

Provable and Verifiable Quantum Advantage in Sample Complexity

Marcello Benedetti^{1,*}, Harry Buhrman^{1,2,3,†} and Jordi Weggemans^{2,4,‡}¹*Quantinuum, London, United Kingdom*²*QuSoft, Amsterdam, The Netherlands*³*University of Amsterdam, Amsterdam, The Netherlands*⁴*CWI, Amsterdam, The Netherlands* (Received 11 August 2025; revised 20 October 2025; accepted 9 December 2025; published 27 January 2026)

Consider a fixed universe of $N = 2^n$ elements and the uniform distribution over elements of some subset of size K . Given samples from this distribution, the task of complement sampling is to provide a sample from the complementary subset. We give a simple quantum algorithm that uses only a single quantum sample—a single copy of the uniform superposition over elements of the subset. When $K = N/2$, we show that the quantum algorithm succeeds with probability 1, whereas any classical algorithm that succeeds with bounded probability of error requires a number of samples of the order of N . This shows that in a sample-to-sample setting, quantum computation can achieve the largest possible separation over classical computation. We show that the same bound can be lifted to prove average-case hardness, paving the way for demonstrations on noisy intermediate-scale quantum (NISQ) computers. It follows that under the assumption of the existence of one-way functions, complement sampling gives provable, verifiable and NISQable quantum advantage in a sample complexity setting.

DOI: [10.1103/q55v-wm7y](https://doi.org/10.1103/q55v-wm7y)

Introduction—In 2019, Google claimed to have demonstrated quantum advantage for the first time; that is, their Sycamore processor performed a computational task in the order of minutes that Google estimated would take the world's most powerful classical computer approximately 10 000 years to complete [1]. Their experiment was based on random circuit sampling where, given a classical description of a quantum circuit, the task is to sample in the computational basis according to the probability distribution associated with the output state. Google's achievement started a flurry of research on both ends of the spectrum: new advantage claims were made based on other sampling experiments [2–4], and new classical simulation techniques [5–9] were designed to efficiently simulate the experiments on a classical computer, challenging whether advantage had been truly achieved.

For a convincing demonstration of quantum advantage, the task at hand should ideally satisfy the following three criteria: (i) it can be solved efficiently in a near-term quantum experiment; (ii) it is provably classically hard to

solve, i.e., takes superpolynomial time as the system size scales; (iii) the solution can be efficiently verified with a classical computer with minimal trust in the experiment.

For random circuit sampling, the main problem is that the verification seems unavoidably tied to the underlying classical hardness. Consequently, it seems to inherently require exponential time to classically verify the output. A recent proposal to address this challenge involves the study of random peaked circuits [10]; however, the current understanding of these circuits is insufficient to determine their viability as a quantum advantage experiment. There exist many other proposals, like Shor's algorithm [11], Yamakawa and Zhandry search [12], variational quantum eigensolvers [13], algorithms for verifying quantumness [14], and more, that all seem to satisfy at most two of the above three points.

However, quantum resources offer more than just potential computational speedups, where ultimately time is the resource that one wants to minimize. For example, it is known (or in some cases strongly conjectured) that superpolynomial advantages exist in communication complexity [15,16], sample complexity [17,18] and space complexity [19,20]. All these advantages stem from the fact that quantum information processing not only provides potential benefits in computation, but also in the way data can be stored, communicated and accessed in a quantum setting.

In this Letter, we present a new approach that combines the two ideas that quantum computers are inherently samplers, and that they can leverage quantum resources in ways classical computers cannot for their classical

*Contact author: marcello.benedetti@quantinuum.com†Contact author: harry.buhrman@quantinuum.com‡Contact author: jrw@cwi.nl

counterparts. We show that in a *sample-to-sample* setting, quantum computing can achieve the largest possible advantage. Proofs and technical discussions are included in Supplemental Material (SM) [21].

Quantum advantage in complement sampling—Let $\omega = \{0, 1\}^n$ be a set of $N = 2^n$ elements. We will focus on the following problem.

Problem 1 (Complement Sampling)—Given (quantum) sampling access to the uniform distribution over a subset $S \subset \omega$ with cardinality K , output an element $y \in \bar{S}$ where $\bar{S}$ is the complement of S .

Complement sampling can be viewed as a sample-to-sample problem where, given samples from some set A according to a distribution μ_A , one has to output something from a related set B , according to a distribution μ_B . This is different than, for example, most problems in classical [29,30] and quantum [31–33] machine learning that use the number of samples as a complexity measure. For a concrete example, [34] studies the sample complexity of the quantum coupon collector, where the task is to fully identify a subset A given quantum sampling access to its elements. In contrast, complement sampling does not require learning anything about the subset A .

In this Letter, we study the power of quantum samples and quantum computing over classical samples in solving complement sampling. Given a probability distribution $D: S \subseteq \omega \rightarrow [0, 1]$, a quantum sample from D is defined as having access to a single copy of the state [17,32]

$$|S_D\rangle = \sum_{x \in S} \sqrt{D(x)} |x\rangle.$$

When D is the uniform distribution, we simply write $|S\rangle$. Measuring copies of $|S_D\rangle$ in the computational basis is identical to classical sampling from the distribution $D(x)$. Thus there would be no advantage to a classical algorithm in having access to the quantum samples if all the classical algorithm can do is computational basis measurements.

From a purely information-theoretic perspective, it is not immediately clear that quantum samples should provide any advantage to a quantum algorithm either. Since the required output for complement sampling is a classical n -bit string, it is natural to compare the amount of classical information received per sample. Holevo’s theorem guarantees that the maximum information content per sample is the same for quantum and classical samples: both yield at most n bits of classical information [35]. Nevertheless, we prove that quantum samples can offer a dramatic advantage over classical samples in complement sampling: our main result is a characterization of the quantum and classical sample complexity in solving this task with certain success probabilities.

Theorem 1 (Informal, from Theorems 5 and 6 in SM [21])—Let $1 \leq K \leq N - 1$ be the cardinality of a subset $S \subset \omega$. Then there is a polynomial-time quantum algorithm

that uses only 1 quantum sample and solves complement sampling with probability

$$\frac{\min\{K, N - K\}}{\max\{K, N - K\}}.$$

Let $\delta \in [0, \frac{1}{2}]$. Then any classical algorithm that solves complement sampling with success probability $\geq \frac{1}{2} + \delta$ needs at least

$$N - \frac{2(N - K)}{2\delta + 1},$$

distinct classical samples.

The classical lower bound is exactly matched by a simple randomized algorithm: it outputs a uniformly random element from the set of all possible elements, excluding those observed in the classical samples. The quantum algorithm, which we describe later, solves complement sampling by doing something even stronger than what is stated in Theorem 1: it (approximately)

$$|S\rangle = \frac{1}{\sqrt{K}} \sum_{x \in S} |x\rangle \longrightarrow |\bar{S}\rangle = \frac{1}{\sqrt{N - K}} \sum_{x \notin S} |x\rangle,$$

from which a sample from $\bar{S}$ is simply obtained by performing a measurement in the computational basis.

Hereafter, we use Landau symbols $\mathcal{O}$ and Ω for asymptotic upper and lower bounds, respectively. For cases where polylogarithmic terms are omitted, we use $\tilde{\mathcal{O}}$ and $\tilde{\Omega}$ instead. Theorem 1 shows that for $K = N/2$, the classical sample complexity is $\Omega(N)$ for any $\delta = \Omega(1)$ whilst the quantum algorithm succeeds with probability 1 using only a single quantum sample. From a theoretical perspective, this is a remarkable feat: it shows that in a sample-to-sample setting, quantum computation can achieve the largest possible separation from classical computation, namely a gap of 1 versus a linear (in N) number of samples. This contrasts with, for example, the case of black-box query complexity, as Aaronson and Ambainis [36] showed that no partial Boolean function has quantum query complexity 1 while requiring a linear number of randomized queries [37].

Using Kolmogorov complexity we also show an explicit circuit lower bound on any classical sampler that uses only a polynomial number of samples and returns samples uniformly at random from the complement (which the quantum algorithm allows one to do).

Theorem 2 (Informal, from Theorem 8 in SM [21])—Let $K = N/2$ and $N = 2^n$. For $n \geq 1$ sufficiently large, any sampler that takes as input $\text{poly}(n)$ classical samples from S and outputs samples from $\bar{S}$ uniformly at random must have circuit complexity at least $\tilde{\Omega}(N)$.

The Kolmogorov complexity argument is particularly insightful of why complement sampling works in a quantum setting, as it highlights the key difference between

classical and quantum samples. Classical samples can be cloned and reused indefinitely. For example, they can be used to construct a circuit that describes S and, in turns, a sampler for $\bar{S}$. In comparison, our quantum algorithm *destroys* the quantum sample in the process of providing a sample from the complement. This means that the quantum algorithm needs a new input for every generated output.

The quantum algorithm—Our initial inspiration comes from the construction by Aaronson *et al.* [39]. The authors show how several notions of complexity on quantum states are related in intricate ways. Consider two arbitrary orthogonal states $|a\rangle$ and $|b\rangle$ spanning a two-dimensional space. The conjugate (or Fourier) basis within this subspace is given by the equal superpositions $|\phi^+\rangle = (|a\rangle + |b\rangle)/\sqrt{2}$ and $|\phi^-\rangle = (|a\rangle - |b\rangle)/\sqrt{2}$. Their main result is that a circuit which (approximately) swaps $|a\rangle$ to $|b\rangle$ and back can be used to build another circuit which (approximately) distinguishes between $|\phi^+\rangle$ and $|\phi^-\rangle$. Vice versa, a distinguisher for $|\phi^+\rangle$ and $|\phi^-\rangle$ can be used to make a swapper for $|a\rangle$ and $|b\rangle$. This implies that the *swapping complexity* of two orthogonal states is related to the *distinguishing complexity* of the conjugate states. In particular, an efficient implementation of one implies an efficient implementation of the other.

How can this be used for the complement sampling problem? A Boolean function f on the domain ω defines the phase state

$$|y_f\rangle = \frac{1}{\sqrt{N}} \sum_{x \in \omega} (-1)^{f(x)} |x\rangle.$$

Consider the case of subsets of cardinality $K = N/2$. It can be observed that any subset state and its complementary state can be expressed using phase states

$$\begin{aligned} |S\rangle &= \frac{1}{\sqrt{2}} (|y_{f_{\text{con}}}\rangle + |y_{f_{\text{bal}}}\rangle), \\ |\bar{S}\rangle &= \frac{1}{\sqrt{2}} (|y_{f_{\text{con}}}\rangle - |y_{f_{\text{bal}}}\rangle), \end{aligned}$$

where f_{con} is a constant function (it maps all inputs $x \in \omega$ to 0) and f_{bal} is a balanced function (half of the inputs in ω get mapped to 0 and the other half get mapped to 1). As the two states are orthogonal, the conjugate states are

$$\begin{aligned} |y_{f_{\text{con}}}\rangle &= \frac{1}{\sqrt{2}} (|S\rangle + |\bar{S}\rangle), \\ |y_{f_{\text{bal}}}\rangle &= \frac{1}{\sqrt{2}} (|S\rangle - |\bar{S}\rangle). \end{aligned}$$

But these are just the postquery states of the Deutsch-Jozsa algorithm [40], and therefore can be efficiently distinguished with probability 1. Figure 1(a) illustrates the distinguisher circuit. Using the Aaronson *et al.* construction, the




FIG. 1. Distinguisher and swappers for complement sampling. (a) Our distinguisher can be thought of as the Deutsch-Jozsa algorithm followed by a copy of the measurement outcome into an ancilla qubit. (b) The complement swapper can be thought of as the Grover diffusion operator $U = 2|+\rangle\langle+| - \mathbb{I}$ up to a global phase. (c) The zero-error swapper uses the controlled- U , the gate $W(q) = e^{i \arccos(\sqrt{q})Y}$, and an ancilla qubit that upon measurement outcome 0 indicates correctness. The parameters are set to $b = 0$, $q = [2(1 - K/N)]^{-1}$ if $K < N/2$, and to $b = 1$, $q = N/(2K)$ if $K \geq N/2$.

Deutsch-Jozsa distinguisher can be turned into a swapper for $|S\rangle$ and $|\bar{S}\rangle$. The swapper is illustrated in Fig. 1(b), and derived in Sec. I of SM [21]. We can now argue that no perfect swapper for $|S\rangle$ and $|\bar{S}\rangle$ exists for values $K \neq N/2$. Indeed, only for $K = N/2$ we have that any two different subsets S_1, S_2 have perfectly distinguishable orthogonal conjugate states (see Proposition 2 in SM [21]). If a perfect swapper would exist for $K \neq N/2$, one could violate the distinguishability limits given by the Helstrom bound [41].

By inspecting the swapper circuit based on the Deutsch-Jozsa distinguisher, we find that, restricted to the register where $|S\rangle$ is held, it implements a reflection around the uniform superposition of basis states. Up to a global phase, this is the Grover diffusion operator $U = 2|+\rangle\langle+| - \mathbb{I}$. With this new insight we generalize the algorithm to $K \neq N/2$ in two different ways. The first, which we simply call *complement swapper*, directly applies U to any $|S\rangle$. For values of K a constant fraction away from $N/2$ the resulting state has constant overlap with the complement $|\bar{S}\rangle$. The second algorithm, which we call *zero-error swapper*, uses an additional ancilla as a flag qubit and a controlled application of U . This is illustrated in Fig. 1(c).

Figure 2 compares the performance of four algorithms as a function of $\beta \in [-1/2, 1/2]$, which we define as the deviation from the ideal subset size ratio $K/N = 1/2 + \beta$. The complement swapper outputs complementary samples with probability $1 - 4\beta^2$ (solid red line), but does not provide information on the correctness of the sample.




FIG. 2. Comparison of complement sampling algorithms given a single sample as input. The vertical axis is the probability of obtaining a sample from the complementary set. The horizontal axis is the deviation from the ideal subset size ratio $K/N = 1/2 + \beta$. When $\beta = 0$ the quantum algorithms based on swappers are exact.

The zero-error swapper outputs complementary samples with probability $[(1 - 2|\beta|)/(1 + 2|\beta|)]$ (dashed blue line), with the guarantee that the output is correct (i.e. it is a *Las Vegas* algorithm). Theorem 5 in SM [21] shows that this is the optimal success probability for any zero-error algorithm when $K \geq N/2$, as well as when no extra auxiliary qubits are allowed (except for the flag qubit) and $K < N/2$. In our proof we construct the hardest instances of S in the sense of maximizing the violation of the conservation of inner products under unitary transformations, which puts an upper bound on the achievable success probability. The quantum algorithm for the coupon collector [34] can also be used for complement sampling (purple dash-dotted line), but it achieves half the probability of success of our complement swapper. Our improvement arises from the coherent reflection around the uniform superposition, whereas the quantum coupon collector performs a measurement $\{|+\rangle\langle+|, \mathbb{I} - |+\rangle\langle+|\}$ on $|S\rangle$, followed by a measurement in the computational basis. The authors [34] additionally consider a stronger input model which gives access to the state preparation circuit for $|S\rangle$. We do not compare against this version of their algorithm because complement sampling does not allow access to the state preparation circuit. Finally, the success probability of the optimal classical algorithm based on random guessing is $[(1/2 - \beta)/(1 - 1/N)] \approx \frac{1}{2} - \beta$ (green dotted line), which is large for negative deviations.

Toward experimental demonstrations—So far we have only discussed worst-case classical hardness results, which are problematic if one also wants to instantiate the problem in an actual experiment. Even worse, one might argue that these hard instances hide all their complexity in the state $|S\rangle$, a comment that we justify in Sec. III of SM [21] by

proving that most states of the form $|S\rangle$ with $K = N/2$ have exponential quantum circuit complexity [42]. To show that complement sampling is indeed capable of demonstrating quantum advantage, we show that the problem remains hard for most choices of S , as long as there is no particular structure to them. We focus on the case where $K = N/2$. Consider the family of subsets $S = \{S^k\}_k$ where each subset is given by $S^k = \{y: y = P^k(x) \text{ and the first bit of } x \in \omega \text{ is } 0\}$ for some strong pseudorandom permutation P^k with a key $k \in \{0, 1\}^\lambda$, and $\lambda \in \mathbb{N}$ is the security parameter. We prove the following result.

Theorem 3 (Informal, from Theorem 7 in SM [21])—Let $K = N/2$ and $N = 2^n$. Let A be a polynomial time algorithm that has sampling access to S . Under the assumption of the existence of one-way functions, for all such A it must hold that

$$\Pr[A \text{ outputs a } y \in \bar{S}] < \frac{1}{2} + \frac{1}{\text{poly}(n)}.$$

Equipped with the above results, we propose an experiment to demonstrate the advantage of quantum over classical resources in a *provable, verifiable* and *NISQable* manner. For a total of r rounds, a player and referee play the following interactive game: (1) The referee picks a pseudorandom permutation P with inverse P^{-1} and defines $S = \{y: y = P(x) \text{ and the first bit of } x \in \omega \text{ is } 0\}$. (2) The (quantum) player asks for j (quantum) samples corresponding to the uniform distribution over elements in S . A quantum player can perform coherent operations on the quantum samples. The player sends back a candidate element $\hat{y}$ supposedly from $\bar{S}$. (3) The referee verifies that $\hat{y} \in \bar{S}$ by checking if $\hat{x} = P^{-1}(\hat{y})$ has the first bit equal to 1.

A sketch of the proposed experiment is provided in Fig. 3. The best strategy for the classical player is to output any sample different from the observed samples (we prove this in Sec. II of SM [21]). For $j = 1$, this random guessing strategy has success probability of $2^{n-1}/(2^n - 1)$, quickly approaching $1/2$ for large n . By repeating the experiment a total of r times, each time with a different pseudorandom permutation, the probability of succeeding in all rounds is approximately $1/2^r$. On the other hand, the quantum player can use our algorithm to obtain $|\bar{S}\rangle$ from $|S\rangle$. This strategy succeeds with probability 1 at each of the r rounds. In reality, the quantum player only needs to succeed with probability $\geq \frac{1}{2} + 1/\text{poly}(n)$ at each round in order to beat the classical player over $r = \text{poly}(n)$ rounds. Thus, the player can tolerate a small rate of noise in the quantum computation. For example, in the simplistic noise model where each gate fails with probability λ , the overall circuit fidelity is of the order of $(1 - \lambda)^m$ where m is the number of gates. As we will see, our quantum circuit only needs $\mathcal{O}(n)$ gates and $\mathcal{O}(\log n)$ circuit depth; the bottleneck being the n -qubit Toffoli gate, which can be implemented in $\mathcal{O}(\log n)$




FIG. 3. A single round of the complement sampling game with j input samples. The referee picks subset S from a family of subsets of fixed cardinality S . The quantum player gets access to j copies of the state $|S\rangle$, and the classical player gets access to j elements sampled according to the uniform distribution over S . Giving the classical player access to measurements of $|S\rangle$ in the computational basis would yield the same input model. Giving the quantum player access to classical samples does not yield any advantage over the classical player. Each player sends back a candidate element $\hat{y}$ and the referee checks whether it belongs to the complementary subset $\bar{S}$.

circuit depth with 1 additional ancilla using the construction from [44].

The preparation of the state $|S\rangle$ by the referee is computationally more expensive as it requires the implementation of pseudorandom permutations. As a near-term toy example, one could implement permutations using the Simplified Advanced Encryption Standard (S-AES) protocol. S-AES uses a block size of 16 bits (as opposed to 128 bits for AES) and a key size of 16 bits as well (128, 192 or 256 for AES). A system size of 16 can display features of quantum advantage: from Theorem 1 with, e.g., $N = 2^{16}$, $K = 2^{15}$ and $\delta = 1/6$, any classical algorithm solving complement sampling requires $2^{14} = 16384$ distinct samples. A quantum circuit implementation of S-AES based on [45] requires only 48 qubits, 168 Toffoli gates, 364 controlled NOT gates, and 75 NOT gates. Beyond $n = 16$ one could resort to *approximate* pseudorandom permutations built from layers of random reversible 3-bit gates. Exponentially small error can be achieved in depth linear in n when using one-dimensional nearest-neighbor layers, and in depth $\sqrt{n}$ when using two-dimensional layers [46]. When noise affects the computation, the referee is allowed to use quantum error detection codes [47] before sending the state $|S\rangle$ to the player. By discarding circuit runs where an error is detected, the referee is expected to increase the fidelity of the state preparation. This comes at the cost of introducing more qubits, gates, and runs, but near-term protocols have shown promising results with low overheads [48].

We expect that in the first complement sampling experiments all computations from both referee and player will be

performed on the same device. Smart choices of families of subsets, and the use of circuit optimizations and error detection, may lead to a NISQable advantage experiment.

Discussion and future work—In this Letter we have shown that in a sample-to-sample setting quantum algorithms provide the largest possible advantage over their classical counterparts. The classical hardness of the problem is very robust: even instances that are generated using strong pseudorandom permutations must be classically hard under the assumption of the existence of one-way functions. On the other hand, the quantum advantage vanishes if classical samples are provided as input instead of the quantum samples. This indicates that the present approach does not provide a purely computational quantum advantage, where both the input and output to a problem are classical.

Our Letter opens up a new direction in demonstrating an advantage of quantum over classical resources. Our back-of-the-envelope resource estimation shows that a proof-of-concept demonstration of complement sampling is NISQable, but further circuit optimizations are likely needed to solve instances that are truly classically intractable. Progress on this matter was recently reported in Ref. [49], where a refined version of the complement sampling game was successfully executed on quantum hardware.

Another topic for future work is the search for other applications of complement sampling. In particular, in the field of cryptography, we believe that there exist applications that can utilize the power of quantum computing for complement sampling.

Acknowledgments—We thank Florian Curchod and Pranav Kalidindi for providing feedback on an earlier version of this paper. J.W. was supported by the Dutch Ministry of Economic Affairs and Climate Policy (EZK), as part of the Quantum Delta NL program.

Data availability—The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

- [1] 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).
- [2] 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).
- [3] 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).

- [4] Q. Zhu, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong *et al.*, Quantum computational advantage via 60-qubit 24-cycle random circuit sampling, *Sci. Bull.* **67**, 240 (2022).
- [5] C. Huang, F. Zhang, M. Newman, J. Cai, X. Gao, Z. Tian, J. Wu, H. Xu, H. Yu, B. Yuan *et al.*, Classical simulation of quantum supremacy circuits, [arXiv:2005.06787](https://arxiv.org/abs/2005.06787).
- [6] B. Barak, C.-N. Chou, and X. Gao, Spoofing linear cross-entropy benchmarking in shallow quantum circuits, in *Proceedings of the 12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, Vol. 185 (Association for Computing Machinery, New York, NY, USA, 2021), p. 30.
- [7] F. Pan, K. Chen, and P. Zhang, Solving the sampling problem of the sycamore quantum circuits, *Phys. Rev. Lett.* **129**, 090502 (2022).
- [8] C. Oh, L. Jiang, and B. Fefferman, Spoofing cross-entropy measure in boson sampling, *Phys. Rev. Lett.* **131**, 010401 (2023).
- [9] 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, *PRX Quantum* **5**, 010334 (2024).
- [10] S. Aaronson and Y. Zhang, On verifiable quantum advantage with peaked circuit sampling, [arXiv:2404.14493](https://arxiv.org/abs/2404.14493).
- [11] 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.
- [12] T. Yamakawa and M. Zhandry, Verifiable quantum advantage without structure, *J. ACM* **71**, 1 (2024).
- [13] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [14] Z. Brakerski, P. Christiano, U. Mahadev, U. Vazirani, and T. Vidick, A cryptographic test of quantumness and certifiable randomness from a single quantum device, *J. ACM* **68**, 1 (2021).
- [15] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, Quantum fingerprinting, *Phys. Rev. Lett.* **87**, 167902 (2001).
- [16] D. Gilboa and J. R. McClean, Exponential quantum communication advantage in distributed learning, [arXiv:2310.07136](https://arxiv.org/abs/2310.07136).
- [17] N. H. Bshouty and J. C. Jackson, Learning DNF over the uniform distribution using a quantum example oracle, in *Proceedings of the Eighth Annual Conference on Computational Learning Theory (Society for Industrial and Applied Mathematics (SIAM), 1995)*, pp. 118–127, [10.1137/S0097539795293123](https://doi.org/10.1137/S0097539795293123).
- [18] A. B. Grilo, I. Kerenidis, and T. Zijlstra, Learning-with-errors problem is easy with quantum samples, *Phys. Rev. A* **99**, 032314 (2019).
- [19] D. Gavinsky, J. Kempe, I. Kerenidis, R. Raz, and R. de Wolf, Exponential separations for one-way quantum communication complexity, with applications to cryptography, in *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2007), pp. 516–525.
- [20] J. Kallaughner, O. Parekh, and N. Voronova, Exponential quantum space advantage for approximating maximum directed cut in the streaming model, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2024), pp. 1805–1815.
- [21] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/q55v-wm7y>, which includes Refs. [22–28].
- [22] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 1996), pp. 212–219.
- [23] C. Zalka, A Grover-based quantum search of optimal order for an unknown number of marked elements, [arXiv:quant-ph/9902049](https://arxiv.org/abs/quant-ph/9902049).
- [24] A. C.-C. Yao, Probabilistic computations: Toward a unified measure of complexity, in *Proceedings of the 18th Annual Symposium on Foundations of Computer Science (sfcs 1977)* (IEEE Computer Society, 1977), pp. 222–227, <https://ieeexplore.ieee.org/document/4567946>.
- [25] A. F. Mendelson, M. A. Zuluaga, B. F. Hutton, and S. Ourselin, What is the distribution of the number of unique original items in a bootstrap sample?, [arXiv:1602.05822](https://arxiv.org/abs/1602.05822).
- [26] M. Luby and C. Rackoff, How to construct pseudorandom permutations from pseudorandom functions, *SIAM J. Comput.* **17**, 373 (1988).
- [27] J. Katz and Y. Lindell, *Introduction to Modern Cryptography*, Chapman & Hall/CRC Cryptography and Network Security Series (CRC Press, London, 2020).
- [28] D. K. Ray-Chaudhuri and R. M. Wilson, On t -designs, *Osaka J. Math.* **12**, 737 (1975).
- [29] D. Haussler, Decision theoretic generalizations of the pac model for neural net and other learning applications, *Inf. Comput.* **100**, 78 (1992).
- [30] S. Hanneke, The optimal sample complexity of PAC learning, *J. Mach. Learn. Res.* **17**, 1 (2016), <http://jmlr.org/papers/v17/15-389.html>.
- [31] S. Aaronson, The learnability of quantum states, *Proc. R. Soc. A* **463**, 3089 (2007).
- [32] S. Arunachalam and R. de Wolf, Guest column: A survey of quantum learning theory, *ACM SIGACT News* **48**, 41 (2017).
- [33] L. Coopmans and M. Benedetti, On the sample complexity of quantum Boltzmann machine learning, *Commun. Phys.* **7**, 274 (2024).
- [34] S. Arunachalam, A. Belovs, A. M. Childs, R. Kothari, A. Rosmanis, and R. de Wolf, Quantum coupon collector, in *Proceedings of the 15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, Vol. 158 (Association for Computing Machinery, New York, NY, USA, 2020), pp. 10:1–10:17.
- [35] A. S. Holevo, Bounds for the quantity of information transmitted by a quantum communication channel, *Prob. Peredachi Inf.* **9**, 3 (1973).
- [36] S. Aaronson and A. Ambainis, Forrelation: A problem that optimally separates quantum from classical computing, in *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC ’15

- (Association for Computing Machinery, New York, 2015), pp. 307–316.
- [37] They proved that the separation can be at most 1 versus $\tilde{\Omega}(\sqrt{N})$, which is achieved by the Forrelation problem. For a constant number of queries, one can achieve $\mathcal{O}_\epsilon(1)$ versus $\Omega(N^{1-\epsilon})$ for any $\epsilon > 0$ [38]. Here $\mathcal{O}_\epsilon(f(n))$ means $\mathcal{O}(f(n))$ where the hidden constant is allowed to depend on ϵ , which is treated as a fixed parameter.
- [38] N. Bansal and M. Sinha, K-forrelation optimally separates quantum and classical query complexity, in *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2021), pp. 1303–1316.
- [39] S. Aaronson, Y. Atia, and L. Susskind, On the hardness of detecting macroscopic superpositions, [arXiv:2009.07450](#).
- [40] D. Deutsch and R. Jozsa, Rapid solution of problems by quantum computation, *Proc. R. Soc. A* **439**, 553 (1992).
- [41] C. W. Helstrom, Quantum detection and estimation theory, *J. Stat. Phys.* **1**, 231 (1969).
- [42] This circuit lower bound was to be expected under the assumption that $\text{QMA} \neq \text{QCMA}$, because if all subset states could be (approximately) constructed with a polynomially-sized circuit then we would have that $\text{QCMA} = \text{QMA}$ by the result of [43].
- [43] A. B. Grilo, I. Kerenidis, and J. Sikora, QMA with subset state witnesses, in *International Symposium on Mathematical Foundations of Computer Science* (Springer, New York, 2015), pp. 163–174.
- [44] J. Nie, W. Zi, and X. Sun, Quantum circuit for multi-qubit Toffoli gate with optimal resource, [arXiv:2402.05053](#).
- [45] Z.-G. Wang, S.-J. Wei, and G.-L. Long, A quantum circuit design of AES requiring fewer quantum qubits and gate operations, *Front. Phys.* **17**, 41501 (2022).
- [46] W. Gay, W. He, N. Kocurek, and R. O’Donnell, Pseudo-randomness properties of random reversible circuits, in *Advances in Cryptology—CRYPTO 2025* (Springer Nature, Cham, 2025), pp. 651–678.
- [47] E. Knill, Fault-tolerant postselected quantum computation: Threshold analysis, [arXiv:quant-ph/0404104](#).
- [48] C. N. Self, M. Benedetti, and D. Amaro, Protecting expressive circuits with a quantum error detection code, *Nat. Phys.* **20**, 219 (2024).
- [49] M. Benedetti, G. Marin-Sanchez, J. Weggemans, M. Rosenkranz, and H. Buhrman, Unconditional and exponentially large violation of classicality, [arXiv:2511.11008](#).