

Representation theory for passive linear optics: anti-concentration, barren plateaus and classical simulability

Hela Mhiri, Léo Monbroussou, Hugo Thomas, Ulysse Chabaud, Zoë Holmes, and Elham Kashefi

This submission is based on Ref. [1] and [2] (in preparation).

Introduction Photonic quantum systems are a promising platform for near-term quantum technologies, combining low noise, high-fidelity linear-optical operations, and scalability to large numbers of optical modes. In particular, passive linear optics provides an experimentally mature model of bosonic quantum computation that has emerged as a candidate for demonstrating quantum advantage through sampling tasks, while increasingly being explored as a hardware-efficient architecture for quantum machine learning and variational quantum algorithms [3–9]. Despite this progress, a general characterization of the statistical properties of passive linear-optical circuits remains elusive. In particular, no general framework currently exists for evaluating second moments of bosonic observables under random interferometers, beyond specific observables or asymptotic regimes.

Second moments play a central role in several aspects of quantum machine learning. They characterize the fluctuations of output probabilities that underpin anti-concentration, a key ingredient in hardness-of-sampling arguments that has remained conjectural beyond the dilute regime. For variational quantum algorithms, they quantify the concentration of expectation values and gradients responsible for barren plateaus. While second-moment analyses have become standard for qubit and fermionic architectures, no comparable theory has been available for passive linear optics in experimentally relevant regimes [10–14].

In this work, we develop the first general framework for computing second moments of particle-number-preserving bosonic observables under Haar-random passive linear-optical transformations. Exploiting representation-theoretic tools for the unitary group, we derive closed-form expressions valid for arbitrary number of photons and modes, thereby avoiding the regime-specific approximations underlying previous analyses [15, 16].

We first show that our formalism provides a unified approach to analysing concentration phenomena in variational photonic algorithms. By classifying observables according to their representation-theoretic structure, we identify broad families of trainable observables that also admit efficient classical simulation. Conversely, for Fock-state inputs, we identify a class of observables, which we call *generalised parity operators*, that can evade concentration while seemingly remaining beyond the reach of currently known simulation techniques, suggesting a regime in which trainability and classical hardness may coexist. Moreover, we further establish anti-concentration of Boson Sampling output probabilities in the experimentally relevant regime, resolving an open conjecture and taking a step closer to a rigorous proof of the hardness of boson sampling [15, 17].

Framework The bosonic creation and annihilation operators act on the m -mode (infinite-dimensional) Fock-Hilbert space $\mathcal{F} = \bigoplus_{n \geq 0} \mathcal{F}_n$, with $\mathcal{F}_n$ the subspace of n photons spread across m modes, satisfying $\dim(\mathcal{F}_n) = \binom{n+m-1}{n}$ as it is spanned by the set of m -tuples of nonnegative integers summing to n , which we denote by $\Phi_{m,n}$. The group of m -mode passive linear optical transformations can be identified with the unitary group $U(m)$ and its action on the n -photon subspace is given by the totally symmetric irrep $\varphi_n(U) := U^{\otimes n}|_{\mathcal{F}_n}$. Our analysis focuses on the space $\mathcal{W} = \bigoplus_{n \geq 0} W_n$ of particle-number-preserving bosonic observables, where each fixed photon-number subspace corresponds to $W_n = \mathcal{F}_n \otimes \mathcal{F}_n^*$. We consider expectation values of such observables under random interferometers, i.e. $U \sim \text{Haar}(U(m))$, of the form $f_U(\rho, O) = \text{Tr}[\rho \varphi_n(U)^\dagger O \varphi_n(U)]$, where $O \in \mathcal{W}$ and either ρ or O is supported on the n -photon subspace W_n . Although the dynamics within each fixed-particle-number sector are described by the representation $\varphi_n(U)$ of the low dimensional Lie group $U(m)$, evaluating expectation values can nevertheless be computationally hard.

A canonical example is provided by Boson Sampling output probabilities, where both the input state and the measurement are projectors onto Fock states [15]. In particular, for an output configuration $S \in \Phi_m^n$ of a boson sampler, we define $p_U(S)$ as the probability of observing outcome $|S\rangle$ after evolving the reference input state $|S_{\text{cf}}\rangle = |1, \dots, 1, 0, \dots, 0\rangle$ under the action of the interferometer U . These probabilities are $\#P$ -hard to

compute exactly in general, which underpins the computational complexity of Boson Sampling and motivates the study of the more general class of bosonic observables for variational photonic models.

Main results We have three main results. First, we develop a representation-theoretic framework that yields *closed-form expressions for second moments of arbitrary particle-number-preserving observables Haar under random interferometers*. Second, we evaluate these expressions for relevant classes of bosonic observables, providing a unified representation-theoretic interpretation of generalized entanglement and locality in the bosonic setting. This allows us to study concentration in photonic quantum machine learning models. We further relate these concentration properties to classical simulation techniques, identifying broad classes of trainable observables that admit efficient classical simulation algorithms. Finally, applying this framework to Fock-state Boson Sampling, we derive a *closed-form expression for the normalized average collision probability* and thereby *prove anti-concentration in the saturated regime*, a conjecture previously only supported by numerical evidence [15, 17] and that is closely associated with hardness claims for sampling.

a. Second moment expression. Our first result gives a general formula for second moments of expectation values of particle-number-preserving observables of the form under Haar-random passive linear-optical transformations. Specifically, by averaging over random interferometers, the second moment expression reduces to computing Hilbert–Schmidt norms of projections into irreducible representation, i.e.

$$\mathbb{E}_{U \sim U(m)} [f_U(\rho, O)^2] = \sum_{\lambda=0}^n \frac{1}{d_{\lambda}^{(n)}} \|P_{\lambda}^{(n)}(\rho)\|_2^2 \|P_{\lambda}^{(n)}(O)\|_2^2, \quad (1)$$

where λ ranges over all irreps of W_n (under the adjoint action of φ_n) and $P_{\lambda}^{(n)}(\cdot)$ projects onto the subspace carrier of λ of dimension $d_{\lambda}^{(n)}$. This is stated formally in Proposition 2 of [1]. Next, we derive a closed-form expression of Eq.(1). To achieve this, we derive intermediary results for irrep projections and their norms, which may be of independent interest. Specifically, we exploit the Howe-duality structure of the operator space to derive expressions for these norms avoiding the computation of Clebsch–Gordan coefficients. We derive a recursive formula for computing the projections $P_{\lambda}^{(n)}(O)$ given in Proposition 3 of [1]. This recursive projection method yields a compact expression of projection norms given in Corollary 1 of [1], that can be efficiently computed for a wide class of bosonic observables of interest.

b. Exponential concentration and classical simulability. Variational quantum algorithms, in which a parametrized quantum circuit is trained by a classical optimizer, have emerged as a popular paradigm for near-term quantum devices [18–20]. Their practical viability, however, is threatened by the Barren Plateau (BP) phenomenon [21]: for wide classes of circuits, the loss function concentrates exponentially around its mean as the system size grows, so that gradients vanish exponentially fast, precluding the efficiency of training strategies [22, 23]. We study the presence of BP in linear optical models, and study their link to classical simulation. We show that concentration is governed by the misalignment of the irrep projections of the input state and the observable, providing a unified representation-theoretic interpretation of generalized entanglement and locality in the bosonic setting. We further relate these concentration properties to existing classical simulation techniques, identifying broad classes of trainable observables that admit efficient classical simulation. Conversely, we show that for Fock-state inputs, generalized parity operator seem to evade concentration while remaining beyond the reach of currently known simulation methods, revealing a regime where trainability and classical hardness may coexist. We summarize our results in the table of Figure 1.

c. Anti-concentration. Our last main result is a proof of anti-concentration of the output distribution of Fock-states Boson Sampling in the *saturated* regime, where the number of modes grows linearly with the number of photons. Anti-concentration is a property of random sampling schemes, related to the *flatness* of the distribution. It is a key ingredient in hardness arguments of approximate sampling, but while proved for many schemes [10–14] remains poorly understood in Boson Sampling. Anti-concentration of Boson Sampling can be formulated as $\Pr_{S,U}[p_U(S) \geq \alpha/|\Phi_m^n|] \geq \gamma$ for constants $\alpha \in [0, 1]$ and $\gamma > 0$. A natural quantity capturing anti-concentration is the so-called *normalized average outcome-collision probability* defined as $P_2(m, n) = |\Phi_m^n| \sum_{S \in \Phi_m^n} \mathbb{E}_{U \sim U(m)} [p_U(S)^2]$. To prove anti-concentration, we use the second moment expression for boson sampling output probabilities and obtain an analytical expression for $P_2(m, n)$ in Theorem 1

$p \backslash O$	$\hat{n}_i^p$		$\hat{n}_{i_1} \dots \hat{n}_{i_p}$		$\frac{\prod \hat{a}_{i_k}^\dagger \hat{a}_{j_k} + \hat{a}_{j_k}^\dagger \hat{a}_{i_k}}{\sqrt{2}}$		$\prod_{i=1}^p (-1)^{\alpha_i \hat{n}_i^2}$		$\prod_{i=1}^p (-1)^{\hat{n}_i}$		$ S \times S $	
	$\mathbb{E}[f_U^2]$	Sim	$\mathbb{E}[f_U^2]$	Sim	$\mathbb{E}[f_U^2]$	Sim	$\mathbb{E}[f_U^2]$	Sim	$\mathbb{E}[f_U^2]$	Sim	$\mathbb{E}[f_U^2]$	Sim
$\Theta(1)$	BP		BP		BP		BP		BP		BP	
$\Theta(\log_2 n)$	unclear		unclear		unclear		BP		unclear		BP	
$\Omega(\sqrt{n})$	BP		BP		BP		BP		BP		BP	

FIG. 1: Exponential concentration and classical simulation for observables in passive linear optics. The different regimes are described by p , with p the number of occupied modes for the last case. All observables considered are normalized by its infinite value in order to make it shot efficient to measure on a quantum computer. We highlight the regime where there are no barren plateaus but also currently no known provably efficient classical simulation methods.

of [1], valid for arbitrary numbers of photons and modes including both dilute and saturated regimes. From this expression we derive its asymptotic scaling and a uniform (in n) upper bound. In the dilute regime, we recover the known linear growth $P_2(m, n) = \Theta(n)$. By contrast, in the regime corresponding to $m = cn^\beta$ with $c \geq 1$ and $1 \leq \beta < 2$, we obtain a different scaling, i.e., $P_2(m, n) = m/n + 1 + o(1)$. Moreover, our uniform upper bound, namely $P_2(m, n) \leq 2(m/n + 1)$ implies via Paley-Zygmund's inequality that standard anti-concentration holds in the *linear* regime $m = cn$, i.e.

$$\Pr_{S,U} \left[p_U(S) \geq \frac{\alpha}{|\Phi_m^n|} \right] \geq \frac{(1-\alpha)^2}{2(c+1)}. \quad (2)$$

The above bound confirms the anti-concentration conjecture of [24] in the experimentally relevant saturated regime and complements recent hardness results for boson sampling beyond the dilute limit. Beyond hardness, the same quantity $P_2(m, n)$ provides (1) the ideal reference value for linear cross-entropy benchmarking [13], linking our analysis to practical validation protocols for photonic experiments (2) an exponential lower-bound on the sample complexity of ε -certification of Boson Sampling [25], showing Boson Sampling does not admit efficient certification procedures, (3) an improved sample complexity upper bound for bosonic filtered randomized benchmarking [26] and (4) rules out classical simulation of Boson Sampling based on the sparsity of the distribution, as this condition is incompatible with anti-concentration [27].

Conclusion and discussion In this work, we develop a representation-theoretic framework for computing second moments of expectation values of bosonic observables under passive linear-optical evolutions. With these tools, we provide a thorough analysis of the exponential-concentration in quantum machine learning models, revealing a regime where efficient trainability and absence of efficient classical simulation methods may coexist. Moreover, by deriving an analytical expression for the normalized average outcome collision probability, we established anti-concentration in Boson Sampling. This provides analytical evidence supporting hardness conjectures in experimentally relevant settings.

Several natural open directions follow from this work. A first question is whether similar representation-theoretic techniques can be extended to higher moments, which would provide access to finer statistical features of boson sampling distributions and further strengthen complexity-theoretic arguments. Another important direction is to relax the fixed particle-number assumption on both observables and input states, requiring to handle representations with multiplicities, and extending the framework to incorporate realistic experimental imperfections, such as photon loss and partial distinguishability. Regarding quantum machine models, our findings on the one hand open the door towards designing efficiently trainable models that escape all known classical simulation methods. This raises the questions of finding broader classes of observables with these properties, and assessing their usefulness in practical scenarios. On the other hand, our findings raise the question of whether new classical simulation methods can be cooked up to efficiently simulate observables in these barren-plateau-free regimes.

-
- [1] H. Mhiri, H. Thomas, L. Monbroussou, U. Chabaud, Z. Holmes, and E. Kashefi, [Boson sampling beyond the dilute regime: second moments and anti-concentration](#) (2026), version Number: 1.
 - [2] L. Monbroussou, H. Thomas, H. Mhiri, U. Chabaud, Z. Holmes, and E. Kashefi, [Classical simulation and model concentration in passive linear optics](#) (2026).
 - [3] N. Killoran, T. R. Bromley, J. M. Arrazola, M. Schuld, N. Quesada, and S. Lloyd, [Physical Review Research](#) **1**, 033063 (2019).
 - [4] T. R. Bromley, J. M. Arrazola, S. Jahangiri, J. Izaac, N. Quesada, A. D. Gran, M. Schuld, J. Swinarton, Z. Zabaneh, and N. Killoran, [Quantum Science and Technology](#) **5**, 034010 (2020).
 - [5] B. Y. Gan, D. Leykam, and D. G. Angelakis, [EPJ Quantum Technology](#) **9**, 16 (2022).
 - [6] G. R. Steinbrecher, J. P. Olson, D. Englund, and J. Carolan, [npj Quantum Information](#) **5**, 60 (2019).
 - [7] A. Kurkin, U. Chabaud, Z. Kolarovszki, B. Bakó, Z. Zimborás, and V. Dunjko, [Universality of Classically Trainable, Quantum-Deployed Boson-Sampling Generative Models](#) (2026), arXiv:2603.11014 [quant-ph].
 - [8] Z. Kolarovszki, B. Bakó, M. Oszmaniec, C. Oh, and Z. Zimborás, [Generative modeling with Gaussian Boson Sampling: classically trainable Bosonic Born Machines](#) (2026), arXiv:2603.11195 [quant-ph].
 - [9] F. Gottlieb, R. Mezher, B. Ventura, S. Mansfield, and A. Salavrakos, [Efficient training of photonic quantum generative models](#) (2026), arXiv:2603.08793 [quant-ph].
 - [10] M. J. Bremner, A. Montanaro, and D. J. Shepherd, [Physical Review Letters](#) **117**, 080501 (2016).
 - [11] M. Oszmaniec, N. Dangniam, M. E. Morales, and Z. Zimborás, [PRX Quantum](#) **3**, 020328 (2022).
 - [12] D. Hangleiter, J. Bermejo-Vega, M. Schwarz, and J. Eisert, [Quantum](#) **2**, 65 (2018).
 - [13] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, [Nature Physics](#) **14**, 595 (2018).
 - [14] A. Bouland, J. F. Fitzsimons, and D. E. Koh, [LIPICs, Volume 102, CCC 2018](#) **102**, 21:1 (2018), arXiv:1709.01805.
 - [15] S. Aaronson and A. Arkhipov, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (ACM, San Jose California USA, 2011) pp. 333–342.
 - [16] S. Nezami, [Permanent of random matrices from representation theory: moments, numerics, concentration, and comments on hardness of boson-sampling](#) (2021).
 - [17] A. Bouland, I. Datta, B. Fefferman, and F. Hernandez, [Exponential improvements to the average-case hardness of bosonsampling](#) (2025).
 - [18] J. Preskill, [Quantum](#) **2**, 79 (2018).
 - [19] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, [Nature Reviews Physics](#) **3**, 625 (2021).
 - [20] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, [Nature](#) **549**, 195 (2017).
 - [21] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, Z. Holmes, and M. Cerezo, [Nature Reviews Physics](#) **7**, 174 (2025).
 - [22] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, [Nature Communications](#) **9**, 4812 (2018).
 - [23] A. Arrasmith, Z. Holmes, M. Cerezo, and P. J. Coles, [Quantum Science and Technology](#) **7**, 045015 (2022).
 - [24] A. Bouland, D. Brod, I. Datta, B. Fefferman, D. Grier, F. Hernandez, and M. Oszmaniec, [Complexity-theoretic foundations of BosonSampling with a linear number of modes](#) (2023).
 - [25] D. Hangleiter and J. Eisert, [Reviews of Modern Physics](#) **95**, 035001 (2023).
 - [26] M. Arienzo, D. Grinko, M. Kliesch, and M. Heinrich, [PRX Quantum](#) **6**, 020305 (2025).
 - [27] H. Pashayan, S. D. Bartlett, and D. Gross, [Quantum](#) **4**, 223 (2020).