I seek to understand the ultimate limits of quantum information processing, primarily in *inference* and *algorithms*. My research identifies fundamental barriers arising in these quantum tasks and the resources that can be used to circumvent them—entanglement [**L**T26], randomness [H**L**W26; H**L**25], and time-reversed access [ACCG**L**26], for example—leading either to more efficient procedures, or else to a theory that rules out further improvements on information-theoretic grounds.

Viewing these problems through the lens of learning theory, with an emphasis on sample and query complexity, has produced valuable insights. That such a perspective should be fruitful for inference-based tasks is hardly surprising. A more striking feature revealed by this lens, however, is that progress in the above two areas is intertwined: inferring quantum properties has algorithmic consequences as well. In my own work, answering a question inspired by quantum sensing also gives tight bounds on robust quantum search [ACCG**L**26].

Going forward, I aim to strengthen these connections between inference and algorithms and further develop my research on quantum advantages in learning problems as well as randomized strategies for quantum and classical algorithms. There is ample work to be done in laying strong foundations for these areas, and the theoretical impetus is compelling in its own right. At the same time, quantum computers are poised to make a significant impact in the world, for example, by breaking encryption. Whether this technology's final legacy is constructive will depend on current research into the full range of possibilities for how these devices could benefit fundamental science, industry, and ultimately society as a whole. My research program is based on both considerations.

Inference. Much of my work revolves around the question: How can we infer properties of a quantum system accurately and efficiently from experiments? This is the quantum analogue (and natural generalization) of the central problem in computational learning theory, replacing distributions with density matrices. It is also of direct relevance to practical settings throughout quantum science and in efforts to build fault-tolerant quantum devices.

For example, in many applications, such as quantum sensing, the goal is to do this using the natural dynamics of the system. To model these experiments in a rigorous manner, we may turn to the well-studied task of Hamiltonian property testing. If one can evolve forward and backward in time, then the Heisenberg limit—evolution time inversely proportional to the desired precision—is known to be achievable [KL25; ST26] on quantum devices, enabling beyond-classical precisions for probing many-body systems; however, access to time-reversed dynamics is not always realistic.

In [ACCG**L**26], we showed *time reversal is necessary for optimal Hamiltonian property testing*. This also proves an algorithmic result; namely, inverse access is necessary to retain any asymptotic speedup in quantum search with bounded-error oracles: the famed Grover speed-up is completely destroyed if the oracle is not perfect, even when it is coherent. Hence, natural, forward-time dynamics is a barrier for probing quantum systems and implementing robust quantum algorithms. Our proof developed a new Fourier-analytic approach to quantum lower bounds which works naturally in continuous time, distinguishes forward from inverse queries, and goes beyond the celebrated compressed oracle method as used in the discrete-query result in [TW26].

Another important inference task on the path to fault-tolerant quantum computing is to determine static properties, such as whether a quantum computation works as intended. Can we estimate such properties without incurring a cost scaling with the dimension, which is exponential in the size of the device? To address this, I helped introduce *random dimension reduction* [**L**T26], which simultaneously reduces the dimensions of many, potentially distinct quantum states while preserving certain symmetric properties[1]. This replaces dimension with rank in quantum learning tasks, giving the best known upper bounds for estimating basic quantities like the trace distance, fidelity, and relative entropies. Moreover,

[1]Those invariant under the tensor power action of an isometry.

the core insight that one may replace dimension with rank in a black-box manner applies just as well to any future improvements to these algorithms.

Remarkably, such dimension independence is impossible in general with single-copy measurements, even for pure states; thus, our work shows ambient dimension is a barrier that can be circumvented using entangling measurements in quantum learning tasks. Thanks to the prior work [PSTW25], dimension reduction also completes an elementary, self-contained analysis of sample-optimal tomography (the quantum analogue of learning a distribution to within some error in total variation distance), avoiding a description in terms of Schur polynomials [OW16; HHJWY17].

By exploiting structure in physically realistic systems, such as ground states of local Hamiltonians, we can connect static properties back to physics. My work on testing tree tensor network states [LL26] uses representation theory to establish nearly tight sample complexity bounds for this task, including nearly linear lower bounds in the number of sites for matrix product states with perfect completeness and sufficiently large bond dimension, resolving questions posed in [SW22; Aar+24]. A key insight, which serves as a precursor to the recently introduced and widely applicable random purification channel [TWZ26], is that one can discard half the registers without loss of generality for learning bipartite entanglement (see also [SW22]).

Algorithms. Apart from algorithmic implications of inference-based tasks, my work also studies how to circumvent barriers in quantum algorithms and classical simulation more directly.

For example, one prominent barrier to fast classical simulation of many-body systems is finite memory. A state-of-the-art approach using matrix product states (MPSs) is to truncate bond dimensions, i.e., keep the top singular values of a matrix, which is optimal among deterministic approximations. Can one surpass this barrier using randomness? It turns out the answer is yes, by as much as a quadratic factor in the accuracy. In [HLW26] we give *efficient algorithms for finding the optimal randomized approximation to a quantum state* in trace distance, and for sampling from the required distributions. Along the way, we developed improved classical algorithms for inference and sampling from maximum-entropy distributions. This enables our method to be used with no asymptotic increase in runtime over the SVD required to form an MPS in the first place. Our ideas have also inspired randomized compilation techniques that reduce gate complexity for state preparation [WZYZ26].

Another example of the utility of randomness lies in quantum algorithms. In many physical platforms, entangling gates between spatially separated qubits can be challenging. In [HL25] we gave *optimal algorithms for circuit cutting*—replacing entangling gates with random local operations—together with matching information-theoretic lower bounds on the resulting overhead. This improves prior methods [PHOW20; CSTWZ21] for "clustered" Hamiltonian simulation using many small quantum devices, and similar techniques have recently been demonstrated on real quantum hardware [Car+24].

Future directions. I now describe my current plans for future work in more detail before situating them in a broader context and discussing my vision for long-term research.

Nature as an oracle. Though historically employed for algorithmic and complexity-theoretic purposes, one may also view oracles as models for probing unknown physical processes, as in learning and sensing. I intend to build on our results [ACCGL26] to explore the landscape of possible quantum advantages in this setting, and the resources required to attain them.

As a concrete example, for general Hamiltonians, we have established [ACCGL26] that scaling with the dimension is necessary to achieve Heisenberg-limited precision in the absence of time-reversed dynamics. By contrast, for local Hamiltonians, forward evolution interleaved with simple quantum control suffices [Blu+26]. Where between these extremes does time reversal become necessary? A natural parameter is the *sparsity* of the Hamiltonian in the Pauli basis: known protocols for simulating negative time evolution from forward dynamics are efficient for Hamiltonians supported on a known, small set of Pauli operators [OKSM24], which suggests that sparsity, rather than dimension, may govern the cost. This was also posed as an open

question for learning in [LTW26]. More generally, I aim to determine which physically relevant structure makes time reversal useful.

Another direction is in trying to push these separations further: super-quadratic advantages from inverse queries are already known for estimating out-of-time-order correlators [CSM23] and for learning the locality structure of shallow circuits [SHH25]. Can such separations be established for natural Hamiltonian learning tasks in the continuous-time setting, and do they persist when other resources, such as control, are available?

Experimental imperfections also motivate the following broader question, which I will continue to explore: when can quantum algorithms be made robust to inexact oracles? Prior work shows the answer is delicate [HMD03; RS08; ACCGL26].

Randomness & quantum algorithms. While we are used to thinking of uniquely quantum phenomena—for example, entanglement—as resources for speeding up classical tasks, several recent examples [GKZ25; Cam19; HL25] flip the script: classical randomness can be a resource for quantum problems.

I will leverage insights from my results on randomized truncation [HLW26] to add to this line of work. Namely, in seeking the best approximations to a quantum object, we are often solving a two-player game, for which mixed strategies can outperform deterministic ones. With the right structure, as we show in [HLW26], this game can be solved, yielding the optimal approximation. The same perspective extends naturally from states to quantum channels, and I will investigate scenarios where these problems can be solved efficiently, yielding optimal randomized implementations and provable reductions in quantum computational resources.

There is also precedent for randomized methods in ground state preparation[2], which is one of the key challenges in efficient end-to-end quantum simulation. A compelling avenue is to optimize over randomly altered Hamiltonians, as recently suggested in [Ans26], and to compare this approach with Gibbs sampling and dissipative ground state preparation. Understanding when such randomized strategies offer an advantage is another promising direction that I plan to pursue.

Broader research goals. The idea that one may reason about information and physics independently has long since been discarded. On the one hand, basic questions in quantum physics are often, in essence, about inference from data: for example, how precisely can we measure a magnetic field? On the other, exploiting quantum phenomena unlocks new capabilities for computation, sensing, and communication. These quantum technologies have already had a significant impact on science—e.g., for detecting gravitational waves—and are marching toward broader applicability. My research will aim to make progress across these areas: inference and algorithms solve complementary problems that belong in the same program.

To do so effectively, I will continue to develop a theory of quantum computing which does not forget about its physical context and the available resources. My work has already shown that this results in not only a more comprehensive and compelling theory, but also one that reveals phenomena relevant to computational models familiar to computer science. For instance, studying physically motivated restrictions on continuous time evolution uncovers fundamental limitations on quantum search.

Despite existing work on quantum algorithms, quantum sensing, classical simulation, and related subjects, opportunities remain to connect these areas through such a systematic study of computational resources. In the future, I hope to foster collaborations between theoretical computer science and quantum physics, leading to insights within both fields. I expect that strengthening these connections will lead to new algorithmic principles, and that these will in turn allow us to recognize new scientific and technological capabilities for quantum computing that we have yet to imagine. This is the ultimate ambition of my research.

[2] For example, randomizing the interpolation path [Far+11] in the adiabatic algorithm, or the evolution times [BKS09].

References

My research

[ACCG**L**26] Richard R. Allen, Matthias C. Caro, Yanlin Chen, Tom Gur, and **Angus Lowe**. *The power of time reversal in Hamiltonian property testing*. 2026. arXiv: 2610.04996 [quant-ph] (pages 1–3).

[H**L**25] Aram W. Harrow and **Angus Lowe**. "Optimal Quantum Circuit Cuts with Application to Clustered Hamiltonian Simulation". In: *PRX Quantum* 6.1 (2025), p. 010316. DOI: 10.1103/PRXQuantum.6.010316 (pages 1–3).

[H**L**W26] Aram W. Harrow, **Angus Lowe**, and Freek Witteveen. *Randomized Truncation of Quantum States*. Jan. 2026. arXiv: 2510.08518 [quant-ph] (pages 1–3).

[L**L**26] Benjamin Lovitz and **Angus Lowe**. "Nearly Tight Bounds for Testing Tree Tensor Network States". In: *IEEE Transactions on Information Theory* 72.5 (2026), pp. 3074–3097. DOI: 10.1109/TIT.2026.3670320 (page 2).

[**L**N25] **Angus Lowe** and Ashwin Nayak. "Lower Bounds for Learning Quantum States with Single-Copy Measurements". In: *ACM Transactions on Computation Theory* 17.1 (2025), pp. 1–42. DOI: 10.1145/3717450.

[**Low**+21] **Angus Lowe**, Max Hunter Gordon, Piotr Czarnik, Andrew Arrasmith, Patrick J. Coles, and Lukasz Cincio. "Unified Approach to Data-Driven Quantum Error Mitigation". In: *Physical Review Research* 3.3 (2021), p. 033098. DOI: 10.1103/PhysRevResearch.3.033098.

[**Low**+23] **Angus Lowe**, Matija Medvidović, Anthony Hayes, Lee J. O'Riordan, Thomas R. Bromley, Juan Miguel Arrazola, and Nathan Killoran. "Fast Quantum Circuit Cutting with Randomized Measurements". In: *Quantum* 7 (2023), p. 934. DOI: 10.22331/q-2023-03-02-934. arXiv: 2207.14734 [quant-ph].

[**L**T26] **Angus Lowe** and Xinyu Tan. *Random Dimension Reduction and Learning Symmetric Properties of Quantum States*. June 2026. arXiv: 2606.23592 [quant-ph] (page 1).

Other references

[Aar+24] Scott Aaronson, Adam Bouland, Bill Fefferman, Soumik Ghosh, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. "Quantum Pseudoentanglement". In: *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Ed. by Venkatesan Guruswami. Vol. 287. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, 2:1–2:21. DOI: 10.4230/LIPIcs.ITCS.2024.2 (page 2).

[Ans26] Anurag Anshu. *An Energy-Based Uncertainty Principle and Low-Energy State Preparation*. Sept. 2026. arXiv: 2603.15495 [quant-ph] (page 3).

[BKS09] Sergio Boixo, Emanuel Knill, and Rolando Somma. "Eigenpath traversal by phase randomization". In: *Quantum Info. Comput.* 9.9 (2009), pp. 833–855 (page 3).

[Blu+26] Andreas Bluhm, Matthias C. Caro, Francisco Escudero Gutiérrez, Junseo Lee, Aadil Oufkir, Cambyse Rouzé, and Myeongjin Shin. *Certifying and Learning Local Quantum Hamiltonians*. Mar. 2026. arXiv: 2603.29809 [quant-ph] (page 2).

[Cam19] Earl Campbell. "Random Compiler for Fast Hamiltonian Simulation". In: *Physical Review Letters* 123.7 (2019). DOI: 10.1103/physrevlett.123.070503 (page 3).

[Car+24] Almudena Carrera Vazquez, Caroline Tornow, Diego Ristè, Stefan Woerner, Maika Takita, and Daniel J. Egger. "Combining Quantum Processors with Real-Time Classical Communication". In: *Nature* 636.8041 (2024), pp. 75–79. DOI: 10.1038/s41586-024-08178-2 (page 2).

[CSM23] Jordan Cotler, Thomas Schuster, and Masoud Mohseni. "Information-theoretic hardness of out-of-time-order correlators". In: *Phys. Rev. A* 108 (6 2023), p. 062608. DOI: 10.1103/PhysRevA.108.062608 (page 3).

[CSTWZ21] Andrew M. Childs, Yuan Su, Minh C. Tran, Nathan Wiebe, and Shuchen Zhu. "Theory of Trotter Error with Commutator Scaling". In: *Physical Review X* 11.1 (2021), p. 011020. DOI: 10.1103/PhysRevX.11.011020 (page 2).

[Far+11] Edward Farhi, Jeffrey Goldstone, David Gosset, Sam Gutmann, Harvey B. Meyer, and Peter Shor. "Quantum adiabatic algorithms, small gaps, and different paths". In: *Quantum Info. Comput.* 11.3 (2011), pp. 181–214 (page 3).

[GKZ25] David Gosset, Robin Kothari, and Chenyi Zhang. *Multi-qubit Toffoli with exponentially fewer T gates*. 2025. arXiv: 2510.07223 [quant-ph] (page 3).

[HHJWY17] Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. "Sample-Optimal Tomography of Quantum States". In: *IEEE Transactions on Information Theory* 63.9 (2017), pp. 5628–5641. DOI: 10.1109/TIT.2017.2719044 (page 2).

[HMD03] Peter Høyer, Michele Mosca, and Ronald De Wolf. "Quantum search on bounded-error inputs". In: *Proceedings of the 30th International Conference on Automata, Languages and Programming*. ICALP'03. Eindhoven, The Netherlands: Springer-Verlag, 2003, pp. 291–299 (page 3).

[KL25] John Kallaugher and Daniel Liang. "Hamiltonian Locality Testing via Trotterized Postselection". In: *20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025)*. Ed. by Bill Fefferman. Vol. 350. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025, 10:1–10:20. DOI: 10.4230/LIPIcs.TQC.2025.10 (page 1).

[LTW26] Laura Lewis, Ewin Tang, and John Wright. *Learning the Structure of Open Quantum Systems*. June 2026. arXiv: 2606.30358 [quant-ph] (page 3).

[OKSM24] Tatsuki Odake, Hlér Kristjánsson, Akihito Soeda, and Mio Murao. "Higher-order quantum transformations of Hamiltonian dynamics". In: *Phys. Rev. Res.* 6 (1 2024), p. L012063. DOI: 10.1103/PhysRevResearch.6.L012063 (page 2).

[OW16] Ryan O'Donnell and John Wright. "Efficient quantum tomography". In: *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '16. Cambridge, MA, USA: Association for Computing Machinery, 2016, pp. 899–912. DOI: 10.1145/2897518.2897544 (page 2).

[PHOW20] Tianyi Peng, Aram W. Harrow, Maris Ozols, and Xiaodi Wu. "Simulating Large Quantum Circuits on a Small Quantum Computer". In: *Physical Review Letters* 125.15 (2020), p. 150504. DOI: 10.1103/PhysRevLett.125.150504 (page 2).

[PSTW25] Angelos Pelecanos, Jack Spilecki, Ewin Tang, and John Wright. *Mixed State Tomography Reduces to Pure State Tomography*. Nov. 2025. arXiv: 2511.15806 [quant-ph] (page 2).

[RS08] Oded Regev and Liron Schiff. "Impossibility of a Quantum Speed-Up with a Faulty Oracle". In: *Proceedings of the 35th International Colloquium on Automata, Languages and Programming - Volume Part I*. ICALP '08. Reykjavik, Iceland: Springer-Verlag, 2008, pp. 773–781. DOI: 10.1007/978-3-540-70575-8_63 (page 3).

[SHH25] Thomas Schuster, Jonas Haferkamp, and Hsin-Yuan Huang. "Random unitaries in extremely low depth". In: *Science* 389.6755 (2025), pp. 92–96. DOI: 10.1126/science.adv8590 (page 3).

[ST26] Myeongjin Shin and Yu Tong. *Hamiltonian Learning and Certification via Eigenphase Engineering*. 2026. arXiv: 2609.26596 [quant-ph] (page 1).

[SW22] Mehdi Soleimanifar and John Wright. "Testing Matrix Product States". In: *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, Jan. 2022, pp. 1679–1701. DOI: 10.1137/1.9781611977073.68 (page 2).

[TW26] Ewin Tang and John Wright. *Amplitude Amplification and Estimation Require Inverses*. Apr. 2026. arXiv: 2507.23787 [quant-ph] (page 1).

[TWZ26] Ewin Tang, John Wright, and Mark Zhandry. *Conjugate Queries Can Help*. May 2026. arXiv: 2510.07622 [quant-ph] (page 2).

[WZYZ26] Yue Wang, Xiao-Ming Zhang, Xiao Yuan, and Qi Zhao. *Efficient Preparation of Quantum States via Randomized Truncation*. Jan. 2026. arXiv: 2510.12247 [quant-ph] (page 2).