quantum GANs, hybrid models have been employed, since they minimize the number of qubits needed, something of great importance in the noisy intermediate-scale quantum era. In addition, architectures that further optimize resources are also included, such as the quantum patch GAN [3] and the PQWGAN [4], which incorporate the patch method, thereby decreasing the resources required. In the case of Quantum Diffusion models, our work has initially focused on the analysis of Q-Dense [5], which is a type of Quantum Denoising Diffusion model that uses highly entangled circuits. Additionally, in this project versions of this model have also been used, including the Q-DenseDirected [5], which guides the process by incorporating the labels, making the process more controlled. Finally, the Fréchet Inception Distance and the Peak Signal-to-Noise Ratio metrics have been used to evaluate the performance of the different architectures and the image quality. Our project allows the comparison of the different models as well as the evaluation of their performance at different image resolutions. Through this analysis, we contribute to a better understanding of the current state of the art, providing a reference point for future research.

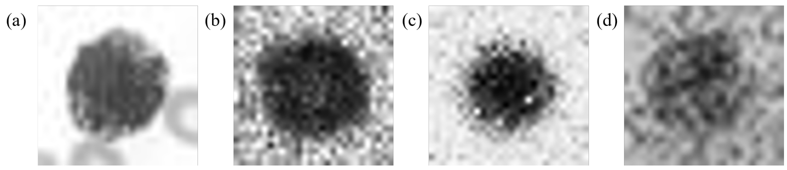

Figure 1: (a) Example of a real image from the Blood MedMNIST dataset [1]. (b) Example of an image generated by the PQWGAN [4]. (c) Example of an image generated by the QC-GAN [6]. (d) Example of an image generated by the Q-Dense [5].

Acknowledgements: This work is part of Maria de Maeztu Units of Excellence Programme CEX2021-001195-M, funded by MICIU/AEI /10.13039/501100011033.

References

[1] Yang, J., et al. (2023). Medmnist v2-a large-scale lightweight benchmark for 2d and 3d biomedical image classification. Scientific Data, 10(1), 41.

[2] Menze, B., et al. (2014). The multimodal brain tumor image segmentation benchmark (BRATS). IEEE transactions on medical imaging, 34(10), 1993-2024.

[3] Huang, H. L., et al. (2021). Experimental quantum generative adversarial networks for image generation. Physical Review Applied, 16(2), 024051.

[4] Tsang, S. L., et al. (2023). Hybrid quantum-classical generative adversarial network for high resolution image generation. IEEE Transactions on Quantum Engineering.

[5] Kölle, M., et al. (2024). Quantum Denoising Diffusion Models. arXiv preprint arXiv:2401.07049.

[6] Shu, R., et al. (2024). Variational Quantum Circuits Enhanced Generative Adversarial Network. arXiv preprint arXiv:2402.01791.