Shallow IQP circuits for generative modeling on NISQ hardware

Oriol Balló-Gimbernat[1,2,3]
Marcos Arroyo-Sánchez[1,2,3]
Paula García-Molina[2]
Adan Garriga[1]
Fernando Vilariño[2,3]
[1] *Eurecat, Centre Tecnològic de Catalunya*
[2] *Centre de Visió per Computador (CVC)*
[3] *Universitat Autònoma de Barcelona (UAB)*

oriol.ballo@eurecat.org

Generative modeling is one of the most promising applications of quantum machine learning, yet training and deploying Quantum Generative Models (QGMs) [1] on near-term hardware remains effectively intractable due to prohibitive gradient estimation and implementation costs. We propose a resource-efficient approach based on shallow Instantaneous Quantum Polynomial-time (IQP) circuits [2] that circumvents these bottlenecks by leveraging efficient classical training while retaining the guarantee of sampling hardness [3, 4]. To validate this approach, we formalize graph generation as a hierarchy of physical correlations, allowing us to map abstract data features such as edge density and bipartiteness directly to the quantum observables required to learn them. We validate our protocol (Fig. 1) through demonstrations both on real hardware (from 28 to 153 qubits) and simulations (28 qubits). Results show that while global structural features exhibit significant degradation beyond 91 qubits, our models achieve high-precision reproduction of local correlations, even up to 153 qubits (Fig. 2). These findings establish shallow IQP circuits as a robust, scalable candidate for generative tasks on current Noisy Intermediate-Scale Quantum (NISQ) devices.

References

[1] C. Zoufal, arXiv, 2111.12738 (2021)
[2] D. Shepherd, M. J. Bremner, Proceedings of the Royal Society A, 465 (2009) 1413
[3] E. Recio-Armengol, S. Ahmed, J. Bowles, arXiv, 2503.02934 (2025)
[4] E. Recio-Armengol, J. Bowles, arXiv, 2501.04776 (2025)

Figures

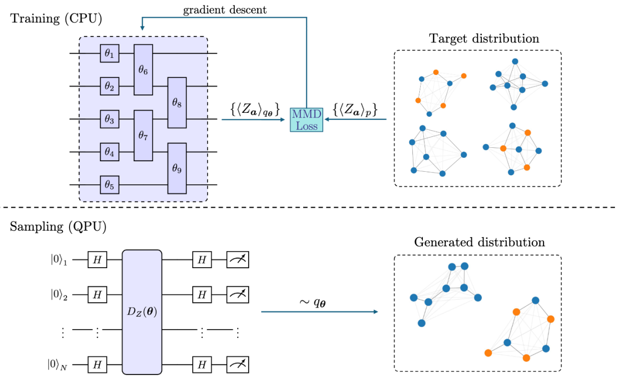

Figure 1: Overview of the proposed workflow. **Top:** Model training is performed entirely on classical hardware via a classical estimation of the MMD loss. **Bottom:** Sampling is carried out on IBM's Aachen quantum processor.

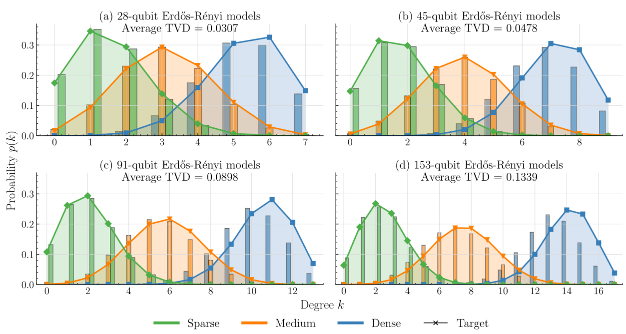

Figure 2: Degree distributions obtained from models from 28 to 153 qubits trained on Erdős-Rényi datasets and executed on NISQ hardware. Bars indicate empirical node-degree frequencies, while solid lines denote theoretical binomial targets. The TVD measures the deviation between generated and target distributions.