

Classically trainable quantum generative modeling with Boson and Fermion Sampling

Bence Bakó, Zoltán Kolarovszki, Michał Oszmaniec, Changhun Oh, and Zoltán Zimborás

Based on Refs. [1, 2] (arXiv:2511.13844, arXiv:2603.11195).

Introduction. Sampling-based quantum advantage schemes currently serve as the leading experimental demonstrations of computational tasks deemed intractable for classical devices. However, their practical significance remains a subject of intense debate as these tasks are synthetically designed to maximize classical hardness rather than to solve meaningful real-world problems. In this contribution, based on Refs. [1, 2], we propose a generative learning application of two such schemes: Fermion Sampling [3] and Gaussian Boson Sampling [4]. We show that, despite the classical hardness of sampling from their output distributions, the generative models built on these schemes (Fig. 1) can nevertheless be trained efficiently by classical means. With these works, we extend the “train on classical, deploy on quantum” generative learning architectures beyond Instantaneous Quantum Polynomial-time (IQP) circuits [5]. We numerically demonstrate their performance and scalability, training fermionic and bosonic models with up to **160 qubits** and **805 photonic modes**, respectively; the latter holding the current parameter record with **1,297,660 trainable parameters**, surpassing the previous record of 500,500 [5].




FIG. 1. **Framework of classically trainable quantum generative models.** The proposed models, Fermionic Born Machines (FBMs) and Gaussian Bosonic Born Machines (GBBMs), are trained classically by minimizing the squared maximum mean discrepancy (MMD^2) between the model and data distributions; sampling from the trained model is then performed on a quantum device.

Classical training of quantum generative models. Quantum generative learning is a promising application of quantum computers, yet it faces several trainability challenges, including the difficulty in experimental gradient estimations. However, for certain structured quantum generative models, expectation values of local observables can be efficiently computed on a classical computer, enabling fully classical training without quantum gradient evaluations. This idea was explored in Refs. [5, 6] for IQP circuits, where direct estimation of probabilities and local expectation values can be performed efficiently on a classical computer.

These works rely on the paradigmatic framework of quantum circuit Born machines, where the model probability distribution is represented using a quantum circuit as $q_{\mathbf{w}}(\mathbf{x}) = |\langle \mathbf{x} | U(\mathbf{w}) | 0 \rangle|^2$, $\mathbf{x} \in \{0, 1\}^n$ and $U(\mathbf{w})$ is the quantum circuit ansatz with trainable parameters $\mathbf{w}$. A highly practical implicit loss function is the squared *maximum mean discrepancy* (MMD^2), which, as shown in Ref. [7], can be formulated using expectation values of Pauli-Z strings as

$$\mathcal{L}_{\text{MMD}^2}(p, q_{\mathbf{w}}) = \mathbb{E}_{i \sim p_K} [(\langle Z_i \rangle_p - \langle Z_i \rangle_{q_{\mathbf{w}}})^2], \quad (1)$$

where $\langle Z_i \rangle_p$ and $\langle Z_i \rangle_{q_w}$ denote the expectation values of the Pauli-Z strings for the target and model probability distribution, respectively. In the above expression, K refers to the choice of the kernel function and p_K denotes the corresponding probability distribution over the possible Pauli-Z strings. A typical choice is the Gaussian kernel $K_\sigma(\mathbf{x}, \mathbf{y}) := \exp(-\|\mathbf{x} - \mathbf{y}\|_2^2/2\sigma)$.

Fermionic Born Machines. Relying on the Fermion Sampling quantum advantage scheme [3], we introduce a classically trainable quantum generative model class called *Fermionic Born Machines* (FBMs) defined as follows:

Definition 1 (Fermionic Born Machine). An FBM on $d = N(k + m)$ fermionic modes is a quantum circuit with the following steps:

1. Initialize the d modes in some (parametrized) quantum state $\rho(\boldsymbol{\alpha})$;
2. Apply parametrized fermionic linear optical (FLO) transformation $U_{\text{FLO}}(\boldsymbol{\theta})$;
3. For each of the N registers, perform particle number measurement on the first k modes.

Based on the hyperparameters and the input state, these models have different expressive power and simulability. Unlike simple Gaussian inputs, the expressivity of the model can be significantly increased by considering parametrized magic input states, which we define as $\rho(\boldsymbol{\alpha}) = |\boldsymbol{\alpha}\rangle\langle\boldsymbol{\alpha}|$, where

$$|\boldsymbol{\alpha}\rangle := \bigotimes_{j=1}^N |\alpha_j\rangle, \quad |\alpha_j\rangle := \cos \alpha_j |0000\rangle + \sin \alpha_j |1111\rangle \quad (2)$$

and $\alpha_j \in [0, 2\pi)$ are trainable parameters. The important and desirable properties of FBMs with these parametrized magic input states can be summarized as follows:

- The output probability distribution can be classically hard to sample from under reasonable complexity theoretic assumptions;
- The expectation value of Pauli-Z strings of length ℓ can be computed exactly in $\mathcal{O}(\ell^3 4^\ell N^{\lceil \ell/2 \rceil})$ time, yielding a classically trainable quantum model;
- The trained model can be sampled efficiently on a qubit-based quantum computer using a linear-depth quantum circuit composed of $\mathcal{O}(d^2)$ matchgates.

To achieve the polynomial runtime discussed above, we introduce a novel technique based on a Gaussian decomposition of the input state. For the specific architecture considered here, our method outperforms state-of-the-art operator backpropagation algorithms [8, 9]. We note, however, that a more efficient approximation algorithm has recently been devised for the passive FLO version of this scheme [10]. In our work, we demonstrate the performance of our model and training framework, using a series of numerical experiments, considering datasets from various application domains and compare our results to classical approximations (Chow-Liu tree approximation, uniformly random samples) and learning models, such as restricted Boltzmann machines (RBMs). We further study important aspects of the framework, such as the effect of overparametrization, or the role of magic in the input state. The test results of our largest benchmark experiment on a gene sequence dataset, using 160 qubits is shown in Figure 2 (left).

Gaussian Bosonic Born Machines. Our second model class is inspired by the *Gaussian Boson Sampling* [4] quantum advantage proposal and generalizes the expectation value-based MMD² formulation from Pauli-Z observables to parity and threshold operators, their analogs in the photonic setting. This model class can be defined as follows.

Definition 2 (Gaussian Bosonic Born Machine). A *Gaussian Bosonic Born Machine* (GBBM) on d bosonic modes is a quantum generative model with the following components:

1. Initial state: a pure Gaussian state, such as the multimode vacuum $|0\rangle$,




FIG. 2. **Benchmark results of our largest experiments, as reflected by the MMD^2 metric with respect to a test set using different Gaussian kernel bandwidths.** (left) FBM (160 qubits) benchmark results on a gene sequence dataset. (right) GBBM (805 photonic modes) test results on a genomic dataset.

2. Variational circuit: a parameterized photonic quantum circuit composed of Gaussian operations, including interferometers, squeezing and displacement gates,
3. Measurement: parity or threshold measurements performed on each mode.

Although most Gaussian Boson Sampling experiments focus on centered Gaussian states, we incorporate displacement gates because they substantially enhance the expressivity of the model. Hence, in total, a single layer of the resulting ansatz contains $2d^2 + 2d$ trainable parameters.

Our efficient classical training strategy operates by tracking the evolution of the first and second statistical moments of the underlying Gaussian quantum state. Since each update involves matrix multiplications of size $2d \times 2d$, the evolution of a Gaussian state ρ can be simulated classically in $O(Ld^3)$ time, where L is the number of layers. Moreover, the expectation value of a parity operator string Π_A , acting on a subset of modes $A \subseteq [d]$ corresponds to

$$\text{Tr}[\Pi_A \rho] = \left(\frac{\pi}{2}\right)^{|A|} W_{\rho_A}(\mathbf{0}), \quad (3)$$

where W_{ρ_A} denotes the Wigner function of the reduced Gaussian state $\rho_A = \text{Tr}_{[d] \setminus A} \rho$. Hence, for a parity string of length $\ell = |A|$, the computational cost of obtaining the expectation value scales as $\mathcal{O}(\ell^3)$, yielding a highly scalable classically trainable quantum model. In practice, $\ell \ll d$ for the parity strings of interest, so this step is typically inexpensive. Instead, the overall computational cost might often be dominated by the construction of the full covariance matrix of the d -mode Gaussian state. As a result, expectation value estimation remains scalable even for bosonic systems with a large number of modes, with the bottleneck shifting from parity evaluation to covariance construction. While we mainly focus on parity expectation values, such observables may be insufficient for balanced datasets; therefore, we also consider threshold operator strings, which avoid this limitation at the cost of a more expensive training procedure.

We demonstrate the scalability of our models through a series of numerical experiments, the largest one using a genomic dataset, with bitstrings of length 805, leading to a GBBM with 805 modes and 1,297,660 trainable parameters. Remarkably, our Python implementation using JAX for automatic differentiation successfully trains this large-scale model using a single Nvidia A100 GPU in under 2 hours. The test results are shown in Figure 2 (right).

Outlook. A promising direction for future work involves extending both approaches beyond the Born machine framework to integrate them into more sophisticated computational pipelines and applications, such as generative models for optimization and quantum chemistry. Another interesting avenue is to supplement the circuits with nonlinear elements, which could enhance the expressivity of the model. Finally, it is worth mentioning that the classical hardness results concerning both quantum advantage schemes are usually formulated in different metrics, such as the total variation distance. Whether these circuits (including IQP) can be effectively spoofed in MMD^2 -related metrics is still an open question.

-
- [1] B. Bakó, Z. Kolarovszki, and Z. Zimborás, Fermionic Born Machines: Classical training of quantum generative models based on Fermion Sampling, arXiv preprint arXiv:2511.13844 (2025).
 - [2] Z. Kolarovszki, B. Bakó, M. Oszmaniec, C. Oh, and Z. Zimborás, Generative modeling with Gaussian Boson Sampling: classically trainable Bosonic Born Machines, arXiv preprint arXiv:2603.11195 (2026).
 - [3] M. Oszmaniec, N. Dangniam, M. E. Morales, and Z. Zimborás, Fermion sampling: a robust quantum computational advantage scheme using fermionic linear optics and magic input states, PRX Quantum **3**, 020328 (2022).
 - [4] C. S. Hamilton, R. Kruse, L. Sansoni, S. Barkhofen, C. Silberhorn, and I. Jex, Gaussian Boson Sampling, Phys. Rev. Lett. **119**, 170501 (2017).
 - [5] E. Recio-Armengol, S. Ahmed, and J. Bowles, Train on classical, deploy on quantum: scaling generative quantum machine learning to a thousand qubits, arXiv preprint arXiv:2503.02934 (2025).
 - [6] S. Kasture, O. Kyriienko, and V. E. Elfving, Protocols for classically training quantum generative models on probability distributions, Physical Review A **108**, 042406 (2023).
 - [7] M. S. Rudolph, S. Lerch, S. Thanasilp, O. Kiss, O. Shaya, S. Vallecorsa, M. Grossi, and Z. Holmes, Trainability barriers and opportunities in quantum generative modeling, npj Quantum Information **10**, 116 (2024).
 - [8] A. Angrisani, A. Schmidhuber, M. S. Rudolph, M. Cerezo, Z. Holmes, and H.-Y. Huang, Classically estimating observables of noiseless quantum circuits, Phys. Rev. Lett. **135**, 170602 (2025).
 - [9] A. Miller, J. Favre, Z. Holmes, Ö. Salehi, R. Chakraborty, A. Nykänen, Z. Zimborás, A. Glos, and G. García-Pérez, Simulation of fermionic circuits using majorana propagation, arXiv preprint arXiv:2503.18939 (2025).
 - [10] C. Oh, M. Oszmaniec, O. Reardon-Smith, and Z. Zimborás, Classical simulation of free-fermionic dynamics and quantum chemistry with magic input, arXiv preprint arXiv:2604.26813 (2026).