

Efficient Learning of Structured Quantum Circuits via Pauli Dimensionality and Sparsity

Given black-box access to an unknown unitary process, how efficiently can one reconstruct its action? This task—known as *unitary process tomography*—is a central problem in quantum information and quantum computation. It has been extensively studied over the past several decades under various models and performance measures (see [20, Section 1.3] for a detailed overview). Recently, Haah, Kothari, O’Donnell, and Tang [20] established that $\Theta(4^n/\varepsilon)$ queries are both necessary and sufficient to learn an unknown n -qubit unitary to ε accuracy in diamond norm.

A complementary challenge is to identify structured subclasses of unitary channels that admit *efficient* learning algorithms in both query complexity and runtime. This direction seeks to characterize which quantum dynamics are tractable to learn, and is practically relevant for the experimental verification of quantum devices. Moreover, efficiently learnable circuit classes cannot be pseudorandom, so such algorithms delineate the boundary of pseudorandomness for quantum circuits, a direction that has received significant recent attention [7, 11, 30–32, 38]. For these reasons, a growing line of work has identified efficiently learnable subclasses of quantum circuits, including constant-depth circuits [24], fermionic matchgates [35], Clifford circuits [29], and other restricted models.

Despite these successes, the existing algorithms appear highly specialized, relying on disparate techniques tailored to individual circuit classes. This leaves open a basic conceptual question:

Is there a unifying structural reason why so many seemingly different classes admit efficient tomography?

In this work, we answer this question by developing a general framework showing that if a generating set of Pauli operators remains low-complexity under conjugation by a unitary, then the unitary can be learned efficiently. This identifies a common structural principle underlying efficient unitary tomography, unifying a broad range of previous results while extending them to substantially more expressive circuit classes. In several cases, our framework also yields improved query and time complexity.

We develop a general framework for unitary tomography based on the structure of a unitary’s *Pauli spectrum*. Using this framework, we obtain polynomial-time learning algorithms for a broad range of circuit classes, including $O(\log \log n)$ -depth circuits, quantum $O(\log n)$ -juntas, near-Clifford circuits, the Clifford hierarchy, fermionic matchgate circuits, and certain compositions thereof. Our results both **unify** previous efficient-learning algorithms under a common structural principle and **generalize** them to substantially more expressive circuit classes than were previously known.

The Unifying Framework

Our framework studies the action of the unknown unitary U on Pauli operators via conjugation, i.e., the map $P \mapsto U^\dagger P U$. Since the Pauli operators form a basis for all n -qubit operators, this action fully determines U . In particular, it suffices to understand the action of U on a *generating set* of Pauli operators. A set G of Pauli operators is said to generate the Pauli group if every Pauli operator can be written as a product of elements of G (up to global phase). Our framework shows that if the operators $\{U^\dagger g U : g \in G\}$ have low-complexity Pauli spectra, then U can be learned efficiently.

We consider two natural notions of complexity of the Pauli spectrum. Any unitary U admits a Pauli expansion $U = \sum_{P \in \{I, X, Y, Z\}^{\otimes n}} \alpha_P P$, where the coefficients $\{\alpha_P\}$ form the Pauli spectrum of U . We say that U has *Pauli dimensionality* k if its spectrum is supported on a subgroup G of size 2^k , and *Pauli sparsity* s if its spectrum is supported on a set S of size s . These notions can be viewed as quantum analogues of Fourier dimensionality and sparsity in Boolean function analysis [12].

Let $\text{dist}_\diamond(\cdot, \cdot)$ denote the diamond distance, i.e., the standard worst-case distance measure between quantum channels.

Theorem 0.1. *Let G be a known generating set of Pauli operators, and let U be an n -qubit unitary such that for every $g \in G$, the operator $U^\dagger g U$ has Pauli dimensionality k (resp. Pauli sparsity s). Then there is an*

algorithm that outputs a unitary V satisfying $\text{dist}_\diamond(U, V) \leq \varepsilon$ with probability at least $1 - \delta$. The algorithm uses $\text{poly}(2^k, n) \cdot \frac{\log(1/\delta)}{\varepsilon}$ (resp. $\text{poly}(s, n) \cdot \frac{\log(1/\delta)}{\varepsilon^2}$) queries.

Thus, [Theorem 0.1](#) reduces unitary learning to identifying a generating set G for which $\{U^\dagger g U : g \in G\}$ have low-complexity Pauli spectra. We therefore obtain learning algorithms for broad classes of quantum circuits by establishing this property.

Among the core technical ingredients underlying [Theorem 0.1](#) are new learning algorithms for learning unitaries with low-complexity Pauli spectra.

Theorem 0.2. *Let U be an n -qubit k -Pauli-dimensional (resp. s -Pauli-sparse) unitary. There is an efficient algorithm that outputs a unitary V satisfying $\text{dist}_\diamond(U, V) \leq \varepsilon$ with high probability using $O(2^k k/\varepsilon)$ (resp. $O(s^2/\varepsilon^2)$) queries. Moreover, any algorithm requires $\Omega(2^k/\varepsilon)$ (resp. $\Omega(s/\varepsilon)$) queries.*

[Theorem 0.2](#) provides the core primitives underlying [Theorem 0.1](#). In particular, the efficiency of the resulting learning algorithms is governed by the complexity of these subroutines, and improvements to these primitives directly translate into improved algorithms for all circuit classes captured by our framework. This gives strong motivation for optimizing the performance of these subroutines.

We make several remarks about [Theorems 0.1](#) and [0.2](#). For unitaries with Pauli dimensionality k , the resulting algorithms are *time efficient* in all cases. For unitaries with Pauli sparsity s , an additional challenge arises: the learning algorithms produce an operator $\hat{A}$ approximating U , which must be rounded to a nearby unitary $\hat{U}$. In general, it is unclear how to perform this rounding while preserving sparsity efficiently, and we leave this as an open problem. Nevertheless, in [??](#), we identify several natural settings in which this rounding step can be implemented efficiently. All circuit classes considered in this work fall into these settings, and hence all of our resulting learning algorithms run in polynomial time.¹

Except for an edge case, the algorithm for learning low-Pauli-dimensional unitaries ([Theorem 0.2](#)) is query-optimal.² When $k = 2n$, so that the Pauli spectrum has full support, [Theorem 0.2](#) recovers the $\Theta(4^n/\varepsilon)$ bound of Haah, Kothari, O’Donnell, and Tang [\[20\]](#). Thus, our result generalizes their optimal tomography algorithm to a broader setting, interpolating between the worst-case regime of arbitrary unitaries and structured subclasses that admit faster learning.

Our work fits into a broader theme in quantum estimation: exploiting structure in the Pauli spectrum of quantum states and circuits. This perspective has led to efficient algorithms for learning and testing quantum states, channels, and unitary processes [\[1–6, 10, 13–16, 18, 21, 22, 25, 27, 33, 34, 41\]](#), and for quantum estimation more generally [\[17, 19\]](#). Our results extend this paradigm to a unified framework for learning structured quantum circuits.

Finally, a recent work due to Honjani and Heidari [\[23\]](#) study the problem of learning an unknown unitary that is promised to be ε -close to an s -Pauli-sparse unitary, where closeness is measured in the ℓ_1 -distance of the Pauli coefficients, i.e., two unitaries $U = \sum_P \alpha_P P$ and $V = \sum_P \beta_P P$ are ε -close if $\sum_P |\alpha_P - \beta_P| \leq \varepsilon$. Their algorithm uses $\tilde{O}(s^6/\varepsilon^4)$ queries. In contrast, our [Theorem 0.2](#) gives more efficient guarantees in the setting of exactly s -Pauli-sparse unitary channels.

Applications

[Theorem 0.1](#) yields efficient learning algorithms for a broad range of quantum circuit classes. At first glance, it may be unclear which circuits satisfy the condition that there exists a generating set G for which $\{U^\dagger g U : g \in G\}$ have low-complexity Pauli spectra. We show that this condition in fact captures a wide variety of natural and well-studied circuit classes, providing a unifying explanation for many previously disparate learning results. Moreover, it extends beyond prior work, yielding efficient algorithms for more expressive classes of circuits than were previously known to be learnable.

¹If one only needs to apply a *quantum channel* close to the unknown unitary, the classical rounding step can be avoided. In particular, given $\hat{A}$, one can approximately implement its polar decomposition using standard techniques, e.g., the algorithm of Quek and Rebentrost [\[37\]](#). This is because $\hat{A}$ is a linear combination of at most s Pauli operators, and thus admits an efficient block-encoding via the Linear Combination of Unitaries (LCU) framework. See [??](#) for details.

²A Pauli subgroup of order 2^k admits a canonical generating set $\{x_1, \dots, x_a, z_1, \dots, z_{a+b}\}$ with $2a + b = k$. The additional factor of k in the $O(2^k k/\varepsilon)$ query complexity of [Theorem 0.2](#) arises only in the regime $a = o(\log b)$. In all other regimes, our algorithm is query optimal. Whether this factor is necessary remains open.

We illustrate this with representative examples and defer full details to the main paper. In particular, [Theorem 0.1](#) recovers prior results on learning arbitrary unitaries [20], the Clifford hierarchy [29], near-Clifford circuits [26, 28], shallow circuits [24], fermionic matchgate circuits [35], the matchgate hierarchy [9], and quantum juntas [6]. Here, “near-Clifford” refers to circuits consisting of Clifford gates together with $O(\log n)$ non-Clifford single-qubit gates, and a quantum k -junta is a unitary acting nontrivially on only k of n qubits.

Importantly, prior works relied on techniques tailored to each circuit class. In contrast, all of these results arise as direct consequences of [Theorem 0.1](#), showing that the low-complexity Pauli spectra condition provides a unifying principle for efficient unitary learning.

Quantum k -juntas An n -qubit quantum k -junta is a unitary that acts nontrivially on only k of the n qubits. We obtain a query-optimal learning algorithm for quantum k -juntas.

Corollary 0.3. *Let U be an n -qubit quantum k -junta. Given query access to U , there is a time-efficient algorithm that outputs V satisfying $\text{dist}_\diamond(U, V) \leq \varepsilon$ with probability at least $1 - \delta$ using $O\left(\frac{4^k}{\varepsilon} \log \frac{1}{\delta}\right)$ queries. This query complexity is optimal up to constant factors.*

Our result yields exponential improvements in both query and time complexity over prior work. Compared to [6], we improve the query complexity quadratically in ε while achieving the stronger guarantee of learning in diamond distance. The algorithm of [24] relies on enumerating all k -local Pauli operators and is therefore efficient only when $k = O(1)$, whereas our algorithm remains efficient for $k = O(\log n)$.

Compositions of shallow and near-Clifford circuits Prior work gives efficient learning algorithms for constant-depth circuits [24] and for Clifford circuits with a small number of non-Clifford gates [28], but these techniques do not extend to their composition.

Using [Theorem 0.1](#), we obtain a unified algorithm that improves on both prior works and extends to compositions of shallow and near-Clifford circuits.

Theorem 0.4. *Let U be an n -qubit unitary of the form $U = QC$ or $U = CQ$, where Q is a depth- d circuit and C is a Clifford circuit augmented with t single-qubit non-Clifford gates. Then U can be learned efficiently for $d = O(\log \log n)$ and $t = O(\log n)$.*

This is the first result that simultaneously handles shallow and near-Clifford circuits, unifying the techniques developed for these settings. Our algorithm is efficient for $d = O(\log \log n)$ and $t = O(\log n)$, which are both provably optimal. In particular, we show that any algorithm requires exponential dependence on t and doubly exponential dependence on d .

Our algorithm remains efficient for depth $O(\log \log n)$ circuits, whereas [24] can only learn constant-depth circuits. Similarly, [28] handles only Clifford circuits augmented with T gates, whereas our algorithm applies to arbitrary single-qubit gates. We note that the class of unitaries covered by [Theorem 0.4](#) includes the first level of the recently introduced *Magic Hierarchy* [36].

Compositions of fermionic matchgates and Clifford circuits. Fermionic matchgates correspond to fermionic Gaussian unitaries under the Jordan–Wigner transformation [39, 40]. Prior work gives efficient learning algorithms for fermionic matchgate circuits [8, 35].

Using [Theorem 0.1](#), we extend these results to compositions with Clifford circuits, analogous to [Theorem 0.4](#).

Theorem 0.5. *Let U be an n -qubit unitary of the form $U = FC$ or $U = CF$, where F is a fermionic matchgate circuit and C is a Clifford circuit. Then U can be learned efficiently using queries $O(n^7/\varepsilon^2)$ and time $O(n^8/\varepsilon^2)$.*

This shows that our framework extends beyond shallow circuits and juntas to other structured models such as fermionic systems, and continues to apply even under composition with Clifford operations.

References

- [1] Srinivasan Arunachalam and Arkopal Dutt. “Polynomial-Time Tolerant Testing Stabilizer States”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 2025. DOI: [10.1145/3717823.3718277](https://doi.org/10.1145/3717823.3718277) (page 2).
- [2] Srinivasan Arunachalam, Arkopal Dutt, Francisco Escudero Gutiérrez, and Carlos Palazuelos. “Learning Low-Degree Quantum Objects”. In: *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*. 2024. DOI: [10.4230/LIPIcs.ICALP.2024.13](https://doi.org/10.4230/LIPIcs.ICALP.2024.13) (page 2).
- [3] Zongbo Bao, Philippe van Dordrecht, and Jonas Helsen. “Tolerant Testing of Stabilizer States with a Polynomial Gap via a Generalized Uncertainty Relation”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 2025. DOI: [10.1145/3717823.3718201](https://doi.org/10.1145/3717823.3718201). URL: <https://doi.org/10.1145/3717823.3718201> (page 2).
- [4] Zongbo Bao and Penghui Yao. “On Testing and Learning Quantum Junta Channels”. In: *Proceedings of Thirty Sixth Conference on Learning Theory*. 2023. URL: <https://proceedings.mlr.press/v195/bao23b.html> (page 2).
- [5] Sitan Chen, Weiyuan Gong, Qi Ye, and Zhihan Zhang. “Stabilizer Bootstrapping: A Recipe for Efficient Agnostic Tomography and Magic Estimation”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 2025. DOI: [10.1145/3717823.3718191](https://doi.org/10.1145/3717823.3718191) (page 2).
- [6] Thomas Chen, Shivam Nadimpalli, and Henry Yuen. “Testing and Learning Quantum Juntas Nearly Optimally”. In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2023. DOI: [10.1137/1.9781611977554.ch43](https://doi.org/10.1137/1.9781611977554.ch43) (pages 2, 3).
- [7] Nai-Hui Chia, Daniel Liang, and Fang Song. “Quantum State and Unitary Learning Implies Circuit Lower Bounds”. In: *Proceedings of Thirty Eighth Conference on Learning Theory*. 2025. URL: <https://proceedings.mlr.press/v291/chia25a.html> (page 1).
- [8] Aria Christensen and Andrew Zhao. *Learning fermionic linear optics with Heisenberg scaling and physical operations*. 2026. arXiv: [2602.05058](https://arxiv.org/abs/2602.05058) [quant-ph]. URL: <https://arxiv.org/abs/2602.05058> (page 3).
- [9] Josh Cudby and Sergii Strelchuk. “Learning Gaussian Operations and the Matchgate Hierarchy”. In: *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 2024. DOI: [10.1109/QCE60285.2024.00026](https://doi.org/10.1109/QCE60285.2024.00026) (page 3).
- [10] Steven T. Flammia and Ryan O’Donnell. “Pauli error estimation via Population Recovery”. In: *Quantum* (2021). DOI: [10.22331/q-2021-09-23-549](https://doi.org/10.22331/q-2021-09-23-549). URL: <https://doi.org/10.22331/q-2021-09-23-549> (page 2).
- [11] Ben Foxman, Natalie Parham, Francisca Vasconcelos, and Henry Yuen. “Random Unitaries in Constant (Quantum) Time”. In: *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*. 2026. DOI: [10.4230/LIPIcs.ITCS.2026.61](https://doi.org/10.4230/LIPIcs.ITCS.2026.61) (page 1).
- [12] Parikshit Gopalan, Ryan O’Donnell, Rocco A. Servedio, Amir Shpilka, and Karl Wimmer. “Testing Fourier Dimensionality and Sparsity”. In: *SIAM Journal on Computing* (2011). DOI: [10.1137/100785429](https://doi.org/10.1137/100785429) (page 1).
- [13] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. “Agnostic Tomography of Stabilizer Product States”. In: *Quantum* (2026). DOI: [10.22331/q-2026-03-13-2027](https://doi.org/10.22331/q-2026-03-13-2027) (page 2).
- [14] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. “Efficient Learning of Quantum States Prepared With Few Non-Clifford Gates”. In: *Quantum* (2025). DOI: [10.22331/q-2025-11-06-1907](https://doi.org/10.22331/q-2025-11-06-1907). URL: <https://doi.org/10.22331/q-2025-11-06-1907> (page 2).
- [15] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. “Improved Stabilizer Estimation via Bell Difference Sampling”. In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 2024. DOI: [10.1145/3618260.3649738](https://doi.org/10.1145/3618260.3649738) (page 2).
- [16] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. “Low-Stabilizer-Complexity Quantum States Are Not Pseudorandom”. In: *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. 2023. DOI: [10.4230/LIPIcs.ITCS.2023.64](https://doi.org/10.4230/LIPIcs.ITCS.2023.64). URL: <https://doi.org/10.4230/LIPIcs.ITCS.2023.64> (page 2).
- [17] Sabee Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. “Lower bounds on the non-Clifford cost of pseudoentanglement”. In: *Phys. Rev. A* (1 2026). DOI: [10.1103/v493-8s1q](https://doi.org/10.1103/v493-8s1q) (page 2).

- [18] David Gross, Sepehr Nezami, and Michael Walter. “Schur–Weyl duality for the Clifford group with applications: Property testing, a robust Hudson theorem, and de Finetti representations”. In: *Communications in Mathematical Physics* (2021). DOI: [10.1007/s00220-021-04118-7](https://doi.org/10.1007/s00220-021-04118-7). URL: <https://doi.org/10.1007/s00220-021-04118-7> (page 2).
- [19] Andi Gu, Salvatore F.E. Oliviero, and Lorenzo Leone. “Magic-Induced Computational Separation in Entanglement Theory”. In: *PRX Quantum* (2 2025). DOI: [10.1103/PRXQuantum.6.020324](https://doi.org/10.1103/PRXQuantum.6.020324) (page 2).
- [20] Jeongwan Haah, Robin Kothari, Ryan O’Donnell, and Ewin Tang. “Query-optimal estimation of unitary channels in diamond distance”. In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. 2023. DOI: [10.1109/FOCS57990.2023.00028](https://doi.org/10.1109/FOCS57990.2023.00028) (pages 1–3).
- [21] Dominik Hangleiter and Michael J. Gullans. “Bell Sampling from Quantum Circuits”. In: *Phys. Rev. Lett.* (2 2024). DOI: [10.1103/PhysRevLett.133.020601](https://doi.org/10.1103/PhysRevLett.133.020601) (page 2).
- [22] Marcel Hinsche and Jonas Helsen. “Single-Copy Stabilizer Testing”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 2025. DOI: [10.1145/3717823.3718169](https://doi.org/10.1145/3717823.3718169). URL: <https://doi.org/10.1145/3717823.3718169> (page 2).
- [23] Zahra Honjani and Mohsen Heidari. *Query Learning Nearly Pauli Sparse Unitaries in Diamond Distance*. 2026. arXiv: [2604.00203](https://arxiv.org/abs/2604.00203) [quant-ph] (page 2).
- [24] Hsin-Yuan Huang, Yunchao Liu, Michael Broughton, Isaac Kim, Anurag Anshu, Zeph Landau, and Jarrod R. McClean. “Learning Shallow Quantum Circuits”. In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 2024. DOI: [10.1145/3618260.3649722](https://doi.org/10.1145/3618260.3649722) (pages 1, 3).
- [25] Vishnu Iyer and Daniel Liang. *Tolerant Testing of Stabilizer States with Mixed State Inputs*. 2025. arXiv: [2411.08765](https://arxiv.org/abs/2411.08765) [quant-ph]. URL: <https://arxiv.org/abs/2411.08765> (page 2).
- [26] Ching-Yi Lai and Hao-Chung Cheng. “Learning Quantum Circuits of Some T Gates”. In: *IEEE Transactions on Information Theory* (2022). DOI: [10.1109/TIT.2022.3151760](https://doi.org/10.1109/TIT.2022.3151760). URL: <https://doi.org/10.1109/TIT.2022.3151760> (page 3).
- [27] Lorenzo Leone, Salvatore F. E. Oliviero, and Alioscia Hamma. “Learning t -doped stabilizer states”. In: *Quantum* (2024). DOI: [10.22331/q-2024-05-27-1361](https://doi.org/10.22331/q-2024-05-27-1361). URL: <https://doi.org/10.22331/q-2024-05-27-1361> (page 2).
- [28] Lorenzo Leone, Salvatore F. E. Oliviero, Seth Lloyd, and Alioscia Hamma. “Learning efficient decoders for quasichotic quantum scramblers”. In: *Phys. Rev. A* (2 2024). DOI: [10.1103/PhysRevA.109.022429](https://doi.org/10.1103/PhysRevA.109.022429) (page 3).
- [29] Richard A. Low. “Learning and Testing Algorithms for the Clifford Group”. In: *Phys. Rev. A* (5 2009). DOI: [10.1103/PhysRevA.80.052314](https://doi.org/10.1103/PhysRevA.80.052314). URL: <https://doi.org/10.1103/PhysRevA.80.052314> (pages 1, 3).
- [30] Chuhan Lu, Minglong Qin, Fang Song, Penghui Yao, and Mingnan Zhao. *Parallel Kac’s Walk Generates PRU*. 2025. arXiv: [2504.14957](https://arxiv.org/abs/2504.14957) [quant-ph] (page 1).
- [31] Fermi Ma and Hsin-Yuan Huang. *A note on pseudorandom unitaries in polylog depth*. 2024. URL: https://hsinyuan-huang.github.io/assets/img/FermiMa_HsinYuanHuang_PolyLogDepthPRUs_against_SubExpAdv.pdf (page 1).
- [32] Fermi Ma and Hsin-Yuan Huang. “How to Construct Random Unitaries”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 2025. DOI: [10.1145/3717823.3718254](https://doi.org/10.1145/3717823.3718254). URL: <https://doi.org/10.1145/3717823.3718254> (page 1).
- [33] Saeed Mehraban and Mehrdad Tahmasbi. “Improved Bounds for Testing Low Stabilizer Complexity States”. In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 2025. DOI: [10.1145/3717823.3718228](https://doi.org/10.1145/3717823.3718228). URL: <https://doi.org/10.1145/3717823.3718228> (page 2).
- [34] Shivam Nadimpalli, Natalie Parham, Francisca Vasconcelos, and Henry Yuen. “On the Pauli Spectrum of QAC^0 ”. In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 2024. DOI: [10.1145/3618260.3649662](https://doi.org/10.1145/3618260.3649662) (page 2).
- [35] Michał Oszmaniec, Ninnat Dangniam, Mauro E.S. Morales, and Zoltán Zimborás. “Fermion Sampling: A Robust Quantum Computational Advantage Scheme Using Fermionic Linear Optics and Magic Input States”. In: *PRX Quantum* (2 2022). DOI: [10.1103/PRXQuantum.3.020328](https://doi.org/10.1103/PRXQuantum.3.020328) (pages 1, 3).
- [36] Natalie Parham. *Quantum circuit lower bounds in the magic hierarchy*. 2025. arXiv: [2504.19966](https://arxiv.org/abs/2504.19966) [quant-ph] (page 3).

- [37] Yihui Quek and Patrick Rebentrost. “Fast algorithm for quantum polar decomposition and applications”. In: *Phys. Rev. Res.* (1 2022). DOI: [10.1103/PhysRevResearch.4.013144](https://doi.org/10.1103/PhysRevResearch.4.013144) (page 2).
- [38] Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang. “Random unitaries in extremely low depth”. In: *Science* (2025). DOI: [10.1126/science.adv8590](https://doi.org/10.1126/science.adv8590) (page 1).
- [39] Barbara M. Terhal and David P. DiVincenzo. “Classical simulation of noninteracting-fermion quantum circuits”. In: *Physical Review A* (2002). DOI: [10.1103/PhysRevA.65.032325](https://doi.org/10.1103/PhysRevA.65.032325). URL: <https://doi.org/10.1103/PhysRevA.65.032325> (page 3).
- [40] Leslie G. Valiant. “Quantum Circuits That Can Be Simulated Classically in Polynomial Time”. In: *SIAM Journal on Computing* (2002). DOI: [10.1137/S0097539700377025](https://doi.org/10.1137/S0097539700377025). URL: <https://doi.org/10.1137/S0097539700377025> (page 3).
- [41] Francisca Vasconcelos and Hsin-Yuan Huang. “Learning shallow quantum circuits with many-qubit gates”. In: *Proceedings of Thirty Eighth Conference on Learning Theory*. 2025. URL: <https://proceedings.mlr.press/v291/vasconcelos25a.html> (page 2).