One T Gate Makes Distribution Learning Hard

M. Hinsche,[1,*] M. Ioannou,[1,†] A. Nietner,[1,‡] J. Haferkamp,[1] Y. Quek,[1] D. Hangleiter,[2] J.-P. Seifert,[3] J. Eisert,[1,4,5] and R. Sweke[1,§]

[1]*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*
[2]*Joint Center for Quantum Information and Computer Science (QuICS), University of Maryland and NIST, College Park, Maryland 20742, USA*
[3]*Department of Electrical Engineering and Computer Science, TU Berlin, 10587 Berlin, Germany*
[4]*Helmholtz-Zentrum Berlin für Materialien und Energie, 14109 Berlin, Germany*
[5]*Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany*

(Received 4 August 2022; revised 10 February 2023; accepted 5 May 2023; published 13 June 2023)

The task of learning a probability distribution from samples is ubiquitous across the natural sciences. The output distributions of local quantum circuits are of central importance in both quantum advantage proposals and a variety of quantum machine learning algorithms. In this work, we extensively characterize the learnability of output distributions of local quantum circuits. Firstly, we contrast learnability with simulatability by showing that Clifford circuit output distributions are efficiently learnable, while the injection of a single T gate renders the density modeling task hard for any depth $d = n^{\Omega(1)}$. We further show that the task of generative modeling universal quantum circuits at any depth $d = n^{\Omega(1)}$ is hard for any learning algorithm, classical or quantum, and that for statistical query algorithms, even depth $d = \omega[\log(n)]$ Clifford circuits are hard to learn. Our results show that one cannot use the output distributions of local quantum circuits to provide a separation between the power of quantum and classical generative modeling algorithms, and therefore provide evidence against quantum advantages for practically relevant probabilistic modeling tasks.

DOI: 10.1103/PhysRevLett.130.240602

Recent years have witnessed a massive increase in the use of machine learning techniques, in both industry and the physical sciences [1]. The subfield of probabilistic modeling is of particular importance, and deep generative models have been successfully applied to solve problems and guide progress in areas as diverse as cosmological structure formation [2], particle shower modeling [3], astrophysical imaging [4], and protein structure prediction [5]. In light of this diverse applicability, including a wide variety of potential applications in the physical sciences, much effort has been dedicated to the development of state-of-the-art models and algorithms for probabilistic modeling.

Simultaneously, the last years have witnessed significant interest in the potential of exploiting quantum devices for machine learning tasks [1,6]. Of particular interest are hybrid quantum-classical schemes, in which parametrized quantum circuits are used as a model class, whose parameters are optimized via classical algorithms [7,8]. In the context of generative modeling, the output distributions of quantum circuits are a particularly natural model class, referred to as *quantum circuit Born machines* (QCBMs) [9–11]. In particular, it is known that this model class is expressive enough to contain many probabilistic graphical models [12,13], while not being classically simulatable [14–16]. These facts, along with a growing body of numerical experiments [17–20], suggest that hybrid quantum-classical algorithms using QCBMs as a model class may offer concrete advantages over state-of-the-art classical generative modeling techniques. However, to date, there are no rigorous results which support this intuition. Given the wide variety of potential applications, understanding concretely the potential of quantum probabilistic modeling techniques is of broad interest.

In this work we provide strong evidence that QCBM based algorithms will not outperform classical algorithms for probabilistic modeling, and in the process provide a variety of fundamental learning-theoretic insights into the properties of local quantum circuits and the many-body quantum systems they model. To do this, we study the learnability of the output distributions of local quantum circuits, within the formal framework of distribution learning [21]. This is motivated by the following observation: If quantum devices offer an advantage for a broad range of physically or practically relevant distribution learning problems, then one would surely believe that they offer advantages for problems which have been fine-tuned in their favor. The task of learning the output distributions of quantum circuits themselves is precisely such a fine-tuned probabilistic modeling problem biased toward quantum learning algorithms. In this work, by proving the hardness of learning the output distributions of local

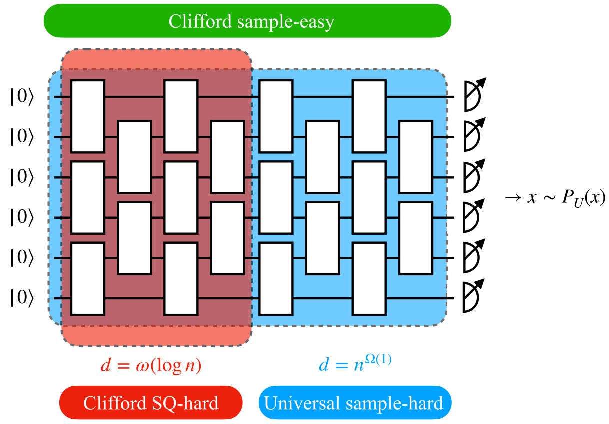

FIG. 1. Regimes of complexity of the task of generator- or evaluator-learning the output distributions of local quantum circuits on n qubits of depth d. In accord with the intuition that deeper circuits generate more complex distributions, the complexity depends on how d scales with n. We find that for any depth, Clifford circuits are efficiently learnable from samples (Theorem 1). By contrast, for $d = \omega(\log n)$, even Clifford circuit output distributions are not efficiently learnable from *statistical queries* (Theorem 4). Moreover, given *sample* access, the output distributions of generic local quantum circuits are not efficiently learnable at any depth $d = n^{\Omega(1)}$, up to standard cryptographic assumptions (Corollary 1 and Theorem 3).

quantum circuits for both quantum and classical algorithms, we prove rigorously the lack of formal quantum advantage in this setting, and therefore provide strong evidence against the existence of a more generally applicable quantum advantage in this important domain (see Fig. 1 for on overview of our results).

Additionally, we note that local quantum circuits, as considered in this Letter, are often considered as a model for a wide variety of many-body quantum systems. In particular, discretization of the time evolution under a local Hamiltonian naturally leads precisely to a local quantum circuit [22,23]. As such, the properties of such circuits are currently of great interest across a broad range of subfields of physics, ranging from condensed matter, across high energy physics, all the way to the theory of black holes [24–27]. In light of this, our results can be naturally viewed as placing rigorous and fundamental limitations on the possibility of efficiently learning certain properties of these wide ranging physical systems.

Finally, we show that, within the context of distribution learning, classical simulatability of a class of quantum circuits *does not* imply efficient learnability. This is in strong contrast to existing conjectures and known results in other related settings [17,28–30]. To do this, we prove that the output distributions of Clifford circuits are efficiently learnable, while adding a single T gate to the circuit renders the learning problem hard. As such, while the complexity of the classical simulation scales with the number of T gates, we find that the addition of a single T gate induces a striking complexity transition in the corresponding distribution learning problem.

Setting.—In this Letter, we are concerned with learning distributions promised to be from a *distribution class* $\mathcal{D}$. In particular, we are interested in the properties of learning algorithms that solve the following problem [31]:

Problem 1 (Distribution learning).—Given a distribution class $\mathcal{D}$, samples from an unknown distribution $P \in \mathcal{D}$, and $\varepsilon, \delta \in (0,1)$, output with probability at least $1-\delta$, a representation of a distribution Q satisfying $\mathrm{TV}(P,Q) \leq \varepsilon$.

We will be concerned with two types of *representations*, namely *generators* and *evaluators*: An evaluator for a distribution Q is a computationally efficient algorithm which, when given some x, outputs the probability $Q(x)$, and a generator for a distribution Q is a computationally efficient algorithm for generating samples from Q.

We note that the problem of distribution learning with respect to an evaluator is often referred to as *density modeling*, while the problem of learning with respect to a generator is often referred to as *generative modeling*. Additionally, we stress that in the case of generative modeling it is *not* sufficient for the learning algorithm to store and later reproduce the samples it received during the learning phase, or to output a larger but still bounded set of samples [32]. Indeed, the learning algorithm is required to output another algorithm—a generator—which can output as many samples as desired, from a distribution which is close in total variation distance to the unknown target distribution.

We are concerned here exclusively with discrete distributions over $\{0,1\}^n$, and denote the set of all such distributions by $\mathcal{D}_n$. Given some $\mathcal{D} \subseteq \mathcal{D}_n$, we say that an algorithm is a computationally (sample) efficient algorithm for learning $\mathcal{D}$ with respect to a particular representation (either generators or evaluators) if it solves the above problem for all $P \in \mathcal{D}$, using $O[\mathrm{poly}(n, 1/\varepsilon, 1/\delta)]$ computational time (samples). If there exists a computationally efficient learning algorithm for $\mathcal{D}$ with respect to a particular representation, then we say that $\mathcal{D}$ is efficiently learnable with respect to that representation. If there does not exist a computationally efficient learning algorithm for some class $\mathcal{D}$ with respect to a particular representation, then we say that $\mathcal{D}$ is hard to learn with respect to that representation.

In this work, we focus on the output distributions of quantum circuits. More specifically, to any unitary U we have the associated probability distribution P_U, with probabilities

$$P_U(x) \coloneqq |\langle x|U|0^{\otimes n}\rangle|^2. \tag{1}$$

We then consider sets of distributions obtained from all unitaries generated by quantum circuits of a specific depth, with gates from a specific gate set. Unless otherwise specified, we consider one-dimensional circuits consisting only of nearest-neighbor gates, which for convenience we refer to as *local* quantum circuits. We are particularly interested in how the complexity of learning depends on both the gate set, and the circuit depth. Our results generalize and extend seminal work on learning the output distributions of classical circuits [21]. We also stress that we consider learning the classical distributions defined by measurement in the computational basis—i.e., we do not assume access to the output state and the ability to measure in different bases.

Throughout the following sections, we only provide intuitive proof sketches and defer the reader to the Supplemental Material [33] for detailed proofs.

Learning Clifford distributions.—We start by studying the learnability of the output distributions of Clifford circuits. This will elucidate the relation between classical simulatability of quantum circuits and their learnability: Famously, by virtue of the Gottesman-Knill theorem, Clifford circuits can be efficiently classically simulated [56,57]. Similarly, it has been found previously that the algebraic structure of the Clifford group also facilitates efficient learning of an unknown stabilizer state [58] or Clifford circuit [29] from few copies of the unknown quantum state. Furthermore, in Aaronson's framework of probably approximately correct (PAC)-learning of quantum states [60], stabilizer states have been found to be efficiently learnable [28,59]. In this setting, Ref. [30] finds a sufficient condition under which the complexity of simulatability and learnability are aligned. Here, we ask whether the alignment in the complexity of classical simulation and learning holds also in the distribution learning setting. Indeed, when studying Clifford circuits, we find that our learning model is no exception.

Theorem 1.—The set $\mathcal{D}_{\mathrm{Cl}}$ of Clifford circuit distributions, for any depth, is efficiently learnable with respect to generators and evaluators.

Proof (sketch).—Clifford circuit output distributions are uniform over affine subspaces of the finite n dimensional vector space $\mathbb{F}_2^n$. Hence using Gaussian elimination on $O(n)$ samples recovers the correct affine subspace, and from this the correct distribution representation, with success rate $1 - \exp[-\Omega(n)]$. ∎

Hardness of learning Clifford + T *distributions.*—Next, we ask whether this alignment of complexity extends even to slightly non-Clifford circuits. In particular, on the simulation side, the run-time of the best-known classical algorithms for simulating T-enriched Clifford circuits will grow exponentially with the number of T gates [61,62]. On the learning side, a first result for learning output states of unknown Clifford + T circuits, from copies of the unknown state, has been obtained in Ref. [29]. They also find an exponential scaling in the number of T gates provided all T gates are applied in a single layer.

Let us now return to the distribution learning setting. We consider the class of output distributions arising from T-enriched Clifford circuits. The following result relies on the *learning parities with noise* (LPN) assumption. It posits that there does not exist an efficient algorithm, quantum or classical, for learning from classical samples the class of Boolean parity functions under the uniform distribution when subject to any constant-rate random classification noise [63]. We note that this is a canonical assumption for many cryptographic schemes [64,65].

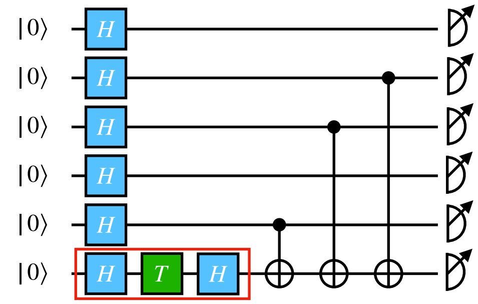

FIG. 2. Example of a circuit used in the proof of Theorem 2. Without the red box, samples from this circuit are of the form $[x, f(x)]$ where x is uniformly random and f is the parity function supported on bits 2,3,5. With the red box, the samples are of the form (x, y) where $y = f(x)$ with probability $1 - \eta$ and $y = \neg f(x)$ with probability η.

Theorem 2.—Under the LPN assumption, the output distributions of local Clifford circuits of depth $d = n^{\Omega(1)}$ enriched with a single T gate are not efficiently learnable with respect to an evaluator.

Proof (Sketch).—Ref. [21] gives a class of distributions such that LPN reduces to evaluator learning this class. Specifically, for each parity function, there is a corresponding distribution. Each such distribution can be realized as the output distribution of the Clifford circuit enriched with a single T gate (see, e.g., Fig. 2). We obtain the stated depth dependence by recompiling the circuit into local gates and using a rescaling argument to trade circuit depth for learning complexity. ∎

The key insight underlying the proof of Theorem 2 is that the LPN noise can be realized by a single T gate. Moreover, it can be seen that, if one relaxes the nearest-neighbour requirement on the Clifford gates, i.e., allowing instead for arbitrary connectivity between qubits, then one obtains the above hardness result in Theorem 2 already for depth $d = \Omega(1)$. The sharp transition in complexity between Theorem 1 and Theorem 2 stands in interesting contrast to the smooth increase in the complexity of classically simulating T-enriched Clifford circuits: In particular, while T-enriched Clifford circuits can be simulated efficiently for up to $O(\log n)$ many T gates [62], a single T gate is enough to make distribution learning with an evaluator at least as hard as LPN. We therefore see that single T-enriched Clifford circuits provide a simple and striking example of a quantum circuit class whose output distributions are hard to learn, while being easy to classically simulate. We note that other such examples could be straightforwardly constructed via classical-to-quantum circuit mappings from hard to learn classes of classical circuits. Additionally, the class of T-enriched local Clifford circuits is a subclass of the class of all local quantum circuits. Hence, the conditional hardness result of Theorem 2 also applies to this more general class:

Corollary 1.—Under the LPN assumption, the output distributions of local quantum circuits of depth $d = n^{\Omega(1)}$ are not efficiently learnable with respect to an evaluator.

Hardness of learning generators.—In the previous section, we saw that Clifford distributions with a single T gate are hard to learn with respect to an evaluator. This leaves open the question of the complexity of learning the output distributions of non-Clifford circuits with respect to a generator. While theoretically interesting, this is also the class of distributions for which one would most expect to obtain a quantum advantage via quantum generative models (QCBMs). As such, our hardness result also leads to practical insights on the potential of QCBMs.

In Ref. [21], it has been shown that the output distributions of polynomially sized classical circuits are not efficiently classically learnable with respect to a generator. In this section, we establish an analogous result for the output distributions of quantum circuits by adapting the proof strategy of Ref. [21]. Our result applies to both quantum and classical learning algorithms. In particular, we show that one can embed *pseudorandom functions* (PRFs) into the output distributions of local quantum circuits. In order to establish hardness for quantum learning algorithms, we use "standard-secure" PRFs—i.e., PRFs secure against quantum adversaries with classical membership queries [66].

Theorem 3.—Assuming the existence of classical-secure (standard-secure) pseudorandom functions, there is no efficient classical (quantum) algorithm for learning the output distributions of depth $d = n^{\Omega(1)}$ local quantum circuits, with gates from any universal gate set.

Proof (sketch).—Instantiating the proof of Theorem 17 in Ref. [21] with a standard-secure PRF yields the following: the output distributions of polynomially sized classical circuits are not efficiently generator learnable, even by quantum learning algorithms. Polynomially sized classical circuits can be realized by polynomially sized local quantum circuits. Therefore, the output distributions of polynomially sized local quantum circuits can also not be learned efficiently with respect to a generator. This result can be extended to any universal gate set by virtue of the Solovay-Kitaev theorem. We obtain the stated depth dependence by use of a rescaling argument trading complexity for depth. ■

Previous work has suggested, and provided numerical evidence, that learning a generator for quantum circuit output distributions is hard for classical learning algorithms [17–20]. Theorem 3 provides a rigorous proof of this and, interestingly, shows that these distributions are hard to learn even using quantum algorithms—including QCBM based learners. As such, one cannot hope to use the output distributions of local quantum circuits to prove a probabilistic modeling separation between QCBM based algorithms and classical algorithms.

We note that our proof technique shares similarities with that of Ref. [67], where it was shown that learning Boolean functions generated by constant depth classical circuits from quantum examples is hard. However, function classes which are hard to learn cannot be generically used to create distribution classes which are hard to learn with respect to a generator [68]. Thus, our results do not follow directly from theirs.

Hardness of learning with statistical query algorithms.—In this section we show that the hardness results of the previous sections can be strengthened to hold for the output distributions of superlogarithmic depth circuits (see Fig. 1), if one considers a restricted—but practically highly relevant—class of learning algorithms.

To understand this restriction, note that in Theorem 1 the algorithm uses individual samples from the unknown target distribution to exploit the algebraic structure of Clifford output distributions. However, in the absence of a strong promise on the structure of the target, it is *a priori* unclear how a learning algorithm should utilize individual samples from the target distribution. As such, most *generic* distribution learning algorithms—i.e., algorithms which are not designed specifically for one particular distribution class—work by using samples from the unknown distribution to estimate statistical averages with respect to that distribution, including many gradient-based algorithms used in practice [10,17,69].

We formally study such learning algorithms by restricting their access to the target distribution to approximate statistical averages. More specifically, we assume that the algorithm has access to a *statistical query oracle*, which when queried with some efficiently computable function $\phi: \{0,1\}^n \to [-1,1]$ returns some v such that $|\mathbb{E}_{x\sim P}[\phi(x)] - v| \leq \tau$ [70]. To relate back to sample hardness we consider at most inverse polynomial accuracy $\tau = \Omega[1/\mathrm{poly}(n)]$ [71].

Theorem 4.—There is no query efficient algorithm for learning from inverse polynomially accurate statistical queries $\mathcal{D}_{\mathrm{Cl}}$ at depth $\omega[\log(n)]$ and $\mathcal{D}_{\mathcal{G}}$ at depth $\omega[\log^k(n)]$ where k is a constant depending on the universal gate set $\mathcal{G}$ (which can be as small as 2), with respect to either generators or evaluators.

Proof (sketch).—As shown in Refs. [73,74], learning parities in the statistical query model is hard. From this, one can prove that the output distributions of parity functions on uniformly random inputs are also hard to learn from statistical queries. We have already shown in the proof of Theorem 2 that the output distributions of parity functions can be realized by linear depth Clifford circuits. Combining these two facts yields the hardness result for linear depth Clifford circuits. We then obtain the first claim by applying a rescaling argument which trades circuit depth for complexity. We obtain the second claim by using robustness properties of the statistical query oracle, coupled with the Solovay-Kitaev theorem to approximate Clifford circuits. ■

A first immediate consequence of the above result is that one cannot hope to use the output distributions of superlogarithmic depth local circuits to prove a practical separation between the power of classical learning algorithms and QCBMs, provided one uses previously proposed QCBM based learning algorithms based on statistical queries [10,17]. Again, as the output distributions of local quantum circuits were the most promising candidate for such a separation, this result provides evidence against such a separation in more practically relevant settings. Additionally, Theorem 3 leaves open the possibility that there exists some efficient learning algorithm for circuits with depth less than $n^{\Omega(1)}$. However, as hardness in the statistical query model is often taken as evidence for hardness in the sample model [70], the above result suggests that Theorem 3 could potentially be strengthened to hold for the output distributions of superlogarithmic depth circuits. At least, any efficient learning algorithm for such circuits must utilize individual samples in a nontrivial way.

Conclusions.—In this Letter, we have provided an extensive characterization of the complexity of learning the output distributions of local quantum circuits. As such circuits provide a model for a wide variety of many-body quantum systems, these results are of fundamental interest in their own right. However, this characterization also contributes to our understanding of the relationship between the learnability and simulatability of local quantum circuit output distributions.

Moreover, our results have multiple implications for the emerging field of *quantum machine learning*. Much research in this vein aims to identify problems which provably separate the power of quantum and classical learning algorithms [75]. Previously, leveraging cryptographic assumptions, highly fine-tuned learning problems were constructed for which fault-tolerant quantum computers can obtain an exponential advantage [76–78]. The output distributions of quantum circuits were a primary candidate for establishing a separation for a natural learning problem. However, our work establishes that this is not possible in a strict sense, and therefore implies the need to identify new strategies for proving practically relevant quantum advantages in machine learning. In particular, our work complements existing results [79,80] that place limitations on the applicability of near-term hybrid quantum-classical learning algorithms, including QCBMs.

There remain many exciting questions. Firstly, are our worst-case bounds tight? In particular, can one exhibit efficient learning algorithms for the circuit depths not covered by our hardness results? Secondly, can one characterize the *sample* complexity of the learning tasks we have considered? Thirdly, in order to gain insight into the performance of heuristic learning algorithms, it is important to understand the *average-case* complexity of learning the output distributions of local quantum circuits. Additionally, it is interesting to study the learnability of other physically motivated distributions, such as those arising from free-fermionic evolutions [81,82]. Finally, to fully characterize the relationship between simulatability and learnability, it is of interest to understand whether hardness of simulation implies hardness of learning. In particular, are there circuit distributions which are hard to classically simulate, while being efficiently learnable?

We are thankful for excellent discussions with Matthias Caro and Hakop Pashayan, and the feedback of an unknown peer reviewer. This work has been funded by the Cluster of Excellence MATH+ (EF1-11), the BMWK (PlanQK, EniQmA), the BMBF (Hybrid, QPIC-1), the DFG (CRC183, EI 519 20-1), the QuantERA (HQCC), the Munich Quantum Valley (K8), the Alexander von Humboldt Foundation and the Quantum Flagship (PasQuans2).

M. H., M. I., A. N., and R. S. have contributed equally to this work.

[*]Corresponding author.
m.hinsche@fu-berlin.de
[†]Corresponding author.
marios.ioannou@fu-berlin.de
[‡]Corresponding author.
a.nietner@fu-berlin.de
[§]Corresponding author.
rsweke@gmail.com

[1] G. Carleo, J. I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, Rev. Mod. Phys. **91**, 045002 2019).

[2] A. C. Rodríguez, T. Kacprzak, A. Lucchi, A. Amara, R. Sgier, J. Fluri, T. Hofmann, and A. Réfrégier, Fast cosmic web simulations with generative adversarial networks, Comput. Astron. and Cosmol. **5**, 2 2018).

[3] M. Paganini, L. de Oliveira, and B. Nachman, CaloGAN: Simulating 3d high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks, Phys. Rev. D **97**, 014021 2018).

[4] K. Schawinski, C. Zhang, H. Zhang, L. Fowler, and G. K. Santhanam, Generative adversarial networks recover features in astrophysical images of galaxies beyond the deconvolution limit, Mon. Not. R. Astron. Soc.: Lett. **467**, L110 (2017).

[5] J. Jumper *et al.*, Highly accurate protein structure prediction with alphafold, Nature (London) **596**, 583 (2021).

[6] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, Nature (London) **549**, 195 (2017).

[7] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum (NISQ) algorithms, Rev. Mod. Phys. **94**, 015004 (2022).

[8] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, Quantum Sci. Technol. **4**, 043001 (2019).

[9] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, A generative modeling approach for benchmarking and training shallow quantum circuits, npj Quantum Inf. **5**, 45 2019).

[10] J.-G. Liu and L. Wang, Differentiable learning of quantum circuit Born machine, Phys. Rev. A **98**, 062324 (2018).

[11] J. Tian, X. Sun, Y. Du, S. Zhao, Q. Liu, K. Zhang, W. Yi, W. Huang, C. Wang, X. Wu, M.-H. Hsieh, T. Liu, W. Yang, and D. Tao, Recent advances for quantum neural networks in generative learning, arXiv:2206.03066.

[12] I. Glasser, R. Sweke, N. Pancotti, J. Eisert, and J. I. Cirac, Expressive power of tensor-network factorizations for probabilistic modeling, Adv. Neural Inf. Process. Syst. **32**, 1498 (2019).

[13] S. Adhikary, S. Srinivasan, J. Miller, G. Rabusseau, and B. Boots, Quantum tensor networks, stochastic processes, and weighted automata, in *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research Vol. 130 (PMLR, 2021), pp. 2080–2088.

[14] M. J. Bremner, R. Jozsa, and D. J. Shepherd, Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy, Proc. R. Soc. A **467**, 459 (2010).

[15] 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).

[16] D. Hangleiter and J. Eisert, Computational advantage of quantum random sampling, arXiv:2206.04079.

[17] B. Coyle, D. Mills, V. Danos, and E. Kashefi, The Born supremacy: Quantum advantage and training of an Ising Born machine, npj Quantum Inf. **6**, 60 (2020).

[18] K. Gili, M. Mauri, and A. Perdomo-Ortiz, Evaluating generalization in classical and quantum generative models, arXiv:2201.08770.

[19] M. S. Rudolph, N. Bashige Toussaint, A. Katabarwa, S. Johri, B. Peropadre, and A. Perdomo-Ortiz, Generation of high-resolution handwritten digits with an ion-trap quantum computer, arXiv:2012.03924.

[20] M. Y. Niu, A. M. Dai, L. Li, A. Odena, Z. Zhao, V. Smelyanskyi, H. Neven, and S. Boixo, Learnability and complexity of quantum samples, arXiv:2010.11983.

[21] M. Kearns, Y. Mansour, D. Ron, R. Rubinfeld, R. E. Schapire, and L. Sellie, On the learnability of discrete distributions, in *Proceedings of the 26th Annual ACM Symposium on Theory of Computing, STOC '94* (ACM, New York, 1994), pp. 273–282.

[22] S. Lloyd, Universal quantum simulators, Science **273**, 1073 (1996).

[23] A. M. Childs and Y. Su, Nearly Optimal Lattice Simulation by Product Formulas, Phys. Rev. Lett. **123**, 050503 (2019).

[24] C. Sünderhauf, D. Pérez-García, D. A. Huse, N. Schuch, and J. I. Cirac, Localization with random time-periodic quantum circuits, Phys. Rev. B **98**, 134204 (2018).

[25] L. Piroli, C. Sünderhauf, and X.-L. Qi, A random unitary circuit model for black hole evaporation, J. High Energy Phys. 04 (2020) 063.

[26] J. Haferkamp, P. Faist, B. T. Kothakonda N, J. Eisert, and N. Yunger Halpern, Linear growth of quantum circuit complexity, Nat. Phys. **18**, 528 (2022).

[27] X. Mi, P. Roushan, C. Quintana, S. Mandrà, J. Marshall, C. Neill, F. Arute, K. Arya, J. Atalaya, R. Babbush *et al.*, Information scrambling in quantum circuit, Science **374**, 1479 (2021).

[28] A. Rocchetto, Stabiliser states are efficiently PAC-learnable, Quantum Inf. Comput. **18**, 541 (2018).

[29] C.-Y. Lai and H.-C. Cheng, Learning quantum circuits of some T gates, IEEE Trans. Inf. Theory **68**,3951 (2022).

[30] M. Yoganathan, A condition under which classical simulability implies efficient state learnability, arXiv:1907.08163.

[31] TV denotes here the total variation distance between two probability distributions.

[32] B. Axelrod, S. Garg, V. Sharan, and G. Valiant, Sample amplification: Increasing dataset size even when learning is impossible, in *Proceedings of the 37th International Conference on Machine Learning*, Proceedings of Machine Learning Research Vol. 119 (PMLR, 2020), pp. 442–451.

[33] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.130.240602 for detailed proofs of the results presented in the main text, which includes Refs. [34–55].

[34] F. G. S. L. Brandáo, W. Chemissany, N. Hunter-Jones, R. Kueng, and J. Preskill, Models of quantum complexity growth, PRX Quantum **2**, 030316 (2021).

[35] C. M. Dawson and M. A. Nielsen, The Solovay-Kitaev algorithm, arXiv:quant-ph/0505030.

[36] A. W. Harrow, B. Recht, and I. L. Chuang, Efficient discrete approximations of quantum gates, J. Math. Phys. (N.Y.) **43**, 4445 (2002).

[37] J. Dehaene and B. De Moor, The Clifford group, stabilizer states, and linear and quadratic operations over GF(2), Phys. Rev. A **68**, 042318 2003).

[38] P. J. S. G. Ferreira, B. Jesus, J. Vieira, and A. J. Pinho, The rank of random binary matrices and distributed storage applications, IEEE Commun. Lett. **17**, 151 (2013).

[39] D. Helmbold, R. Sloan, and M. K. Warmuth, Learning integer lattices, SIAM J. Comput. **21**, 240 (1992).

[40] P. Auer and N. Cesa-Bianchi, On-line learning with malicious noise and the closure algorithm, in *Algorithmic Learning Theory* (Springer Berlin Heidelberg, Berlin, Heidelberg, 1994), Vol. 872, pp. 229–247.

[41] A. Blum, A. T. Kalai, and H. Wasserman, Noise-tolerant learning, the parity problem, and the statistical query model, J. ACM **50**, 506 (2003).

[42] S. Bravyi and D. Maslov, Hadamard-free circuits expose the structure of the Clifford group, IEEE Trans. Inf. Theory **67**, 4546 (2021).

[43] A. Bogdanov and A. Rosen, Pseudorandom functions: Three decades later, in *Tutorials on the Foundations of Cryptography* (Springer, New York, 2017), pp. 79–158.
[44] O. Goldreich, *Foundations of Cryptography: Volume 2, Basic Applications* (Cambridge University Press, Cambridge, 2009).
[45] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, 2011).
[46] R. Jozsa and M. van den Nest, Classical simulation complexity of extended Clifford circuits, Quantum Inf. Comput. **14**, 633 (2014).
[47] S. Anders and H. J. Briegel, Fast simulation of stabilizer circuits using a graph-state representation, Phys. Rev. A **73**, 022334 (2006).
[48] A. Mari and J. Eisert, Positive Wigner Functions Render Classical Simulation of Quantum Computation Efficient, Phys. Rev. Lett. **109**, 230503 (2012).
[49] V. Veitch, C. Ferrie, D. Gross, and J. Emerson, Negative quasi-probability as a resource for quantum computation, New J. Phys. **14**, 113011 (2012).
[50] R. Raussendorf, J. Bermejo-Vega, E. Tyhurst, C. Okay, and M. Zurel, Phase-space-simulation method for quantum computation with magic states on qubits, Phys. Rev. A **101**, 012350 (2020).
[51] H. Pashayan, J. J. Wallman, and S. D. Bartlett, Estimating Outcome Probabilities of Quantum Circuits Using Quasiprobabilities, Phys. Rev. Lett. **115**, 070501 (2015).
[52] S. Bravyi, G. Smith, and J. A. Smolin, Trading Classical and Quantum Computational Resources, Phys. Rev. X **6**, 021043 (2016).
[53] L. Kocia, Improved strong simulation of universal quantum circuits, arXiv:2012.11739v4.
[54] H. Qassim, H. Pashayan, and D. Gosset, Improved upper bounds on the stabilizer rank of magic states, Quantum **5**, 606 (2021).
[55] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, Nature (London) **549**, 172 (2017).
[56] D. Gottesman, The Heisenberg representation of quantum computers, arXiv:quant-ph/9807006.
[57] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, Phys. Rev. A **70**, 052328 2004).
[58] A. Montanaro, Learning stabilizer states by Bell sampling, arXiv:1707.04012.
[59] A. Gollakota and D. Liang, On the hardness of PAC-learning stabilizer states with noise, Quantum **6**, 640 (2022).
[60] S. Aaronson, The learnability of quantum states, Proc. R. Soc. A **463**, 3089 (2007).
[61] S. Bravyi and D. Gosset, Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates, Phys. Rev. Lett. **116**, 250501 (2016).
[62] S. Bravyi, D. Browne, Padraic Calpin, E. Campbell, D. Gosset, and M. Howard, Simulation of quantum circuits by low-rank stabilizer decompositions, Quantum **3**, 181 (2019).
[63] We note that a similar hardness result based on the LPN assumption can be obtained for output distributions of Clifford circuits subject to single-qubit depolarizing noise.
[64] O. Regev, On lattices, learning with errors, random linear codes, and cryptography, J. ACM **56**, 84 2009).
[65] K. Pietrzak, Cryptography from learning parity with noise, in *SOFSEM 2012: Theory and Practice of Computer Science* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2012), pp. 99–114.
[66] M. Zhandry, How to construct quantum random functions, in *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundation of Compuatational Science* (IEEE, New York, 2012), pp. 679–687.
[67] S. Arunachalam, A. B. Grilo, and H. Yuen, Quantum statistical query learning, arXiv:1903.02840.
[68] D. Xiao, Learning to create is as hard as learning to appreciate, in *Proceedings of the COLT 2020 23rd Conference on Learning Theory, Haifa, Israel, 2010*, edited by A. T. Kalai and M. Mohri (Omnipress, 2010), pp. 516–528.
[69] S. Mohamed and B. Lakshminarayanan, Learning in implicit generative models, arXiv:1610.03483.
[70] V. Feldman, A general characterization of the statistical query complexity, in *Proceedings of the 2017 Conference on Learning Theory* (PMLR, 2017), pp. 785–830.
[71] Note that for inverse polynomial accuracy the statistical query oracle can be efficiently simulated from samples. Thus, query-complexity lower bounds with respect to statistical queries yield computational complexity lower bounds with respect to sample queries [72].
[72] I. Diakonikolas, D. M. Kane, and A. Stewart, Statistical query lower bounds for robust estimation of high-dimensional Gaussians and Gaussian mixtures, in *Proceedings of the 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, New York, 2017), pp. 73–84.
[73] M. Kearns, Efficient noise-tolerant learning from statistical queries, J. ACM **45**, 983 (1998).
[74] A. Blum, M. Furst, J. Jackson, M. Kearns, Y. Mansour, and S. Rudich, Weakly learning DNF and characterizing statistical query learning using Fourier analysis, in *Proceedings of the 26th Annual ACM Symposium on Thery of Computing (STOC) '94* (ACM Press, 1994), pp. 253–262.
[75] S. Arunachalam and R. de Wolf, Guest column: A survey of quantum learning theory, ACM SIGACT News **48**, 41 (2017).
[76] Y. Liu, S. Arunachalam, and K. Temme, A rigorous and robust quantum speed-up in supervised machine learning, Nat. Phys. **17**, 1013 (2021).
[77] R. Sweke, J.-P. Seifert, D. Hangleiter, and J. Eisert, On the quantum versus classical learnability of discrete distributions, Quantum **5**, 417 (2021).
[78] S. Jerbi, C. Gyurik, S. Marshall, H. Briegel, and V. Dunjko, Parametrized quantum policies for reinforcement learning, in *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)* (Curran Associates, Inc., New York, 2021), Vol. 34, pp. 28362–28375.

[79] D. Stilck França and R. García-Patrón, Limitations of optimization algorithms on noisy quantum devices, Nat. Phys. **17**, 1221 (2021).

[80] E. R. Anschuetz and B. T. Kiani, Beyond barren plateaus: Quantum variational algorithms are swamped with traps, Nat. Commun. **13**, 7760 (2022).

[81] S. Aaronson and S. Grewal, Efficient learning of non-interacting fermion distributions, arXiv:2102.10458.

[82] S. Aaronson, Shtetl-optimized: Yet more mistakes in papers, https://www.S. aaronson.com/blog/?p=5706, 2021, accessed: 2021-09-27.