

Instantaneous Quantum Polynomial Circuit Born Machines: Trainability and Non-Simulability

Kevin Shen^{1,2,3}, Susanne Pielawa³, Vedran Dunjko^{1,2}, Hao Wang^{1,2}

¹ $\langle aQa^L \rangle$ Applied Quantum Algorithms, Leiden University, The Netherlands

²LIACS, Leiden University, Niels Bohrweg 1, 2333 CA, Leiden, The Netherlands

³BMW Group, 80788 München, Germany

1 Introduction

Instantaneous Quantum Polynomial Quantum Circuit Born Machines (IQP-QCBM) [11, 16] are emerging quantum generative models that have attracted increasing attention. IQP-QCBMs represent probability distributions over bitstrings through Born-rule measurements of parameterized IQP circuits. IQP-QCBMs are distinguished from many other quantum machine learning (QML) models in two ways: First, the models are conjectured to inherit the classical sampling hardness from IQP circuits under standard complexity-theoretic assumptions [1, 2]. Second, as Recio-Armengol et al. [14] showed, by exploiting the special circuit architectures of IQP circuits, the models can be trained on a classical computer using the maximum mean discrepancy (MMD) loss [8], while deployment on a quantum computer is required for sample generation.

Despite the growing body of exciting application-oriented work on IQP-QCBMs, their scalability depends on a central unresolved question: their trainability. A major obstruction to trainability in QML is the barren plateau (BP) phenomenon, in which the variance of the loss gradient with respect to the parameter distribution decays exponentially with the number of qubits n [4, 9, 10, 12]. To our knowledge, the occurrence of BPs in IQP-QCBMs has not been formally studied. *The fact that the MMD loss and its gradients for IQP-QCBMs can be evaluated classically does not imply trainability; it only addresses the cost of computing the loss signal, rather than this signal being significant enough for effective training.*

Another important criterion for QML architectures is that they possess representational or computational power beyond efficient classical simulation, thereby providing the possibility of quantum advantage. Accordingly, promising IQP-QCBM architectures must be assessed not only for trainability, but also for whether their output distribution families remain classically hard to simulate. Existing classical sampling-hardness results for IQP circuits have been established only under specific architectural assumptions [1–3]. Indeed, certain restricted IQP architectures are provably classically simulable; for example, planar IQP circuits are connected to exactly solvable Ising models and can be efficiently simulated in relevant regimes [7]. This issue is further highlighted by recent work on supervised QML, which argues that many architectures with provable absence of barren plateaus are also efficiently classically simulable, because the mechanisms that avoid gradient concentration often restrict the computation to polynomially sized subspaces [5]. It is therefore natural to ask whether an analogous trainability–simulability tradeoff arises in quantum generative modeling.

In this work, we show that whether an IQP-QCBM exhibits a BP depends on two practical design choices: the IQP circuit architecture and the kernel used in the MMD loss. We derived the closed-form expression of variance of the MMD loss gradient, and we identified the “critical rank” of the IQP-QCBM model, the dimension of the space spanned by the anti-commuting generators w.r.t. the observable, as the deciding factor for the presence or absence of BP. We summarize the main results of the critical rank in Table 1. Our analytical work provides guidance for avoiding architectures that are provably susceptible to BPs and narrows the design space of trainable IQP-QCBMs.

Specifically, we show that IQP-QCBMs whose architecture connectivity is generated by a sparse Erdős–Rényi graph with $\mathcal{O}(n \log n)$ edges are trainable, while their output distribution family remains hard to approximate classically, applying theoretical arguments from Bremner et al. [3].

The full presentation will also include the characterization of IQP-QCBM’s expressivity and experiment results for material science applications.

2 Background

Definition 2.1 (IQP-QCBM). An n -qubit IQP-QCBM is a variational quantum generative model defined by a commuting diagonal circuit $U_{\theta} = \prod_{j=1}^D \exp(i\theta_j Z^{s^j})$, $\theta_j \in [0, 2\pi)$, $s^j \in \mathbb{F}_2^n$, where $D \in \mathcal{O}(\text{poly}(n))$. The model prepares $U_{\theta} |+\rangle^{\otimes n}$ and measures all qubits in the Pauli- X basis. By the Born rule, sampling from the model follows the distribution: For each $x \in \mathbb{F}_2^n$, $q_{\theta}(x) = \text{Tr} \left[|x\rangle\langle x| H^{\otimes n} U_{\theta} \rho U_{\theta}^{\dagger} H^{\otimes n} \right]$, where $\rho = |+\rangle\langle +|^{\otimes n}$. We refer to $S = \{s^j\}_j$ the generator set, which determines the model architecture.

Proposition 2.2 (MMD loss in IQP-QCBM). *Consider an IQP model and its output distribution q_θ defined in Theorem 2.1. Given a target distribution p over $\mathbb{F}_2^n$ and a kernel k , the MMD loss between p and q_θ admits the following form [14, 15]: $\mathcal{L}(\theta) := \text{MMD}_k^2(p, q_\theta) = \mathbb{E}_{\mathbf{a} \sim \Lambda} \left[(C_p(\mathbf{a}) - C_\theta(\mathbf{a}))^2 \right]$, where $C_p(\mathbf{a}) = \mathbb{E}_{\mathbf{x} \sim p} [(-1)^{\mathbf{x} \cdot \mathbf{a}}]$ and $C_\theta(\mathbf{a}) = \text{Tr}(X^\mathbf{a} U_\theta \rho U_\theta^\dagger)$ and Λ is the Fourier transform of k .*

In QML, loss functions expressed in terms of expectation values often suffer from the barren plateau (BP) phenomenon, whereby gradients vanish exponentially with the number of qubits n [12]. Ragone et al. [13] proposed a general algebraic framework for bounding gradient variance using the dynamical Lie algebra (DLA) generated by the circuit, and Fontana et al. [6] obtained related results using Lie group representation theory.

We note, however, that the DLA-based approach of Ragone et al. [13] does not directly apply to the BP analysis of the MMD loss in IQP-QCBMs. In IQP-QCBMs, the gate generators are Pauli- Z strings, whereas the observables appearing in the MMD loss are Pauli- X strings. These operators are orthogonal under the Hilbert–Schmidt inner product. Consequently, the required nonzero projection of either the input state or the observable onto the circuit’s DLA is absent. We therefore develop our analytical results from first principles.

3 Main results

Theorem 3.1 (Critical rank determines the variance of the characteristic function’s partial derivative). *Consider an n -qubit IQP-QCBM in Theorem 2.1 with the generator set $S = \{\mathbf{s}^j\}_{j=1}^D$, and let the reference measure on the parameter space be uniform, namely $\mu_n = \mathcal{U}([0, 2\pi)^D)$. For a frequency $\mathbf{a} \in \mathbb{F}_2^n \setminus \{\mathbf{0}\}$, define $S^\mathbf{a} := \{\mathbf{s}^j \in S \mid \mathbf{a} \cdot \mathbf{s}^j = 1\}$ and define the critical rank $r^\mathbf{a} := \dim \text{span}_{\mathbb{F}_2}(S^\mathbf{a})$. Then*

$$\text{Var}_{\theta \sim \mu_n} [C_\theta(\mathbf{a})] = 2^{-r^\mathbf{a}}. \quad (1)$$

For $\ell \in [D]$, if $\mathbf{s}^\ell \notin S^\mathbf{a}$, then $\frac{\partial C_\theta(\mathbf{a})}{\partial \theta_\ell} = 0$. If $\mathbf{s}^\ell \in S^\mathbf{a}$, then

$$\text{Var}_{\theta \sim \mu_n} \left[\frac{\partial C_\theta(\mathbf{a})}{\partial \theta_\ell} \right] = 2^{2-r^\mathbf{a}}. \quad (2)$$

The critical rank $r^\mathbf{a}$ depends on the frequency $\mathbf{a}$, as well as on the type and number of generators in the model. In contrast to other QML models, where gradient variances are often bounded only asymptotically as functions of circuit depth, IQP-QCBMs allow an exact characterization: once the generator set is specified, the critical ranks and therefore the variances, can be computed exactly. We apply Theorem 3.1 to IQP-QCBMs of different architectures.

Table 1: We consider four architectures for the IQP circuit (Figure 1), and report the scaling of the critical rank $r^\mathbf{a}$ (Theorem 3.1) – the key property of the circuit determining the scaling of MMD’s derivative. “no B.P.” indicates the absence of barren plateau in the MMD loss. “A.C.” means L^2 anti-concentration of the IQP-QCBM’s output distribution, needed to claim classical non-simulability.

Architecture	$r^\mathbf{a}$ ($ \mathbf{a} = \mathcal{O}(\log n)$)	no B.P.	A.C.
Product	$\mathcal{O}(\log n)$	✓	×
2D lattice	$\mathcal{O}(\log n)$	✓	×
Sparse ER	$\mathcal{O}(\text{polylog}(n))$	✓	✓
Complete	n	×	✓

Example 3.2 (Sparse Erdős–Rényi graph). *Consider a generator set containing all single-qubit generators, and a random subset of two-qubit generators. Specifically include each two-qubit generator $e_u + e_v$ independently with probability $p = c \ln(n)/n$ for some constant $c = \Theta(1)$. Equivalently, let $G = (V, E)$ be an Erdős–Rényi graph with edge probability $c \ln(n)/n$, and include $e_u + e_v$ iff $(u, v) \in E$. We can compute that $\mathbb{E}_G[r^\mathbf{a}] \in \Theta(n)$ if $|\mathbf{a}| \in \Theta(n)$. By Chernoff’s bound, with high probability, for a single Erdős–Rényi graph, low Hamming-weight Fourier frequencies with $|\mathbf{a}| \in \Theta(1)$ will have polynomial decay, ($|\mathbf{a}| \in \mathcal{O}(\log(n))$) gives subexponential but not*

polynomial), while high Hamming-weight frequencies with $|\mathbf{a}| \in \Theta(n)$ will have exponential decay variances. Therefore, by choosing a kernel with sufficiently large weights on the low Hamming-weight part of the spectrum, one can avoid barren plateaus.

We identified IQP-QCBMs of sparse-ER architecture as potentially relevant, in the sense that they are trainable in the low-weight moments while remaining non-dequantizable as a sampling family. Nevertheless, whether the learning task remains classically hard, when given access to data, remains as an open question.

Theorem 3.3 (Trainable but not classically-simulable, for low moments). *Under the average-case hardness assumption for discrete-parameter sparse IQP circuits [3] and noncollapse of polynomial hierarchy, for every fixed body order k_0 , there exists a high-probability sparse-IQP family $\mathcal{T} = \{\mathcal{T}_n\}_n$ such that:*

1. every $C \in \mathcal{T}_n$ has inverse-polynomial gradient variance for all Fourier moments of Hamming weight at most k_0 , and hence a kernel focusing on low-weight frequencies grants trainability (no BP).
2. no single uniform polynomial-time classical sampler can approximately sample all circuits in $\mathcal{T}_n$.

We provide two experimental validations of our theoretical predictions. First, we verify the scaling of MMD’s partial derivative w.r.t. the qubit number on synthetic datasets of size $n = 20, 40, \dots, 200$, whose marginals and correlations approximately match the Fourier spectrum of binarized MNIST. We test the five IQP architectures considered in the paper: Product States, 2D Lattice, Sparse ER with $c = 0.5$, Sparse ER with $c = 2$, and (complete) all-to-all-connected 2D graph. We consider three kernels supported respectively on frequencies of Hamming-weight 2, $\log(n)$, and $n/4$. The results are shown in Figure 1 (top row). Across all kernels and architectures, under uniform distribution of parameters, the observed gradient scaling is either polynomial or exponential in n , in agreement with the predictions from Theorem 3.1. Second, we study MMD loss trajectories to test whether the predicted presence or absence of barren plateaus matches practical trainability. We show the training MMD trajectories in Figure 1 (bottom row). The 2D Lattice and Sparse ER architectures exhibit improved training stability and convergence as the sample size for gradient estimation increases. Moreover, the 2D Lattice architecture consistently shows a faster convergence than Sparse ER, matching its smaller critical rank. In contrast, the (complete) all-to-all-connected graph architecture remains untrainable even at the largest sample size, consistent with the predicted barren plateau.




Figure 1: Top: Scaling of MMD’s derivative w.r.t. the qubit number, measured on MNIST-inspired datasets with the mean and standard deviation reported over 50 repetitions. Bottom: Trajectories of the training MMD loss for learning the 784-qubit binarized MNIST dataset, with increasing sample sizes for estimating the MMD loss.

References

- [1] Michael J. Bremner, Richard Jozsa, and Dan J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, August 2010. ISSN 1471-2946. doi: 10.1098/rspa.2010.0301.
- [2] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. Average-case complexity versus approximate simulation of commuting quantum computations. *Physical Review Letters*, 117(8), August 2016. ISSN 1079-7114. doi: 10.1103/physrevlett.117.080501.
- [3] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. Achieving quantum supremacy with sparse and noisy commuting quantum computations. *Quantum*, 1:8, April 2017. ISSN 2521-327X. doi: 10.22331/q-2017-04-25-8.
- [4] M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature Communications*, 12(1), March 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-21728-w.
- [5] M. Cerezo, Martin Larocca, Diego García-Martín, N. L. Diaz, Paolo Braccia, Enrico Fontana, Manuel S. Rudolph, Pablo Bermejo, Aroosa Ijaz, Supanut Thanasilp, Eric R. Anschuetz, and Zoë Holmes. Does provable absence of barren plateaus imply classical simulability? *Nature Communications*, 16:7907, 2025. doi: 10.1038/s41467-025-63099-6.
- [6] Enrico Fontana, Dylan Herman, Shouvanik Chakrabarti, Niraj Kumar, Romina Yalovetzky, Jamie Heredge, Shree Hari Sureshbabu, and Marco Pistoia. Characterizing barren plateaus in quantum ansätze with the adjoint representation. *Nature Communications*, 15(1):7171, 2024.
- [7] Keisuke Fujii and Tomoyuki Morimae. Commuting quantum circuits and complexity of ising partition functions. *New Journal of Physics*, 19(3):033003, March 2017. ISSN 1367-2630. doi: 10.1088/1367-2630/aa5fdb.
- [8] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012.
- [9] Zoë Holmes, Kunal Sharma, M. Cerezo, and Patrick J. Coles. Connecting ansatz expressibility to gradient magnitudes and barren plateaus. *PRX Quantum*, 3(1), January 2022. ISSN 2691-3399. doi: 10.1103/prxquantum.3.010313.
- [10] Martín Larocca, Supanut Thanasilp, Samson Wang, Kunal Sharma, Jacob Biamonte, Patrick J. Coles, Lukasz Cincio, Jarrod R. McClean, Zoë Holmes, and M. Cerezo. Barren plateaus in variational quantum computing. *Nature Reviews Physics*, 7(4):174–189, March 2025. ISSN 2522-5820. doi: 10.1038/s42254-025-00813-9.
- [11] Jin-Guo Liu and Lei Wang. Differentiable learning of quantum circuit born machines. *Physical Review A*, 98(6), December 2018. ISSN 2469-9934. doi: 10.1103/physreva.98.062324.
- [12] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812, 2018.
- [13] Michael Ragone, Bojko N. Bakalov, Frédéric Sauvage, Alexander F. Kemper, Carlos Ortiz Marrero, Martín Larocca, and M. Cerezo. A lie algebraic theory of barren plateaus for deep parameterized quantum circuits. *Nature Communications*, 15(1), August 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-49909-3.
- [14] Erik Recio-Armengol, Shahnawaz Ahmed, and Joseph Bowles. Train on classical, deploy on quantum: scaling generative quantum machine learning to a thousand qubits. *arXiv preprint arXiv:2503.02934*, 2025.
- [15] Manuel S Rudolph, Sacha Lerch, Supanut Thanasilp, Oriel Kiss, Oxana Shaya, Sofia Vallecorsa, Michele Grossi, and Zoë Holmes. Trainability barriers and opportunities in quantum generative modeling. *npj Quantum Information*, 10(1):116, 2024.

- [16] Dan Shepherd and Michael J. Bremner. Temporally unstructured quantum computation. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465(2105):1413–1439, 02 2009. ISSN 1364-5021. doi: 10.1098/rspa.2008.0443.