

# Classically trained, quantum-deployed generative models: universality and kernel-adaptive training

Andrii Kurkin<sup>1,2</sup>, Kevin Shen<sup>1,2,3</sup>, Ulysse Chabaud<sup>4</sup>, Zoltán Kolarovszki<sup>5,6</sup>, Bence Bakó<sup>5,6</sup>,  
Susanne Pielawa<sup>3</sup>, Hao Wang<sup>1,2</sup>, Zoltán Zimborás<sup>7,8</sup>, and Vedran Dunjko<sup>1,2</sup>

<sup>1</sup> $\langle aQa^L \rangle$  Applied Quantum Algorithms, Leiden Univ., NL <sup>2</sup>LIACS, Leiden Univ., NL <sup>3</sup>BMW Group,  
München, DE <sup>4</sup>DIENS, ENS/PSL/CNRS/INRIA, Paris, FR, <sup>5</sup>HUN-REN Wigner RCP, Budapest, HU  
<sup>6</sup>Eötvös Loránd Univ., Budapest, HU <sup>7</sup>Univ. Helsinki, FI <sup>8</sup>Algorithmiq Ltd, Helsinki, FI

## 1 Overview

A fundamental bottleneck holds back quantum machine learning: heuristic quantum methods cannot yet be meaningfully tested, since current hardware cannot run them at interesting scales. What one would ideally want is a method that can be meaningfully benchmarked *without* being run on a quantum device. This seems impossible, yet the *train-classically, deploy-quantumly* paradigm [1] achieves it: the training objective is evaluated efficiently on a classical computer, while sampling from the trained model may remain classically intractable. In this work we show that truly universal quantum generative models can be achieved in this paradigm, using two very different restricted models of computation with sampling-advantage guarantees: the instantaneous quantum polynomial (IQP) circuit Born machine (QCBM) and the Boson Sampling Born Machine (BSBM). This extended abstract is based on two manuscripts<sup>1</sup>, where technical details and proofs can be found.

For IQP-QCBM models, previous work had already shown that while such models do not allow classical sampling under standard assumptions [2], certain loss functions remain classically computable, and the model can be classically trained with respect to them [1, 3]. Nonetheless, two critical issues are identified: the model is known to be *not* universal i.e., not capable of representing arbitrary distributions; and the fixed Gaussian kernel for which this efficient training was established [1] can restrict the discriminative power of the MMD loss and cause training issues. For the former, we give the first universality result for IQP-QCBM: adding  $n$  *hidden* qubits to an  $n$ -qubit IQP generator makes the model universal in the worst case. This result plays the role that universal-approximation theorems play in classical machine learning: without it, one cannot tell whether the architecture is merely a useful heuristic sampler or a principled model class capable of representing any target distribution. For the latter, we propose a *kernel-adaptive* training method in which the kernel is adversarially trained. We prove that such adaptive kernels have greater discriminative power than Gaussian fixed kernel, and that the kernel-adaptive MMD is consistent with convergence in distribution of the generator. We validate the method empirically on a parity-check benchmark and binarized MNIST.

We then ask whether the same paradigm can be realized with a very different model of computation. Motivated by sampling-hardness results for linear optics [4], and building on prior explorations of photonic generative modeling [5], we focus on Boson Sampling as the basis for a generative model, which we make rigorous and call the Boson Sampling Born Machine (BSBM): a boson sampler with trainable optical parameters whose Fock-basis measurement defines the model distribution. Mirroring the IQP-QCBM results, we show that BSBMs are classically trainable for broad families of MMD losses. As with IQP, the bare BSBM is *not* universal, though the structural obstructions are different. Nonetheless, we resolve non-universality by extending the model with a fixed postprocessing map on the measurement outcomes. Finally, we prove the existence of a sequence of models with increasing expressivity that achieves universality in the limit while retaining hardness in the intermediate regime.

---

<sup>1</sup>[arxiv.org/abs/2510.08476](https://arxiv.org/abs/2510.08476), [arxiv.org/abs/2603.11014](https://arxiv.org/abs/2603.11014)

## 2 IQP-QCBM

IQP-QCBM is a generative model, which consists of a parameterized IQP circuit acting on  $n$  qubits and preparing the state  $|\psi(\theta)\rangle$ , followed by measurement in the computational basis. The measurement produces bitstrings  $x \in \{0,1\}^n$  according to the probability distribution  $q_\theta(x) = |\langle x | \psi(\theta) \rangle|^2$ . To address the non-universality of this model, we extend the original IQP-QCBM architecture by incorporating hidden qubits.

**Definition 1** (IQP-QCBM with hidden qubits). *Let the full system have  $m+n$  qubits, where the final  $n$  qubits are designated as visible and the first  $m$  as hidden. An IQP-QCBM with hidden qubits is an IQP-QCBM on  $m+n$  qubits in which the hidden subsystem is traced out prior to measurement. The output distribution on the visible subsystem is  $q_\theta(x) = \text{Tr}(\text{Tr}_{\text{hidden}}(\rho) |x\rangle\langle x|)$ , where  $\rho$  is the state produced by the parametrized IQP circuit.*

**Theorem 1** (Exact universality for IQP-QCBM). *For any target probability distribution  $p$  over  $\{0,1\}^n$ , there exists an IQP-QCBM acting on  $n$  visible qubits and  $m = n$  hidden qubits whose output distribution  $q_\theta$  equals  $p$  exactly.*

**Remark 1.** *While  $n$  ancillas suffice, their worst-case optimality is open: an explicit Sidon-set construction suggests a lower bound of  $\lfloor n/2 \rfloor$ , so the worst-case optimum lies in  $[\lfloor n/2 \rfloor, n]$ . For specific targets, far fewer ancillary qubits suffice, especially if extremely high accuracy is not required. As an example, we exhibit distributions unreachable by the visible-only model yet reproduced exactly with a single hidden qubit.*

Although the Gaussian-kernel MMD enables classical training of IQP-QCBMs, the fixed Gaussian choice has limited discriminative power: there exist pairs of distributions that are far apart in TV distance but on which the Gaussian-kernel MMD is exponentially small in  $n$ , while an alternative characteristic kernel keeps it bounded below.

**Lemma 1** (Adaptive vs. Gaussian discriminative power). *For every  $n \in \mathbb{N}$ , there exist a target distribution  $p$  on  $\{0,1\}^n$ , an IQP circuit on  $n$  qubits with a parameter setting  $\theta$ , and the corresponding distribution  $q_\theta$ , such that: (1)  $\text{TV}(p, q_\theta) \in \Omega(1)$ ; (2) for any Gaussian kernel  $k_\sigma$ ,  $\text{MMD}_{k_\sigma}^2(p, q_\theta) \in \mathcal{O}(2^{-n})$ ; (3) there exists a constant  $C > 0$  such that for every  $n$  some characteristic kernel  $\kappa$  achieves  $\text{MMD}_\kappa^2(p, q_\theta) \geq C$ .*

**Theorem 2** (Generalized kernel MMD). *For any stationary and bounded kernel  $k$  over  $\{0,1\}^n$ , with Fourier transform  $G$ , it holds that:  $\text{MMD}_G^2(p, q_\theta) = \mathbb{E}_{\alpha \sim G} (\langle Z_\alpha \rangle_p - \langle Z_\alpha \rangle_{q_\theta})^2$*

This theorem recovers the Gaussian-kernel result of [1] and the expression admits efficient estimation.




Figure 1: Adaptive kernel vs. fixed Gaussian.

We use this result to introduce kernel-adaptive training, which allows the kernel to be chosen during training. We parametrize the spectral measure  $G_\gamma$  with a classical critic network and train adversarially, defining  $\mathcal{L}(p, q_\theta) := \max_\gamma \text{MMD}_{G_\gamma}^2(p, q_\theta)$  subject to  $\sum_\alpha G_\gamma(\alpha) = 1$  and

$G_\gamma(\alpha) > 0$ . The critic concentrates spectral mass where  $\langle Z_\alpha \rangle_p$  and  $\langle Z_\alpha \rangle_{q_\theta}$  differ most, and an unbiased estimator of  $\mathcal{L}$  and its gradients is available classically. We empirically validate the kernel-adaptive method on two datasets; results are presented in Figure 1. For the *parity-check* dataset, the adaptive kernel achieves lower total variation error. For *binarized MNIST* ( $n = 784$ ), both methods minimize their respective MMD training losses, which converge in each case. Yet only the adaptive kernel keeps the data-informed second-moment error decreasing, ending well below the fixed Gaussian, which briefly lowers it early on before rising and saturating an order of magnitude higher.

### 3 Boson Sampling Born Machine

**Definition 2** (Boson Sampling Born Machine). *A BSBM is a generative model specified by the number of optical modes  $m$ , the linear-optical network described by the unitary  $U \in \mathcal{U}(m)$ , and a fixed initial state of  $k$  photons. By a photon-number measurement of all output modes, postselecting on the no-collision case, it represents a family of distributions  $\mathcal{Q}(m, k) := \{q_{m,k,U}\}_{U \in \mathcal{U}(m)}$  with sample space  $\Omega_m^k := \{\mathbf{s} \in \{0, 1\}^m : \|\mathbf{s}\|_1 = k\}$ , where  $q_{m,k,U}$  is the probability distribution on  $\Omega_m^k$  given by  $q_{m,k,U}(\mathbf{s}) = |\text{Per}(U_{S(\mathbf{s}),[k]})|^2$  and  $S(\mathbf{s})$  indexes the occupied modes. The interferometer  $U$  is parametrized, and its parameters are the trainable parameters of the model.*

First, we show that a BSBM can be trained w.r.t. the MMD loss efficiently on a classical computer. We apply Theorem 2, with the only difference that instead of  $Z_\alpha$  we use the parity word  $\Pi_\alpha = \bigotimes_j (-1)^{\alpha_j \hat{n}_j}$  supported on  $\alpha$ . On the model side,  $\langle \Pi_\alpha \rangle_{q_\theta} = \langle \psi_{\text{in}} | \hat{U}^\dagger(\theta) \Pi_\alpha \hat{U}(\theta) | \psi_{\text{in}} \rangle$  equals the permanent of a  $k \times k$  submatrix of  $U^\dagger(\theta) \text{diag}(1 - 2\alpha) U(\theta)$ , which Gurvits’ algorithm [6, 4] approximates to additive error efficiently. Efficient estimation extends to its gradients, so the kernel-adaptive training procedure, introduced above for IQP-QCBM, applies here as well.

Second, the bare BSBM is not universal over  $\{0, 1\}^m$ . There are two restrictions: (1) the generated strings have fixed Hamming weight  $k$ ; (2) the interferometer has only  $O(m^2)$  parameters against the  $2^m - 1$  degrees of freedom of an arbitrary distribution. We restore universality by introducing the *extended* BSBM (EBSBM), which augments the bare model with a fixed postprocessing map  $f : \Omega_m^k \rightarrow \{0, 1\}^n$ , producing the pushforward distribution  $q_{m,k,U,f}^*(x) = \sum_{\mathbf{s}: f(\mathbf{s})=x} q_{m,k,U}(\mathbf{s})$ . This already yields a simple universal construction: fix  $n \in \mathbb{N}$  and take the EBSBM with a single photon ( $k = 1$ ),  $m = 2^n$  modes, and a postprocessing map  $f : \Omega_{2^n}^1 \rightarrow \{0, 1\}^n$  returning the position of the measured photon in binary. In practice, however, we seek intermediate regimes that use fewer parameters yet still cover rich classes of distributions while retaining sampling-hardness potential; we establish this possibility below.

**Theorem 3** (Expressivity sequence). *For each target output size  $n$ , there exists a sequence of EBSBM models indexed by a single capacity-increasing hyperparameter  $j$  with properties: (1) expressivity increases monotonically in  $j$ ; (2) universality over  $\{0, 1\}^n$  in the large-capacity limit; and (3) inheritance of worst-case boson-sampling hardness for  $\text{poly}(n)$ -sized model descriptions.*

### 4 Outlook

We advanced the train-classically, deploy-quantumly paradigm on two models of computation: IQP-QCBM and BSBM. We resolve universality in both, and introduce kernel-adaptive training with empirical validation up to 784 qubits. Several questions remain open: the gap between MMD-based training and total-variation hardness at inference, genuine trainability beyond efficient loss estimation, and the compatibility of trainability with sampling hardness, for which initial progress exists in the case of IQP-QCBMs [7]. Realizing the paradigm’s promise ultimately calls for an end-to-end deployment of these classically trained models on quantum hardware at scale.

## References

- [1] Erik Recio-Armengol, Shah Nawaz 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.
- [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), 2016.
- [3] Erik Recio-Armengol and Joseph Bowles. Iqpopt: Fast optimization of instantaneous quantum polynomial circuits in jax. *arXiv preprint arXiv:2501.04776*, 2025.
- [4] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 333–342, 2011.
- [5] Tigran Sedrakyan and Alexia Salavrakos. Photonic quantum generative adversarial networks for classical data. *Optica Quantum*, 2(6):458–467, 2024.
- [6] Leonid Gurvits. On the complexity of mixed discriminants and related problems. In *International Symposium on Mathematical Foundations of Computer Science*, pages 447–458. Springer, 2005.
- [7] Kevin Shen, Susanne Pielawa, Vedran Dunjko, and Hao Wang. Characterizing trainability of instantaneous quantum polynomial circuit born machines, 2026.