Spoofing Cross-Entropy Measure in Boson Sampling

Changhun Oh,[1] Liang Jiang,[1] and Bill Fefferman[2]
[1]*Pritzker School of Molecular Engineering, University of Chicago, Chicago, Illinois 60637, USA*
[2]*Department of Computer Science, University of Chicago, Chicago, Illinois 60637, USA*

(Received 28 January 2023; accepted 9 June 2023; published 6 July 2023)

Cross-entropy (XE) measure is a widely used benchmark to demonstrate quantum computational advantage from sampling problems, such as random circuit sampling using superconducting qubits and boson sampling (BS). We present a heuristic classical algorithm that attains a better XE than the current BS experiments in a verifiable regime and is likely to attain a better XE score than the near-future BS experiments in a reasonable running time. The key idea behind the algorithm is that there exist distributions that correlate with the ideal BS probability distribution and that can be efficiently computed. The correlation and the computability of the distribution enable us to postselect heavy outcomes of the ideal probability distribution without computing the ideal probability, which essentially leads to a large XE. Our method scores a better XE than the recent Gaussian BS experiments when implemented at intermediate, verifiable system sizes. Much like current state-of-the-art experiments, we cannot verify that our spoofer works for quantum-advantage-size systems. However, we demonstrate that our approach works for much larger system sizes in fermion sampling, where we can efficiently compute output probabilities. Finally, we provide analytic evidence that the classical algorithm is likely to spoof noisy BS efficiently.

DOI: 10.1103/PhysRevLett.131.010401

Quantum computers are believed to efficiently solve problems that classical counterparts cannot, such as integer factoring [1] and quantum simulation [2]. Whereas scalability and fault tolerance are required to solve hard and practical problems, currently available devices are of noisy intermediate-scale quantum (NISQ) type. Hence, much attention has been paid to demonstrating quantum advantage by exploiting NISQ devices. Sampling problems are particularly promising for the demonstration thanks to rigorous evidence that classical computers cannot efficiently solve them [3–6]. Indeed, we have recently seen the first claims of quantum advantage using random circuit sampling (RCS) with superconducting qubits [7–9] and Gaussian boson sampling (GBS) [10–13]. However, the apparent limitation of current experiments is that they are not scalable because uncorrected noise decays quantum signals as the system size grows [14–29]. Thus, finding appropriate-size experiments is crucial, where the system size is sufficiently large so that classical computers cannot efficiently simulate them, but not too large for noise to annihilate quantum signals.

Finding an appropriate regime is also important to enable classical verification because verification techniques such as cross-entropy benchmarking (XEB) cost exponential time [30]. The state-of-the-art method is XEB, which is sample efficient for RCS and GBS [31]; scoring a large cross entropy (XE) is considered evidence of quantum advantage. The premise behind the benchmarking is that appropriate-size experiments maintain sufficient quantum signals so that the experimental score cannot be attained by classical devices within reasonable costs due to the remaining quantum signals. For RCS [7,8], there have been extensive studies to support the premise [7,30,32,33]. While we have some limited understanding of what XE measures in RCS [7,30,32–35], with many interesting debates of XEB [34,35], our understanding of XE in boson sampling (BS) is much more limited than RCS. Crucially, there is no theoretical evidence that attaining a high XE is classically hard, to the best of our knowledge. Despite this, the XEB has been used in many state-of-the-art quantum advantage experiments [10–13]. XEB is particularly attractive since the benchmark does not depend on an adversarial mock-up distribution. This "adversary independence" feature allows us to use XEB to benchmark noisy experiments in situations in which we are ignorant of the best possible classical spoofing algorithm, as is the case with GBS experiments [36].

In this Letter, we provide a heuristic classical algorithm that scores better than the current intermediate-scale GBS (in a verifiable regime) and is likely to score better than the near-future BS experiments for XEB. First, we numerically demonstrate, using a small-scale GBS, that our algorithm selectively generates heavy outcomes of BS and scores better than the ideal distribution. For larger systems in a quantum advantage regime, due to the inefficiency of estimating XE and a large computational cost, the frequently used method is to analyze an intermediate-size experiment in a classically verifiable regime, instead of the largest experiment. Following this, we demonstrate that the XE of the proposed sampler achieves a significantly larger

0031-9007/23/131(1)/010401(7)

XE score than the intermediate-size GBS of the most recent experiments in Refs. [11,12]. To predict its behavior for large system sizes, we analyze its performance for fermion sampling (FS), which is efficiently simulable and verifiable, and provide analytical evidence of efficient spoofing of noisy BS. Therefore, our classical algorithm is expected to score better than the near-future BS experiments in a reasonable time. We finally discuss other existing spoofing methods and benchmarking.

XE.—The (log) XE of $q_U(x)$ with respect to the ideal probability $p_U(x)$ is defined and estimated as [7,30]

$$\mathrm{XE} \equiv \sum_x q_U(x) \log p_U(x) \approx \frac{1}{N_s} \sum_{i=1}^{N_s} \log p_U(x^{(i)}), \qquad (1)$$

where N_s is the number of samples from $q_U(x)$, $q_U(x)$ is an experimental probability distribution or a mock-up probability distribution, and the sum is taken over all measurement outcomes x. Since there are exponentially many different outcomes for BS, the XE of an experimental or a mock-up distribution with respect to an ideal probability distribution cannot be efficiently computed; thus it is estimated by sampling $\{x^{(i)}\}_{i=1}^{N_s}$ from $q_U(x)$ in practice (even estimation is inefficient because the cost for computing probabilities is exponential in the system size.). Unlike RCS, due to the lack of understanding of the ideal score of BS, the XE was used as a relative quantity against mock-up distributions [12]. Specifically, if an experimental XE is larger than the mock-up distribution's, it implies that the former generates heavier outcomes, considered as evidence of quantum advantage for BS. Such a method was used as evidence of quantum computational advantage in the recent GBS experiment [12].

Heavy outcome generation.—Now, we present a classical algorithm attaining a large XE. For BS, there are probability distributions that can be efficiently sampled and correlated with the ideal probability distribution. For example, fully distinguishable BS has nonzero correlation [37]. Another example is a probability distribution having the same low-order marginal distributions, which was used to spoof a GBS experiment [25]. Nevertheless, the correlations are too small to obtain a larger XE than the experiments [12]. We now provide how to increase the XE from the small correlation. Note that our scheme is not limited to BS in principle.

Let us denote a probability distribution for sampling as $q_U(x)$ and a distribution for postselection as $h_U(x)$. We now present the algorithm, illustrated in Fig. 1. Step 1: Choose an efficient classical sampler $q_U(x)$. Step 2: Generate $k \times N_s$ samples from $q_U(x)$. Step 3: Compute $h_U(x)$ for each sample. In this procedure, we require $h_U(x)$ *efficiently computable* and *correlated* to $p_U(x)$. Step 4: Postselect and output N_s samples whose $h_U(x)$'s are largest out of $k \times N_s$ samples [38].

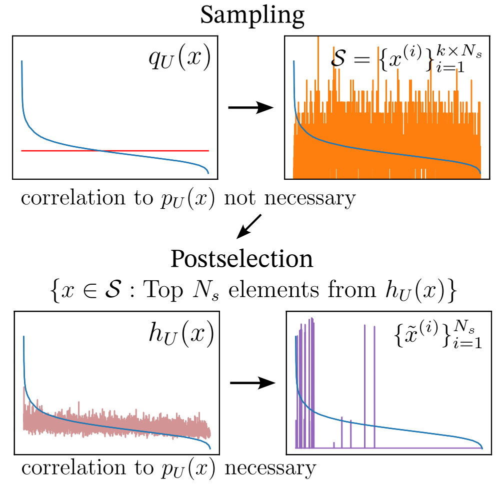

FIG. 1. Proposed algorithm for spoofing XEB. The blue curve represents the ideal distribution $p_U(x)$. See the main text for the detailed procedures.

The principle behind the postselection is that, due to the correlation between $p_U(x)$ and $h_U(x)$, we obtain samples likely to have a larger probability with respect to the ideal distribution $p_U(x)$ by selecting the samples with larger $h_U(x)$. Thus, the most crucial step is to find $h_U(x)$ correlated with the ideal distribution, but easy to compute. Together with our numerical result below, the existence of such indicators might be related to hardness of heavy outcome generation in that the indicators enable us to generate heavy outcomes, where the heaviness means that its probability is larger than the median of probabilities [39]. Note that we do *not* compute the ideal probability in the entire procedure for sampling. We provide more discussions about the choice of $q_U(x)$ and $h_U(x)$ in the Supplemental Material [40].

Spoofing XEB in GBS experiments.—To illustrate how the spoofing procedure operates in practice, we analyze the method for a small-size GBS circuit with the number of modes $M = 16$ and the number of photons $N = 4$ [12]. Here, we choose uniform distribution $q_U(x)$ and first-order-marginal-based distribution $h_U(x) \equiv \prod_{i=1}^{M} p_U(x_i)$, where $p_U(x_i)$ is the ideal marginal probability of the ith mode. Here, the marginal probabilities can be easily computed since the reduced state is a single-mode Gaussian state whose covariance matrix is a submatrix of the full covariance matrix [46]. Because the chosen $h_U(x)$ perfectly recovers the first-order marginals of the ideal distribution $p_U(x)$ by definition, it correlates with the ideal distribution. After computing all the ideal probabilities $p_U(x)$, we sort the outcomes in descending order. We also compute $q_U(x)$ and $h_U(x)$ and sort them in the same order as the ideal case. Although we use log XE for comparison with experiments in Refs. [11,12], a similar result using linear XE is provided in the Supplemental Material [40].

As shown in Figs. 2(a) and 2(b), the ideal probability $p_U(x)$ is likely to be large if $h_U(x)$ is large, which clearly

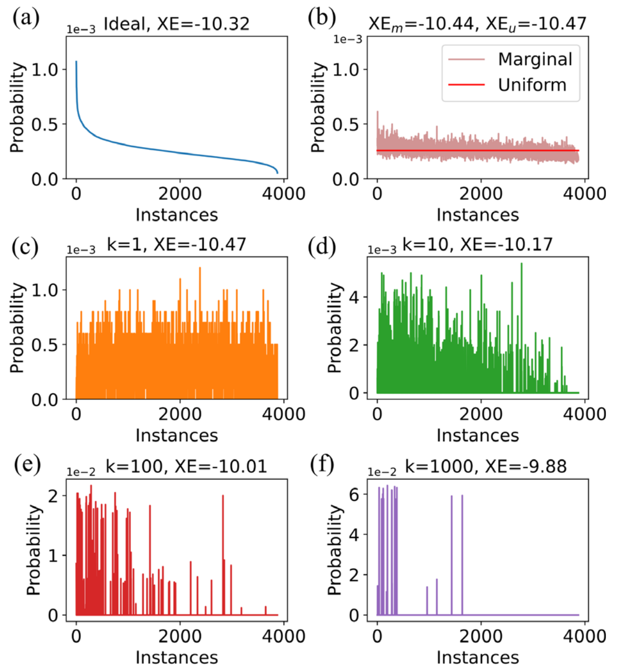

FIG. 2. (a) Ideal distribution $p_U(x)$ and (b) the distribution $h_U(x)$ are obtained by computing all probabilities of N photon outcomes, where $h_U(x)$ and $q_U(x)$ are sorted in the same order of $p_U(x)$. Here, XE_m and XE_u represent the XE of $h_U(x)$ and $q_U(x)$, respectively. (c)–(f) The spoofing with different k with $N_s = 10^4$.

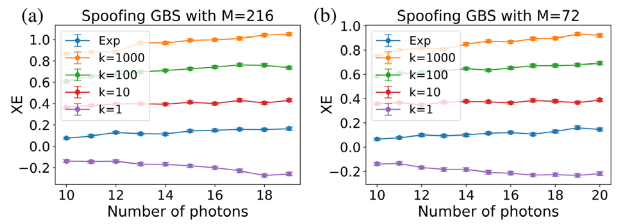

FIG. 3. (a) 216 modes with 216 squeezed states with $r \approx 0.89$. (b) 72 modes with 72 squeezed states with $r \approx 0.53$. The error bar represents the standard deviation of $\log p_U(x)$ of obtained samples divided by $\sqrt{N_s}$, and $N_s = 10^3$.

shows their correlation, while each XE of $q_U(x)$ and $h_U(x)$ is smaller than the ideal distribution. However, as the postselection rate k increases, the samples are concentrated on large probabilities of the ideal distribution, which clearly shows that our procedure selectively generates heavy outcomes. Especially for $k = 10^2$ and 10^3, the samples are highly concentrated on the heavy outcomes. Consequently, the spoofer's XE is larger than the ideal probability. Therefore, one can attain a large XE by sampling heavy outcomes without directly simulating the desired circuit. Here, the XE might not monotonically increase as the postselection rate because one might end up selecting less heavy outcomes. We also analyze the same procedure for Fock-state BS (FBS) using different h_U in Ref. [40].

Now, we consider intermediate-size circuits used for demonstrating quantum advantage in Ref. [12], where the experimental samples attain a larger XE than various mock-up distributions. Since XE is not an efficient verification method, the result based on the intermediate-scale circuits is claimed to be evidence of the quantum advantage of their largest circuit. Following Ref. [12], we normalize as $\mathrm{XE} = \sum_{i=1}^{N_s} \log[p_U(x^{(i)})/\mathcal{N}]/N_s$, where $\mathcal{N} \equiv \Pr(N)/\binom{N+M-1}{N}$ and $\Pr(N)$ is the probability of obtaining N photons from $p_U(x)$.

We emphasize that the XEB we employ is applied for each photon number sector instead of the entire sample set, consistent with the previous methods in experiments [10–12]. Therefore, it is not sensitive to the total photon number distribution. If we conduct the benchmarking for the entire set, the trivial way can spoof it is by manipulating the total photon number distribution of a mock-up sampler so that the sampler generates samples from a particular sector, having a larger probability of each outcome than other sectors on average. Meanwhile, we show in the Supplemental Material [40] that we can adjust our spoofer's total number distribution to be consistent with the ideal case.

Figure 3 again shows that without postselection, the XE is much smaller than the experimental samples' XE from Ref. [12]. Meanwhile, for $k = 10$, the spoofer's XE is much larger than the experiment for the wide range of photon numbers. Additionally, as the postselection rate increases further, the difference becomes more significant. Based on the trend, we expect that such a gap may not close for the large photon number sectors. Similar results for different experiments [11] are provided in Ref. [40].

Prediction for larger systems.—Since XEB is computationally inefficient, it is demanding to check if the spoofer works even for larger systems. To predict its performance for large-size circuits and how the postselection overhead scales, we consider a particular type of FS [47], a variant of FBS (see Ref. [37] and the Supplemental Material [40]). Since the FS is easy to simulate and its probabilities are easy to compute [37], it does not provide quantum advantage. Nevertheless, its similar structure to BS may enable us to predict a larger BS, since the key idea of our algorithm is to postselect heavy outcomes without computing the probability of the ideal distribution; thus it does not directly rely on the easiness of FS.

Figure 4 shows the result. We implement the same spoofer for FS with larger system sizes with postselection rate k differently to analyze the postselection overhead and investigate the XE difference: $\Delta\mathrm{XE} \equiv \mathrm{XE}_{\mathrm{spoofer}} - \mathrm{XE}_{\mathrm{id}}$. Although the XE from different choices of the postselection rate varies, all the different cases' XE difference gradually become negative as the system size grows. Such a trend may be caused because the first-order marginals' correlation to the ideal distribution is not sufficiently large for larger system sizes. It implies that, as the system size increases, the postselection rate might have to increase, superpolynomially at worst. Even if this is true, since we are comparing the XE score against the *ideal* case and noise significantly decreases the XE of experiments (see below),

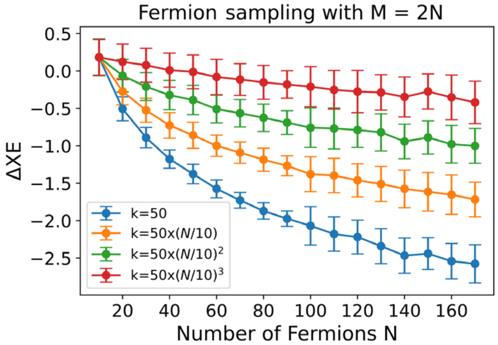

FIG. 4. XE of FS. $N_s = 10^3$ for $N \leq 90$ and $N_s = 10^2$ for $N \geq 10^2$. We used 10^3 different random unitary matrices for circuit configuration. The error bar represents the standard deviation of XE difference.

it would be extremely difficult to surpass the score from our spoofer in the near future.

Evidence of efficiency of spoofing noisy BS.—We now provide evidence that our algorithm can efficiently spoof *noisy* BS experiments. To this end, we study the *linear* XE, $\mathrm{XE} \equiv \sum_x q_U(x) p_U(x)$, of ideal and noisy FBS, assuming a collision-free regime, i.e., $M = \omega(N^5)$. For the ideal case, i.e., $q_U(x) = p_U(x) = |\mathrm{Per} U_x|^2$ [3], the average XE for Haar-random unitary U is given by

$$\mathbb{E}_U[\mathrm{XE}_{\mathrm{id}}] = \sum_x \mathbb{E}_U[|\mathrm{Per} U_x|^4] \approx \frac{\sum_x \mathbb{E}_Z[|\mathrm{Per} Z|^4]}{M^{2N}} \approx \frac{(N+1)!}{M^N}. \quad (2)$$

Here, we approximated submatrices of large Haar-random unitary matrices by random Gaussian matrices, $U_x \approx Z/\sqrt{M}$ with Z following the complex normal distribution, and used $\mathbb{E}_Z[|\mathrm{Per} Z|^4] = N!(N+1)!$ [3] and $\binom{M}{N} \approx M^N/N!$. Meanwhile, the XE of q_U independent of p_U is given by $\mathbb{E}_U[\mathrm{XE}_{\mathrm{idp}}] = N!/M^N$ [40], where one can clearly see the additional factor $(N+1)$ for the ideal case.

We now analyze the behavior of the XE under partial distinguishability [40,49,50], one of the most important noise models in practice, describing partial overlaps of wave functions of different photons (see the Supplemental Material [40] and Refs. [49,50] for more details of the model). The XE of BS under partial distinguishability $0 \leq \rho < 1$ with respect to the ideal distribution, normalized by the XE of an independent distribution, is bounded by [40]

$$\frac{\mathbb{E}_U[\mathrm{XE}_{\mathrm{pd}}]}{\mathbb{E}_U[\mathrm{XE}_{\mathrm{idp}}]} \leq \frac{e^2(1-\rho^{N+1})}{1-\rho}. \quad (3)$$

Observe that even for constant ρ, the normalized XE converges to $e^2/(1-\rho)$, which only provides an additional constant factor $e^2/(1-\rho)$, while the ideal score's factor increases as $(N+1)$. Clearly, the XE significantly decreases under experimental noise, suggesting that the experimental XE will be much smaller than the ideal case.

We now provide analytical evidence that the noisy XE might be attainable using our method. Since analyzing our algorithm's output probability distribution is difficult due to postselection, we consider a different probability distribution $\propto [h_U(x)]^s$, which contains our algorithm's core idea, as it tends to generate heavy outcomes from $p_U(x)$. Effectively, a similar effect to postselection is expected for $s > 1$, because the resulting probability favors large-probability outcomes from $h_U(x)$ (we do not expect that we can sample from the distribution). Thus, the power s is associated with postselection overhead k. We show for linear XE that the distribution with a multinomial distribution $h_U(x) = N^{-N} \prod_{i=1}^N \sum_{j=1}^N |U_{j,x_i}|^2$ provides [40]

$$\frac{\mathbb{E}_U[\mathrm{XE}(p_U, h_U^s)]}{\mathbb{E}_U[\mathrm{XE}_{\mathrm{idp}}]} \equiv \frac{1}{\mathbb{E}_U[\mathrm{XE}_{\mathrm{idp}}]} \mathbb{E}_U\left[\frac{\sum_x p_U(x) h_U^s(x)}{\sum_x h_U^s(x)}\right] \approx e^s, \quad (4)$$

where the approximation holds for small s and large N. Since it suffices to achieve a constant factor due to noise from Eq. (3) and there are other types of noise, such as loss [40], we expect that choosing a constant power s, which may be interpreted as constant postselection overhead k, might be sufficient to spoof noisy BS unless the noise can be highly suppressed.

Comparison to existing spoofing methods.—Most existing methods for spoofing GBS work by attempting to model noisy experiments. It is often the case that noisy experiments converge to classically easy distribution [15,18,23–25,50,51]. These algorithms exploit this by sampling from this easy distribution. For example, Kalai and Kindler claimed that noisy experimental boson sampler's probabilities are approximable by low-degree polynomials [15]. Since then, there have been many subsequent proposals to take advantage of noise [18,50]. Similarly, the algorithm in Ref. [25] aims to reproduce low-order marginals without recovering high-order marginals. Another method is to use the classical state because noise often transforms the output state to be close to classical states [20,23,51,52].

Unlike the existing methods, our method's goal is spoofing XEB instead of approximate simulation. Since XE increases by generating heavy outcomes, a large XE can be achieved without simulation. Thus, our approach does not necessarily work for other benchmarking that does not rely on heavy outcome generation, such as the Bayesian test and correlation functions [53]. The Bayesian test employed in current GBS experiments as another evidence of quantum advantage [10–12] defines the score as

$$\mathrm{score} = \frac{1}{N_s} \sum_{i=1}^{N_s} \log \frac{p_U(x^{(i)})}{q_U(x^{(i)})}, \quad (5)$$

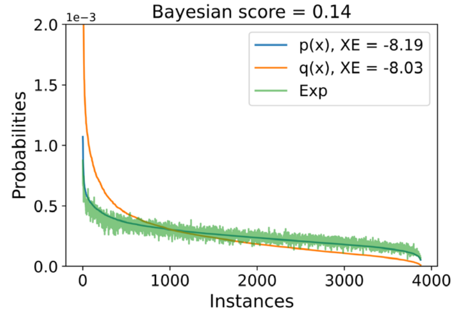

FIG. 5. Bayesian test using $M = 16$ and $N = 4$ case in Ref. [12]. The green curve is the experimental data from Ref. [12]. The XE score is different from Fig. 2 because the probability is normalized over the sector for this figure.

where $\{x^{(i)}\}_{i=1}^{N_s}$ is a sample set from *experiment*. A positive score implies that the experimental samples are more likely to be sampled from the ideal distribution $p_U(x)$ than the mock-up distribution $q_U(x)$. Because the test requires the profile of the mock-up distribution and computing the postselected distribution's probability is difficult, it is not easy to perform Bayesian testing against our method. Nevertheless, because our method favorably generates heavy outcomes, the output distribution may be highly concentrated. Thus, even if we can compute the probability distribution, our method is unlikely to pass the Bayesian test.

As an example that generates heavy outcomes but fails to pass the Bayesian test, consider a mock-up distribution $q_U(x) \propto p_U(x)^2$ (without postselection). Thus, $q_U(x)$ generates heavier outcomes than $p_U(x)$ with respect to $p_U(x)$, which guarantees to spoof the XE test, as shown in Fig. 5. However, since the mock-up distribution is farther than the ideal distribution to empirical experimental distribution, the Bayesian score becomes positive, implying that $q_U(x)$ fails to pass the Bayesian test, although it generates heavy outcomes.

Finally, our results spoofing the XE test do not imply that the experiments in Refs. [10–12] are easy to simulate because our algorithm's goal is to spoof the test without simulation. Instead, the implication is that the XE scores may not be a proper measure as evidence of quantum advantage of BS. Our results open many questions about the verification of sampling tasks. First, making our method analytical is crucial to predict the asymptotic behavior of the method precisely. Second, it would also be interesting to apply the same method to other sampling tasks. Especially for RCS, finding a distribution that is correlated with the ideal distribution and easy to compute would be an important first step to applying the presented method. We emphasize again that although we chose h_U relying on marginals, other various possible quantities may not necessarily rely on marginals. Also, since our algorithm is specialized to spoof the XE test, we do not expect it to pass other tests, such as the Bayesian test. Hence, it would also be interesting to analyze the Bayesian test to see if it can be spoofed.

We thank Benjamin Villalonga and Sergio Boixo for interesting and fruitful discussions. We acknowledge support from the ARO MURI (W911NF-16-1-0349, W911NF-21-1-0325), AFOSR MURI (FA9550-19-1-0399, FA9550-21-1-0209), DOE Q-NEXT, NSF (OMA-1936118, ERC-1941583, OMA-2137642), NTT Research, and the Packard Foundation (2020-71479). B. F. acknowledges support from AFOSR (FA9550-21-1-0008). This material is based upon work partially supported by the National Science Foundation under Grant No. CCF-2044923 (CAREER) and by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers (Q-NEXT). This material is also supported in part by the DOE QuantISED Award No. DE-SC0020360. The authors are also grateful for the support of the University of Chicago Research Computing Center for assistance with the numerical simulations carried out in this Letter.

[1] P. W. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (IEEE, New York, 1994), pp. 124–134.
[2] S. Lloyd, Universal quantum simulators, Science **273**, 1073 (1996).
[3] S. Aaronson and A. Arkhipov, The computational complexity of linear optics, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (2011), pp. 333–342.
[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] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, On the complexity and verification of quantum random circuit sampling, Nat. Phys. **15**, 159 (2019).
[6] A. Deshpande, A. Mehta, T. Vincent, N. Quesada, M. Hinsche, M. Ioannou, L. Madsen, J. Lavoie, H. Qi, J. Eisert *et al.*, Quantum computational advantage via high-dimensional Gaussian boson sampling, Sci. Adv. **8**, eabi7894 (2022).
[7] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, Nature (London) **574**, 505 (2019).
[8] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan *et al.*, Strong Quantum Computational Advantage Using a Superconducting Quantum Processor, Phys. Rev. Lett. **127**, 180501 (2021).
[9] A. Morvan, B. Villalonga, X. Mi, S. Mandrà, A. Bengtsson, P. Klimov, Z. Chen, S. Hong, C. Erickson, I. Drozdov *et al.*, Phase transition in random circuit sampling, arXiv: 2304.11119.
[10] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu *et al.*, Quantum computational advantage using photons, Science **370**, 1460 (2020).

[11] H.-S. Zhong, Y.-H. Deng, J. Qin, H. Wang, M.-C. Chen, L.-C. Peng, Y.-H. Luo, D. Wu, S.-Q. Gong, H. Su *et al.*, Phase-Programmable Gaussian Boson Sampling Using Stimulated Squeezed Light, Phys. Rev. Lett. **127**, 180502 (2021).

[12] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins *et al.*, Quantum computational advantage with a programmable photonic processor, Nature (London) **606**, 75 (2022).

[13] Y.-H. Deng, Y.-C. Gu, H.-L. Liu, S.-Q. Gong, H. Su, Z.-J. Zhang, H.-Y. Tang, M.-H. Jia, J.-M. Xu, M.-C. Chen *et al.*, Gaussian boson sampling with pseudo-photon-number resolving detectors and quantum computational advantage, arXiv:2304.12240.

[14] D. Aharonov, M. Ben-Or, R. Impagliazzo, and N. Nisan, Limitations of noisy reversible computation, arXiv:quant-ph/9611028.

[15] G. Kalai and G. Kindler, Gaussian noise sensitivity and boson sampling, arXiv:1409.3093.

[16] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Achieving quantum supremacy with sparse and noisy commuting quantum computations, Quantum **1**, 8 (2017).

[17] X. Gao and L. Duan, Efficient classical simulation of noisy quantum computation, arXiv:1810.03176.

[18] J. Renema, V. Shchesnovich, and R. Garcia-Patron, Classical simulability of noisy boson sampling, arXiv:1809.01953.

[19] V. S. Shchesnovich, Noise in boson sampling and the threshold of efficient classical simulatability, Phys. Rev. A **100**, 012340 (2019).

[20] R. García-Patrón, J. J. Renema, and V. Shchesnovich, Simulating boson sampling in lossy architectures, Quantum **3**, 169 (2019).

[21] K. Noh, L. Jiang, and B. Fefferman, Efficient classical simulation of noisy random quantum circuits in one dimension, Quantum **4**, 318 (2020).

[22] Y. Takahashi, Y. Takeuchi, and S. Tani, Classically simulating quantum circuits with local depolarizing noise, Theor. Comput. Sci. **893**, 117 (2021).

[23] H. Qi, D. J. Brod, N. Quesada, and R. García-Patrón, Regimes of Classical Simulability for Noisy Gaussian Boson Sampling, Phys. Rev. Lett. **124**, 100502 (2020).

[24] C. Oh, K. Noh, B. Fefferman, and L. Jiang, Classical simulation of lossy boson sampling using matrix product operators, Phys. Rev. A **104**, 022407 (2021).

[25] B. Villalonga, M. Y. Niu, L. Li, H. Neven, J. C. Platt, V. N. Smelyanskiy, and S. Boixo, Efficient approximation of experimental Gaussian boson sampling, arXiv:2109.11525.

[26] D. Aharonov, X. Gao, Z. Landau, Y. Liu, and U. Vazirani, A polynomial-time classical algorithm for noisy random circuit sampling, *STOC 2023: Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (2023), pp. 945–957, 10.1145/3564246.3585234.

[27] C. Oh, L. Jiang, and B. Fefferman, On classical simulation algorithms for noisy boson sampling, arXiv:2301.11532.

[28] M. Liu, C. Oh, J. Liu, L. Jiang, and Y. Alexeev, Complexity of Gaussian boson sampling with tensor networks, arXiv:2301.12814.

[29] C. Oh, M. Liu, Y. Alexxev, B. Fefferman, and L. Jiang, Tensor network algorithm for simulating experimental Gaussian boson sampling, arXiv:2306.03709.

[30] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, Nat. Phys. **14**, 595 (2018).

[31] XEB is sample efficient in the sense that it does not require an exponentially large number of samples, while it is computationally inefficient because computing the ideal output probability requires exponential time.

[32] S. Aaronson and L. Chen, Complexity-theoretic foundations of quantum supremacy experiments, arXiv:1612.05903.

[33] S. Aaronson and S. Gunn, On the classical hardness of spoofing linear cross-entropy benchmarking, arXiv: 1910.12085.

[34] X. Gao, M. Kalinowski, C.-N. Chou, M. D. Lukin, B. Barak, and S. Choi, Limitations of linear cross-entropy as a measure for quantum advantage, arXiv:2112.01657.

[35] F. Pan, K. Chen, and P. Zhang, Solving the Sampling Problem of the Sycamore Quantum Supremacy Circuits, Phys. Rev. Lett. **129**, 090502 (2022).

[36] Other benchmarking methods employed [10–13], such as the Bayesian test and correlator methods, also do not have rigorous evidence for hardness. Also, the Bayesian test is genuinely implemented against adversarial samplers; thus, it cannot rule out all possible classical algorithms. We emphasize that the current GBS experiments employ XEB against certain adversarial samplers because we do not know the ideal score.

[37] S. Aaronson and A. Arkhipov, Boson sampling is far from uniform, arXiv:1309.7460.

[38] Another possible way is to generate k samples and post-select a single sample out of k samples and iterate N_s times, instead of generating $k \times N_s$ and postselecting N_s samples at once. While we do not expect a significant difference, we employ the latter in this Letter consistently.

[39] We note that generating outcomes whose probabilities are larger than the median is sufficient to spoof the XEB because even running a quantum device can only generate heavy outcomes in this sense slightly better than the light outcomes [33].

[40] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.131.010401 for more discussion about the spoofing algorithm and more numerical results for different experiments and linear cross-entropy measure, which includes Refs. [41–45].

[41] C. Oh, Y. Lim, B. Fefferman, and L. Jiang, Classical Simulation of Boson Sampling Based on Graph Structure, Phys. Rev. Lett. **128**, 190501 (2022).

[42] A. Popova and A. Rubtsov, Cracking the quantum advantage threshold for Gaussian boson sampling, in *Quantum 2.0* (Optica Publishing Group, 2022), pp. QW2A–15.

[43] P. D. Drummond, B. Opanchuk, A. Dellios, and M. D. Reid, Simulating complex networks in phase space: Gaussian boson sampling, Phys. Rev. A **105**, 012427 (2022).

[44] D. A. Ivanov and L. Gurvits, Complexity of full counting statistics of free quantum particles in product states, Phys. Rev. A **101**, 012303 (2020).

[45] M. C. Tichy, K. Mayer, A. Buchleitner, and K. Mølmer, Stringentand Efficient Assessment of Boson-Sampling Devices, Phys. Rev. Lett. **113**, 020502 (2014).

[46] A. Serafini, *Quantum Continuous Variables: A Primer of Theoretical Methods* (CRC Press, Boca Raton, 2017).

[47] While the fermionic sampling we consider is easy, there exists different types of fermionic sampling that is proven to be hard under plausible assumptions [48].

[48] M. Oszmaniec, N. Dangniam, M. E. S. 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).

[49] M. C. Tichy, Sampling of partially distinguishable bosons and the relation to the multidimensional permanent, Phys. Rev. A **91**, 022316 (2015).

[50] J. J. Renema, A. Menssen, W. R. Clements, G. Triginer, W. S. Kolthammer, and I. A. Walmsley, Efficient Classical Algorithm for Boson Sampling with Partially Distinguishable Photons, Phys. Rev. Lett. **120**, 220502 (2018).

[51] J. Martńez-Cifuentes, K. Fonseca-Romero, and N. Quesada, Classical models are a better explanation of the Jiuzhang Gaussian boson samplers than their targeted squeezed light models, arXiv:2207.10058.

[52] M. Oszmaniec and D. J. Brod, Classical simulation of photonic linear optics with lost particles, New J. Phys. **20**, 092002 (2018).

[53] D. S. Phillips, M. Walschaers, J. J. Renema, I. A. Walmsley, N. Treps, and J. Sperling, Benchmarking of Gaussian boson sampling using two-point correlators, Phys. Rev. A **99**, 023836 (2019).