Classical training of quantum generative models based on Fermion Sampling

Bence Bakó [1,2], Zoltán Kolarovszki [1,2] and Zoltán Zimborás [1,3,4]

[1] *HUN-REN Wigner Research Centre for Physics, Budapest, Hungary*
[2] *Eötvös Loránd University, Budapest, Hungary*
[3] *Algorithmiq Ltd, Kanavakatu 3C 00160 Helsinki, Finland*
[4] *University of Helsinki, Yliopistonkatu 4 00100 Helsinki, Finland*

Quantum generative learning offers immense potential as the natural machine learning application of quantum computers, but it faces several trainability challenges. However, certain restricted and structured quantum generative models have a potential to overcome these trainability issues, through the efficient classical estimation of the expectation values and gradients of local observables. These estimates can be used to train the quantum model completely classically, thus also eliminating the need for quantum gradient computation. Furthermore, some classes of quantum models enable this type of classical training, while requiring a quantum device for efficient sampling [1,2].

The classical trainability of such models builds on the results of Ref. [3], where the authors showed that the squared maximum mean discrepancy loss function can be reformulated using expected values, instead of samples. The training framework of Ref. [4] is formulated around this loss function and introduces the class of parametrized instantaneous quantum polynomial time (IQP) circuits, as classically trainable generative models that are also classically hard to sample from in the final inference step. Similarly to IQP circuits, the Fermion Sampling quantum advantage scheme [2] also permits the classical estimation of local expectation values, while showing rigorous classical hardness in sampling. Furthermore, these circuits exhibit favourable trainability properties through the lack of barren plateaus [5].

In this work, we introduce fermionic Born machines as classically trainable quantum generative models and the corresponding efficient training scheme. The model consists of parametrized magic states and fermionic linear optical (FLO) transformations with trainable weights. The training algorithm relies on the decomposition of the parametrized magic states into Gaussian operators, enabling efficient expectation value estimation. Finally, through the Jordan-Wigner transformation, these FLO circuits can be directly implemented on qubit-based architectures to sample from the output probability distribution in the inference stage.

Through a series of numerical experiments, we showcase the performance of our model and training framework. We start by considering small-scale examples to show that this model can actually approach the target probability distribution, as measured in the total variational distance, not only match its finite moments. Furthermore, we demonstrate the scalability of our approach on larger problems of 100 qubits and beyond.

[1] M. J. Bremner, R. Jozsa, and D. J. Shepherd, Proc. R. Soc. A 467, 459 (2011).
[2] M. Oszmaniec, N. Dangniam, M. E. Morales, and Z. Zimborás, PRX Quantum 3, 020328 (2022).
[3] M. S. Rudolph, S. Lerch, S. Thanasilp, O. Kiss, O. Shaya, S. Vallecorsa, and Z. Holmes, npj Quantum Inf. 10, 116 (2024).
[4] E. Recio-Armengol, S. Ahmed, and J. Bowles, arXiv:2503.02934 (2025).
[5] N. L. Diaz, D. García-Martín, S. Kazi, M. Larocca, and M. Cerezo, arXiv:2310.11505 (2023).