Ewin Tang. Research statement.

Broadly, my guiding question is: *What tasks can be performed efficiently in a universe governed by quantum mechanics?* This question has implications for both physics and computation. For the latter, quantum computers are poised to reshape the landscape of computation, with downstream effects across society. For the former, many fundamental questions in physics ask what we can do in a quantum world, making them algorithmic at heart: can nature produce exotic quantum phenomena? How can we witness this, either by experiment or by simulation?

Concretely, *I work on quantum algorithms, investigating applications of quantum computing*: my PhD work is on machine learning, and my recent work is on many-body quantum systems. Such systems—think large molecules, superconducting materials, anything where entanglement comes into the picture—are a central subject of computational scrutiny across physics and chemistry. Both applications present bold visions of paradigm-shifting speedups over our usual, 'classical' computers, but proving the existence of such speedups is surprisingly tricky [A15; L+23]. My goal is to pin down formal evidence for what we can actually expect from the quantum computers of the future. Some insights I've made include the following.

1. **Showing that a large class of quantum machine learning algorithms do not admit exponential speedups** [T19; T21; CGLLTW20; CGLLTW22].[1] At the time, a growing body of work had suggested that quantum computers could achieve massive speedups for a broad class of tasks in linear algebra. I created a model of 'quantum-inspired' classical algorithms which matches the performance of these quantum algorithms, showing that these hopes were in fact false. Through this work, we now understand that quantum speedups in machine learning are likely limited to highly specialized data.

2. **Giving the first efficient algorithm for learning the parameters of an unknown quantum system at thermal equilibrium** [BLMT24b].[2] It had been conjectured that no efficient algorithm existed [AA23], roughly meaning that a scientist wishing to identify the underlying mechanics of a system in their lab must perform an exponentially long computation, an intractable barrier. My work shows that the scientist can rest easy, at least in theory.

3. **Showing that 'heat destroys entanglement': quantum systems above a constant temperature exhibit *zero* entanglement** [BLMT24c]. This work started with the algorithmic question of when we can efficiently *simulate* the kind of systems from the aforementioned learning result [BLMT24b], but through our algorithm we uncovered this deeper structural 'law' describing the relationship between entanglement and temperature. This result cuts against the intuitions from a large body of work on entanglement at thermal equilibrium.

Moving forward, I intend to continue working on designing algorithmic approaches for doing science. This kind of work excites me, as it features a rich interplay between fields: from science comes a wealth of impactful questions, and from computer science comes a wealth of algorithmic tools to resolve them. Both sides benefit from the exchange: the right algorithm leads to new discoveries, and scientific inquiry fosters beautiful CS theory. I have found this to be the case for my research in quantum algorithms: working towards this aim has elicited understanding and, sometimes, surprising implications for both physics and computer science. In what follows, I describe my work and future plans in more detail.

[1] My first two papers [T19; T21] were jointly awarded a plenary talk and best student paper at QIP 2020.
[2] This work [BLMT24b] was awarded a plenary talk at QIP 2024.

1 Quantum speed-ups in machine learning?

I began my career studying *quantum machine learning* (QML). The pitch for QML proceeds as follows. By the rules of quantum mechanics, nature implicitly runs massive matrix computations to evolve even simple systems. We might then hope to trick nature into using its linear algebraic processor for our own purposes, like, say, linear regression [HHL09] (or principal component analysis [LMR14], semi-definite programming [BS17], etc.) far faster than we can on classical computers.

QML took off in the 2010's, with a range of proposals along these lines appearing to achieve an *exponential speedup*, the largest speedup possible for quantum computation [BWPRWL17; P18]. If true, these speedups would be far more general than any other kind of speedup in its class. The omnipresence of linear algebra in industry means that such a seismic speedup could revolutionize society in the same way classical computation did. However, analyzing these claims is not straightforward, as these QML algorithms typically assume that input vectors and matrices are already 'loaded' into the quantum computer as quantum states [A15], and so do not take into account the cost of loading. In essence, these proposals are assuming that their data is in a sort of data structure, and then not counting the (otherwise ruinously large) pre-processing cost of creating the data structure. It wasn't known how to compare this to classical algorithms, leaving us in the dark on the true nature of these speedups. The most promising QML algorithm was one for low-rank approximation from Kerenidis and Prakash [KP17], as it made its loading step explicit enough that one could conceivably prove a genuine exponential speedup.

Results. I failed to prove a speedup; instead, I discovered that none exists: the algorithm of Kerenidis and Prakash [KP17] admits *no* exponential speedup [T19]. This led to a sea change in the field. I then invented a framework for comparing QML algorithms to classical algorithms [T21], by defining a data structure for classical computers which is analogous to quantum computers receiving their input preloaded into quantum states. In this and subsequent joint work with collaborators [CGLLTW20; CGLLTW22; BT24; T22], I showed that this data structure made classical computers about as strong as quantum in many settings, ultimately concluding that nine of the most notable QML algorithms do not admit exponential speedups. This work radically reshaped our understanding of exponential speedups in quantum computing, changing what could have been a wide lane for advantage into mostly a dead-end [A22; D+23; P22], with remaining proposals ekeing out advantage only under exceptional, specialized circumstances [KLLP19; YSSK20; B+24b].

This work leverages a crucial issue in QML: though quantum computers can manipulate data extremely efficiently, extracting features from the computer is difficult. Often, QML proposals assume that input data is low-rank, since then there is a signal large enough to extract; my work shows that this signal must also be large enough to be present in extremely small subsamples of the data. The natural choice of classical data structure allows for efficient subsampling, specifically what's called ℓ^2 importance sampling in the classical 'sketching' literature [FKV04; KV17]. I developed a quantum-inspired theory of importance sampling to simulate the techniques in QML.

My work also provides insight on classical algorithms, demonstrating that importance sampling possesses unexpected strength. We show that, if given a low-rank matrix A in the aforementioned data structure, we can simulate the same data structure for $p(A)$, a polynomial applied to A, in essentially *constant time*, much faster than reading the full input. This follows from natural compositional properties of the data structure, that it is, in a sense, closed under multiplication and addition. These properties are very similar to the QML data structures which form the basis of the well-known meta-algorithm, 'quantum singular value transformation' (QSVT) [GSLW19]. In the sketching literature, compositionality is a familiar but perhaps under-appreciated property; my work presents a unique perspective by showing that importance sampling, unlike the other tools

common in that literature, enjoys a vast compositional reach in merely constant time.

2 Learning and simulating big quantum systems

My other line of work is on another promising application of quantum computing: understanding many-body quantum systems. This is the arena in which Feynman originally dreamed quantum computers could find success. The demand for algorithmic advances here is larger than ever, as this type of computation is the bottleneck for resolving many basic questions across physics and chemistry. Determining the behavior of electrons in the reactions studied in quantum chemistry [C24], mapping out the phase transitions studied in condensed matter [L+15], simulating the Standard Model physics which occur in particle accelerators [D+23]—all these are versions of this task, and all still elude our best supercomputers.

For these systems, we demand efficient algorithms for answering the questions of the scientist who studies them. Such a scientist seeks to understand the relationship between the macroscopic behavior of a system and the underlying mechanisms which drive it. So, we ask the complementary questions of learning and simulation: given measurements of an unknown system in a lab, can I efficiently learn its mechanics; and given a description of a system's mechanics, can I efficiently simulate it on a quantum computer?

Beyond being relevant to practice, these directions comprise a new algorithmic perspective for tackling questions in basic physics: instead of asking what is possible, we ask what is possible *efficiently*. Everything we observe, from the chemicals in our beakers to the black holes in our telescopes, is the output of a universe which was given fairly little time to prepare it [L02]. It follows, then, that any phenomenon we hope to discover or reproduce must be one which can be produced efficiently [A05]. In this sense, simulation and learning ask, *what kinds of systems can nature prepare*; and *what properties of systems are "feelable", in the sense that we can detect them through experiment?* I have thought about this question across a variety of basic systems [HKOT23; BLMT24c; B+24a; PTTW24], but in this statement I'd like to highlight one of the most natural versions of this problem, where we suppose the system has reached equilibrium at a certain temperature, forming a 'Gibbs state'.

A quantum system is modeled by a Hamiltonian H, which is an exponentially large matrix encoding local interaction forces.[3] This Hamiltonian governs everything about a system, from its dynamics, i.e. how it evolves from an initial state, to various types of equilibrium states which are fixed under the dynamics, including its Gibbs state. For applications, we are interested in how the *macroscopic* properties of the Gibbs state, like its energy and phase of matter, changes as we vary *atomic-scale* interaction strengths. Simulation and learning of Gibbs states is about going between the two; both tasks are faced by experimentalists [KBEVZ21].

Results. On the learning side, my work culminates in an algorithm for learning a Hamiltonian from copies of Gibbs states at any temperature [HKT22; BLMT24b]. Importantly, we can still efficiently learn the forces driving the quantum systems of greatest interest to physicists—ones which are cold enough to have long-range entanglement and exhibit exotic quantum phenomena. Though this task is easy for classical systems, the presence of entanglement prevents classical algorithms from extending naturally. In fact, some had suggested that entanglement makes learning, and the physicist's job, inherently computationally difficult [AA23]. To the contrary, we show that the

[3] For example, in models like the Ising model, there are n sites which correspond to atoms in a lattice, and the $2^n \times 2^n$-sized Hamiltonian describes something like a magnetic force, which dictates that atoms which are nearby in the lattice want to be magnetically aligned. For the purposes of this statement, we only consider systems which are locally interacting, meaning that the Hamiltonian only contains interactions between nearby particles.

entanglement can be tamed with a deft application of ideas from quantum information and convex optimization.

On the simulation side, my work [BLMT24a] shows that Gibbs states above a critical temperature can be simulated efficiently; and, surprisingly, our proof that the algorithm succeeds also proves that such states have *zero* entanglement. This result came as a shock. Prior intuition, informed by a large body of literature on entanglement in Gibbs states [A23], proceeds as follows. Entanglement is a (quantum) type of correlation, and we would expect a system at high temperature to be fairly uncorrelated. Indeed, we can prove bounds on the amount of correlation at high temperature. But these systems always have at least some correlation, so we would expect that the amount of entanglement is also never zero.

Nevertheless, we demonstrate the 'death' of entanglement at high temperature. This is a striking instance in which a structural result about physics was discovered via an algorithmic technique. Our novel technique is an adaptation of the sampling-to-counting reduction [SJ89], the classic method for sampling from an n-bit distribution by sampling one bit at a time, recursively splitting it into a marginal and a posterior. A priori, using this seems hopeless, since such a splitting cannot be done in quantum systems due to entanglement. However, we attempt the strategy anyways, and find that we can get a strong enough control over correlations between parts of the system to make this approach succeed, proving lack of entanglement in the process. We anticipate that this approach will have purchase for understanding entanglement in other settings.

3 Future directions

The limits of scientific inquiry are often computational: when experiments fail, we turn to simulation to understand the world. My sights are set toward developing algorithms to advance science. In addition to quantum algorithms, these include computational methods based on classical numerical linear algebra [S11; M19], machine learning, and statistics [D08; RSS19]. Though I can't predict what specific directions will bear fruit in the future, I expect that looking at science through the lens of computation will always generate insight.

The following are specific directions I intend to pursue.

More to do in big quantum systems. I'm excited to continue investigating problems in many-body quantum systems. In particular, I think simulation is the most promising application of quantum computing. No-go results had stalled prior investigations of this, but recent results have opened a new angle of attack: Monte Carlo-style methods for preparing Gibbs states, i.e. quantum systems at thermal equilibrium [CKBG23]. A flurry of recent results demonstrates that a variety of standard algorithms for the classical version of this problem lift to corresponding quantum algorithms. However, we have no idea if these will actually work, since we have almost no bounds on the mixing times of these algorithms for interesting systems. So, I would like to build a theory of *mixing times of quantum systems*.

A first, tractable step would be to 'quantize' the classical Markov chain theory to whatever extent possible, but my aim is to go beyond this theory to tackle uniquely quantum phenomena. I will describe just one such question of several: the physics of *quantum thermalization*. Quantum mechanics is reversible, meaning that time-reversed quantum mechanics is still quantum mechanics. However, we model large systems with thermodynamics, which is irreversible, since entropy increases with time. How do we deal with the fundamental incongruity between these two theories? My work on high-temperature Gibbs states has already been used to give partial answers to this question [H24; PC24]; however, we have far from a complete picture. I hope to build mathematical frameworks to reconcile this incongruity, and show that an arrow of time can emerge from pure quantum mechanics.

Another open question: can we develop an algorithm, quantum or classical, to find high-temperature superconductors? Our current understanding of this problem is very minimal [HKTAP22], but it is more within grasp than it might seem. My past work is on relating Hamiltonians to their Gibbs states; the difference is here, we want to optimize over a class of Hamiltonians to find a Gibbs state satisfying some efficiently checkable property, i.e. superconductivity.

Complexity theory in a quantum world. To build a theory of what is possible with quantum computation, we need an idea of what is impossible with quantum computation. Recent breakthroughs have made clear that we still lack understanding of even basic questions about quantum hardness: are "fully quantum" problems fundamentally harder than any classical problem [LMW24]? Can we build cryptography from the hardness of random quantum circuits, even when P = NP [KQST23]?

My particular interest is in *Hamiltonian complexity*, characterizing the hardness of Hamiltonian problems. The landmark question in this domain is the quantum PCP conjecture [AAV13], which conjectures that solving Hamiltonian problems approximately is, in the worst case, as hard as solving them exactly. However, I'm interested broadly in how to robustly embed properties in the solutions of Hamiltonian problems. Certain kinds of robust properties are what physicists call phases of matter, and so have import on condensed matter physics and related engineering questions, like on the existence of self-correcting quantum memory. I'd love to explore these connections further, to understand what is possible in a quantum world.

References

My research

[B+24a] A. Bakshi, J. Bostanci, W. Kretschmer, Z. Landau, J. Li, A. Liu, R. O'Donnell, and E. Tang. "Learning the closest product state". In submission. 2024. arXiv: 2411.04283 [quant-ph].

[BLMT24a] A. Bakshi, A. Liu, A. Moitra, and E. Tang. "High-temperature Gibbs states are unentangled and efficiently preparable". 2024. arXiv: 2403.16850 [quant-ph].

[BLMT24b] A. Bakshi, A. Liu, A. Moitra, and E. Tang. "Learning quantum Hamiltonians at any temperature in polynomial time". In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 2024. arXiv: 2310.02243 [quant-ph].

[BLMT24c] A. Bakshi, A. Liu, A. Moitra, and E. Tang. "Structure learning of Hamiltonians from real-time evolution". 2024. arXiv: 2405.00082 [quant-ph].

[BT24] A. Bakshi and E. Tang. "An improved classical singular value transformation for quantum machine learning". In: *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2024. arXiv: 2303.01492 [quant-ph].

[CGLLTW20] N.-H. Chia, A. Gilyén, H.-H. Lin, S. Lloyd, E. Tang, and C. Wang. "Quantum-inspired algorithms for solving low-rank linear equation systems with logarithmic dependence on the dimension". In: *31st International Symposium on Algorithms and Computation (ISAAC 2020)*. Published version of [CLW18] and [GLT20]. 2020.

[CGLLTW22] N.-H. Chia, A. P. Gilyén, T. Li, H.-H. Lin, E. Tang, and C. Wang. "Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning". In: *Journal of the ACM* (2022). arXiv: 1910.06151 [cs.DS].

[GLT20] A. Gilyén, S. Lloyd, and E. Tang. "Quantum-inspired low-rank stochastic regression with logarithmic dependence on the dimension". 2020. arXiv: 1811.04909 [cs.DS].

[HKOT23] J. Haah, R. Kothari, R. O'Donnell, and E. Tang. "Query-optimal estimation of unitary channels in diamond distance". In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. 2023. arXiv: 2302.14066 [quant-ph].

[HKT22] J. Haah, R. Kothari, and E. Tang. "Optimal learning of quantum Hamiltonians from high-temperature Gibbs states". In: *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. 2022. arXiv: 2108.04842 [quant-ph].

[PTTW24] A. Pelecanos, X. Tan, E. Tang, and J. Wright. *Beating full state tomography for unentangled spectrum estimation*. In submission. 2024.

[T19] E. Tang. "A quantum-inspired classical algorithm for recommendation systems". In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing - STOC 2019*. 2019. arXiv: 1807.04271 [cs.IR].

[T21] E. Tang. "Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions". In: *Physical Review Letters* (6 2021). arXiv: 1811.00414 [cs.IR].

[T22] E. Tang. "Dequantizing algorithms to understand quantum advantage in machine learning". In: *Nature Reviews Physics* (2022).

Other references

[A05] S. Aaronson. "Guest column: NP-complete problems and physical reality". In: *ACM SIGACT News* (2005). arXivh: quant-ph/0502072.

[A15] S. Aaronson. "Read the fine print". In: *Nature Physics* (2015).

[A22] S. Aaronson. "How much structure is needed for huge quantum speedups?" 2022. arXiv: 2209.06930 [quant-ph].

[AAV13] D. Aharonov, I. Arad, and T. Vidick. "Guest column: the quantum PCP conjecture". In: *ACM SIGACT News* (2013). arXiv: 1309.7495 [quant-ph].

[A23] Á. M. Alhambra. "Quantum many-body systems in thermal equilibrium". In: *PRX Quantum* (2023). arXiv: 2204.08349 [quant-ph].

[AA23] A. Anshu and S. Arunachalam. "A survey on the complexity of learning quantum states". In: *Nature Reviews Physics* (2023). arXiv: 2305.20069 [quant-ph].

[B+24b] D. W. Berry, Y. Su, C. Gyurik, R. King, J. Basso, A. D. T. Barba, A. Rajput, N. Wiebe, V. Dunjko, and R. Babbush. "Analyzing prospects for quantum advantage in topological data analysis". In: *PRX Quantum* (2024). arXiv: 2209.13581 [quant-ph].

[BWPRWL17] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd. "Quantum machine learning". In: *Nature* (2017). arXiv: 1611.09347.

[BS17] F. G. Brandao and K. M. Svore. "Quantum speed-ups for solving semidefinite programs". In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. 2017. arXiv: 1609.05537 [quant-ph].

[C24] G. K.-L. Chan. "Quantum chemistry, classical heuristics, and quantum advantage". In: (2024). arXiv: 2407.11235 [quant-ph].

[CKBG23] C.-F. Chen, M. J. Kastoryano, F. G. S. L. Brandão, and A. Gilyén. "Quantum thermal state preparation". 2023. arXiv: 2303.18224 [quant-ph].

[CLW18] N.-H. Chia, H.-H. Lin, and C. Wang. "Quantum-inspired sublinear classical algorithms for solving low-rank linear systems". 2018. arXiv: 1811.04852 [cs.DS].

[D+23] A. M. Dalzell, S. McArdle, M. Berta, P. Bienias, C.-F. Chen, A. Gilyén, C. T. Hann, M. J. Kastoryano, E. T. Khabiboulline, A. Kubica, G. Salton, S. Wang, and F. G. S. L. Brandão. "Quantum algorithms: a survey of applications and end-to-end complexities". 2023. arXiv: 2310.03011 [quant-ph].

[D08] P. Diaconis. "The Markov chain Monte Carlo revolution". In: *Bulletin of the American Mathematical Society* (2008).

[FKV04] A. Frieze, R. Kannan, and S. Vempala. "Fast Monte-Carlo algorithms for finding low-rank approximations". In: *Journal of the ACM* (2004).

[GSLW19] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe. "Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics". In: *Proceedings of the 51st ACM Symposium on the Theory of Computing (STOC)*. 2019. arXiv: 1806.01838.

[HHL09] A. W. Harrow, A. Hassidim, and S. Lloyd. "Quantum algorithm for linear systems of equations". In: *Physical Review Letters* (15 2009).

[HKTAP22] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill. "Provably efficient machine learning for quantum many-body problems". In: *Science* (2022). arXiv: 2106.12627 [quant-ph].

[H24] Y. Huang. "Random product states at high temperature equilibrate exponentially well". In: (2024). arXiv: 2409.08436 [cond-mat.stat-mech].

[KV17] R. Kannan and S. Vempala. "Randomized algorithms in numerical linear algebra". In: *Acta Numerica* (2017).

[KLLP19] I. Kerenidis, J. Landman, A. Luongo, and A. Prakash. "q-means: a quantum algorithm for unsupervised machine learning". In: *Advances in Neural Information Processing Systems*. 2019. arXiv: 1812.03584.

[KP17] I. Kerenidis and A. Prakash. "Quantum recommendation systems". In: *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*. 2017. arXiv: 1603.08675.

[KBEVZ21] C. Kokail, R. van Bijnen, A. Elben, B. Vermersch, and P. Zoller. "Entanglement Hamiltonian tomography in quantum simulation". In: *Nature Physics* (2021). arXiv: 2009.09000 [quant-ph].

[KQST23] W. Kretschmer, L. Qian, M. Sinha, and A. Tal. "Quantum cryptography in Algorithmica". In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 2023. arXiv: 2212.00879 [quant-ph].

[L+15] J. P. F. LeBlanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, E. Kozik, X.-W. Liu, A. J. Millis, N. V. Prokof'ev, M. Qin, G. E. Scuseria, H. Shi, B. V. Svistunov, L. F. Tocchio, I. S. Tupitsyn, S. R. White, S. Zhang, B.-X. Zheng, Z. Zhu, and E. Gull. "Solutions of the two-dimensional Hubbard model: benchmarks and results from a wide range of numerical algorithms". In: *Physical Review X* (2015). arXiv: 1505.02290 [cond-mat.str-el].

[L+23] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu, M. Kastoryano, R. Babbush, J. Preskill, D. R. Reichman, E. T. Campbell, E. F. Valeev, L. Lin, and G. K.-L. Chan. "Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry". In: *Nature Communications* (2023). arXiv: 2208.02199 [physics.chem-ph].

[L02] S. Lloyd. "Computational capacity of the universe". In: *Physical Review Letters* (2002). arXiv: quant-ph/0110141.

[LMR14] S. Lloyd, M. Mohseni, and P. Rebentrost. "Quantum principal component analysis". In: *Nature Physics* (2014). arXiv: 1307.0401 [quant-ph].

[LMW24] A. Lombardi, F. Ma, and J. Wright. "A one-query lower bound for unitary synthesis and breaking quantum cryptography". In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 2024. arXiv: 2310.08870 [quant-ph].

[M19] P.-G. Martinsson. *Fast direct solvers for elliptic PDEs*. 2019.

[P22] A. Parshall. "After a Quantum Clobbering, One Approach Survives Unscathed". In: *Quanta Magazine* (2022). URL: https://www.quantamagazine.org/after-a-classical-clobbering-a-quantum-advantage-remains-20221207/.

[PC24] S. Pilatowsky-Cameo and S. Choi. "Quantum thermalization of translation-invariant systems at high temperature". 2024. arXiv: 2409.07516 [quant-ph].

[P18] J. Preskill. "Quantum computing in the NISQ era and beyond". In: *Quantum* (2018). arXiv: 1801.00862.

[RSS19] P. Raghavendra, T. Schramm, and D. Steurer. "High dimensional estimation via sum-of-squares proofs". In: *Proceedings of the International Congress of Mathematicians (ICM 2018)*. 2019. arXiv: 1807.11419 [cs.DS].

[S11] U. Schollwöck. "The density-matrix renormalization group in the age of matrix product states". In: *Annals of Physics* (2011).

[SJ89] A. Sinclair and M. Jerrum. "Approximate counting, uniform generation and rapidly mixing Markov chains". In: *Information and Computation* (1989).

[YSSK20] H. Yamasaki, S. Subramanian, S. Sonoda, and M. Koashi. "Learning with optimized random features: Exponential speedup by quantum machine learning without sparsity and low-rank assumptions". In: *Advances in Neural Information Processing Systems*. 2020. arXiv: 2004.10756 [quant-ph].