Quantum advantage in learning from experiments

Hsin-Yuan Huang, 1, 2, * Michael Broughton, 3 Jordan Cotler, 4, 5 Sitan Chen, 6, 7 Jerry Li, 8 Masoud Mohseni, 3 Hartmut Neven, 3 Ryan Babbush, 3 Richard Kueng, 9 John Preskill, 1, 2, 10 and Jarrod R. McClean^{3, †}

¹Institute for Quantum Information and Matter, Caltech, Pasadena, CA, USA
²Department of Computing and Mathematical Sciences, Caltech, Pasadena, CA, USA
³Google Quantum AI, 340 Main Street, Venice, CA 90291, USA
⁴Harvard Society of Fellows, Cambridge, MA 02138 USA
⁵Black Hole Initiative, Cambridge, MA 02138 USA
⁶Department of Electrical Engineering and Computer Science,
University of California, Berkeley, Berkeley, CA, USA
⁷Simons Institute for the Theory of Computing, Berkeley, CA, USA
⁸Microsoft Research AI, Redmond, WA 98052, USA
⁹Institute for Integrated Circuits, Johannes Kepler University Linz, Austria
¹⁰AWS Center for Quantum Computing, Pasadena, CA 91125, USA
(Dated: December 3, 2021)

Quantum technology has the potential to revolutionize how we acquire and process experimental data to learn about the physical world. An experimental setup that transduces data from a physical system to a stable quantum memory, and processes that data using a quantum computer, could have significant advantages over conventional experiments in which the physical system is measured and the outcomes are processed using a classical computer. We prove that, in various tasks, quantum machines can learn from exponentially fewer experiments than those required in conventional experiments. The exponential advantage holds in predicting properties of physical systems, performing quantum principal component analysis on noisy states, and learning approximate models of physical dynamics. In some tasks, the quantum processing needed to achieve the exponential advantage can be modest; for example, one can simultaneously learn about many noncommuting observables by processing only two copies of the system. Conducting experiments with up to 40 superconducting qubits and 1300 quantum gates, we demonstrate that a substantial quantum advantage can be realized using today's relatively noisy quantum processors. Our results highlight how quantum technology can enable powerful new strategies to learn about nature.

I. Introduction

Humans learn about nature by doing experiments, but up until now our ability to acquire knowledge has been hampered by viewing the quantum world through a classical lens. The rapid advance of quantum technology portends an opportunity to observe the world in a fundamentally different and more powerful way. Instead of measuring physical systems and then processing the classical measurement outcomes to infer properties of the physical systems, quantum sensors [1] will eventually be able to transduce [2] quantum information in physical systems directly to a quantum memory [3, 4], where it can be processed by a quantum computer. Fig. 1(a) illustrates the distinction between conventional and quantum-enhanced experiments. For example, in a quantum-enhanced experiment, multiple photons might be captured and stored coherently at each node of a quantum network, and then processed coherently to extract an informative signal [5–7].

Recently, mathematical analyses done by some of the authors show that there exist properties of an *n*-qubit

*Electronic address: hsinyuan@caltech.edu

system that a quantum machine can learn efficiently, while the required number of conventional experiments to achieve the same task is exponential in n [8, 9]. This exponential advantage contrasts sharply with the quadratic advantage achieved in many previously proposed strategies for improving sensing using quantum technology [1]. In this article, we propose and analyze three classes of learning tasks with exponential quantum advantage, and report on proof-of-principle experiments using up to 40 qubits on a Google Sycamore processor [10]. These experiments confirm that a substantial quantum advantage can be realized even when the quantum memory and processor are both noisy.

To be more concrete, suppose that each experiment generates an n-qubit state ρ , and our goal is to learn some property of ρ ; see Fig. 1. We depict conventional and quantum-enhanced experiments for this scenario in Fig. 1(b). In conventional experiments, each copy of ρ is measured separately, the measurement data is stored in a classical memory, and a classical computer outputs a prediction for the property after processing the classical data. In quantum-enhanced experiments, each copy of ρ is stored in a quantum memory, and then the quantum machine outputs the prediction after processing the quantum data in the quantum memory. We prove that for some tasks, the number of experiments needed to learn a desired property is exponential in n

[†]Electronic address: jmcclean@google.com

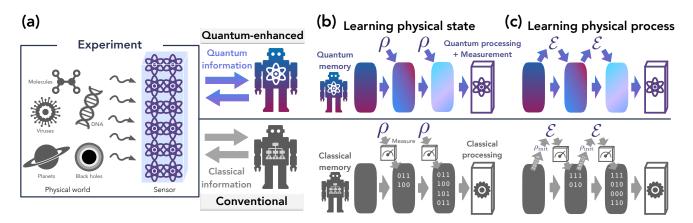


Figure 1: (a) Quantum-enhanced experiments versus conventional experiments. Quantum-enhanced/conventional experiments interface with a quantum/classical machine running a quantum/classical learning algorithm that can store and process quantum/classical information. (b) Learning physical state ρ . Each experiment produces a physical state ρ . In the conventional setting, we measure each ρ to obtain classical data (the measurement could depend on prior measurement outcomes) and store the data in a classical memory. In the quantum-enhanced setting, ρ can coherently alter the quantum information stored in the memory of the quantum machine (illustrated by the change in color). With large enough quantum memory, the quantum machine can simply store each copy of ρ . After multiple rounds of experiments, quantum processing followed by a measurement is performed on the quantum memory. (c) Learning physical process \mathcal{E} . Each experiment is an evolution under \mathcal{E} . In the conventional setting, the classical machine specifies the input state to \mathcal{E} using a classical bitstring, and obtains classical measurement data [11]. In the quantum-enhanced setting, the evolution \mathcal{E} coherently alters the memory of the quantum machine: the input state to \mathcal{E} is entangled with the quantum memory in the quantum machine and the output state is retrieved coherently by the quantum machine.

using the conventional strategy, but only polynomial in n using the quantum-enhanced strategy. For suitably defined tasks, we can achieve exponential quantum advantage using a protocol as simple as storing two copies of ρ in quantum memory and performing an entangling measurement. We also show that quantum-enhanced experiments have a similar exponential advantage in a related scenario shown in Fig. 1(c), in which the goal is to learn about a quantum process $\mathcal E$ rather than a quantum state ρ .

Building on observations in [8, 12] we prove that for a task that entails acquiring information about a large number of non-commuting observables, quantumenhanced experiments can have an exponential advantage even when the measured quantum state is unentangled. By performing experiments with up to 40 superconducting qubits, we show that this quantum advantage persists even when using currently available quantum processors. We also demonstrate quantum advantage in learning the symmetry class of a physical evolution operator, inspired by recent theoretical advances [9, 12]. Finally, in a theoretical contribution, we rigorously prove that quantum-enhanced experiments have an exponential advantage in learning about the principal component of a noisy state, as previously indicated in [13].

In our proof-of-principle experiments, we directly execute the state preparation or process to be learned within the quantum processor. In an actual application, the quantum data analyzed by the learning algorithm

might be produced by an analog quantum simulator or a gate-based quantum computer. We also envision future applications in which quantum sensors equipped with quantum processors interact coherently with the physical world. The robustness of quantum advantage with respect to noise, validated by our experiments using a noisy superconducting device, boosts our confidence that the quantum-enhanced strategies described here can be exploited someday to achieve a substantial advantage in realistic applications.

II. Provable quantum advantage

Here we present three classes of learning tasks and the associated quantum-enhanced experiments, each yielding a provable exponential advantage over conventional experiments. Each result is encapsulated by a theorem which we state informally. Precise statements and proofs are presented in the Supplemental Material. Our experimental demonstrations are discussed in Sec. III. The proofs proceed by representing a classical algorithm with a decision tree depicted at the center of the gray robot in Fig. 1. The tree representation encodes how the classical memory changes as we obtain more experimental data. We then analyze how the transitions on the tree differ for distinct measured physical systems to provide rigorous information-theoretic lower bounds. A general mathematical framework building on [12] is given in Appendix C.

The first task concerns learning about a physical system described by an *n*-qubit state ρ . We suppose that each experiment generates one copy of ρ . In the conventional setting, we measure each copy of ρ to obtain classical data. The procedure can be adaptive, that is, each measurement can depend on the data obtained in earlier measurements. In the quantum-enhanced setting, a quantum computer can store each copy of ρ in a quantum memory, and act jointly on multiple copies of ρ . In both scenarios, we require all quantum data to be measured at the end of the learning phase of the procedure, so that only classical data survives. After the learning is completed, the learner is asked to provide an accurate prediction for the expectation value of one observable drawn from a set $\{O_1, O_2, \dots\}$ where the number of observables in the set is exponentially large in n. The observables in the set can be highly incompatible, that is, each observable may fail to commute with many others in the set.

In prior work by some of the authors [8, 12], we required the learner to predict exponentially many observables, which is not possible in practice if the system size is large. In order to demonstrate the advantage in an actual device, we prove that predicting just the absolute value of one observable requires exponentially many copies in the *conventional* scenario. In contrast, predicting the entire set of observables can be achieved with a polynomial number of copies in the *quantumenhanced* scenario. We thereby establish the following constant versus exponential separation. The proof is given in Appendix D.

Theorem 1 (Predicting observables). There exists a distribution over n-qubit states and a set of observables such that in the conventional scenario, at least order 2^n experiments are needed to predict the absolute value of one observable selected from the set, while a constant number of experiments suffice in the quantum-enhanced scenario.

The exponential quantum advantage can occur even if the state ρ is unentangled. For example, in our experiments we consider $\rho \propto (I + \alpha P)$ where P is an *n*-qubit Pauli operator and $\alpha \in (-1,1)$. This state can be realized as a probabilistic ensemble of product states, each of which is an eigenstate of P with eigenvalue α . Even if the state is known to be of this form, but P, α are unknown, the exponential separation between conventional and quantum-enhanced experiments persists. Moreover, the quantum advantage can be achieved by performing simple entangling measurements on pairs of copies of ρ . That the quantum advantage applies even when correlations among the n qubits are classical leads us to believe that the quantum-enhanced strategy will be beneficial in a broad class of sensing applications. In Appendix G, we extend this theorem, showing that a sufficiently large quantum memory is needed to achieve this task in the quantum-enhanced scenario.

Our second machine learning task with a quantum advantage is quantum principal component analysis (PCA) [13]. In this task, each experiment produces one copy of ρ , and our goal is to predict properties of the (first) principal component of ρ , namely the eigenstate $|\psi\rangle$ of ρ with the largest eigenvalue. For example, we may want to predict the expectation values of a few observables in the state $|\psi\rangle$. This task may become a valuable ingredient in future quantum-sensing applications. If an imperfect quantum sensor transduces a detected quantum state into quantum memory, the state is likely to be corrupted by noise. But it is reasonable to expect that properties of the principal component are relatively robust with respect to noise [14], and therefor highly informative about the uncorrupted state. To perform quantum PCA, a learning algorithm was introduced in [13] based on phase estimation which requires fault-tolerant quantum computers. One can also obtain information about the principal component of ρ using more near-term algorithms, such as virtual cooling [15] and virtual distillation [16].

While the quantum PCA algorithm in [13] is exponentially faster than known algorithms based on conventional experiments, this advantage was not proven against *all* possible algorithms in the conventional scenario. Here, we rigorously establish the exponential quantum advantage for performing quantum PCA. The exponential quantum advantage also holds in the nearterm proposals [15, 16]. The proofs are in Appendix E.

Theorem 2 (Performing quantum PCA). In the conventional scenario, at least order $2^{n/2}$ experiments are needed to learn a fixed property of the principal component of an unknown n-qubit quantum state, while a constant number of experiments suffice in the quantum-enhanced scenario.

It is worth commenting on recent results in Refs. [17, 18] showing that quantum PCA can be achieved by polynomial-time classical algorithms, which may seem to contradict Theorem 2. Those works assume the ability to access any entry of the exponentially large matrix ρ to exponentially high precision in polynomial time. Achieving such a high precision requires measuring exponentially many copies of ρ , which takes an exponential number of experiments and exponential time. Hence, the assumptions of [17, 18] do not hold here. See [19] which provides a detailed exposition of these matters.

Another core task in quantum mechanics is understanding physical processes rather than states. Here, each experiment implements a physical process \mathcal{E} , and we can interface with \mathcal{E} through a quantum/classical machine in the quantum-enhanced/conventional setting; see Fig. 1(c). We show that a quantum machine can learn an approximate model of any polynomial-time quantum process \mathcal{E} from only a polynomial number of experiments. Given a distribution on input states, the approximate model can predict the output state from \mathcal{E}

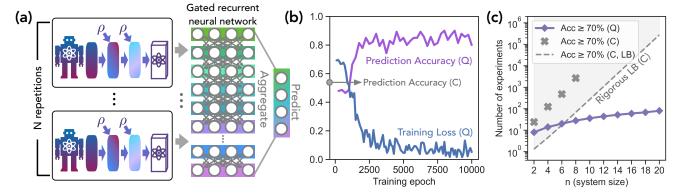


Figure 2: Quantum advantage in learning physical states. (a) Supervised machine learning (ML) model based on quantum-enhanced experiments. N repetitions of quantum-enhanced experiments are performed and the data is fed into a gated recurrent neural network (GRU) [20, 21]. The neurons in the GRU are aggregated to predict an output. (b) Training process of the supervised ML model. We train the supervised ML model to determine which of two n-qubit Pauli operators has a larger magnitude for the expectation value in an unknown state ρ using noiseless simulation for small system sizes (n < 8). We consider the cross entropy [22] as the training loss. Then we use the supervised ML model to make predictions using data from noisy quantum-enhanced experiments running on the Sycamore processor [10] for larger system sizes ($8 \le n \le 20$). We consider the probability to predict correctly as the prediction accuracy. Random guessing yields a prediction accuracy of 0.5. (c) Quantum advantage in the number of experiments needed to achieve $\ge 70\%$ accuracy. Here, (Q) corresponds to results running the supervised ML model based on quantum-enhanced experiments and (C) corresponds to results running the best known conventional strategy. The dotted line is a lower bound for any conventional strategy (C, LB) as proven in Appendix D.4. Even running on a noisy quantum processor, quantum-enhanced experiments are seen to vastly outperform the best theoretically achievable conventional results (C, LB).

accurately on average. In contrast, we would need an exponential number of experiments to achieve the same task in the conventional setting. The proof for general quantum processes is given in Appendix F.

Theorem 3 (Learning quantum processes). Suppose we are given a polynomial-time physical process \mathcal{E} acting on n qubits and a probability distribution over n-qubit input states. In the conventional scenario, at least order 2^n experiments are needed to learn an approximate model of \mathcal{E} that predicts output states accurately on average, while a polynomial number of experiments suffice in the quantum-enhanced scenario.

III. Demonstrations of quantum advantage

The exponential quantum advantage captured by Theorems 1, 2, and 3 applies no matter how much classical processing power is leveraged in the conventional experiments. The conventional strategy fails because there is just no way to access enough classical data to perform the specified tasks, if the number of experiments is subexponential in n. But these exponential separations apply in an idealized setting where quantum states are stored and processed perfectly. Will access to quantum memory unlock a substantial quantum advantage under more realistic conditions?

For two different tasks, we have investigated the robustness of the quantum advantage by conducting experiments using a superconducting quantum processor.

The first task we studied pertains to Theorem 1. The task is to approximately estimate the magnitude for the expectation value of Pauli observables. The unknown state is an unentangled n-qubit state $\rho = 2^{-n} (I + \alpha P)$, where $\alpha = \pm 0.95$, P is a Pauli operator, and both α, P are unknown. After all measurements are completed and learning is terminated, two distinct Pauli operators Q_1 and Q_2 are announced, one of which is P and the other of which is not equal to P. We ask the machine to determine which of $|\operatorname{tr}(Q_1\rho)|$ and $|\operatorname{tr}(Q_2\rho)|$ is larger.

In the conventional scenario, where copies of ρ are measured one by one, the best known strategy is to use randomized Clifford measurements requiring an exponential number of copies to achieve the task with reasonable success probability [8, 24]. In the quantumenhanced scenario, copies of ρ are deposited in quantum memory two at a time, and a Bell measurement across the two copies is performed to extract a snapshot of the state. We then feed the snapshots into a supervised machine learning (ML) model based on a gated recurrent neural network [20, 21, 25] to make a prediction, as depicted in Fig. 2(a). We train the neural network using noiseless simulation data for small system sizes (n < 8). Then we use the neural network to make predictions when we are provided with experimental data for large system sizes $(8 \le n \le 20)$. We report the prediction accuracy, which is equal to the probability for correctly answering whether $|\operatorname{tr}(Q_1\rho)|$ or $|\operatorname{tr}(Q_2\rho)|$ is larger. Fig. 2(b) shows the performance of the ML model as we train the neural network. Fig. 2(c) depicts,

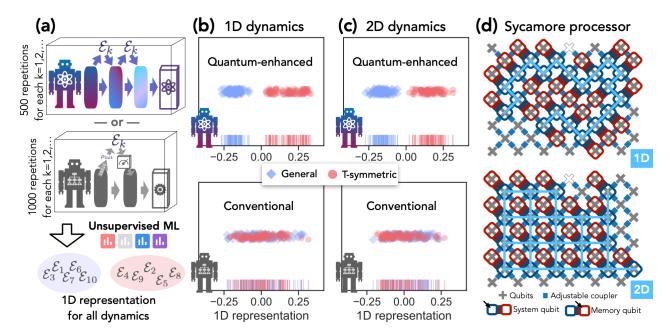


Figure 3: Quantum advantage in learning physical dynamics. (a) Unsupervised machine learning (ML) model. We perform 500 repetitions of quantum-enhanced experiments (each accessing \mathcal{E}_k twice) for every physical process \mathcal{E}_k , and feed the data into an unsupervised ML model (Gaussian kernel PCA [23]) to learn a one-dimensional representation for describing distinct physical dynamics $\mathcal{E}_1, \mathcal{E}_2, \ldots$ Similarly, we also consider applying unsupervised ML to data obtained from 1000 repetitions of the best known conventional experiments (each accessing \mathcal{E}_k once) for every physical process \mathcal{E}_k . (b) Representation learned by unsupervised ML for 1D dynamics. Each point corresponds to a distinct physical process \mathcal{E}_k . The vertical line at the bottom shows the exact 1D representation of each \mathcal{E}_k . Half of the processes satisfy time-reversal symmetry (blue diamonds) while the other half of them do not (red circles). When fed with data from quantum-enhanced experiments, the ML model accurately discovers the underlying symmetry pattern. In contrast, the ML model fails to do so when fed with data from conventional experiments. (c) Representation learned by unsupervised ML for 2D dynamics. (d) The geometry implemented on the Sycamore processor [10]. We consider two different classes of connectivity geometry for implementing 1D (top) and 2D (bottom) dynamics.

as a function of the system size n, the number of experiments needed in each scenario to achieve 70% prediction accuracy. We show the experimental results when using conventional and quantum-enhanced experiments, along with a theoretical lower bound on the number of experiments needed in the conventional scenario as proven in Appendix D.4. Despite the noisy storage and processing in the experimental device, we observe a significant quantum advantage.

The second task we studied, which pertains to Theorem 3, is inspired by the recent observation that quantum-enhanced experiments can efficiently identify the symmetry class of a quantum evolution operator, while conventional experiments cannot $[9,\ 12]$. An unknown n-qubit quantum evolution operator is presented, drawn either from the class of all unitary transformations, or from the class of time-reversal-symmetric unitary transformations (i.e., real orthogonal transformations). We consider whether an unsupervised ML can learn to recognize the symmetry class of the unknown evolution operator based on data obtained from either quantum-enhanced experiments or conventional experiments. An illustration is shown in Fig. 3(a).

In the conventional scenario, we repeatedly apply the unknown evolution operator to the initial state $|0\rangle^{\otimes n}$, and then measure each qubit of the output state in the Y-basis. Under T-symmetric evolution the output state has purely real amplitudes; hence the expectation value of any purely imaginary observable, such as the Pauli-Y operator, is always zero. In contrast, the expectation value of Y after general unitary evolution is generically nonzero, but may be exponentially small and hence hard to distinguish from zero. In the quantum-enhanced scenario, we make use of n additional memory qubits. We prepare an initial state in which the n system qubits are entangled with the n memory qubits, evolve the system qubits under the unknown evolution operator, swap the system and memory qubits, evolve the system qubits again, and finally perform n Bell measurements, each acting on one system qubit and one memory qubit.

Each evolution operator is a one-dimensional or two-dimensional n-qubit quantum circuit as shown in Fig. 3(d). After sampling many different evolution operators from both symmetry classes (and obtaining data from each sampled evolution multiple times), we use an unsupervised ML model (kernel PCA [23]) to find a one-

dimensional representation of the evolution operators. The representations learned by the unsupervised ML model are shown in Fig. 3(b, c). Using the quantum-enhanced data, the ML model discovers a clean separation between the two symmetry classes, while there is no discernable separation into classes when using data from conventional experiments. The signal from the quantum-enhanced experiments is strong enough that the two classes are easily recognized without access to any labeled training data.

IV. Outlook

We have investigated how quantum technology can enhance our ability to discover new phenomena occurring in nature. For a variety of tasks, we proved that quantum-enhanced strategies using quantum memory and quantum processing can predict properties of physical systems using exponentially fewer experiments than conventional strategies. This exponential advantage is achievable even if the amount of classical processing used in the conventional strategies is unlimited, and even when the physical system exhibits only classical correlations. While many previous studies of quantum advantage have focused on computational tasks with known inputs, our work focuses instead on learning tasks, where the goal is to learn about an a priori unknown physical system. This work provides a new approach to understanding and achieving quantum advantage in quantum machine learning [26, 27] and quantum sensing [1].

Our experiments with up to 40 qubits in a superconducting quantum processor show that a substantial quantum advantage is already evident when using today's noisy intermediate-scale quantum platforms [28]. These experiments demonstrate that supervised and unsupervised machine learning models [25, 29] employing data obtained from quantum-enhanced experiments can predict properties and discover underlying structure in physical systems that are beyond the scope of conventional experiments.

We envision that future quantum sensing systems will be able to transduce detected quantum data to a quantum memory and then process the stored data using a quantum computer. Lacking for now suitably advanced sensors and transducers, we have conducted proof-of-concept experiments in which quantum data is directly planted in our quantum processor. Nevertheless, the robust quantum advantage we have validated highlights the potential for advancing quantum platforms to unlock facets of nature that would otherwise remain concealed.

Acknowledgments:

The quantum hardware used for this experiment was developed by the Google Quantum AI hardware team, under the direction of Anthony Megrant, Julian Kelly and Yu Chen. Methods for device calibrations were developed by the physics team led by Vadim Smelyanskiy. Data was collected via cloud access through Google's Quantum Computing Service. HH is supported by a Google PhD Fellowship. JC is supported by a Junior Fellowship from the Harvard Society of Fellows, the Black Hole Initiative, as well as in part by the Department of Energy under grant DE-SC0007870. SC is supported by the National Science Foundation under Award 2103300 and was visiting the Simons Institute for the Theory of Computing while part of this work was completed. JP acknowledges funding from the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, (DE-NA0003525, DE-SC0020290), and the National Science Foundation (PHY-1733907). The Institute for Quantum Information and Matter is an NSF Physics Frontiers Center.

^[1] C. L. Degen, F. Reinhard, and P. Cappellaro, "Quantum sensing," Rev. Mod. Phys. 89, 035002 (2017).

^[2] N. Lauk, N. Sinclair, S. Barzanjeh, J. P. Covey, M. Saffman, M. Spiropulu, and C. Simon, "Perspectives on quantum transduction," Quantum Science and Technology 5, 020501 (2020).

^[3] A. I. Lvovsky, B. C. Sanders, and W. Tittel, "Optical quantum memory," Nature photonics **3**, 706 (2009).

^[4] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, "Topological quantum memory," Journal of Mathematical Physics 43, 4452 (2002).

^[5] D. Gottesman, T. Jennewein, and S. Croke, "Longer-baseline telescopes using quantum repeaters," Phys. Rev. Lett. 109, 070503 (2012).

^[6] J. Bland-Hawthorn, M. J. Sellars, and J. G. Bartholomew, "Quantum memories and the double-slit

experiment: implications for astronomical interferometry," Journal of the Optical Society of America B **38**, A86 (2021).

^[7] V. Giovannetti, S. Lloyd, and L. Maccone, "Advances in quantum metrology," Nature photonics 5, 222 (2011).

^[8] H.-Y. Huang, R. Kueng, and J. Preskill, "Information-theoretic bounds on quantum advantage in machine learning," Phys. Rev. Lett. 126, 190505 (2021).

^[9] D. Aharonov, J. Cotler, and X.-L. Qi, "Quantum algorithmic measurement," arXiv:2101.04634 (2021).

^[10] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandão, D. A. Buell, et al., "Quantum supremacy using a programmable superconducting processor," Nature 574, 505 (2019).

^[11] M. Mohseni, A. T. Rezakhani, and D. A. Lidar,

- "Quantum-process tomography: Resource analysis of different strategies," Phys. Rev. A 77, 032322 (2008).
- [12] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, "Exponential separations between learning with and without quantum memory," to appear in FOCS (2021), arXiv:2111.05881.
- [13] S. Lloyd, M. Mohseni, and P. Rebentrost, "Quantum principal component analysis," Nat. Phys. 10, 631 (2014).
- [14] B. Koczor, "The dominant eigenvector of a noisy quantum state," arXiv:2104.00608 (2021).
- [15] J. Cotler, S. Choi, A. Lukin, H. Gharibyan, T. Grover, M. E. Tai, M. Rispoli, R. Schittko, P. M. Preiss, A. M. Kaufman, et al., "Quantum virtual cooling," Phys. Rev. X 9, 031013 (2019).
- [16] W. J. Huggins, S. McArdle, T. E. O'Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, "Virtual distillation for quantum error mitigation," arXiv:2011.07064 (2020).
- [17] E. Tang, "Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions," Phys. Rev. Lett. 127, 060503 (2021).
- [18] N.-H. Chia, A. 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 *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing* (2020) pp. 387–400.
- [19] J. Cotler, H.-Y. Huang, and J. R. McClean, "Revisiting dequantization and quantum advantage in learning tasks," to appear (2021).
- [20] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," arXiv:1412.3555 (2014).
- [21] D. Tang, B. Qin, and T. Liu, "Document modeling with gated recurrent neural network for sentiment classification," in *Proceedings of the 2015 conference on empir*ical methods in natural language processing (2015) pp. 1422–1432.
- [22] K. P. Murphy, Machine learning: a probabilistic perspective (MIT press, 2012).
- [23] B. Schölkopf, A. Smola, and K.-R. Müller, "Nonlinear component analysis as a kernel eigenvalue problem," Neural computation 10, 1299 (1998).
- [24] H.-Y. Huang, R. Kueng, and J. Preskill, "Predicting many properties of a quantum system from very few measurements," Nature Physics 16, 1050 (2020).
- [25] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning (The MIT Press, 2016).
- [26] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," Nature 549, 195 (2017).
- [27] M. Broughton, G. Verdon, T. McCourt, A. J. Martinez, J. H. Yoo, S. V. Isakov, P. Massey, R. Halavati, M. Y. Niu, A. Zlokapa, E. Peters, O. Lockwood, A. Skolik, S. Jerbi, V. Dunjko, M. Leib, M. Streif, D. V. Dollen, H. Chen, S. Cao, R. Wiersema, H.-Y. Huang, J. R. McClean, R. Babbush, S. Boixo, D. Bacon, A. K. Ho, H. Neven, and M. Mohseni, "Tensorflow quantum: A software framework for quantum machine learning," (2021), arXiv:2003.02989 [quant-ph].
- [28] J. Preskill, "Quantum computing in the NISQ era and

- beyond," Quantum 2, 79 (2018).
- [29] M. Mohri, A. Rostamizadeh, and A. Talwalkar, Foundations of machine learning (The MIT Press, 2018).
- [30] C. Developers, "Cirq," (2021), See full list of authors on Github: https://github.com/quantumlib/Cirq/graphs/contributors.
- [31] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, "Provably efficient machine learning for quantum many-body problems," arXiv:2106.12627 (2021).
- [32] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," Neural Computation 9, 1735 (1997).
- [33] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural in*formation processing systems (2017) pp. 5998–6008.
- [34] G. Koch, R. Zemel, R. Salakhutdinov, et al., "Siamese neural networks for one-shot image recognition," in *ICML deep learning workshop*, Vol. 2 (Lille, 2015).
- [35] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," arXiv:1412.6980 (2014).
- [36] L. Buitinck, G. Louppe, M. Blondel, F. Pedregosa, A. Mueller, O. Grisel, V. Niculae, P. Prettenhofer, A. Gramfort, J. Grobler, R. Layton, J. VanderPlas, A. Joly, B. Holt, and G. Varoquaux, "API design for machine learning software: experiences from the scikitlearn project," in ECML PKDD Workshop: Languages for Data Mining and Machine Learning (2013) pp. 108– 122.
- [37] S. Bubeck, S. Chen, and J. Li, "Entanglement is necessary for optimal quantum property testing," in 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS) (IEEE, 2020) pp. 692–703.
- [38] M. Mohseni and D. A. Lidar, "Direct characterization of quantum dynamics: General theory," Phys. Rev. A 75, 062331 (2007).
- [39] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, "Sample-optimal tomography of quantum states," IEEE Trans. Inf. Theory 63, 5628 (2017).
- [40] S. Chen, S. Zhou, A. Seif, and L. Jiang, "Quantum advantages for pauli channel estimation," arXiv:2108.08488 (2021).
- [41] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, "A hierarchy for replica quantum advantage," arXiv:2111.05874 (2021).
- [42] A. Anshu, Z. Landau, and Y. Liu, "Distributed quantum inner product estimation," arXiv:2111.03273 (2021).
- [43] M. Coudron and S. Menda, "Computations with greater quantum depth are strictly more powerful (relative to an oracle)," in *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing* (2020) pp. 889–901.
- [44] N.-H. Chia, K.-M. Chung, and C.-Y. Lai, "On the need for large quantum depth," in *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Com*puting (2020) pp. 902–915.
- [45] S. Aaronson, "Shadow tomography of quantum states," SIAM Journal on Computing 49, STOC18 (2019).
- [46] S. Aaronson and G. N. Rothblum, "Gentle measurement of quantum states and differential privacy," in STOC (2019) pp. 322–333.
- [47] C. Bădescu and R. O'Donnell, "Improved quantum data

- analysis," arXiv:2011.10908 (2020).
- [48] B. Yu, "Assouad, fano, and le cam," in Festschrift for Lucien Le Cam (Springer, 1997) pp. 423–435.
- [49] Z. Ji, Y.-K. Liu, and F. Song, "Pseudorandom quantum states," in *Annual International Cryptology Conference* (Springer, 2018) pp. 126–152.
- [50] Z. Brakerski and O. Shmueli, "(Pseudo) Random Quantum States with Binary Phase," in *Theory of Cryptography Conference* (Springer, 2019) pp. 229–250.
- [51] Z. Brakerski and O. Shmueli, "Scalable pseudorandom quantum states," in Annual International Cryptology

- Conference (Springer, 2020) pp. 417-440.
- [52] R. Vershynin, High-Dimensional Probability: An Introduction with Applications in Data Science (Cambridge University Press, 2018).
- [53] J. Watrous, The Theory of Quantum Information (Cambridge University Press, 2018).
- [54] M. C. Caro, H.-Y. Huang, M. Cerezo, K. Sharma, A. Sornborger, L. Cincio, and P. J. Coles, "Generalization in quantum machine learning from few training data," arXiv:2111.05292 (2021).

Supplementary information

Α.	Experimental details
	General description • Experiments on learning physical states • Experiments on learning physical dynamics • Additional experimental results • Performance and characterization data
В.	A brief review on quantum information theory
	Definition and properties of quantum processes • Definition and properties of POVMs
C.	Mathematical framework for proving exponential advantage
	Tree representation \bullet Many-versus-one distinguishing tasks \bullet Many-versus-many distinguishing task \bullet Partially-revealed many-versus-one distinguishing task \bullet Presence of noise \bullet Related works
D.	Predicting highly-incompatible observables
Ε.	Performing quantum principal component analysis
	An exponential lower bound for conventional experiments $ullet$ An exponential lower bound using pseudorandomness
F.	Learning a polynomial-time quantum process
	Problem setting \bullet Rigorous statements \bullet Proof of polynomial upper bound in Theorem 11
G.	Predicting highly-incompatible observables with bounded quantum memory
	Background and statement of results • Review of learning tree framework for bounded quantum memories • Hardness result for small quantum memories

A. Experimental details

In this section, we present the details for the physical experiments ran on the superconducting processor as well as the supervised/unsupervised machine learning models used to analyze the data.

A.1. General description

The experiments were performed on a Google Sycamore processor containing up to 53 superconducting transmon qubits. The largest error source in the Sycamore processor [10] is qubit readout error, which ranges from 3% to 7%. The second largest error source is the two-qubit gate with an error around 0.5% to 1.5%. Single-qubit gates have the smallest error around 0.05% to 0.5%. The Sycamore chip was introduced in Ref. [10], where additional details concerning the hardware implementation and performance can be found. The Sycamore chip was controlled remotely using an internal cloud interface programmed using Cirq [30] and TensorFlow Quantum [27]. The layout of the chip including connectivity is depicted in Supp. Fig. 1(a).

For all of our experiments on learning about states and dynamics, the total number of qubits was varied from 4 to 40 qubits (where in each case half of the qubits are used to simulate a physical system). As the system size was varied, the subset of qubits used for the experiments was varied in order to maximize experimental performance and minimize any overhead related to the 2D connectivity. That is, the largest contiguous patches with low gate and measurement error rates were selected, and swap operations were used to meet connectivity requirements when necessary. The experimental requirements for learning states, 1D dynamics, and 2D dynamics differ considerably in their experimental complexity.

In all experiments, the implementations of the unknown states or dynamics are performed in the quantum processor, where the learning algorithms do not know about them. While this is only an emulation of the process of data collection from a physical system in an actual sensing experiment, it allows us to examine the proposed pipeline for quantum data processing in a situation where data collection is imperfect.

A.2. Experiments on learning physical states

We separate this subsection into the concrete procedure for generating the unknown states, the conventional experiments we run, the quantum-enhanced experiments we run, and the supervised neural network model for making prediction based on data from quantum-enhanced experiments.

A.2.a. Procedure for generating the class of unknown states

The state preparation we consider is relatively simple in that the unknown state ρ is unentangled, but has strong non-local classical correlations. In the experimental demonstration we consider states of the form $\rho = 2^{-n} (I + \alpha P)$, where $\alpha \in \{-0.95, 0.95\}$, $P = \bigotimes_{i=1}^{n} P_i$ is an n-qubit Pauli operator, and both α and P are unknown. The state $\rho = 2^{-n} (I + \alpha P)$ is prepared by a randomized constant-depth circuit described in the following. To generate one copy of ρ , we introduce a parameter $\eta = \text{sign}(\alpha)$. Then for each qubit $i = 1, \ldots, n$, we do the following.

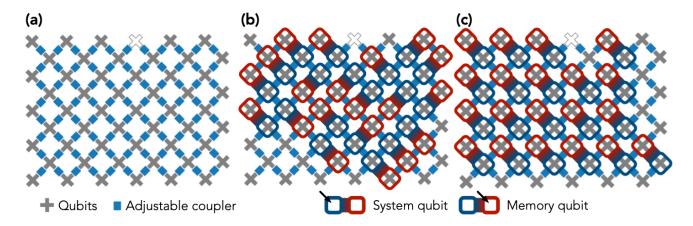
- 1. If $P_i = I$, then we set qubit i to be $|0\rangle\langle 0|$ with probability 1/2, and be $|1\rangle\langle 1|$ with probability 1/2.
- 2. If $P_i \neq I$ and there exists j > i such that $P_j \neq I$, then we set qubit i to be one of the two eigenstates of P_i with equal probability. We multiply η by the eigenvalue (+1 or -1) of the selected eigenstate of P_i .
- 3. If $P_i \neq I$ and there does not exist j > i such that $P_j \neq I$, then we use the following procedure.
 - (a) With probability 0.05, we set qubit i to be either $|0\rangle\langle 0|$ or $|1\rangle\langle 1|$ with equal probability.
 - (b) With probability 0.95, we set qubit i to be the positive eigenstate of P_i if $\eta = +1$, and set qubit i to be the negative eigenstate of P_i if $\eta = -1$.

By construction, the density operator prepared by this procedure is realized as an ensemble of pure states, where each pure state is a tensor product of Pauli operator eigenstates. Therefore, there is no quantum entanglement across different qubits. Furthermore, step 3 is designed to assure that $|\operatorname{tr}(P\rho)| = 0.95$.

A.2.b. Conventional experiments

In the conventional setting, the optimal strategy (up to logarithmic factors) for estimating expectation values of high-weight n-qubit Pauli observables uses classical shadow tomography based on randomized Clifford measurements [24]. Using this strategy, in each experiment we randomly sample a unitary transformation from the Clifford group, apply the sampled transformation to the unknown state ρ , and then measure in the computational basis. Although such randomized Clifford measurements can be executed using quantum circuits of polynomial size, the required circuits are too large to be performed accurately with today's noisy quantum devices except for quite modest values of n. Furthermore, the classical post-processing of the measurement results has complexity exponential in n.

In our conventional experiments, because randomized Clifford measurements are infeasible we instead use classical shadow tomography based on randomized Pauli measurements [24]. Using this strategy, in each experiment, we randomly sample from depth-1 Clifford circuits, apply the sampled circuit to ρ , and then measure in the computational basis. That is, for each of the n qubits, we decide uniformly at random to measure one of the three



Supplementary Figure 1: (a) Layout of a Google Sycamore processor. There is a total of 53 superconducting transmon qubits (the qubit corresponding to an empty cross is out of order). The blue rectangles show the adjustable couplers that can apply the entangling two-qubit gate SYC to neighboring qubits. (b) Layout used for learning states and for learning 1D dynamics. We partition the 40 qubits into 20 system qubits and 20 memory qubits. Either an unknown state of the system qubits is prepared, or an unknown process is applied to the system qubits. (c) Layout used for learning 2D dynamics.

Pauli observables X, Y, or Z. Many such measurements are performed, each time on a new copy of ρ , and the classical data collected is post-processed to predict expectation values of observables in the state ρ .

Classical shadow tomography based on randomized Pauli measurements is a powerful technique that enables classical ML models to predict quantum many-body ground states and quantum phases of matter with rigorous guarantees [31]. However, for the task of estimating the expectation value of an n-qubit Pauli observable that is announced after all measurements are completed, both randomized Clifford and randomized Pauli measurements require a number of experiments that scale exponentially in n, as we have proven in Appendix D.4. Likewise, exponentially many experiments are needed to perform Task 1 defined in Appendix D.4 with high success probability. In Fig. 2(b), we report the average prediction accuracy (the probability of performing Task 1 successfully) over system sizes n = 8, 10, 12, 14, 16, 18, 20. For each n-qubit state ρ , we conduct 1000 experiments to obtain the measurement data. The prediction accuracy is indicated by the gray point shown on the vertical axis, which is only slightly better than random guessing (0.5). For system sizes $n \geq 10$, the prediction accuracy is very close to 0.5. Classical shadow tomography based on randomized Pauli measurements is a statistical estimation procedure that has no training phase. Therefore, the prediction accuracy for conventional experiments is a single point Fig. 2(b). The training epoch on the horizontal axis in Fig. 2(b) is only for quantum-enhanced strategy. In Fig. 2(c), we report the number of experiments required to achieve a prediction accuracy of at least 70% for different system sizes. We consider a maximum of 5000 experiments. For system size $n \ge 10$, we are unable to achieve at least 70\% prediction accuracy with 5000 experiments. Hence we only show system size n=2,4,6,8 for conventional experiments in Fig. 2(c).

A.2.c. Quantum-enhanced experiments

Quantum-enhanced experiments are executed by performing an entangling Bell measurement across two copies of ρ . We prepare the state ρ on the system qubits (marked blue in Supp. Fig. 1), swap the state to the memory qubits (marked red), prepare another state ρ on the system qubits, then perform an entangling Bell measurement across the two copies of ρ . Note that every preparation of ρ generates a random product state according to a classical probability distribution described in Appendix A.2.a.

For each system size $n=2,\ldots,20$, we choose n qubits from among the 20 pairs of qubits shown in Supp. Fig. 1(b); these pairs are selected to minimize errors in the state preparations, gates, and measurements. Because the state ρ is not entangled, no entangling gates are used during the state preparation; therefore there is no advantage in choosing the pairs of qubits to be in proximity to one another. For each system size n, we use the same qubits for the conventional experiments as for the quantum-enhanced experiments, except that in conventional experiments we prepare the unknown state ρ only on the system qubits, and then perform a randomized Pauli measurement of the system immediatly after the state preparation.

While a Bell measurement on a pair of qubits can be performed via a simple circuit containing one Hadamard

gate, one CNOT gate, and two Z-basis measurements, we instead compile these operations into operations better suited for the Sycamore processor. In particular, the native two-qubit entangling gate is the Sycamore Gate, which has a unitary matrix representation given by

$$SYC = iSWAP^{\dagger} CPHASE(-\pi/6) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & -i & 0 & 0 \\ 0 & 0 & 0 & e^{-i\pi/6} \end{pmatrix}.$$
(A1)

Using this gate, the Bell state measurements may be performed by the gate sequence needed to unprepare a Bell state, which when compiled to Sycamore's native gates is expressed as a product of SYC gates and phased XZ gates ($PhXZ_i$). The phased XZ gate $PhXZ_i$ on the *i*-th qubit is a native gate on the Sycamore device that can be expressed as

$$PhXZ(a, x, z)_i = Z_i^z Z_i^a X_i^x Z_i^{-a}.$$
(A2)

where X_i and Z_i are the standard Pauli operators acting on qubit i, and the exponents a, x, z are real numbers. The particular angles (a, x, z) for each of the gates used in our experiments were compiled numerically via a variational optimization. As the inverse Sycamore is not a native gate of the architecture, the compilations to hardware for inverse gates have to compensate for this difference, which we do numerically. The full decomposition of all the gates and circuits we reference are provided as Cirq circuits in additional supplemental material.

Each quantum-enhanced experiment generates a classical bitstring of size 2n. We collect the bitstrings from all experiments and feed them into a neural network model. In addition to the experimental data, the neural network model also takes in two Pauli strings Q_1, Q_2 and predicts which one of $|\operatorname{tr}(Q_1\rho)|$ and $|\operatorname{tr}(Q_2\rho)|$ is larger. In Fig. 2(b), we repeat each quantum-enhanced experiments 500 times. This provides a fair comparison with conventional experiments because two copies of the unknown state ρ are used in each quantum-enhanced experiment; therefore a total of 1000 copies are consumed in both our conventional and quantum-enhanced experiments. In Fig. 2(c), we repeat quantum-enhanced experiments for a maximum of 500 times over different system sizes from n=2,4,6,8,10,12,14,16,18,20.

A.2.d. A supervised neural network model using data from quantum-enhanced experiments

We train a supervised neural network model using noiseless simulation data from small system sizes. Then we use the trained neural network model on the noisy experimental data obtained from performing quantum-enhanced experiments. The neural network model has three layers. Each of the outer layers runs the preceding inner layer multiple times. In the following, we describe each layer of the neural network model.

- 1. The *inner layer* is a recurrent neural network based on gated recurrent unit (GRU) [20, 21, 25]. The recurrent neural network takes in a size-2n bitstring, corresponding to the measurement outcome from a single quantum-enhanced experiment, and an n-qubit Pauli operator Q, which can be represented as a size-2n bitstring. The recurrent neural network outputs a two-dimensional real vector. Other popular choices of recurrent units, such as long short-term memory (LSTM) [32] or transformer [33], could be used instead of GRU.
- 2. The intermediate layer is an aggregation layer. This layer runs the inner layer for all the bitstrings obtained from each of the quantum-enhanced experiments. For example, if we run the quantum-enhanced experiments for 100 times, we would obtain 100 size-2n bitstring and we would run the inner layer for 100 times over each bitstring. The intermediate layer outputs the average of the two-dimensional output vectors from the multiple runs of the inner layers.
- 3. The outer layer is inspired by a Siamese neural network (twin neural network) [34]. This layer runs the intermediate layer twice, one for each of the two Pauli operators Q_1 and Q_2 . Each intermediate layer generates a real-valued vector x of dimension two, which we map to a single real value by considering $x_1 x_2$. The outer layer compares the real value from the two intermediate layers and outputs Q_1 or Q_2 based on which one of them has a higher real value.

The specific details of the above neural network structure is given in the accompanying code repository¹.

¹ https://github.com/quantumlib/ReCirq/tree/master/recirq/qml_lfe

In the inner layer, we create a recurrent neural network with an encoding layer that maps an integer between 0 and 15 to a vector of dimension 30, a GRU with 30 neurons, and a decoding layer that maps 30 neurons to 2 neurons. Only the inner layer contains trainable parameters. The intermediate layer and the outer layer are both fixed operations based on outputs from the inner layer, which will not be updated.

Next, we discuss the process for training the neural network model. We use noiseless simulation data (for small system sizes n < 8) to train the recurrent neural network. During training, we pick a state $\rho = 2^{-n} (I + \alpha P)$ where $\alpha \in \{-0.95, 0.95\}$ and P is an n-qubit Pauli operator, and pick an n-qubit Pauli operator Q that is equal to P with probability 1/2 and is not equal to P with probability 1/2. We encode the training data into two tensors, inp and target. The encoding is defined by the following.

- The tensor inp is of size $b \times n$, where b is the number of quantum-enhanced experiments we performed, n is the number of qubits, and each entry of inp is an integer from 0 to 15. The (t, i)-th entry of inp encodes the component of Q on qubit i (a choice of 4 for I, X, Y, Z) and the Bell measurement outcome on qubit i from the t-th quantum-enhanced experiment (also a choice of 4). Each entry takes a total of 16 possible values.
- The tensor target is of size 1. The entry in target is equal to 1 if P=Q, and is equal to 0 if $P\neq Q$.

We update the neural network model once using inp of size $b \times n$ and target of size 1. We are using the cross entropy loss and employ the Adam optimizer [35], which is a gradient-based optimization algorithm that adaptively estimates lower-order moments. We generate multiple different states ρ and Q corresponding to different inp and target to train the neural network model.

During the training process, we are not using the *outer layer*. Also, we simultaneously run the *b* repetitions of the inner layer for each outcome from a single quantum-enhanced experiment by leveraging parallel computing. Then, we average over the *b* repetitions of the inner layer. Also, the output of the neural network model is a two-dimensional real vector, denoted as $v = (v_0, v_1)$. When target is $a \in \{0, 1\}$, the loss function is given by

$$-\log\left(\frac{\mathrm{e}^{v_a}}{\mathrm{e}^{v_0} + \mathrm{e}^{v_1}}\right). \tag{A3}$$

The two real values v_0 , v_1 are combined to produce a probability distribution

$$\frac{e^{v_0}}{e^{v_0} + e^{v_1}} = 1 - \frac{1}{e^{v_0 - v_1} + 1}, \quad \frac{e^{v_1}}{e^{v_0} + e^{v_1}} = \frac{1}{e^{v_0 - v_1} + 1}, \tag{A4}$$

indicating which of a=0 and a=1 is more likely. If v_0-v_1 is large, then a=0 corresponding to $P\neq Q$ is more likely. On the other hand, if v_0-v_1 is small, then a=1 corresponding to P=Q is more likely. We compute the gradient through back-propagation and update the model using the Adam optimizer [35].

Finally, we discuss the prediction process in the neural network model. Due to the significant amount of measurement errors, we employ a form of measurement error mitigation. We first characterize the measurement errors for every qubit assuming the zero state preparations and X-gates are perfect. For each qubit i, we obtain a 2×2 matrix specifying the probability to measure 0 or 1 if the qubit is in $|0\rangle\langle 0|$ or $|1\rangle\langle 1|$. We store that as a list of 2×2 matrices called calib_2x2. We then expand the data, referred to as data in the pseudo-code, obtained from the quantum-enhanced experiments, which is a two-dimensional array of size $b \times (2n)$. Basically, we expand each measurement to 20 measurements with a real-valued coefficient associated to each of the expanded measurements. Therefore, data_expanded is a two-dimensional array of size $(20b) \times (2n)$ and coefficients is a one-dimensional array of size 20b.

```
def noise_inversion(data, calib_2x2, inverse_cnt=20):
    Set data_expanded as an empty array
    Set list_of_coefficients as an empty array

for t from 0 to b-1:
    for r from 0 to inverse_cnt-1:
        Set single_data as an empty array
        Set coefficient as 1.0

    for i from 0 to 2n-1:
        Set p as calib_2x2[i][1, 1] if data[t][i] = 0
        Set p as calib_2x2[i][0, 0] if data[t][i] = 1
```

```
With probability 1—p do:
    single_data.append(1—data[t][i])
    coefficient *= -1

Else do:
    single_data.append(data[t][i])

Append single_data to data_expanded
    Append coefficient to list_of_coefficients

return data_expanded, list_of_coefficients
```

After obtaining data_expanded, we construct two tensors inp1 and inp2 corresponding to the same experimental data data_expanded, but different Pauli operators Q_1, Q_2 . Both inp1 and inp2 are tensors of size $(20b) \times n$, where b is the number of quantum-enhanced experiments we performed, n is the number of qubits, and each entry of inp1 and inp2 is an integer from 0 to 15 similar to the training process. Then the neural network make a prediction using the two input tensors inp1 and inp2. In the outer layer, the neural network model runs the intermediate layer (as well as the multiple repetitions of inner layer) for each of the two input tensors to obtain two 2D vectors denoted as $u = (u_0, u_1), v = (v_0, v_1)$. From the discussion given around Eq. (A4). If $u_0 - u_1$ is small, then it is more likely that $P = Q_1$. If $v_0 - v_1$ is small, then it is more likely that $P = Q_1$. The neural network hence compare $u_0 - u_1$ and $v_0 - v_1$ to predict whether $P = Q_1$ or $P = Q_2$.

A.3. Experiments on learning physical dynamics

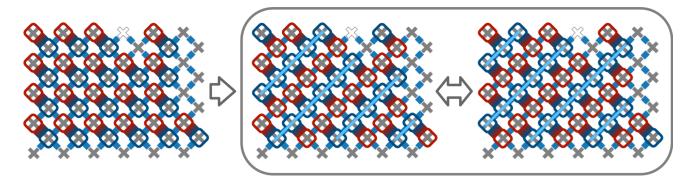
For the task of learning about physical dynamics in 1D and 2D, we considered unitary transformations implemented by 1D and 2D random quantum circuits. We generated many random circuits, half of which are time-reversal symmetric (i.e., real orthogonal), and half of which are general unitary circuits without any symmetry. For each of these circuits, we performed both conventional and quantum-enhanced experiments to generate classical measurement data. This data was fed to an unsupervised machine learning model to learn a low-dimensional classical representation of the physical dynamics. We wished to see whether the unsupervised ML model could recognize the difference between time-reversal symmetric dynamics and general dynamics. The results summarized in Fig. 3 were obtained in experiments analyzing 180 different circuits in each of the the two classes, using methods described below. The largest quantum circuits we ran on the Sycamore processor are presented in Table I.

	Number of qubits	Number of gates	Circuit depth
1D dynamics	40	842	40
2D dynamics	40	1388	54

Table I: Circuit information for the experiments on learning physical dynamics.

In [9], a restricted subclass of conventional strategies was shown to require an exponential number of experiments to distinguish between general unitary dynamics and time-reversal-symmetric dynamics. In [12], some of the authors of the present work have shown that an exponential number of experiments are required for this task even when arbitrary conventional strategies are allowed. Furthermore, it is plausible that under appropriate cryptographic assumptions, the superpolynomial difficulty of characterizing quantum dynamics in conventional experiments would persist even for psuedo-random dynamical processes that can be efficiently generated on a quantum computer. However, at present, explicit constructions of cryptographically-secure pseudo-random unitaries are not known. In light of this, in our experiments we resort to studying random quantum circuits similar to those used for demonstrating quantum computational supremacy [10].

In this subsection, we provide further details regarding how our samples of 1D dynamics and 2D dynamics are generated, how our experiments are conducted, and how our unsupervised machine learning model works. As we will see, the unsupervised ML successfully learns to classify quantum circuits into symmetry classes when provided with data from quantum-enhanced experiments, but not when provided with data from conventional experiments.



Supplementary Figure 2: To implement the 2D dynamics, we first move from the layout on the left to the layout in the middle (by swapping some pairs of system and memory qubits). Then, we iterate between the layout in the middle and the layout in the right (by swapping all pairs of system and memory qubits).

A.3.a. 1D dynamics

We use the layout provided in Supp. Fig. 1(b). A 1D circuit is implemented along the 1D line connecting all the qubits circled in blue. For system size n < 20, we consider a contiguous region on the 1D line with the smallest gate/measurement error. For general unitary dynamics, we use a 1D version of the quantum supremacy circuit [10]. The quantum supremacy circuit in 1D interleaves between a layer of random single-qubit gates and a layer of two-qubit entangling gates, namely SYC gates applied to neighboring qubits. We alternate the partitioning of the two-qubit entangling gates, e.g., $(1,2), (3,4), (5,6) \longleftrightarrow (2,3), (4,5)$ for n=6.

For T-symmetric (time-reversal-symmetric) dynamics, the single-qubit gates are real orthogonal 2×2 matrices of the form e^{-itY} , where Y is the Pauli-Y matrix and t is a randomly chosen real number. In addition, we replace the two-qubit entangling gate SYC by a T-symmetric two-qubit entangling gate

$$V = (U_3 \otimes U_4) SYC(U_1 \otimes U_2), \tag{A5}$$

where U_1, U_2, U_3, U_4 are appropriately chosen single-qubit gates. In order to find a suitable choice of U_1, \ldots, U_4 such that the two-qubit gate V is time-reversal symmetric, we employ a numerical optimization. We parameterize each of U_i as $\exp(\mathrm{i}(a_iX+b_iY+c_iZ))$, where $a_i, b_i, c_i \in \mathbb{R}$ are initialized randomly. There are a total of 12 variables. Next, we define the loss function to be equal to the Frobenius norm of the imaginary part of the matrix V (after fixing the top left entry of V to be real). We then perform gradient descent to minimize the loss function and terminate once we have found that the loss function is below 10^{-9} . We then replace SYC by V.

In the conventional experiment, we begin with $|0^n\rangle\langle 0^n|$ on the system qubits, evolve under the 1D dynamics, and measure in the Y-basis. We also considered using randomized Pauli measurement at the end, but the performance for measuring in the Y-basis is slightly better. The rationale is that the output state under T-symmetric evolution has purely real amplitudes; hence the expectation value of any purely imaginary observable, such as the Pauli-Y operator, is always zero. In contrast, the expectation value of Y after a general unitary evolution is non-zero. T-symmetric unitaries are nevertheless hard to distinguish from general unitaries in the conventional setting because the expectation value of Y is exponentially small for general unitaries; theefore an exponentially large number of experiments are needed to discern its nonzero value. Indeed, the result in [9, 12] shows that conventional strategies require an exponential number of experiments for distinguishing T-symmetric unitaries from general unitaries.

require an exponential number of experiments for distinguishing T-symmetric unitaries from general unitaries. In the quantum-enhanced experiment, we prepare a Bell state $\frac{1}{\sqrt{2}}(|00\rangle+|11\rangle)$ for every pair of system and memory qubits. Then we evolve the system qubits under the unknown dynamics. After the evolution, we swap the system and the memory qubits. Then we evolve the system qubits under the unknown dynamics again. Finally, we measure every pair of system and memory qubits in the Bell basis. Each quantum-enhanced experiment generates a 2n-bit string. We perform gradient descent to find our implementation of the Bell state preparation, swap operation, and Bell measurement using the native gates in the Sycamore processor.

A.3.b. 2D dynamics

For our 2D circuits we use the layout provided in Supp. Fig. 1(c) which is also shown as the leftmost layout in Supp. Fig. 2. In the leftmost layout, none of the system qubits (circled blue) are connected to one another. In order

to implement 2D dynamics, we first swap some pairs of the system and memory qubits to obtain the layout shown in the middle. In the middle layout, we can see that many of the system qubits are connected (the light blue line). We implement a depth-4 1D random quantum circuit for each light blue line. Each depth-4 circuit corresponds to 1 layer of single-qubit gates, 1 layer of two-qubit gates, 1 layer of single-qubit gates, and 1 layer of two-qubit gates. The partitioning of the two layers of two-qubit gates are different. Then, we swap all pairs of qubits to obtain the layout shown on the right. The right-most layout connects a different set of system qubits (the light blue line). We again implement a depth-4 1D random quantum circuit for each light blue line. After that, we move back to the middle layout and repeat for multiple rounds. After sufficiently many repetitions, the n qubits become globally entangled.

A.3.c. An unsupervised machine learning model

For each circuit, we create a feature vector by obtaining statistics for each bit in the measurement outcome bitstring. In conventional experiments, each experiment produces a n-bit measurement outcome. In quantum-enhanced experiments, each experiment produces a 2n-bit measurement outcome. In the following, we consider $\ell = n$ or 2n depending on either we are running conventional or quantum-enhanced experiments. For an ℓ -bit measurement outcome, we obtain a feature vector of size 2ℓ including the first and second moment of each bit. After constructing a feature vector for each circuit, we map the feature vector to an infinite-dimensional reproducing kernel Hilbert space (corresponding to a Gaussian kernel) that includes all the polynomial expansions of the feature vector. Then, we find a low-dimensional subspace in the infinite-dimensional Hilbert space using principal component analysis (PCA) [23]. The entire procedure can be performed efficiently using kernel PCA [23]. Kernel PCA is implemented using scikit-learn [36].

In Fig. 3 of the main text, we show a one-dimensional subspace found by the unsupervised ML model, for both 1D and 2D random quantum circuits. We can use this one-dimensional representation to classify the circuits into two classes (by splitting the one-dimensional representation in the middle). Then, we can evaluate the accuracy of the unsupervised ML model by checking the percentage of circuits that are correctly classified as general circuits or as T-symmetric circuits. A two-dimensional subspace found by the unsupervised ML model, and an assessment of classification accuracy, are discussed in Appendix A.4

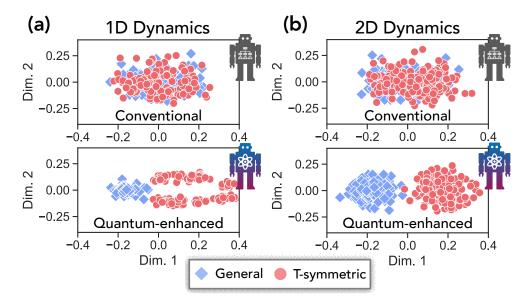
A.4. Additional experimental results

In Supp. Fig. 3, we provide the two-dimensional representations learned by unsupervised ML for the various random quantum circuits investigated in our conventional and quantum-enhanced experiments. (One-dimensional representations are presented in Fig. 3 in the main text.) We see that in the second dimension found by unsupervised ML using quantum-enhanced experiments for 1D dynamics, T-symmetric dynamics are clustered into two groups. Further inspection shows that the unsupervised ML model has learned substructure corresponding to the parity of the depth of the evolution (recall that the depth is always an integer). In principle, the unsupervised ML model should be able to learn a wide variety of structures in the dynamics. Notably, we see that it places the distinction between general unitary dynamics and T-symmetric dynamics as the major axis (the first dimension), and places less prominent structure as the second major axis (the second dimension).

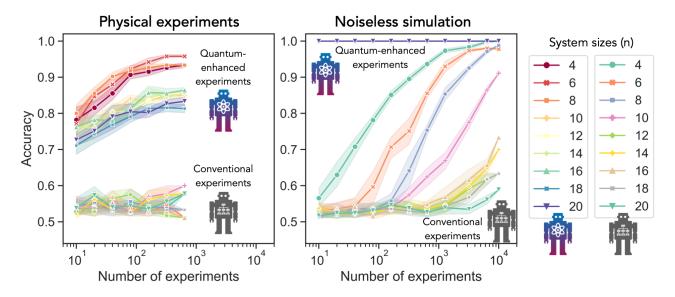
In Supp. Fig. 4, we provide the accuracy of the unsupervised ML model for distinguishing between general unitary dynamics and T-symmetric dynamics. We see a substantial advantage for using the quantum-enhanced strategy in both the physical experiments and the noiseless simulation. We perform brute-force noiseless simulation for conventional experiments because the system size is at most 20. The noiseless simulation for quantum-enhanced experiments uses the fact that $(U \otimes U) \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |ii\rangle = (UU^T \otimes I) \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |ii\rangle$, hence we can effectively reduce the simulation to a system size at most 20.

A.5. Performance and characterization data

The performance of the device was characterized before each run. The measurement data is collected explicitly and used for measurement error mitigation in the prediction process of the supervised neural network model (see discussion in the last part of Appendix A.2.d). The task of learning quantum states is largely limited by the qubit measurement fidelities. A representative sample of the data from the device is reported in Fig. 5, where one can see that typical readout errors (conflated with errors in preparing zero and one states) range from 3% to 7%.



Supplementary Figure 3: Two-dimensional representation learned by unsupervised ML for (a) 1D dynamics and (b) 2D dynamics. Each point in the two-dimensional plane corresponds to a distinct physical process. Half of the processes have time-reversal symmetry (blue diamonds) while the other half do not (red circles). When fed with data from quantum-enhanced experiments, the ML model accurately discovers the underlying symmetry pattern. In contrast, the ML model fails to do so when fed with data from conventional experiments.



Supplementary Figure 4: Accuracy of the unsupervised ML model for classifying general unitary and T-symmetric dynamics. For each system size, we generate 100 different circuits for each of the two classes (general and T-symmetric). The one-dimensional representation found by the unsupervised ML model is used to classify the 200 circuits into two classes. We consider both physical experiments and noiseless simulations. Accuracy is plotted as a function of the number of experiments in both the conventional and quantum-enhanced settings.

For transmons, the $|0\rangle$ preparation has a small error; hence Supp. Fig. 5(a) is dominated by the readout error. Furthermore, the single-qubit gate error (shown in Supp. Fig. 6) is much smaller than the error shown in Fig. 5(b), hence the error shown in Supp. Fig. 5(b) is mostly due to readout errors rather than gate errors. During the actual run of the experiments, we avoid using qubits with the worst readout errors by checking the measurement errors before the experiment and selecting the layout accordingly.

The task of learning quantum dynamics involves circuits of higher complexity and hence is limited by both

measurement errors and errors in two-qubit gates. For these experiments, we report in Supp. Fig. 6 the quality of the single-qubit and two-qubit gates across the device. This data was obtained via parallel cross-entropy benchmarking and single-qubit randomized benchmarking. The typical single-qubit gate error is around 0.001 to 0.005, while the typical two-qubit gate error is around 0.01 to 0.05.

B. A brief review on quantum information theory

In this section we review some relevant definitions and basic results in quantum information theory which are leveraged throughout our problem statements and proofs. Specifically, we will discuss quantum processes, which are a general mathematical formalism for describing physical processes, and positive operator-valued measures (POVMs), which encompass all possible physical measurements. Readers familiar with these concepts can skip this section.

B.1. Definition and properties of quantum processes

For concreteness, let us consider a Hilbert space $\mathcal{H}_S \simeq \mathbb{C}^d$. Here the subscript S stands of 'system', since the Hilbert space describes the space of states of some particular system we wish to study. Given a density matrix ρ on this Hilbert space, we might ask: how can it evolve in time? The Schrödinger equations tells us that a state can evolve via unitary time evolution, and as such a density matrix can evolve by $\rho \mapsto U\rho U^{\dagger}$. However, there is a more general type of time evolution allowed by quantum mechanics. Suppose that we append our Hilbert space by another $\mathcal{H}_E \simeq \mathbb{C}^{d'}$ which describes an external environment. The joint Hilbert space is then the tensor product $\mathcal{H}_S \otimes \mathcal{H}_E$. We can imagine having an initial state $\rho_S \otimes \rho_E$ which factorizes between the system and environment, and then evolving the state by a unitary on the joint Hilbert space which couples the system and environment: $\rho_S \otimes \rho_E \mapsto U_{SE}(\rho_S \otimes \rho_E)U_{SE}^{\dagger}$. If we only have access to \mathcal{H}_S , then our knowledge of $U_{SE}(\rho_S \otimes \rho_E)U_{SE}^{\dagger}$ is described by performing a partial trace over the environment, namely $\operatorname{tr}_E\left(U_{SE}(\rho_S \otimes \rho_E)U_{SE}^{\dagger}\right)$. As such, if we are only aware of the initial density matrix ρ_S on \mathcal{H}_S , then only having access to \mathcal{H}_S the time evolution would appear to be

$$\rho_S \longmapsto \operatorname{tr}_E \left(U_{SE}(\rho_S \otimes \rho_E) U_{SE}^{\dagger} \right) .$$
(B1)

Note that, viewed as time evolution on ρ_S alone, the above map is not unitary. This is because information in ρ_S can leak into the environment, and similarly information from the environment can influence the state on our system S of interest. This mapping is an example of a quantum process, which can be more compactly notated as $\rho_S \mapsto \mathcal{C}[\rho_S]$. Here, \mathcal{C} is a map from density matrices on \mathcal{H}_S to (other) density matrices on \mathcal{H}_S . We visualize this dynamical process in Supp. Fig. 7.

Our quantum process \mathcal{C} has two properties that are worth highlighting:

1. C is trace-preserving. This means that $\operatorname{tr}(C[\rho_S]) = \operatorname{tr}(\rho_S)$. The equality follows from the definition of $C[\rho]$ via the right-hand side of (B1), since

$$\operatorname{tr}(\mathcal{C}[\rho_S]) = \operatorname{tr}_S\left(\operatorname{tr}_E\left(U_{SE}(\rho_S \otimes \rho_E)U_{SE}^{\dagger}\right)\right) = \operatorname{tr}\left(U_{SE}(\rho_S \otimes \rho_E)U_{SE}^{\dagger}\right)$$
(B2)

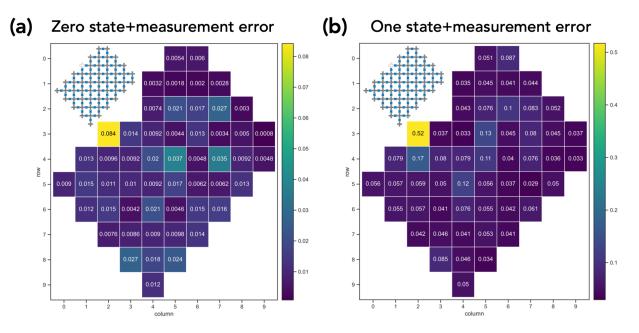
$$= \operatorname{tr}(\rho_S \otimes \rho_E) = \operatorname{tr}(\rho_S) \operatorname{tr}(\rho_E) = \operatorname{tr}(\rho_S), \tag{B3}$$

where we have used the cyclicity of the trace to cancel U_{SE} with U_{SE}^{\dagger} , and have also leveraged $\operatorname{tr}(\rho_E) = 1$.

2. C is completely positive. Suppose we append our system Hilbert space \mathcal{H}_S by ancillas \mathcal{H}_A to arrive at the joint Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_S$. Then complete positivity means that for any density matrix ρ_{AS} on this joint system (and for any choice of ancilla Hilbert space), $(\mathrm{Id}_A \otimes \mathcal{C})[\rho_{AS}]$ is positive-semidefinite; here Id_A acts as the identity on the ancillas. To see why this property holds, we can write out $(\mathrm{Id}_A \otimes \mathcal{C})[\rho_{AS}]$ more explicitly:

$$(\mathrm{Id}_A \otimes \mathcal{C})[\rho_{AS}] = \mathrm{tr}_E\Big((I_A \otimes U_{SE})(\rho_{AS} \otimes \rho_E)(I_A \otimes U_{SE}^{\dagger})\Big) \ .$$

Since the right-hand side is merely performing a unitary transformation on the density matrix $\rho_{AS} \otimes \rho_E$ and then tracing out a subsystem (i.e. the environment subsystem), positive semi-definiteness is preserved.

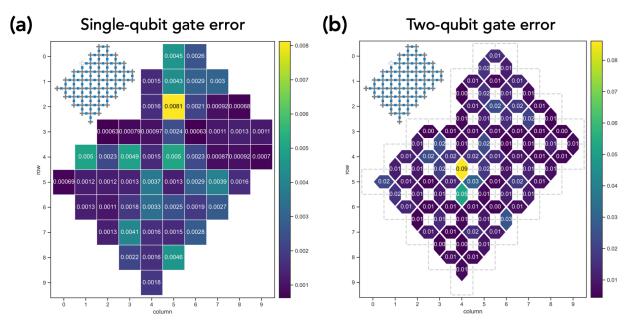


Supplementary Figure 5: Sycamore state preparation and measurement error data.

(a) $|0\rangle$ state preparation and measurement error. We prepare a noisy zero state $|0\rangle$ and measure in the computational basis using the noisy single-qubit readout. We show the probability of measuring $|1\rangle$ in the qubit readout.

(b) $|1\rangle$ state preparation and measurement error. We prepare a noisy one state $|1\rangle$ and measure in the computational basis using the noisy single-qubit readout. We show the probability of measuring $|0\rangle$ in the qubit readout.

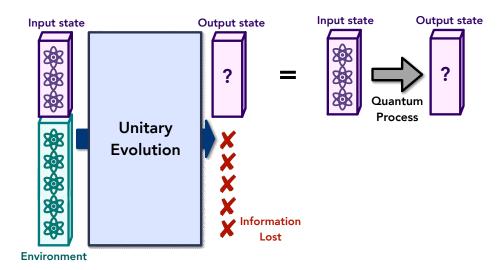
While these values change over time, we present here a representative sample of the error. One can see that in accordance with physical expectations based on T_1 errors, the readout in the physical 1 state is substantially higher than the 0 state.



Supplementary Figure 6: Sycamore single- and two-qubit gate error data.

(a) Single-qubit gate error. The figure shows the error of single-qubit gates across the chip using parallel single-qubit randomized benchmarking.

(b) Two-qubit gate error. The figure shows the error across the chip of two-qubit gates being executed in parallel, as to account for errors that occur during simultaneous operation of qubits. We can see that the distribution of errors varies across the chip couplers, showing the extent to which performance is non-uniform.



Supplementary Figure 7: *Illustration of quantum process: a formalism for describing physical processes.* Quantum process is also known as quantum operation or quantum dynamical map, and is often referred to as quantum channel in quantum communication theory.

We have thus shown that our \mathcal{C} is a completely positive, trace-preserving (CPTP) linear map from density matrices on \mathcal{H}_S to density matrices on \mathcal{H}_S . Henceforth, when we refer to a map as being CPTP, we will implicitly suppose that the map is linear. Moreover, we will interchangeably call a CPTP map a quantum process.

What is not immediately obvious is the following fact:

Theorem 4 (Stinespring dilation). Any CPTP map C taking density matrices on $\mathcal{H}_S \simeq \mathbb{C}^d$ to density matrices on $\mathcal{H}_S \simeq \mathbb{C}^d$ can be written in the form

$$\mathcal{C}[\rho_S] = \operatorname{tr}_E \left(U_{SE}(\rho_S \otimes \rho_E) U_{SE}^{\dagger} \right)$$

for any ρ_S , where U_{SE} , ρ_E , and the dimension of the environment d' are all fixed.

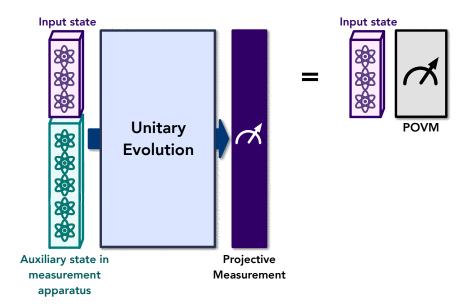
This theorem means that any CPTP map on a density matrix can be realized as a unitary operation on a larger system, i.e. coupling the density matrix to an appropriate environment and evolving the joint state and ultimately tracing out the environment. In this sense, a quantum process is the most general form of evolution of a density matrix. Note that a special case of a quantum process is simply a unitary channel, i.e. $\mathcal{C}[\rho] = U\rho U^{\dagger}$. A way of summarizing the above Theorem is that a quantum process that is not a unitary channel can be thought of as implementing open system dynamics.

B.2. Definition and properties of POVMs

The most conventional way to measure a quantum state $|\psi\rangle$ is by decohering it with respect to a complete orthonormal basis. More specifically, suppose that $|\psi\rangle\in\mathbb{C}^d$ and we choose some complete orthonormal basis $\{|i\rangle\}_{i=0}^{d-1}$ of \mathbb{C}^d . Then upon measuring $|\psi\rangle$ with respect to this basis, we will measure $|\psi\rangle$ to be in the state $|i\rangle$ with probability $\operatorname{Prob}(i) = |\langle i|\psi\rangle|^2$. Analogously for a density matrix ρ on the same Hilbert space, if we measure it with respect to the same orthonormal basis we will measure the state to be $|i\rangle\langle i|$ with probability $\operatorname{Prob}(i) = \operatorname{tr}(|i\rangle\langle i|\rho)$.

There is a nice way of conceptualizing measurements which will admit useful generalizations. First, let us develop some notation. We define $\Pi_i = |i\rangle\langle i|$ which is the projector onto state $|i\rangle$, and will speak of the collection of projectors $\{\Pi_i\}_{i=0}^{d-1}$. It is readily seen that $\sum_{i=0}^{d-1}\Pi_i=I$ since this is just a resolution of the identity. Observe that each Π_i is Hermitian and positive semi-definite. Now suppose we append to our Hilbert space another copy \mathbb{C}^d . Then we can define a unitary on both copies which acts by

$$U(|\psi\rangle \otimes |0\rangle) = \sum_{i=0}^{d-1} \Pi_i |\psi\rangle \otimes |i\rangle$$
 (B4)



Supplementary Figure 8: *Illustration of POVM: a formalism encompassing all physical measurements.* POVM considers a composition of the input state with an auxiliary state in the measurement apparatus (which can be thought of as a set of ancilla qubits) that undergoes an unitary evolution, followed by a projective measurement.

for any $|\psi\rangle$. Note that, as required of a unitary,

$$\left(\langle \psi | \otimes \langle 0 | \right) U^{\dagger} U \left(| \psi \rangle \otimes | 0 \rangle \right) = \sum_{i,j=0}^{d-1} \langle \psi | \Pi_i \Pi_j | \psi \rangle \langle i | j \rangle = \sum_{i=0}^{d-1} \langle \psi | \Pi_i^2 | \psi \rangle = \sum_{i=0}^{d-1} \langle \psi | \Pi_i | \psi \rangle = 1$$

on account of $\Pi_i^2 = \Pi_i$ and $\sum_{i=0}^{d-1} \Pi_i = I$. Given the right-hand side of (B4), we can make a measurement on the appended Hilbert space in the $\{|i\rangle\}_{i=0}^{d-1}$ basis; we will then measure the appended register to be in the state $|i\rangle$ with probability

$$Prob(i) = (\langle \psi | \otimes \langle 0 |) U^{\dagger} (I \otimes |i\rangle \langle i|) U(|\psi\rangle \otimes |0\rangle) = tr(\Pi_i |\psi\rangle \langle \psi|) = |\langle i|\psi\rangle|^2.$$
 (B5)

Similarly, if we consider $\rho \otimes |0\rangle\langle 0|$, conjugate by U, and then measure the state of the ancilla, the probability of measuring the ancilla to be $|i\rangle$ is $\text{Prob}(i) = \text{tr}(\Pi_i \rho) = \text{tr}(|i\rangle\langle i|\rho)$.

We can think about the above in terms of the following procedure. First we prepare a state $|\psi\rangle$; then we bring in an ancilla $|0\rangle$ and cause the two states to interact such that the ancilla goes into a state $|i\rangle$ upon coupling with the $|i\rangle$ -component of $|\psi\rangle$. This leads to the right-hand side of (B4). The ancilla can be thought of as a proxy for the readout of a measurement apparatus: upon reading off the value of $|i\rangle$, we are informed that the state $|\psi\rangle$ has been projected into its $|i\rangle$ -component.

This type of procedure can be generalized as follows. Suppose we have a set of N $d \times d$ operators $\{M_i\}_{i=0}^{N-1}$ satisfying the completeness relation $\sum_{i=0}^{N-1} M_i^{\dagger} M_i = I$. Let us append to our Hilbert space \mathbb{C}^d and ancillary Hilbert space \mathbb{C}^N with complete orthonormal basis $\{|i\rangle\}_{i=0}^{N-1}$. Then we can consider a unitary map

$$U(|\psi\rangle \otimes |0\rangle) = \sum_{i=0}^{N-1} M_i |\psi\rangle \otimes |i\rangle.$$
 (B6)

The fact that $(\langle \psi | \otimes \langle 0 |) U^{\dagger} U(|\psi \rangle \otimes |0 \rangle) = 1$ can be checked using the completeness relation $\sum_{i=0}^{N-1} M_i^{\dagger} M_i = I$. Now if we measure the ancilla with respect to the $\{|i\rangle\}_{i=0}^{N-1}$ basis, then we will measure the ancilla to be in the state $|i\rangle$ with probability $\operatorname{Prob}(i) = |M_i|\psi\rangle|^2$. If we performed an analogous procedure at the level of density matrices, namely starting with a state $\rho \otimes |0\rangle\langle 0|$, conjugating both sides by U, and then measuring the ancilla in the $\{|i\rangle\}_{i=0}^{N-1}$ basis, we would measure the ancilla to be $|i\rangle$ with probability $\operatorname{Prob}(i) = \operatorname{tr}(M_i^{\dagger} M_i \rho)$. We visualize the above procedure in Supp. Fig. 8.

We can abstract this procedure into what is called a positive operator-valued measure (POVM):

Definition 1 (POVM). A POVM is a set Hermitian, positive semi-definite operators $\{F_i\}_{i=0}^{N-1}$ on \mathbb{C}^d satisfying the completeness relation $\sum_{i=0}^{N-1} F_i = I$. A POVM measurement is a procedure in which, given a state ρ on \mathbb{C}^d , an analysis of the procedure of the proced ancillary measurement apparatus registers the index i with probability $tr(F_i\rho)$.

This relates to our previous procedure as follows. We simply decompose $F_i = M_i^{\dagger} M_i$ (say, by a Cholesky decomposition) and perform the procedure previously stated with the M_i 's.

We remark that the term 'measure' is used above in two distinct ways. When we speak of a POVM, the M means measure in the sense of measure theory, since we can think of $\{F_i\}_{i=0}^{N-1}$ as comprising a type of discrete measure on the space of operator on \mathbb{C}^d . Otherwise, we use 'measure' in the sense of measurement.

A useful fact is that given a POVM $\{F_i\}_{i=0}^{N-1}$, we can refine it into another, larger POVM $\{F_{i,j}\}_{i=0,j=0}^{N-1,d-1}$ such that (1) each $F_{i,j}$ is rank-1, and (2) a POVM measurement of $\{F_{i,j}\}_{i=0,j=0}^{N-1,d-1}$ can simulate a POVM measurement of $\{F_i\}_{i=0}^{N-1}$. Let us explain this construction. Since each F_i is a positive semi-definite Hermitian operator, we can diagonalize each operator as $F_i = \sum_{j=0}^{d-1} \lambda_j^{(i)} |v_j^{(i)}\rangle\langle v_j^{(i)}|$. Then let $F_{i,j} := \lambda_j^{(i)} |v_j^{(i)}\rangle\langle v_j^{(i)}|$ which is manifestly positive semi-definite, Hermitian, and rank-1; it is also clear that $\sum_{i=0}^{N-1} \sum_{j=0}^{d-1} F_{i,j} = \sum_{i=0}^{N-1} F_i = I$. We can use a POVM measurement of $\{F_{i,j}\}_{i=0,j=0}^{N-1,d-1}$ to simulate a POVM measurement of $\{F_i\}_{i=0}^{N-1}$ by simply summing measurement

$$\sum_{j=0}^{d-1} \operatorname{tr}(F_{i,j}\rho) = \operatorname{tr}(F_i\rho). \tag{B7}$$

Accordingly, we can without loss of generality choose to work with rank-1 POVMs, since we can use these to simulate any other POVMs.

Mathematical framework for proving exponential advantage

One of the central ingredients for establishing exponential advantage is to prove an exponential lower bound for any learning algorithm with only external classical memory. In this section, we present the basic framework for proving such lower bounds. This framework enables us to establish a suite of exponential advantage in various physically-relevant tasks, which will be presented after this section. The purpose of this section is to provide the readers with the essential tools to prove quantum advantage in the tasks they want to study.

The basic tools include the tree representation of learning algorithms (Appendix C.1), reduction to distinguishing tasks and information-theoretic lower bounds (Appendix C.2.a and C.3). Many of these techniques were introduced and leveraged in [12]. We also present a novel partially-revealed many-versus-one distinguishing task in Appendix C.4 that is crucial for realizing the advantage in practice. Then, in Appendix C.5, we discuss how having noise in the unknown physical states and dynamics only makes the lower bounds for conventional experiments larger. This result on the presence of noise is simple to establish but also crucial in practice because there is often noise in the unknown physical states and dynamics. There are some other techniques presented in a theory paper written by some of the authors [12], such as a multi-linear tensor analysis on the learning tree, which may be of interest to some readers.

C.1. Tree representation

We begin by presenting the tree representation for analyzing algorithms with only classical memory [12]. The key idea is to track changes in the classical memory state in the algorithm using a graph, which we can take to be a rooted tree. We consider each node u of the graph to be a classical memory state. Based on the memory state, the algorithm performs an experiment to obtain a measurement outcome s.

C.1.a. Experiments for learning physical world

To motivate the definitions in the sequel, we separately describe the two types of experimental setups that we focus on in this work: one on learning an unknown physical state and the other on learning an unknown process [8, 9, 12].

• Learning an unknown physical state: A physical state is represented by a density matrix ρ . An algorithm leveraging the classical memory state u measures the physical system ρ using a rank-1 POVM $\{w_s^u|\phi_s^u\rangle\langle\phi_s^u|\}$ with $\sum_s w_s^u|\phi_s^u\rangle\langle\phi_s^u|=I$. Note that from the discussion in Appendix B.2, we can always consider rank-1 POVMs only. The measurement outcome s occurs with probability

$$w_s^u \langle \phi_s^u | \rho | \phi_s^u \rangle$$
. (C1)

Here, the rank-1 POVM $\{w_s^u | \phi_s^u \rangle \langle \phi_s^u | \}$ depends on the classical memory state u.

• Learning an unknown physical process: A physical process is represented by a quantum process \mathcal{E} (equivalently, a CPTP map). An algorithm leveraging the memory state u prepares an initial state $|\psi^u\rangle$, feeds it into the physical evolution \mathcal{E} , and measures the output state $\mathcal{E}(|\psi^u\rangle\langle\psi^u|)$ with a rank-1 POVM $\{w_s^u|\phi_s^u\rangle\langle\phi_s^u|\}$ with $\sum_s w_s^u|\phi_s^u\rangle\langle\phi_s^u|=I$. The outcome s is obtained from the experiment with probability

$$w_s^u \langle \phi_s^u | \mathcal{E}(|\psi^u\rangle\langle\psi^u|) | \phi_s^u \rangle$$
 (C2)

In this case, both the initial state and the measurement depend on the classical memory state u.

We consider the initial state $|\psi^u\rangle$ to be an (n+n')-qubit state, where \mathcal{E} acts on the first n qubits. The rank-1 POVM $\{w_s^u|\phi_s^u|\phi_s^u|\phi_s^u\}$ is on an (n+n')-qubit state.

C.1.b. Dynamics of the learning algorithm

The classical memory state of the learning algorithm is initialized in a certain state, which we represent by the root node r. The memory state of the algorithm begins at the root r. Each measurement outcome s resulting from a single experiment causes the algorithm to transition to a node neighbor of r. Whenever the algorithm obtains a measurement outcome s, the memory state changes. This is represented by a transition from a node s to another node s,

$$u \xrightarrow{s} v.$$
 (C3)

The directed edge e = (u, v, s) from u to v represents the transition of the memory in an algorithm when we receive the measurement outcome s. We illustrate the transition under a single experiment in Supp. Fig. 9(a).

If a different s leads us to the same node, then the algorithm is not retaining full information of the measurement outcome. An example is given in Supp. Fig. 9(b). Since we do not limit the size of the classical memory, there is no need to lose (or forget) information. Hence, all the outgoing edges of the root node r indexed by the measurement outcome s will point to distinct nodes. The same argument holds for any node in the graph. More precisely, every outgoing edge from a node u will connect to a node v, such that v has exactly one incoming edge (the edge is from u). The only node in the graph without an incoming edge is the root r. This is exactly the definition of a directed rooted tree \mathcal{T} . We will focus on an algorithm that performs T experiments. This means the depth of the tree, namely the number of edges in any root-to-leaf path, will be T. The tree representation is shown in Supp. Fig. 9(c).

When we execute the algorithm to achieve a certain task (such as to verify entanglement, or learn a model of the physical system), the entire dynamic process of how the memory state changes will be represented by a path from the root r to a leaf ℓ in the tree \mathcal{T} ,

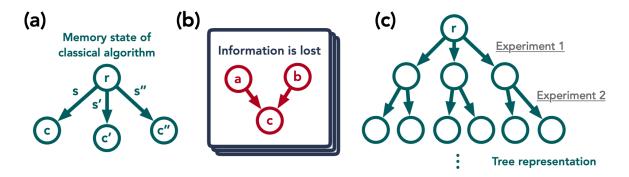
$$u_0 = r \xrightarrow{s_1} u_1 \xrightarrow{s_2} u_2 \xrightarrow{s_3} \dots \xrightarrow{s_{T-1}} u_{T-1} \xrightarrow{s_T} u_T = \ell. \tag{C4}$$

To establish a lower bound against any learning algorithm for a particular task, we need to analyze each such path along with the probability that the path is taken.

C.2. Many-versus-one distinguishing tasks

C.2.a. Reduction

In a learning task, we often want the learning algorithm to be able to make accurate predictions about some properties of the unknown physical system or dynamics. We will have a set of states (the mathematical representation of a physical system) or a set of channels (the mathematical representation of physical dynamics) to which we assume the unknown system or dynamics belong. The basic technique we employ in all of our proofs is to pick out one of the states/channels as the null hypothesis, and consider all the rest as the alternative hypothesis [8, 12]. Let \mathcal{X} denote the set of possible states/channels.



Supplementary Figure 9: Illustration of the tree representation for a learning algorithm. (a) Dynamics of memory. The memory state changes based on the measurement outcome s. (b) No cycles. If two memory states a, b transition into the same memory state c, then some information is lost. (c) Tree representation of an algorithm. When no information is lost, the transition graph of the memory states must be a directed tree. Each layer of the tree corresponds to one experiment. After T experiments, the memory state is represented by a node in the T-th layer.

- Null hypothesis: The unknown state/channel is an element $X_0 \in \mathcal{X}$. To establish a tight lower bound, we should choose an X_0 that we think is close to every other state/channel in \mathcal{X} .
- Alternative hypothesis: The unknown state/channel is a random element in $\mathcal{X} \setminus \{X_0\}$.

Furthermore, we need to choose X_0 such that the desired property we would like to learn enables us to distinguish between X_0 and the entire set of $\mathcal{X} \setminus \{X_0\}$.

To prove a lower bound against any classical algorithm, we try to answer the following question.

How hard is it to distinguish the alternative hypothesis from the null hypothesis?

Because the alternative hypothesis consists of many elements and the null hypothesis consists of only one element, we refer to this distinguishing task as the many-versus-one distinguishing task.

C.2.b. Information-theoretic lower bound

In order to establish a lower bound for the many-versus-one distinguishing task, we need to first discuss the leaf probability distribution in the tree representation of the learning algorithm [8, 12]. Recall that depending on the unknown state/process, the transition probabilities among the memory states in the learning algorithm will be different. This is because the outcome probability for each experiment differs when the unknown state/process differs. Therefore, the probability to traverse a certain path in the tree representing an execution of the learning algorithm will change according to the unknown state/process. Hence, the probability to arrive at a particular leaf node in the depth-T tree will change. An illustration is given in Supp. Fig. 10.

For each element X in \mathcal{X} , the set of all admissible states/channels, we write the probability distribution over leaves as

$$p_X(\ell), \quad \ell : \text{leaf node of the tree.}$$
 (C5)

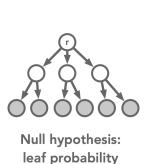
The probability distribution over the leaves ℓ for the null hypothesis and for the alternative hypothesis are respectively

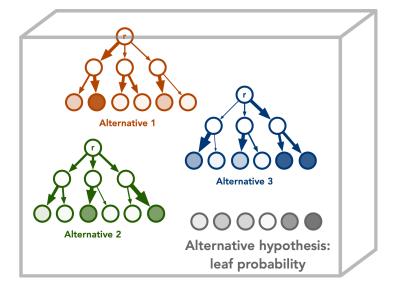
$$p_{X_0}(\ell)$$
 and $\underset{X \in \mathcal{X} \setminus \{X_0\}}{\mathbb{E}} p_X(\ell)$. (C6)

The probability distribution over $X \in \mathcal{X} \setminus \{X_0\}$ in the expectation $\mathbb{E}_{X \in \mathcal{X} \setminus \{X_0\}}$ is arbitrary. We should choose the probability distribution that yields the largest lower bound.

Suppose that the null hypothesis and the alternative hypothesis are true with probability 1/2 each. If we want to use the memory state of the learning algorithm to distinguish between the null hypothesis and the alternative hypothesis, then the success probability of any procedure is upper bounded by

$$\frac{1}{2} + \frac{1}{2} \text{TV}\left(p_{X_0}, \underset{X \in \mathcal{X} \setminus X_0}{\mathbb{E}} p_X\right) = \frac{1}{2} + \frac{1}{4} \sum_{\ell} \left| p_{X_0}(\ell) - \underset{X \in \mathcal{X} \setminus X_0}{\mathbb{E}} p_X(\ell) \right|, \tag{C7}$$





Supplementary Figure 10: *Illustration for the leaf probability distribution*. The leaf probability distribution depends on the unknown physical state/process and the learning algorithm. In the null hypothesis, we have a single state/process, which gives rise to a probability distribution over leaves. In the alternative hypothesis, there are multiple possible states/processes. Each state/process produces a different leaf probability distribution. The leaf probability distribution for the alternative hypothesis is the average of all the leaf probability distributions.

which is also known as LeCam's two-point method. $TV(p_0, p_1)$ is the total variation distance between the two probability distributions p_0, p_1 .

Intuitively, as we perform more experiments, the depth of the tree increases and the total variation distance between the leaf probability distribution increases. If we want to achieve a prediction accuracy of $p \ge \frac{1}{2}$, then we need the total variation distance to be lower bounded by,

$$\operatorname{TV}\left(p_{X_0}, \underset{X \in \mathcal{X} \setminus X_0}{\mathbb{E}} p_X\right) \ge 2p - 1. \tag{C8}$$

On the other hand, the total variation distance can be upper bounded by a monotonically increasing function of the number of experiments T equal to the depth of the tree. Altogether, this allows us to lower bound the number of experiments T by a function of the success probability p.

C.3. Many-versus-many distinguishing task

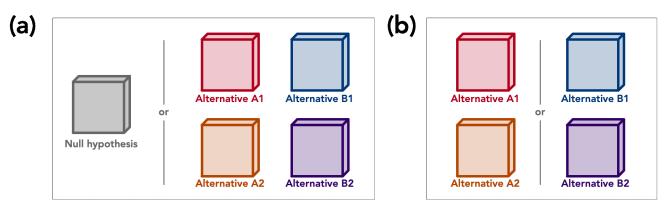
Sometimes, it is easier to first reduce the learning task to a many-versus-many distinguishing task before reducing to a many-versus-one task. This technique is used in Appendix E to prove exponential advantage for quantum principal component analysis. Consider \mathcal{X} to be the set of allowed states/channels. We consider a subset $\mathcal{A} \subseteq \mathcal{X}$, and define $\mathcal{B} = \mathcal{X} \setminus \mathcal{A}$. Here, we consider the following two hypotheses.

- Hypothesis A: The unknown state/channel is a random element in A.
- Hypothesis B: The unknown state/channel is a random element in \mathcal{B} .

Assume each hypothesis happens with probability 1/2. The goal is to distinguish which hypothesis is true. Because \mathcal{A}, \mathcal{B} can contain many elements in \mathcal{X} , we refer to this as the many-versus-many distinguishing task. In Supp. Fig. 11, we visualize the difference between the *many-versus-many distinguishing task* and the other tasks.

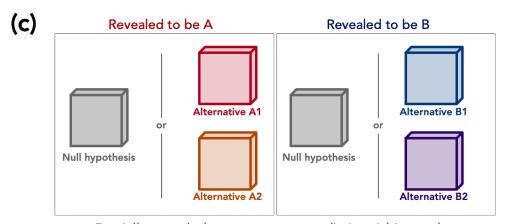
Following a similar derivation as the many-versus-one distinguishing task, for any learning algorithm in the conventional setting, we represent the algorithm as a learning tree. Given a tree representation \mathcal{T} , the success probability for any procedure to distinguish hypothesis A and B using the final memory state of the learning algorithm is upper bounded by

$$\frac{1}{2} + \frac{1}{2} \text{TV} \left(\underset{X \in \mathcal{A}}{\mathbb{E}} p_X, \underset{X \in \mathcal{B}}{\mathbb{E}} p_X \right), \tag{C9}$$



Many-versus-one distinguishing task

Many-versus-many distinguishing task



Partially-revealed many-versus-one distinguishing task

Supplementary Figure 11: Visualization of the different distinguishing tasks. (a) In the many-versus-one distinguishing task, we are distinguishing between the null hypothesis and the alternative hypothesis (which can be one of many alternatives). (b) In the many-versus-many distinguishing task, we want to distinguish between two hypotheses (each of which can be one of many alternatives). (c) In the partially-revealed many-versus-one distinguishing task, some information about the alternatives is revealed, which makes the distinguishing task easier.

where p_X is a probability distribution over the leaf nodes of the tree \mathcal{T} when the unknown state/channel is X. Hence, if we want to achieve a prediction accuracy of $p \geq \frac{1}{2}$, then we need the total variation distance to be lower bounded by

$$\operatorname{TV}\left(\underset{X\in\mathcal{A}}{\mathbb{E}}p_X, \underset{X\in\mathcal{B}}{\mathbb{E}}p_X\right) \ge 2p-1.$$
 (C10)

This last inequality will be important for establishing the lower bound on the number of experiments T.

C.4. Partially-revealed many-versus-one distinguishing task

In some tasks, it can be challenging to verify if an algorithm has learned accurately without revealing some information to the learning algorithm. Here, we consider a setting where after learning, the algorithm is additionally given some partial information about the underlying state/process.

The set \mathcal{X} of all admissible states/processes can be represented as follows,

$$X = (\xi, \chi) \in \mathcal{X},\tag{C11}$$

where ξ is the information that will be revealed during prediction and χ remains hidden. After performing all experiments, the algorithm can obtain ξ^* , such that the unknown state/process X is guaranteed to be either

$$X_0 \quad \text{or} \quad (\xi^*, \chi) \in \mathcal{X} \setminus \{X_0\},$$
 (C12)

i.e. the null hypothesis or an element of the alternative hypothesis. In Supp. Fig. 11, we visualize the difference between the *partially-revealed many-versus-one distinguishing task* and other tasks. Due to the additional information revealed to the learning algorithm, the distinguishing task becomes easier. However, in many examples, we show that revealing a significant amount of information to a learning algorithm that only has external classical memory will not significantly help its distinguishing power.

Suppose that after revealing the information ξ^* to the learning algorithm, the conditional probability for whether the unknown state/process X is X_0 (null hypothesis) or one of $(\xi^*, \chi) \in \mathcal{X} \setminus \{X_0\}$ (alternative hypothesis) is still uniform, i.e., 1/2 and 1/2. Then similar to when the information is not revealed, the success probability of any procedure to distinguish between null and alternative hypothesis is upper bounded by

$$\frac{1}{2} + \frac{1}{2} \text{TV} \left(p_{X_0}, \underset{(\xi^*, \chi) \in \mathcal{X} \setminus \{X_0\}}{\mathbb{E}} p_{(\xi^*, \chi)} \right) = \frac{1}{2} + \frac{1}{4} \sum_{\ell} \left| p_{X_0}(\ell) - \underset{(\xi^*, \chi) \in \mathcal{X} \setminus \{X_0\}}{\mathbb{E}} p_{(\xi^*, \chi)}(\ell) \right|. \tag{C13}$$

When ξ^* is chosen randomly, the average success probability is upper bounded by

$$\frac{1}{2} + \frac{1}{2} \underset{\xi^*}{\mathbb{E}} \operatorname{TV} \left(p_{X_0}, \underset{(\xi^*, \chi) \in \mathcal{X} \setminus \{X_0\}}{\mathbb{E}} p_{(\xi^*, \chi)} \right). \tag{C14}$$

Again, similar to the discussion before, in order to achieve a prediction accuracy of $p \ge 1/2$, we need to satisfy the inequality

$$\mathbb{E}_{\xi^*} \text{TV}\left(p_{X_0}, \mathbb{E}_{(\xi^*, \chi) \in \mathcal{X} \setminus \{X_0\}} p_{(\xi^*, \chi)}\right) \ge 2p - 1. \tag{C15}$$

The left-hand side of the above inequality can be upper bounded by a monotonically increasing function of T, hence we can obtain a lower bound on T. Note that by Jensen's inequality, the left hand side of the above inequality is larger than the left hand side in Eq. (C8), so we will obtain a weaker lower bound on T. This makes sense because making accurate predictions with partially-revealed information is easier.

C.5. Presence of noise

So far, we have considered protocols for learning quantum states or quantum processes in the absence of noise. There are several forms of noise we can consider: (i) noise on input states, (ii) noise on the POVMs which are measured, and (iii) noise on the quantum process (if there is one). Let us prove the following result:

Theorem 5 (Noise cannot decrease the lower bound). If the upper bound

$$\operatorname{TV}\left(\mathbb{E}_{X\in\mathcal{A}}p_X, \mathbb{E}_{X\in\mathcal{B}}p_X\right) \le 2p - 1 \tag{C16}$$

holds for all learning protocols with a classical memory, then this same bound holds for all learning protocols with a classical memory in the presence of noise. Because the upper bound on total variation distance applies when noise is present, so does the lower bound on the number of experiments needed to achieve the distinguishing task.

Proof. Consider first the setting of learning an unknown physical state ρ . Suppose we have a learning protocol with a classical memory described by a learning tree \mathcal{T} . At node u in the protocol, we measure the state ρ with the POVM $\{F_s^u\}_s$. We will measure the sth outcome with probability

$$\operatorname{tr}(F_s^u \rho)$$
. (C17)

If there is noise on ρ , we can use $\rho \mapsto \mathcal{N}[\rho]$ for some noise quantum process \mathcal{N} . Likewise if there is noise on the POVM, we can use $F_s^u \mapsto \mathcal{M}^{\dagger}[F_s^u]$ for a noise quantum process \mathcal{M} . Then the probability of the sth outcome is instead

$$\operatorname{tr}(\mathcal{M}^{\dagger}[F_s^u]\mathcal{N}[\rho]) = \operatorname{tr}((\mathcal{N}^{\dagger} \circ \mathcal{M}^{\dagger})[F_s^u]\rho). \tag{C18}$$

But $\{(\mathcal{N}^{\dagger} \circ \mathcal{M}^{\dagger})[F_s^u]\}_s$ also forms a POVM. We can apply this same argument to each node in the tree; note that the noise channels can be node-dependent. The result is that we simply get a new learning tree with classical

memory, with POVM's augmented by the noise channels. But since by hypothesis (C16) holds for all learning protocols with a classical memory, the bound evidently still holds in the noisy setting.

In the setting where we are learning a physical process \mathcal{E} , the argument is similar. Given a learning tree \mathcal{T} for learning the physical process, at node u we (i) prepare the state ρ^u , (ii) apply the physical process \mathcal{E} , and (iii) measure with the POVM $\{F_s^u\}_s$ and obtain outcome s with probability

$$\operatorname{tr}(F_s^u \mathcal{E}[\rho^u]). \tag{C19}$$

If the initial state is noisy, we can implement this by a channel mapping $\rho^u \mapsto \mathcal{N}[\rho^u]$. If \mathcal{E} is noisy, this can be implemented by $\mathcal{E} \mapsto \mathcal{D} \circ \mathcal{E} \circ \mathcal{F}$. Finally, if the POVM is noisy, we can implement this by $F_s^u \mapsto \mathcal{M}^{\dagger}[F_s^u]$. In these circumsntances, the probability of the sth outcome is instead

$$\operatorname{tr}(\mathcal{M}^{\dagger}[F_s^u](\mathcal{D} \circ \mathcal{E} \circ \mathcal{F})[\mathcal{N}[\rho^u]]) = \operatorname{tr}((\mathcal{D}^{\dagger} \circ M^{\dagger})[F_s^u]\mathcal{E}[\mathcal{F} \circ \mathcal{N}[\rho^u]]). \tag{C20}$$

But $\{(\mathcal{D}^{\dagger} \circ M^{\dagger})[F_s^u]\}_s$ also forms a POVM, and $(\mathcal{F} \circ \mathcal{N})[\rho^u]$ is also a valid choice of input state. The same argument can be used for noise channels applied at every node in the tree, and moreover the noise can be node-dependent. The result is that we just get a modified learning tree with classical memory, which by assumption satisfies (C16), as desired.

In summary, we have shown that if a task is hard for all learning protocols with classical memory, then the task is still just as hard (if not harder) in the presence of noise.

C.6. Related works

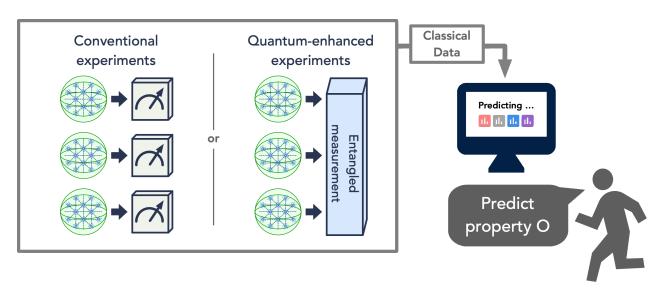
We begin with existing works that study the separation between conventional and quantum-enhanced strategies for learning physical systems and dynamics. In [37], they establish a polynomial separation between conventional and quantum-enhanced strategies for testing if a state is maximally mixed or not. In [8, 9, 12], exponential separations between conventional and quantum-enhanced strategies are established for tasks regarding the learning of physical systems and dynamics.

We also mention some relevant works that study other classes of strategies for learning or characterizing physical systems and dynamics. It was shown in Ref. [38] that the dynamics of open quantum systems with dimension d^n , where d is a prime, can be fully reconstructed with a quadratically fewer experiments over conventional quantum process tomography, with a quantum-enhanced strategy consist of n auxiliary systems of same dimensions d and performing generalized Bell-sate preparations and generalized Bell-state measurements. The results in [39] give a polynomial separation between a restricted class of conventional strategies and quantum-enhanced strategies for learning the complete description of a quantum state. The results in [24] give an exponential separation between a restricted class of conventional strategies and quantum-enhanced strategies for learning to predict properties of a quantum state. [40] establish an exponential separation between ancilla-free strategies and ancilla-assisted strategies for learning the eigenvalues in Pauli channels. [41] give an exponential separation between restricted quantum-enhanced strategies and quantum states. [42] consider a problem on learning two spatially separated quantum states using local quantum learning algorithms and give an exponential separation between having a quantum or a classical communication channel between the local quantum learning algorithms. In [43, 44], an exponential separation between two bounded-depth quantum learning algorithms are given for learning about an exponential-time quantum process.

D. Predicting highly-incompatible observables

The first task we study using the framework of the previous section involves learning about a physical system represented by an n-qubit state ρ . We provide an illustration of the task in Supp. Fig. 12.

- In conventional experiments, we consider algorithms that can measure each copy of ρ one at a time. The algorithm can choose to perform any POVM measurement on each copy, where the POVM measurement can be chosen adaptively based on the outcomes of previous experiments.
- In quantum-enhanced experiments, we consider algorithms that can use a quantum computer to act collectively on multiple copies of ρ to obtain entangled measurement data.



Supplementary Figure 12: Illustration for the task of predicting highly-incompatible observables. The unknown quantum state ρ is represented by the green sphere. Conventional experiments measure each copy of state ρ individually, and the measurements can depend adaptively on previous measurements. Quantum-enhanced experiments store many copies of ρ in a quantum memory, process the copies with a quantum computer, and produce an entangled measurement outcome. The classical data obtained from the experiments are used to predict a property of ρ .

In both scenarios, we consider all quantum data to be used during the learning phase, and we are left only with classical measurement data. After this learning phase, the learner is then asked to provide accurate predictions for the expectation value of an observable O, using the classical data obtained from the experiments. The observable O is selected from an exponentially large set $\{O_1, O_2, \ldots O_M\}$, where O_1, \ldots, O_M may not be mutually commuting and M is exponential in n.

Note that when the observables in the set are not mutually commuting, it is impossible to measure all of them simultaneously. Hence, a naïve algorithm in the conventional scenario would be to measure the exponential number of observables individually, which would result in exponential sample complexity.

D.1. Exponential advantage in predicting absolute value of a single observable

We will prove that even predicting the absolute value of just a single observable requires exponentially many copies in the conventional scenario. In contrast, an algorithm with quantum memory can predict the expectation values for M arbitrary observables from only $\mathcal{O}(n \log(M)/\epsilon^4)$ copies of ρ through the procedure known as shadow tomography [45–47]. Hence, even if we would like to predict an exponential number of observables, an algorithm with quantum memory only needs a polynomial number of copies.

In fact, for certain natural instances, we can show an even more dramatic separation. Specifically, for the following states and observables, we will show how to achieve an exponential versus *constant* separation.

Definition 2 (Separation instance). Consider a distribution \mathcal{D} over n-qubit state ρ and observable O.

- 1. With probability 1/2, the state is $\rho = I/2^n$ and $O \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}$ is chosen uniformly at random.
- 2. With probability 1/2, the state is $\rho = (I+0.9sP)/2^n$ and O = P, where $s = \{\pm 1\}$ with equal probability and $P \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}$ is chosen uniformly at random.²

The *n*-qubit state ρ considered in the above definition does not contain any quantum entanglement. The state ρ can be written as a classical probability distribution over tensor products of single-qubit states. Despite the

² While 0.9 is used in the definition $\rho = (I + 0.9sP)/2^n$, any constant value smaller than 1 is sufficient to obtain the exponential separation. A technical difficulty arises when we consider $(I + sP)/2^n$, and it is unclear whether this difficulty is fundamental.

lack of quantum entanglement, we can still achieve an exponential versus constant separation. This result is a substantial improvement over the result established in [8]. In [8], some of the authors showed an $\Omega(2^{n/3})$ versus $\mathcal{O}(n)$ separation between conventional and quantum-enhanced strategies, but the task was to predict an exponential number of observables, which can only be verified using an exponential amount of time.

Theorem 6 (Exponential advantage in predicting highly-incompatible observables). We sample an n-qubit state ρ and an observable O according to D given in Definition 2; both of these are unknown to the algorithm. The algorithm then learns about ρ through conventional or quantum-enhanced experiments. After the learning phase, we ask the learning algorithm to predict $|\operatorname{tr}(O\rho)|$.

- Upper bound: There is an algorithm in the quantum-enhanced scenario using only $\mathcal{O}(1)$ copies of ρ to predict up to 0.25 additive error with probability at least 0.8.
- Lower bound: For any algorithm in the **conventional** scenario, it needs at least $\Omega(2^n)$ copies of ρ to predict up to 0.25 additive error with probability at least 0.8.

Here, we are using the standard Big- \mathcal{O} and Big- Ω notations: $f = \Omega(g)$ if there is an $n_0, C > 0$ such that $\forall n > n_0$, $f(n) \geq Cg(n)$; and $f = \mathcal{O}(g)$ if there is an $n_0, M > 0$ such that $\forall n > n_0, |f(n)| \leq Mg(n)$. We separate the proof of Theorem 6 into the following two subsections. In Appendix D.2, we prove a constant upper bound for quantum-enhanced experiments for this task. In Appendix D.3, we prove an exponential lower bound for conventional experiments for the same task.

D.2. A constant upper bound for quantum-enhanced experiments

The learning algorithm in the quantum-enhanced scenario builds on results presented in [8]. We separate the protocol into the learning phase, where entangled measurements are performed, and the prediction phase, where we predict the desired properties.

D.2.a. Learning phase

Consider N_Q rounds of two-copy entangled measurements. In round $t \in \{1, ..., N_Q\}$, for every $k \in \{1, ..., n\}$ we measure the k-th qubit from the first and second copies of ρ in the Bell basis to obtain

$$S_k^{(t)} \in \left\{ |\Psi^+\rangle\langle\Psi^+|, |\Psi^-\rangle\langle\Psi^-|, |\Phi^+\rangle\langle\Phi^+|, |\Phi^-\rangle\langle\Phi^-| \right\}, \tag{D1}$$

where the Bell basis encompasses four maximally entangled two-qubit states. Here, $|\Omega\rangle = \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle)$ is the Bell state, and we additionally have

$$\begin{split} |\Psi^{+}\rangle &= I \otimes I \, |\Omega\rangle = \frac{1}{\sqrt{2}} \left(|00\rangle + |11\rangle \right), \\ |\Phi^{+}\rangle &= I \otimes X \, |\Omega\rangle = \frac{1}{\sqrt{2}} \left(|00\rangle + |11\rangle \right), \\ |\Phi^{-}\rangle &= \mathrm{i} I \otimes X \, |\Omega\rangle = \frac{1}{\sqrt{2}} \left(|01\rangle + |10\rangle \right), \end{split}$$

Then, we efficiently store the measurement data $S_k^{(t)}, \forall k = 1, \dots, n, \forall t = 1, \dots, N_Q$ in a classical memory with $2nN_Q$ classical bits.

D.2.b. Prediction phase

Given an observable O drawn from $\{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}$, we can use the block of classical memory obtained in the learning phase to estimate $|\operatorname{tr}(O\rho)|$. First let us consider the case where ρ is a single-qubit state. When we measure $\rho \otimes \rho$ in the Bell basis, the measurement outcome S is a projector onto one of the four Bell states given in Eq. (D1). Let $\sigma \in \{I, X, Y, Z\}$ be any Pauli matrix. Each Bell state is an eigenstate of $\sigma \otimes \sigma$ with an eigenvalue ± 1 . The probability that the Bell measurement outcome S is an eigenstate of $\sigma \otimes \sigma$ with eigenvalue +1 is

$$Prob(+) = \frac{1}{2} tr ((I \otimes I + \sigma \otimes \sigma)(\rho \otimes \rho)), \qquad (D2)$$

while the -1 eigenvalue occurs with probability

$$Prob(-) = \frac{1}{2} \operatorname{tr} \left((I \otimes I - \sigma \otimes \sigma)(\rho \otimes \rho) \right). \tag{D3}$$

Therefore, we have

$$\mathbb{E}\left[\operatorname{tr}\left((\sigma\otimes\sigma)S\right)\right] = \operatorname{Prob}(+) - \operatorname{Prob}(-) = \operatorname{tr}\left((\sigma\otimes\sigma)(\rho\otimes\rho)\right) = |\operatorname{tr}(\sigma\rho)|^{2},\tag{D4}$$

where \mathbb{E} denotes the expectation with respect to the probability distribution over Bell measurement outcomes. We see that the entangling Bell measurement enables us to estimate the absolute value $|\operatorname{tr}(\sigma\rho)|$ for any Pauli matrix $\sigma \in \{I, X, Y, Z\}$.

We can generalize this observation to the case where ρ is an n-qubit state, and each pair of qubits in $\rho \otimes \rho$ is measured in the Bell basis to yield the outcomes $\{S_k, k=1,2,\ldots,n\}$. If $O=\sigma_1\otimes\cdots\otimes\sigma_n$ is a Pauli observable, then as in the n=1 case the Bell state S_k is an eigenstate of $\sigma_k\otimes\sigma_k$ with eigenvalue ± 1 for each k. This implies that $\bigotimes_{k=1}^n S_k$ is an eigenstate of $O\otimes O$ with an eigenvalue ± 1 . In particular, let us consider the product

$$\prod_{k=1}^{n} \operatorname{tr} \left((\sigma_k \otimes \sigma_k) S_k \right) = \pm 1. \tag{D5}$$

This product is equal to +1 when the tensor product of the Bell measurement outcomes $\bigotimes_{k=1}^{n} S_k$ is an eigenstate of $O \otimes O$ with eigenvalue +1, and it is -1 when $\bigotimes_{k=1}^{n} S_k$ is an eigenstate of $O \otimes O$ with eigenvalue -1. We conclude that

$$\mathbb{E}\left[\prod_{k=1}^{n} \operatorname{tr}\left((\sigma_{k} \otimes \sigma_{k})S_{k}\right)\right] = \mathbb{E}\left[\operatorname{tr}\left((O \otimes O) \bigotimes_{k=1}^{n} S_{k}\right)\right]$$

$$= \operatorname{Prob}(O \otimes O = +1) - \operatorname{Prob}(O \otimes O = -1)$$

$$= \operatorname{tr}\left((O \otimes O)(\rho \otimes \rho)\right)$$

$$= |\operatorname{tr}(O\rho)|^{2}, \tag{D6}$$

where \mathbb{E} denotes the expectation with respect to the probability distribution of Bell measurement outcomes. The above derivation shows that the *n*-qubit entangling Bell measurement enables us to estimate the absolute value $|\operatorname{tr}(O\rho)|$ for any O considered in Definition 2.

Because Equation (D6) relates the probability distribution of Bell measurement outcomes to the absolute value $|\operatorname{tr}(O\rho)|$, we can estimate $|\operatorname{tr}(O\rho)|$ accurately by repeatedly making entangling Bell measurements on successive pairs of copies of ρ sufficiently many times. Specifically, in the learning phase, we perform the entangling Bell measurement on $N_{\rm Q}$ pairs of copies of ρ , and collect the measurement data $\{S_k^{(t)}\}$ in the classical memory, where $k=1,2,\ldots,n$ labels the qubit pairs, and $t=1,2,\ldots,N_{\rm Q}$ labels the different rounds of measurements. For any given n-qubit Pauli observable $O=\sigma_1\otimes\cdots\otimes\sigma_n$, we consider the following estimator

$$\hat{a}(O) = \frac{1}{N_Q} \sum_{t=1}^{N_Q} \prod_{k=1}^n \operatorname{tr}\left((\sigma_k \otimes \sigma_k) S_k^{(t)} \right), \tag{D7}$$

which can be computed efficiently in time $\mathcal{O}(nN_{\rm O})$.

Using the expectation evaluated in Equation (D6), we can apply Hoeffding's inequality to show that the estimate $\hat{a}(O)$ is equal to the expectation value $\operatorname{tr}((O \otimes O)(\rho \otimes \rho)) = |\operatorname{tr}(O\rho)|^2$ up to a small error with high probability. The formal statement is given below.

Lemma 1. Given $N_Q = \Theta(\log(1/\delta)/\epsilon^2)$. For any observable O considered in Definition 2, we have

$$\left|\hat{a}(O) - |\operatorname{tr}(O\rho)|^2\right| \le \epsilon,\tag{D8}$$

with probability at least $1 - \delta$.

To obtain an estimate for the absolute value $|\operatorname{tr}(O\rho)|$, we consider the estimate

$$\hat{b} = \sqrt{\max(0, \hat{a})}. (D9)$$

We can show the inequalities

$$|\operatorname{tr}(O\rho)|^2 - \epsilon \le \hat{a} \le |\operatorname{tr}(O\rho)|^2 + \epsilon \implies \max(0, \sqrt{|\operatorname{tr}(O\rho)|^2} - \sqrt{\epsilon}) \le \hat{b} \le \sqrt{|\operatorname{tr}(O\rho)|^2} + \sqrt{\epsilon}, \tag{D10}$$

using the fact that $\sqrt{x+y} \le \sqrt{x} + \sqrt{y}$.

In the final step of the upper bound proof, we use Lemma 1 to obtain the following result. As long as $N_Q = \mathcal{O}(1)$, we can estimate the absolute value of $\operatorname{tr}(O\rho)$ for any observable O given in Definition 2 to an error 0.25 with probability at least 0.8.

Corollary 1. Let $N_Q = \Theta(1)$. For any observable O considered in Definition 2, we have

$$\left|\hat{b} - |\operatorname{tr}(O\rho)|\right| \le 0.25,\tag{D11}$$

with probability at least 0.8.

This concludes the constant upper bound for quantum-enhanced experiments in Theorem 6.

D.3. An exponential lower bound for conventional experiments

The proof begins with a reduction to the partially-revealed many-versus-one distinguishing task followed by bounding the total variation distance.

D.3.a. Reduction to partially-revealed many-versus-one distinguishing task

We consider the following partially-revealed many-versus-one distinguishing task discussed in Appendix C.4, namely where:

- The null hypothesis is $I/2^n$.
- The alternative hypothesis is $(I + 0.9sP)/2^n$.

The partially revealed information is the Pauli operator P. Recall the following from Definition 2,

- 1. With probability 1/2, the state is $\rho = I/2^n$ and $O \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}$ is sampled uniformly at random. (Corresponds to the null hypothesis)
- 2. With probability 1/2, the state is $\rho = (I + 0.9sP)/2^n$ and O = P, where $s = \{\pm 1\}$ with equal probability and $P \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}$ uniformly. (Corresponds to the alternative hypothesis)

For $\rho = I/2^n$, we have $|\operatorname{tr}(O\rho)| = 0$. For $\rho = (I + 0.9sP)/2^n$, we have $|\operatorname{tr}(O\rho)| = 0.9$. Therefore, if an algorithm could predict $|\operatorname{tr}(O\rho)|$ to 0.25 error with probability at least $1 - \delta$, it could be used to distinguish between the null and alternative hypotheses with success probability at least $1 - \delta$.

D.3.b. Total variation distance

From the information-theoretic lower bound for partially-revealed many-versus-one distinguishing task given in Appendix C.4, if we let $p_{\rho}(\ell)$ be the leaf probability distribution under ρ , then

$$\mathbb{E}_{P \in \{I, X, Y, Z\} \otimes n \setminus \{I \otimes n\}} \operatorname{TV} \left(p_{I/2^n}, \mathbb{E}_{s \in \{\pm 1\}} p_{(I+0.9sP)/2^n} \right) \ge 1 - 2\delta.$$
(D12)

For each leaf node ℓ , we consider the path from the root to ℓ ,

$$u_0 = r \xrightarrow{s_1} u_1 \xrightarrow{s_2} u_2 \xrightarrow{s_3} \dots \xrightarrow{s_{T-1}} u_{T-1} \xrightarrow{s_T} u_T = \ell. \tag{D13}$$

At each node u, we perform a POVM measurement $\{w_s^u|\phi_s^u\rangle\langle\phi_s^u|\}_s$ on ρ to obtain an outcome s with probability

$$w_s^u \langle \phi_s^u | \rho | \phi_s^u \rangle$$
. (D14)

Hence, we can write down the probability to arrive at the leaf ℓ as

$$p_{\rho}(\ell) = \prod_{t=1}^{T} w_{s_{t}}^{u_{t-1}} \left\langle \phi_{s_{t}}^{u_{t-1}} \middle| \rho \middle| \phi_{s_{t}}^{u_{t-1}} \right\rangle. \tag{D15}$$

Recalling the definition of total variation distance, note that for any probability distributions p_A, p_B for which $p_A(\ell) > 0$ whenever $p_B(\ell) > 0$,

$$TV(p_A, p_B) = \frac{1}{2} \sum_{\ell} |p_A(\ell) - p_B(\ell)| = \sum_{\ell} \max(0, p_A(\ell) - p_B(\ell)) = \sum_{\ell} p_A(\ell) \cdot \max\left(0, 1 - \frac{p_B(\ell)}{p_A(\ell)}\right), \quad (D16)$$

where the last equality follows from $\max(ax, ay) = a \max(x, y)$ for all $x, y \in \mathbb{R}$ and all $a \ge 0$. Observe that for leaf ℓ ,

$$\frac{\mathbb{E}_{s\in\{\pm 1\}} p_{(I+0.9sP)/2^n}(\ell)}{p_{I/2^n}(\ell)} = \frac{\mathbb{E}_{s\in\{\pm 1\}} \prod_{t=1}^T w_{s_t}^{u_{t-1}} \langle \phi_{s_t}^{u_{t-1}} | \frac{I+0.9sP}{2^n} | \phi_{s_t}^{u_{t-1}} \rangle}{\prod_{t=1}^T w_{s_t}^{u_{t-1}} \langle \phi_{s_t}^{u_{t-1}} | \frac{I}{2^n} | \phi_{s_t}^{u_{t-1}} \rangle}$$
(D17)

$$= \mathbb{E}_{s \in \{\pm 1\}} \prod_{t=1}^{T} \left(1 + 0.9s \left\langle \phi_{s_t}^{u_{t-1}} \middle| P \middle| \phi_{s_t}^{u_{t-1}} \right\rangle \right). \tag{D18}$$

Combining (D16) and (D18), we can express the total variation distance inside the expectation in (D12) as

$$\text{TV}\left(p_{I/2^n}, \underset{s \in \{\pm 1\}}{\mathbb{E}} p_{(I+0.9sP)/2^n}\right) = \sum_{\ell} p_{I/2^n}(\ell) \, \max\left(0, \, 1 - \underset{s \in \{\pm 1\}}{\mathbb{E}} \prod_{t=1}^{T} \left(1 + 0.9s \left\langle \phi_{s_t}^{u_{t-1}} \middle| P \middle| \phi_{s_t}^{u_{t-1}} \right\rangle\right)\right) \quad \text{(D19)}$$

D.3.c. Upper bound for total variation distance

We analyze one of the terms in the total variation distance using Jensen's inequality (note that $\exp(x)$ is a convex function in x).

$$\mathbb{E}_{s \in \{\pm 1\}} \prod_{t=1}^{T} \left(1 + 0.9s \left\langle \phi_{s_t}^{u_{t-1}} \middle| P \middle| \phi_{s_t}^{u_{t-1}} \right\rangle \right) \tag{D20}$$

$$= \mathbb{E}_{s \in \{\pm 1\}} \exp \left[\sum_{t=1}^{T} \log \left(1 + 0.9s \left\langle \phi_{s_t}^{u_{t-1}} \middle| P \middle| \phi_{s_t}^{u_{t-1}} \right\rangle \right) \right]$$
 (D21)

$$\geq \exp\left[\mathbb{E}_{s\in\{\pm 1\}} \sum_{t=1}^{T} \log\left(1 + 0.9s \left\langle \phi_{s_{t}}^{u_{t-1}} \middle| P \middle| \phi_{s_{t}}^{u_{t-1}} \right\rangle\right)\right]$$
(D22)

$$= \exp\left[\sum_{t=1}^{T} \frac{1}{2} \log\left(1 - 0.81 \left\langle \phi_{s_t}^{u_{t-1}} \middle| P \middle| \phi_{s_t}^{u_{t-1}} \right\rangle^2\right)\right]$$
 (D23)

$$= \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_t}^{u_{t-1}} \right| P \left| \phi_{s_t}^{u_{t-1}} \right\rangle^2} \,. \tag{D24}$$

We can then upper bound the total variation distance as

$$\operatorname{TV}\left(p_{I/2^n}, \underset{s \in \{\pm 1\}}{\mathbb{E}} p_{(I+0.9sP)/2^n}\right) \tag{D25}$$

$$\leq \sum_{\ell} p_{I/2^{n}}(\ell) \, \max\left(0, \, 1 - \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| P \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}} \right) \tag{D26}$$

$$= \sum_{\ell} p_{I/2^n}(\ell) \left(1 - \prod_{t=1}^T \sqrt{1 - 0.81 \left\langle \phi_{s_t}^{u_{t-1}} \right| P \left| \phi_{s_t}^{u_{t-1}} \right\rangle^2} \right). \tag{D27}$$

The last equality follows from the fact that all eigenvalues of P are ± 1 , hence $1 \ge \prod_{t=1}^T \sqrt{1 - 0.81 \left\langle \phi_{s_t}^{u_{t-1}} \middle| P \middle| \phi_{s_t}^{u_{t-1}} \middle\rangle^2}$.

D.3.d. Lower bound for the number of measurements

We can combine Eq. (D27) and Eq. (D12) to find

$$\mathbb{E}_{P \in \{I, X, Y, Z\} \otimes n \setminus \{I^{\otimes n}\}} \sum_{\ell} p_{I/2^{n}}(\ell) \left(1 - \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| P \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}} \right) \ge 1 - 2\delta. \tag{D28}$$

By linearity of expectation, we have

$$\sum_{\ell} p_{I/2^n}(\ell) \left(1 - \mathbb{E}_{P \in \{I, X, Y, Z\} \otimes n \setminus \{I \otimes n\}} \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_t}^{u_{t-1}} \middle| P \middle| \phi_{s_t}^{u_{t-1}} \right\rangle^2} \right) \ge 1 - 2\delta. \tag{D29}$$

We analyze the expectation value term in the summand as follows:

$$\mathbb{E}_{P \in \{I, X, Y, Z\} \otimes n \setminus \{I \otimes n\}} \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_t}^{u_{t-1}} \right| P \left| \phi_{s_t}^{u_{t-1}} \right\rangle^2}$$
 (D30)

$$= \underset{P \in \{I, X, Y, Z\} \otimes n \setminus \{I^{\otimes n}\}}{\mathbb{E}} \exp \left[\frac{1}{2} \sum_{t=1}^{T} \log \left(1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| P \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2} \right) \right]$$
(D31)

$$\geq \exp\left[\frac{1}{2} \sum_{t=1}^{T} \underset{P \in \{I, X, Y, Z\} \otimes n \setminus \{I^{\otimes n}\}}{\mathbb{E}} \log\left(1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \middle| P \middle| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}\right)\right]$$
(D32)

$$\geq \exp\left[\frac{1}{2} \sum_{t=1}^{T} \underset{P \in \{I,X,Y,Z\} \otimes n \setminus \{I \otimes n\}}{\mathbb{E}} -1.701 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| P \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}\right] \tag{D33}$$

$$= \exp\left[-0.8505 \sum_{t=1}^{T} \frac{1}{2^n + 1}\right] = \exp\left(-\frac{0.8505T}{2^n + 1}\right). \tag{D34}$$

The second line follows from Jensen's inequality because $\exp(x)$ is convex in x. The third line uses $\log(1-x) \ge -2.1x, \forall x \in [0, 0.82]$. The fourth line uses the fact that

$$\underset{P \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}}{\mathbb{E}} P \otimes P = \frac{2^n \text{SWAP} - I \otimes I}{4^n - 1}, \tag{D35}$$

hence $\mathbb{E}_{P \in \{I,X,Y,Z\}^{\otimes n} \setminus \{I^{\otimes n}\}} \left< \phi_{s_t}^{u_{t-1}} \right| P \left| \phi_{s_t}^{u_{t-1}} \right>^2 = \frac{2^n - 1}{4^n - 1} = \frac{1}{2^n + 1}$.

Combining the analysis with Eq. (D29), we find that

$$\sum_{\ell} p_{I/2^n}(\ell) \left(1 - \exp\left(-\frac{0.8505T}{2^n + 1} \right) \right) \ge 1 - 2\delta.$$
 (D36)

Because $\sum_{\ell} p_{I/2^n}(\ell) = 1$, we have

$$1 - \exp\left(-\frac{0.8505T}{2^n + 1}\right) \ge 1 - 2\delta. \tag{D37}$$

Together, the following lower bound on the number of experiments can be obtained:

$$T \ge \frac{2^n + 1}{0.8505} \log \left(\frac{1}{2\delta} \right). \tag{D38}$$

After setting $\delta = 0.2$ (corresponding to a success probability of at least 0.8), we conclude the exponential lower bound for conventional experiments in Theorem 6.

D.4. An exponential lower bound for comparing absolute values

In the physical experiment presented in the main text, we considered a slightly different task that also yields an exponential lower bound for conventional experiments. This slightly different task has two main differences when compared with the task described in Appendix D.3. First, we do not consider the maximally mixed state $I/2^n$. Second, we ask the learner to predict which of two observables O_1, O_2 has a larger absolute value. The task description is given below.

Task 1 (Comparing absolute values). There is an unknown state $\rho = (I + 0.9sP)/2^n$ where $s = \{\pm 1\}$ and $P \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}$ are both sampled uniformly at random. The algorithm learns about ρ through conventional or quantum-enhanced strategies. The algorithm transforms all quantum data to classical data. After learning, we present the learning algorithm with

$$O_1 = P, O_2 = Q \quad \text{or} \quad O_1 = Q, O_2 = P,$$
 (D39)

with equal probability, where $Q \neq P \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}$ is sampled uniformly. The learning algorithm succeeds if it correctly classifies whether $|\operatorname{tr}(O_1\rho)| > |\operatorname{tr}(O_2\rho)|$ or $|\operatorname{tr}(O_1\rho)| < |\operatorname{tr}(O_2\rho)|$.

Using the procedure presented in Appendix D.2, it is not hard to show that quantum-enhanced strategies could accomplish the above task with classification accuracy (i.e., the probability that the classification is correct) at least $1-\delta$ from only $\mathcal{O}(\log(1/\delta))$ experiments. In contrast, we have the following theorem for conventional experiments.

Theorem 7 (Exponential lower bound for Task 1). A learning algorithm in the conventional setting (without quantum memory) requires at least

$$\frac{(2^n+1)}{0.8505}\log\left(\frac{2}{1+2\delta}\right) \tag{D40}$$

experiments to accomplish Task 1 with an accuracy of $1 - \delta$, for a given $\delta > 0$.

D.4.a. Lower bound for total variation distance

Here we begin the proof of Theorem 7. Task 1 is closely related to the partially-revealed many-versus-one distinguishing task, but is not exactly the same. We will utilize a slightly different information-theoretic bound for this task. Let us define the following notation

$$\rho_{sP} \equiv \frac{I + 0.9sP}{2^n}.\tag{D41}$$

We consider a learning algorithm in the conventional setting. We consider the probability distribution $p_{\rho}(\ell)$ over the leaf node ℓ when the underlying state is ρ . Recall that the leaf node ℓ is the final memory state of the learning algorithm. Any procedure that makes the prediction based on the final memory state of the learning algorithm must have a classification accuracy upper bounded by

$$\frac{1}{(4^n - 1)(4^n - 2)} \sum_{P \neq Q \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}} \left[\frac{1}{2} + \frac{1}{2} \text{TV} \left(\underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sP}}(\ell), \underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sQ}}(\ell) \right) \right]. \tag{D42}$$

To understand why the above inequality holds, consider a fixed $P \neq Q$. There is an equal probability that the underlying state is $\rho_{+P}, \rho_{-P}, \rho_{+Q}$, or ρ_{-Q} because $s \in \{\pm 1\}$ with equal probability and $O_1 = P, O_2 = Q$ or $O_1 = Q, O_2 = P$ with probability 1/2. In order to distinguish the event of ρ_{+P}, ρ_{-P} from the event ρ_{+Q}, ρ_{-Q} based on the leaf node ℓ , we need the two distributions $\mathbb{E}_{s \in \{\pm 1\}} p_{\rho_{sP}}(\ell)$ and $\mathbb{E}_{s \in \{\pm 1\}} p_{\ell}(\rho_{sQ})$ to be sufficiently distinct. Formally, one can show that the success probability is upper bounded by

$$\frac{1}{2} + \frac{1}{2} \text{TV} \left(\underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sP}}(\ell), \underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sQ}}(\ell) \right)$$
(D43)

using LeCam's two-point method, see e.g. Lemma 1 in [48]. To achieve the above success probability, one can use the maximum likelihood protocol that outputs P if $\mathbb{E}_{s\in\{\pm 1\}} p_{\rho_{sP}}(\ell) > \mathbb{E}_{s\in\{\pm 1\}} p_{\rho_{sQ}}(\ell)$ and outputs Q otherwise.

(D53)

Because P, Q are both chosen uniformly at random (but distinct), the average classification accuracy is given in Eq. (D42). If the learning algorithm could achieve an accuracy of $1 - \delta$, we would have

$$(1 - \delta) \le \frac{1}{2} + \frac{1}{2(4^n - 1)(4^n - 2)} \sum_{P \ne Q \in \{I, X, Y, Z\} \otimes n \setminus \{I \otimes n\}} \text{TV}\left(\underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sP}}(\ell), \underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sQ}}(\ell) \right). \tag{D44}$$

This implies that

 $=2-2\exp\left(-\frac{0.8505T}{2^{n}+1}\right).$

$$1 - 2\delta \le \frac{1}{(4^n - 1)(4^n - 2)} \sum_{P \ne Q \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}} \text{TV}\left(\underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sP}}(\ell), \underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sQ}}(\ell) \right). \tag{D45}$$

D.4.b. Upper bound for total variation distance

We now perform triangle inequalities and reuse inequalities in Appendix D.3 to upper bound the total variation distance:

$$\operatorname{TV}\left(\underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sP}}(\ell), \underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sQ}}(\ell)\right) \leq \operatorname{TV}\left(p_{I/2^{n}}(\ell), \underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sP}}(\ell)\right) + \operatorname{TV}\left(p_{I/2^{n}}(\ell), \underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sQ}}(\ell)\right) \tag{D46}$$

$$\leq \sum_{\ell} p_{I/2^{n}}(\ell) \left(1 - \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| P \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}}\right)$$

$$+ \sum_{\ell} p_{I/2^{n}}(\ell) \left(1 - \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| Q \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}}\right)$$

$$= 1 - \sum_{\ell} p_{I/2^{n}}(\ell) \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| P \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}}$$

$$+ 1 - \sum_{\ell} p_{I/2^{n}}(\ell) \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| Q \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}}.$$
(D48)

The first line is triangle inequality. The second inequality uses Eq. (D27) for both P and Q. The equality step uses $\sum_{\ell} p_{I/2^n}(\ell) = 1$. Therefore, we have

$$\frac{1}{(4^{n}-1)(4^{n}-2)} \sum_{P \neq Q \in \{I,X,Y,Z\} \otimes n \setminus \{I \otimes n\}} \text{TV} \left(\underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sP}}(\ell), \underset{s \in \{\pm 1\}}{\mathbb{E}} p_{\rho_{sQ}}(\ell) \right) \tag{D49}$$

$$\leq 1 - \frac{1}{(4^{n}-1)(4^{n}-2)} \sum_{P \neq Q \in \{I,X,Y,Z\} \otimes n \setminus \{I \otimes n\}} \left[\sum_{\ell} p_{I/2^{n}}(\ell) \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| P \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}} \right]$$

$$+ 1 - \frac{1}{(4^{n}-1)(4^{n}-2)} \sum_{P \neq Q \in \{I,X,Y,Z\} \otimes n \setminus \{I \otimes n\}} \left[\sum_{\ell} p_{I/2^{n}}(\ell) \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| Q \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}} \right]$$

$$= 1 - \sum_{\ell} p_{I/2^{n}}(\ell) \sum_{P \in \{I,X,Y,Z\} \otimes n \setminus \{I \otimes n\}} \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| P \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}}$$

$$+ 1 - \sum_{\ell} p_{I/2^{n}}(\ell) \sum_{Q \in \{I,X,Y,Z\} \otimes n \setminus \{I \otimes n\}} \prod_{t=1}^{T} \sqrt{1 - 0.81 \left\langle \phi_{s_{t}}^{u_{t-1}} \right| Q \left| \phi_{s_{t}}^{u_{t-1}} \right\rangle^{2}}$$

$$= 1 - \sum_{\ell} p_{I/2^{n}}(\ell) \exp\left(-\frac{0.8505T}{2^{n}+1}\right) + 1 - \sum_{\ell} p_{I/2^{n}}(\ell) \exp\left(-\frac{0.8505T}{2^{n}+1}\right)$$
(D51)

In the above inequalities, the first inequality uses Eq. (D48). The equality thereafter uses the following analysis,

$$\frac{1}{(4^n - 1)(4^n - 2)} \sum_{P \neq Q \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}} f(P)$$
(D54)

$$= \frac{1}{4^n - 1} \sum_{P \in \{I, X, Y, Z\} \otimes n \setminus \{I \otimes n\}} f(P) \left(\frac{1}{4^n - 2} \sum_{\substack{Q \in \{I, X, Y, Z\} \otimes n \setminus \{I \otimes n\} \\ \text{s.t., } O \neq P}} 1 \right)$$
(D55)

$$=\frac{1}{4^n-1}\sum_{P\in\{I,X,Y,Z\}^{\otimes n}\setminus\{I^{\otimes n}\}}f(P)$$
(D56)

$$= \underset{P \in \{I, X, Y, Z\} \otimes n \setminus \{I^{\otimes n}\}}{\mathbb{E}} f(P), \tag{D57}$$

where $f(P) = \sum_{\ell} p_{I/2^n}(\ell) \prod_{t=1}^T \sqrt{1 - 0.81 \left\langle \phi_{s_t}^{u_{t-1}} \right| Q \left| \phi_{s_t}^{u_{t-1}} \right\rangle^2}$, as well as linearity of expectation, i.e.

$$\underset{P\in\{I,X,Y,Z\}\otimes^n\backslash\{I^{\otimes n}\}}{\mathbb{E}} \sum_{\ell} p_{I/2^n}(\ell) = \sum_{\ell} p_{I/2^n}(\ell) \underset{P\in\{I,X,Y,Z\}\otimes^n\backslash\{I^{\otimes n}\}}{\mathbb{E}}.$$
 (D58)

The third step in Eq. (D52) uses Eq. (D30) to (D34), and the final step uses $\sum_{\ell} p_{I/2^n}(\ell) = 1$.

D.4.c. Combining upper and lower bounds

We can combine the lower bound obtained in Eq. (D45) and the upper bound in Eq. (D53) to find

$$1 - 2\delta \le 2 - 2\exp\left(-\frac{0.8505T}{2^n + 1}\right). \tag{D59}$$

Basic algebraic manipulations give

$$\frac{0.8505T}{2^n + 1} \ge -\log\left(\frac{1}{2} + \delta\right) = \log\left(\frac{2}{1 + 2\delta}\right). \tag{D60}$$

We have thus concluded the desired lower bound

$$T \ge \frac{(2^n + 1)}{0.8505} \log \left(\frac{2}{1 + 2\delta} \right) \tag{D61}$$

stated in Theorem 7.

E. Performing quantum principal component analysis

The (first) principal component of a nonnegative Hermitian matrix A is the eigenvector of A with the largest eigenvalue. Here, we consider a well-known task called quantum principal component analysis (PCA), which can be achieved efficiently using the quantum algorithm given in [13]. The formal definition of the quantum PCA task is given in the following definition.

Task 2 (Quantum principal component analysis task). Let ρ be an unknown n-qubit mixed state whose top eigenvector $|\phi\rangle$ has eigenvalue larger than all other eigenvalues by a constant factor independent of n. Given a fixed observable O, we would like to predict $\langle \phi | O | \phi \rangle$ up to a small additive error.

We can accomplish this task using the quantum PCA algorithm in [13]. In this algorithm, multiple copies of ρ are used in a protocol that approximates the unitary operator $\sum_t |t\rangle\langle t| \otimes \exp(-i\rho t)$, where t is the "time" stored in an auxiliary register. By applying this conditional $\exp(-i\rho t)$ operation to the initial state $\rho = \sum_i \lambda_i |\phi_i\rangle\langle\phi_i|$, performing the quantum Fourier transform and measuring the auxiliary register, we read out an eigenvalue λ_i and prepare the corresponding eigenstate $|\phi_i\rangle$ with probability λ_i . The eigenvalue can be measured with constant

accuracy, and the eigenstate prepared with constant fidelity, using a constant number of copies of ρ . Once $|\phi_i\rangle$ has been prepared, we can measure the observable O in this state.

By assumption, the largest eigenvalue of ρ is a constant independent of n, and furthermore is greater than all other eigenvalues by a constant. Hence by repeating the above procedure a constant number of times, we can estimate $\langle \phi | O | \phi \rangle$ to constant accuracy, where $| \phi \rangle$ is the eigenstate of ρ with the largest eigenvalue. In contrast, in Appendix E.1 we show that algorithms that can only learn about ρ through conventional experiments require an exponential number of copies of ρ . Bringing these arguments all together, we establish the following theorem.

Theorem 8 (Exponential advantage for quantum principal component analysis). Let ρ be an unknown n-qubit mixed state (for n > 1) whose top eigenvector $|\phi\rangle$ has eigenvalue larger than all the other eigenvalues by a constant, and let Z_1 be an observable which is equal to the Pauli-Z operator on the first qubit. Algorithms learn about ρ through conventional or quantum-enhanced experiments.

- Upper bound: There is an algorithm in the **quantum-enhanced** scenario using only $\mathcal{O}(1)$ copies of ρ to predict $\langle \phi | Z_1 | \phi \rangle$ up to 0.25 error with probability at least 0.8.
- Lower bound: Any algorithm in the **conventional** scenario needs at least $\Omega(2^{n/2})$ copies of ρ to predict $\langle \phi | Z_1 | \phi \rangle$ up to 0.25 error with probability at least 0.8.

Instead of estimating $\langle \phi | Z_1 | \phi \rangle$, we can also consider the near-term proposals [15, 16] to obtain some information about the principal component $|\psi\rangle$ of an unknown state ρ as follows. The proposals consider $\rho^M/\operatorname{tr}(\rho^M)$, which approaches $|\psi\rangle\langle\psi|$ when M is large. In particular, [16] shows that one can efficiently estimate $\operatorname{tr}(Z_1\rho^2)/\operatorname{tr}(\rho^2)$ by performing entangling Bell measurements over at most two copies of ρ at a time. By the analysis in [16], if the eigenvalue associated to the principal component of ρ is a constant, then $\operatorname{tr}(Z_i\rho^2)/\operatorname{tr}(\rho^2)$ can be estimated to any constant error by performing quantum-enhanced experiments over a constant number of copies of ρ . In contrast, we show that if one can only measure a single copy of ρ at a time, exponentially many copies are necessary to estimate $\operatorname{tr}(Z_i\rho^2)/\operatorname{tr}(\rho^2)$.

Theorem 9 (Exponential advantage for near-term quantum principal component analysis). Suppose we are given an observable Z_1 which is equal to the Pauli-Z operator on the first qubit, as well as an n-qubit state ρ (for n > 1) where an eigenvector $|\phi\rangle$ of ρ has an eigenvalue that is larger than all the other eigenvalues by a constant. We consider algorithms which learn about ρ through conventional or quantum-enhanced experiments. Then we have the following bounds:

- Upper bound: There is an algorithm in the quantum-enhanced scenario using only $\mathcal{O}(1)$ copies of ρ to predict $\operatorname{tr}(Z_1\rho^2)/\operatorname{tr}(\rho^2)$ up to 0.25 error with probability at least 0.8.
- Lower bound: Any algorithm in the **conventional** scenario needs at least $\Omega(2^{n/2})$ copies of ρ to predict $\operatorname{tr}(Z_1\rho^2)/\operatorname{tr}(\rho^2)$ up to 0.25 error with probability at least 0.8.

E.1. An exponential lower bound for conventional experiments

The lower bound proofs for both Theorem 8 and Theorem 9 are essentially the same. We first reduce quantum PCA (or near-term analogs thereof) to a many-versus-many distinguishing task. Then we bound the total variation distance to arrive at the exponential lower bound.

E.1.a. Reduction to many-versus-many distinguishing task

We begin by considering a many-versus-many distinguishing task, as discussed in Appendix C.3. The two hypotheses are given below.

• Hypothesis A: The unknown n-qubit state ρ is given by

$$\rho_A(|\psi\rangle) = \frac{1}{2}|0\rangle\langle 0| \otimes |\psi\rangle\langle \psi| + \frac{1}{2}|1\rangle\langle 1| \otimes \frac{I}{2^{n-1}},\tag{E1}$$

where $|\psi\rangle$ is an fixed (n-1)-qubit pure state, sampled at the outset from the Haar measure.

• Hypothesis B: The unknown n-qubit state ρ is given by

$$\rho_B(|\psi\rangle) = \frac{1}{2}|1\rangle\langle 1| \otimes |\psi\rangle\langle \psi| + \frac{1}{2}|0\rangle\langle 0| \otimes \frac{I}{2^{n-1}}, \tag{E2}$$

where $|\psi\rangle$ is again a fixed (n-1)-qubit pure state, sampled at the outset from the Haar measure.

It is not hard to see that in hypothesis A, the principal component (largest eigenvector) is $|\phi\rangle = |0\rangle \otimes |\psi\rangle$. On the other hand, in hypothesis B, the principal component (largest eigenvector) is $|\phi\rangle = |1\rangle \otimes |\psi\rangle$. Hence, $\langle \phi | Z_1 | \phi \rangle = 1$ in hypothesis A, but $\langle \phi | Z_1 | \phi \rangle = -1$ in hypothesis B. If an algorithm in the conventional scenario can predict $\langle \psi | Z_1 | \psi \rangle$ up to 0.25 error with probability at least 0.8, then we can use the output from the algorithm to distinguish between hypotheses A and B with a success probability of at least 0.8.

Similarly, we have $\operatorname{tr}(Z_1\rho^2)/\operatorname{tr}(\rho^2)=(2^{n-1}-1)/(2^{n-1}+1)$ in hypothesis A and $\operatorname{tr}(Z_1\rho^2)/\operatorname{tr}(\rho^2)=-(2^{n-1}-1)/(2^{n-1}+1)$ in hypothesis B. If an algorithm in the conventional scenario can predict $\operatorname{tr}(Z_1\rho^2)/\operatorname{tr}(\rho^2)$ up to 0.25 error with probability at least 0.8, then we can use the output from the algorithm to distinguish between hypothesis A and B with a success probability of at least 0.8.

Together, a lower bound for distinguishing hypotheses A and B using conventional experiments immediately implies a lower bound for both Theorem 8 and Theorem 9.

E.1.b. Total variation distance

As in previous sections, let $p_{\rho}(\ell)$ denote the probability to arrive at the leaf node ℓ using the learning algorithm in the conventional setting when the unknown state is ρ . If the algorithm can distinguish between hypotheses A and B with success probability 0.8, then using Eq. (C10) we have

$$TV\left(\underset{|\psi\rangle}{\mathbb{E}} p_{\rho_A(|\psi\rangle)}, \underset{|\psi\rangle}{\mathbb{E}} p_{\rho_B(|\psi\rangle)}\right) \ge 0.6.$$
 (E3)

From the triangle inequality, we have

$$TV\left(\underset{|\psi\rangle}{\mathbb{E}} p_{\rho_A(|\psi\rangle)}, p_{I/2^n}\right) + TV\left(\underset{|\psi\rangle}{\mathbb{E}} p_{\rho_B(|\psi\rangle)}, p_{I/2^n}\right) \ge 0.6.$$
 (E4)

For each leaf node ℓ , we consider the path from the root to ℓ ,

$$u_0 = r \xrightarrow{s_1} u_1 \xrightarrow{s_2} u_2 \xrightarrow{s_3} \dots \xrightarrow{s_{T-1}} u_{T-1} \xrightarrow{s_T} u_T = \ell.$$
 (E5)

At each node u, we perform a POVM measurement $\{w_s^u|\phi_s^u\rangle\langle\phi_s^u|\}_s$ on ρ to obtain an outcome s with probability

$$w_s^u \langle \phi_s^u | \rho | \phi_s^u \rangle$$
. (E6)

Hence, we can write down the probability to arrive at the leaf ℓ as

$$p_{\rho}(\ell) = \prod_{t=1}^{T} w_{s_t}^{u_{t-1}} \left\langle \phi_{s_t}^{u_{t-1}} \middle| \rho \middle| \phi_{s_t}^{u_{t-1}} \right\rangle. \tag{E7}$$

We will use $\rho(|\psi\rangle)$ to denote either $\rho_A(|\psi\rangle)$ or $\rho_B(|\psi\rangle)$. Then recalling (D16), we have

$$\operatorname{TV}\left(\mathbb{E}_{|\psi\rangle} p_{\rho(|\psi\rangle)}, p_{I/2^n}\right) = \sum_{\ell} p_{I/2^n}(\ell) \max\left(0, 1 - \mathbb{E}_{|\psi\rangle} \prod_{t=1}^T 2^n \left\langle \phi_{s_t}^{u_{t-1}} \middle| \rho(|\psi\rangle) \middle| \phi_{s_t}^{u_{t-1}} \right\rangle\right). \tag{E8}$$

E.1.c. Upper bound for total variation distance

The central quantity to control in our bound on the total variation distance is

$$\mathbb{E}_{|\psi\rangle} \prod_{t=1}^{T} 2^n \left\langle \phi_{s_t}^{u_{t-1}} \middle| \rho(|\psi\rangle) \middle| \phi_{s_t}^{u_{t-1}} \right\rangle. \tag{E9}$$

Without loss of generality, let us consider $\rho(|\psi\rangle) = \rho_A(|\psi\rangle) = \frac{1}{2}|0\rangle\langle 0| \otimes |\psi\rangle\langle \psi| + \frac{1}{2}|1\rangle\langle 1| \otimes \frac{I}{2^{n-1}}$. Suppose each $|\phi_{s_t}^{u_{t-1}}\rangle$ takes the form

$$|\phi_{s_t}^{u_{t-1}}\rangle = \alpha_{s_t}^{u_{t-1}}|0\rangle \otimes |\phi_{s_t}^{u_{t-1},0}\rangle + \beta_{s_t}^{u_{t-1}}|1\rangle \otimes |\phi_{s_t}^{u_{t-1},1}\rangle,$$
 (E10)

where $\alpha_{s_t}^{u_{t-1}}, \beta_{s_t}^{u_{t-1}} \in \mathbb{C}$ and $\left|\alpha_{s_t}^{u_{t-1}}\right|^2 + \left|\beta_{s_t}^{u_{t-1}}\right|^2 = 1$. Then we have

$$\mathbb{E}_{|\psi\rangle} \prod_{t=1}^{T} 2^{n} \left\langle \phi_{s_{t}}^{u_{t-1}} | \rho(|\psi\rangle) | \phi_{s_{t}}^{u_{t-1}} \right\rangle = \mathbb{E}_{|\psi\rangle} \prod_{t=1}^{T} \left(2^{n-1} \left| \alpha_{s_{t}}^{u_{t-1}} \right|^{2} \left| \left\langle \psi | \phi_{s_{t}}^{u_{t-1},0} \right\rangle \right|^{2} + \left| \beta_{s_{t}}^{u_{t-1}} \right|^{2} \right)$$
(E11)

$$= \sum_{S \subseteq \{1, \dots, T\}} \prod_{t \notin S} \left| \beta_{s_t}^{u_{t-1}} \right|^2 \left[\mathbb{E} \prod_{t \in S} 2^{n-1} \left| \alpha_{s_t}^{u_{t-1}} \right|^2 \left| \langle \psi | \phi_{s_t}^{u_{t-1}, 0} \rangle \right|^2 \right].$$
 (E12)

We need to lower bound the above quantity in order to upper bound the total variation distance. In order to do so, we utilize the following lemma. The proof of the lemma is based on Haar integration. For readers unfamiliar with Haar integration, we would suggest skipping the proof of this lemma.

Lemma 2 (High moment bound for Haar-random state). Consider any m-qubit pure states $|\phi_1\rangle, \ldots, |\phi_K\rangle$ and an m-qubit pure state $|\psi\rangle$ sampled from the Haar measure, we have

$$\mathbb{E}_{|\psi\rangle} \prod_{k=1}^{K} |\langle \psi | \phi_k \rangle|^2 \ge \frac{1}{(2^m + K - 1) \dots (2^m + 1)(2^m)}.$$
(E13)

Proof. The Haar integration over states shows that

$$\underset{|\psi\rangle}{\mathbb{E}}|\psi\rangle\langle\psi|^{\otimes K} = \frac{1}{(2^m + K - 1)\dots(2^m + 1)(2^m)} \sum_{\pi \in \mathcal{S}_K} \pi, \tag{E14}$$

where S_K is the permutation group of K items, and π is the permutation operator over the K tensor-product space. From Lemma 5.12 in [12], we have

$$\sum_{\pi \in \mathcal{S}_K} \operatorname{tr}\left(\pi \bigotimes_{k=1}^K |\phi_k\rangle \langle \phi_k|\right) \ge 1. \tag{E15}$$

Therefore, we find

$$\mathbb{E}_{|\psi\rangle} \prod_{k=1}^{K} |\langle \psi | \phi_k \rangle|^2 \ge \frac{1}{(2^m + K - 1) \dots (2^m + 1)(2^m)}.$$
 (E16)

This concludes the proof.

We apply this lemma with

$$m \equiv n - 1, \quad K \equiv |S|, \quad |\phi_k\rangle \equiv |\phi_{s_t}^{u_{t-1},0}\rangle$$
 (E17)

to obtain the following lower bound,

$$\mathbb{E}_{|\psi\rangle} \prod_{t \in S} 2^{n-1} \left| \alpha_{s_t}^{u_{t-1}} \right|^2 \left| \langle \psi | \phi_{s_t}^{u_{t-1}, 0} \rangle \right|^2 \ge \prod_{t \in S} \frac{\left| \alpha_{s_t}^{u_{t-1}} \right|^2}{\left(1 + \frac{|S| - 1}{2^{n-1}} \right) \dots \left(1 + \frac{1}{2^{n-1}} \right) (1)}$$
(E18)

$$\geq \prod_{t \in S} \frac{\left| \alpha_{s_t}^{u_{t-1}} \right|^2}{\left(1 + \frac{|S| - 1}{2^{n-1}} \right)^{|S| - 1}} \tag{E19}$$

$$\geq \left(1 + \frac{T - 1}{2^{n-1}}\right)^{-(T-1)} \prod_{t \in S} \left| \alpha_{s_t}^{u_{t-1}} \right|^2.$$
 (E20)

Combining with Eq. (E12), we have

$$\mathbb{E}_{|\psi\rangle} \prod_{t=1}^{T} 2^{n} \left\langle \phi_{s_{t}}^{u_{t-1}} \middle| \rho(|\psi\rangle) \middle| \phi_{s_{t}}^{u_{t-1}} \right\rangle \ge \left(1 + \frac{T-1}{2^{n-1}}\right)^{-(T-1)} \sum_{S \subset \{1, \dots, T\}} \prod_{t \notin S} \left| \beta_{s_{t}}^{u_{t-1}} \right|^{2} \prod_{t \in S} \left| \alpha_{s_{t}}^{u_{t-1}} \right|^{2} \tag{E21}$$

$$= \left(1 + \frac{T-1}{2^{n-1}}\right)^{-(T-1)} \prod_{t=1}^{T} \left(\left|\beta_{s_t}^{u_{t-1}}\right|^2 + \left|\alpha_{s_t}^{u_{t-1}}\right|^2\right)$$
 (E22)

$$= \left(1 + \frac{T-1}{2^{n-1}}\right)^{-(T-1)}.$$
(E23)

Next we leverage Eq. (E8) to obtain

$$\operatorname{TV}\left(\mathbb{E}_{|\psi\rangle} p_{\rho(|\psi\rangle)}, p_{I/2^n}\right) \le \sum_{\ell} p_{I/2^n}(\ell) \max\left(0, \ 1 - \left(1 + \frac{T-1}{2^{n-1}}\right)^{-(T-1)}\right) \tag{E24}$$

$$=1-\left(1+\frac{T-1}{2^{n-1}}\right)^{-(T-1)}. (E25)$$

The second line follows from $\sum_{\ell} p_{I/2^n}(\ell) = 1$ because $p_{I/2^n}(\ell)$ is a probability distribution.

E.1.d. Lower bound for the number of measurements

We can now utilize the lower bound on the total variation distance given in Eq. (E4) and the upper bound obtained above to find

$$2\left(1 - \left(1 + \frac{T-1}{2^{n-1}}\right)^{-(T-1)}\right) \ge \text{TV}\left(\underset{|\psi\rangle}{\mathbb{E}} p_{\rho_A(|\psi\rangle)}, p_{I/2^n}\right) + \text{TV}\left(\underset{|\psi\rangle}{\mathbb{E}} p_{\rho_B(|\psi\rangle)}, p_{I/2^n}\right) \ge 0.6. \tag{E26}$$

Hence, we have the inequality

$$0.7 \ge \left(1 + \frac{T - 1}{2^{n - 1}}\right)^{-(T - 1)},\tag{E27}$$

$$\implies (T-1)\log\left(1+\frac{T-1}{2^{n-1}}\right) \ge \log(10/7). \tag{E28}$$

Because $\log(1+x) \le x$ for all x > -1, we have

$$(T-1)^2 \ge 2^{n-1} \log(10/7) \implies T \ge 1 + \sqrt{\frac{\log\left(\frac{10}{7}\right)}{2}} 2^{n/2}.$$
 (E29)

Finally, we have established the lower bound $T = \Omega(2^{n/2})$ stated in Theorem 8.

E.2. An exponential lower bound using pseudorandomness

In our exponential lower bound in the previous subsection, we relied on a state that was in part constructed using a Haar-random (n-1)-qubit state $|\psi\rangle$. However, preparing a Haar-random state from a simple initial state (say, a product state) requires circuit depth exponential in n. As such, we can not prepare Haar-random states in practice. Accordingly we cannot prepare either $\rho_A(|\psi\rangle)$ or $\rho_B(|\psi\rangle)$ in realistic circumstances.

However, we could instead efficiently construct pseudorandom states $|\psi\rangle$ which are (very plausibly) indistinguishable from Haar-random states if we probe with any POVM instantiated by a poly(n)-time quantum algorithm [49]. We will elaborate on this shortly. The parenthetical caveat 'very plausibly' is due to the fact that the construction we use relies on cryptographic assumptions which are not proven, but are widely believed. In particular, we need to suppose the existence of quantum-secure one-way functions [49]; these are functions which are efficient to evaluate but are hard to invert even with a quantum computer. Making such an assumption is standard practice in computational complexity theory, and so we proceed apace.

Let us recall the definition of a pseudorandom quantum state:

Definition 3 (Pseudorandom quantum states; paraphrased from Definition 3 of [49]). Given a set K, a family of pseudorandom quantum states on n qubits is a family of states $\{|\phi_k\rangle\}_{k\in\mathcal{K}}$ and a probability distribution \mathcal{D} over K such that:

- There is a poly(n)-time quantum algorithm that samples a single element k from K according to D and generates the corresponding state $|\phi_k\rangle$;
- For any polynomial t(n) and any poly(n)-time quantum algorithm A with outputs in $\{0,1\}$, we have

$$\left| \Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A}(|\phi_k\rangle^{\otimes t(n)}) = 1 \right] - \Pr_{|\psi\rangle \leftarrow \text{Haar}} \left[\mathcal{A}(|\psi\rangle^{\otimes t(n)}) = 1 \right] \right| \le \text{negl}(n).$$
 (E30)

Here negl(n) is a function such that for all constants c > 0 we have $negl(n) < n^{-c}$ for n sufficiently large.

We can interpret the above definition as follows. Eq. (E30) says that if we are given a polynomial t(n) number of copies of either (i) a fixed pseudorandom state, or (ii) a fixed Haar random state, any polynomial time quantum algorithm with binary outputs cannot distinguish between the two cases. Since an exponential depth quantum algorithm can distinguish between the two cases, (E30) describes a notion of *computational indistinguishability*, i.e. we cannot distinguish with polynomial time computational resources.

A key question is: do pseudorandom quantum states exist? We recall the following, contingent result.

Lemma 3 (Existence of pseudorandom quantum states [49]). If there exist quantum-secure one-way functions, then they can be used to construct pseudorandom quantum states.

Explicit details of the construction can be found in [49]; there have also been refinements and generalizations in follow-up work (see e.g. [50, 51]).

Before stating the main result of this section, we require the following definition:

Definition 4 (Polynomial-time algorithms in the conventional scenario). A polynomial-time algorithm \mathcal{A} in the conventional scenario is constructed as follows. We consider a learning tree \mathcal{T} in the **conventional scenario** with leaves ℓ , and require that the protocol described by the learning tree can be implemented by an at most poly(n)-time quantum algorithm. (As such, the depth of \mathcal{T} is at most polynomial in n.) Then let \mathcal{D} be a poly(n)-time classical algorithm which, given the transcript of measurements encoded into the leaves ℓ of \mathcal{T} , provides a binary output 0 or 1. We let \mathcal{A} be a map from n-qubit density matrices ρ to $\{0,1\}$, corresponding to instantiating the learning tree \mathcal{T} on copies of ρ , followed by using \mathcal{D} on the measurement transcript to determine a binary outcome.

We can now leverage the putative pseudorandom states and the above definition to establish the following result:

Theorem 10 (Lower bound many-versus-many distinguishing task using pseudorandom states). Let $\{|\phi_k\rangle\}_{k\in\mathcal{K}}$ be a family of pseudorandom states on n-1 qubits. Then any polynomial-time algorithm \mathcal{A} in the **conventional** scenario with binary output cannot distinguish $\rho_A(|\phi_k\rangle)$ from $\rho_B(|\phi_k\rangle)$ for k sampled from the probability distribution over \mathcal{K} . That is:

$$\left| \Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A}(\rho_A(|\phi_k\rangle)) = 1 \right] - \Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A}(\rho_B(|\phi_k\rangle)) = 1 \right] \right| \le \text{negl}(n). \tag{E31}$$

Proof. Using the triangle inequality several times we have

$$\begin{aligned} \left| \operatorname{Pr}_{k \leftarrow \mathcal{K}} \left[\mathcal{A}(\rho_{A}(|\phi_{k}\rangle)) = 1 \right] - \operatorname{Pr}_{k \leftarrow \mathcal{K}} \left[\mathcal{A}(\rho_{B}(|\phi_{k}\rangle)) = 1 \right] \right| \\ &\leq \left| \operatorname{Pr}_{|\psi\rangle \leftarrow \operatorname{Haar}} \left[\mathcal{A}(\rho_{A}(|\psi\rangle)) = 1 \right] - \operatorname{Pr}_{|\psi\rangle \leftarrow \operatorname{Haar}} \left[\mathcal{A}(\rho_{B}(|\psi\rangle)) = 1 \right] \right| \\ &+ \left| \operatorname{Pr}_{k \leftarrow \mathcal{K}} \left[\mathcal{A}(\rho_{A}(|\phi_{k}\rangle)) = 1 \right] - \operatorname{Pr}_{|\psi\rangle \leftarrow \operatorname{Haar}} \left[\mathcal{A}(\rho_{A}(|\psi\rangle)) = 1 \right] \right| \\ &+ \left| \operatorname{Pr}_{k \leftarrow \mathcal{K}} \left[\mathcal{A}(\rho_{B}(|\phi_{k}\rangle)) = 1 \right] - \operatorname{Pr}_{|\psi\rangle \leftarrow \operatorname{Haar}} \left[\mathcal{A}(\rho_{B}(|\psi\rangle)) = 1 \right] \right| . \end{aligned}$$
(E32)

Let \mathcal{T} be the learning tree corresponding to \mathcal{A} , and let \mathcal{D} be the binary decision function mapping leaves of \mathcal{T} to $\{0,1\}$. The depth T of \mathcal{T} is necessarily at most polynomial in n; let us denote the depth by T(n). Then the first term on the right-hand side of (E32) is upper bounded by

$$\begin{aligned} & \left| \Pr_{|\psi\rangle \leftarrow \text{Haar}} \left[\mathcal{A}(\rho_A(|\psi\rangle)) = 1 \right] - \Pr_{|\psi\rangle \leftarrow \text{Haar}} \left[\mathcal{A}(\rho_B(|\psi\rangle)) = 1 \right] \right| \\ & = \left| \sum_{\ell \in \text{leaf}(\mathcal{T}) : \mathcal{D}(\ell) = 1} \left(\underset{|\psi\rangle}{\mathbb{E}} p_{\rho_A(|\psi\rangle)}(\ell) - \underset{|\psi\rangle}{\mathbb{E}} p_{\rho_B(|\psi\rangle)}(\ell) \right) \right| \end{aligned}$$
(E33)

$$\leq \sum_{\ell \in \text{leaf}(\mathcal{T})} \left| \underset{|\psi\rangle}{\mathbb{E}} p_{\rho_{A}(|\psi\rangle)}(\ell) - \underset{|\psi\rangle}{\mathbb{E}} p_{\rho_{B}(|\psi\rangle)}(\ell) \right| \\
\leq 2 \,\text{TV} \left(\underset{|\psi\rangle}{\mathbb{E}} p_{\rho_{A}(|\psi\rangle)}, p_{I/2^{n}} \right) + 2 \,\text{TV} \left(\underset{|\psi\rangle}{\mathbb{E}} p_{\rho_{B}(|\psi\rangle)}, p_{I/2^{n}} \right) \\
\leq 4 \left(1 - \left(1 + \frac{T(n) - 1}{2^{n-1}} \right)^{-(T(n) - 1)} \right) \tag{E34}$$

where the last inequality comes from (E26).

Next we turn to bounding the second term on the right-hand side of the inequality in (E32), namely

$$\left| \Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A}(\rho_A(|\phi_k\rangle)) = 1 \right] - \Pr_{|\psi\rangle\leftarrow \text{Haar}} \left[\mathcal{A}(\rho_A(|\psi\rangle)) = 1 \right] \right|.$$

First, we observe that every \mathcal{A} takes in T(n) copies of ρ_A . Accordingly, there is a polynomial-time algorithm $\widetilde{\mathcal{A}}$ such that $\widetilde{\mathcal{A}}(\rho_A^{\otimes T(n)}) = \mathcal{A}(\rho_A)$ for all inputs ρ_A ; this just amounts to a slightly different way of notating the domain of the algorithm \mathcal{A} . Then we can rewrite our term of interest as

$$\left| \Pr_{k \leftarrow \mathcal{K}} \left[\widetilde{\mathcal{A}}(\rho_A(|\phi_k\rangle)^{\otimes T(n)}) = 1 \right] - \Pr_{|\psi\rangle \leftarrow \text{Haar}} \left[\widetilde{\mathcal{A}}(\rho_A(|\psi\rangle)^{\otimes T(n)}) = 1 \right] \right|. \tag{E35}$$

But now observe that for any state $|\omega\rangle$, there is a polynomial-time quantum algorithm which takes $|\omega\rangle$ as input and produces $\rho_A(|\omega\rangle)$ as output. Accordingly, repeating this algorithm on T(n) copies of $|\omega\rangle$, we produce T(n) copies of $\rho_A(|\omega\rangle)$; let us denote this T(n)-copy algorithm by \mathcal{B} . We have $\mathcal{B}(|\omega\rangle^{\otimes T(n)}) = \rho_A(|\omega\rangle)^{\otimes T(n)}$, where \mathcal{B} runs in polynomial time (recalling that T(n) is polynomial in n). Then we can write (E35) as

$$\left| \Pr_{k \leftarrow \mathcal{K}} \left[(\widetilde{\mathcal{A}} \circ \mathcal{B}) (|\phi_k\rangle^{\otimes T(n)}) = 1 \right] - \Pr_{|\psi\rangle \leftarrow \text{Haar}} \left[(\widetilde{\mathcal{A}} \circ \mathcal{B}) (|\psi\rangle^{\otimes T(n)}) = 1 \right] \right| . \tag{E36}$$

Since $\widetilde{\mathcal{A}} \circ \mathcal{B}$ is itself a polynomial-time quantum algorithm, Definition 3 tells us that (E36) above is upper bounded by $\operatorname{negl}(n)$. So in summary, we find

$$\left| \Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A}(\rho_A(|\phi_k\rangle)) = 1 \right] - \Pr_{|\psi\rangle \leftarrow \text{Haar}} \left[\mathcal{A}(\rho_A(|\psi\rangle)) = 1 \right] \right| \le \text{negl}(n). \tag{E37}$$

In an identical manner we can show that the third term on the right-hand side of (E32) satisfies the bound

$$\left| \Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A}(\rho_B(|\phi_k\rangle)) = 1 \right] - \Pr_{|\psi\rangle \leftarrow \text{Haar}} \left[\mathcal{A}(\rho_B(|\psi\rangle)) = 1 \right] \right| \le \text{negl}(n).$$
 (E38)

Putting together the inequalities in (E32), (E33), (E37) and (E38), as well as observing that T(n) is at most polynomially large in n, we achieve the desired bound.

The above Theorem immediately implies that the exponential advantage for quantum principal component analysis in Theorem 8 has a counterpart for pseudorandom states. We note that in the pseudorandom context the advantage is not strictly exponential; rather, the advantage holds for arbitrary *polynomial*-time quantum learning algorithms in the quantum-enhanced scenario versus in the conventional scenario.

F. Learning a polynomial-time quantum process

Here we consider the problem of learning a polynomial-time quantum process, and provide a rigorous exponential separation between the conventional and quantum-enhanced learning settings.

F.1. Problem setting

We consider an unknown quantum process \mathcal{E} on n qubits generated as follows.

- An *n*-qubit state σ is accompanied by an *m* ancillary qubits initialized at $|0^n\rangle\langle 0^n|$.
- p unknown two-qubit unitary gates are applied on the (n+m)-qubit system $\rho \otimes |0^n\rangle\langle 0^n|$.

• The ancillary qubits are hidden, resulting in an n-qubit mixed state.

When m, p = poly(n), we refer to \mathcal{E} as a polynomial-time quantum process. Next, we consider an input probability distribution \mathcal{D} over n-qubit mixed states σ . The goal is to learn an approximate model $\tilde{\mathcal{E}}$ of \mathcal{E} , such that we can accurately predict the output state on average:

$$\mathbb{E}_{\sigma \sim \mathcal{D}} \left\| \tilde{\mathcal{E}}(\sigma) - \mathcal{E}(\sigma) \right\|_{1} \le \epsilon. \tag{F1}$$

In the above, $||X||_1 = \max_{O:||O||_{\infty} \le 1} |\operatorname{tr}(OX)|$ is the trace norm.

F.2. Rigorous statements

We have the following theorem for quantum-enhanced experiments showing that they can efficiently learn a polynomial-time quantum process. We will prove the theorem later in Appendix F.3.

Theorem 11 (Approximate learning of quantum processes – polynomial upper bound). For any distribution \mathcal{D} and any $\epsilon, \delta > 0$, there exists a learning algorithm in the quantum-enhanced setting that can learn an approximate model $\tilde{\mathcal{E}}$ such that with probability at least $1 - \delta$,

$$\mathbb{E}_{\sigma \sim \mathcal{D}} \left\| \tilde{\mathcal{E}}(\sigma) - \mathcal{E}(\sigma) \right\|_{1} \le \epsilon \tag{F2}$$

from at most $\widetilde{\mathcal{O}}(\operatorname{poly}(n)\log(1/\delta)/\epsilon^4)$ accesses to \mathcal{E} , where $\widetilde{\mathcal{O}}(\cdot)$ hides factors of $\log(1/\epsilon)$.

In contrast, our hardness results for predicting properties of physical states in the conventional setting (see Theorem 6 in Appendix D.1) immediately implies the following exponential lower bound.

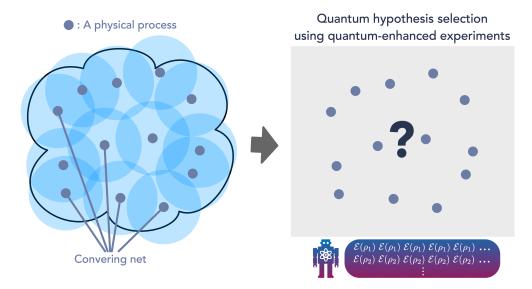
Corollary 2 (Approximate learning of quantum processes – exponential lower bound). Let \mathcal{D} be any distribution and \mathcal{E} be the quantum process that always generates a state ρ considered in Def. 2. Any algorithm in the conventional setting that learns an approximate model $\tilde{\mathcal{E}}$ such that

$$\mathbb{E}_{\sigma \circ \mathcal{D}} \left\| \tilde{\mathcal{E}}(\sigma) - \mathcal{E}(\sigma) \right\|_{1} \le 0.25, \tag{F3}$$

must use at least $\Omega(2^n)$ accesses to \mathcal{E} .

Proof. We consider m=2n. The two-qubit gates swap the input state σ to the first n ancillary qubits. Then we use the rest of the n ancillary qubits and the n system qubits (i.e. the qubits in the input state ρ) to prepare a state ρ considered in Def. 2. To prepare the maximally mixed state $I/2^n$, we entangle each of the system qubits with the corresponding ancillary qubit to prepare a Bell state $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. To prepare the alternative state $(I+0.9sP)/2^n$, we perform the following procedure.

- 1. For qubit i, where P_i is not the last non-identity Pauli operator (i.e. last in terms of having the largest index i), we entangle the i-th system qubit with the corresponding ancillary qubit to prepare a Bell state $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$.
- 2. For the last remaining qubit i, which corresponds to the index i such that P_i is the last non-identity Pauli operator, we apply a sequence of two-qubit gates entangling qubit i to qubit j with $P_j \neq I$. The sequence of two-qubit gates stores the parity (or 1 parity) of all qubits j with $P_j \neq I$. After this step, when we trace over the ancillary qubits, we have generated $(I + sP^{(Z)})/2^n$, where $P^{(Z)} = \bigotimes_{i=1}^n F(P_i)$ and F(I) = I, F(X) = Z, F(Y) = I, F(Z) = I. Then, we rotate the corresponding ancillary qubit for qubit i from $|0\rangle$ to $\sqrt{0.95} |0\rangle + \sqrt{0.05} |1\rangle$ and apply a controlled-not gate from the ancillary qubit (control) to qubit i. The system qubits are now in the state $(I + 0.9sP^{(Z)})/2^n$.
- 3. Finally, for each qubit i with $P_i = X$, we rotate the system qubit from $|0\rangle$ to $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and from $|1\rangle$ to $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle |1\rangle)$. For each qubit i with $P_i = Y$, we rotate from 0 to y+ and from 1 to y-. Tracing over the ancillary qubits, the system qubits are now in the state $(I+0.9sP)/2^n$.



Supplementary Figure 13: Illustration for the proof of Theorem 11 on learning polynomial-time quantum processes. We first form a covering net (all dark blue dots) for the space of all polynomial-time quantum processes (the cloud shape). Any polynomial-time quantum process is close to an element in the covering net. Then we perform quantum hypothesis selection [47] using a quantum dataset stored in the quantum memory to find the approximate physical process.

We can see that the number of gates p is $\mathcal{O}(n)$. Furthermore, no matter what the input state σ is, the above quantum process always produces a state ρ considered in Def. 2.

When an algorithm in the conventional setting has learned an approximate model $\tilde{\mathcal{E}}$ with

$$\mathbb{E}_{\sigma \sim \mathcal{D}} \left\| \tilde{\mathcal{E}}(\sigma) - \mathcal{E}(\sigma) \right\|_{1} \le 0.25, \tag{F4}$$

we can use Jensen's inequality to conclude

$$\left\| \underset{\sigma \sim \mathcal{D}}{\mathbb{E}} \tilde{\mathcal{E}}(\sigma) - \rho \right\|_{1} \leq \underset{\sigma \sim \mathcal{D}}{\mathbb{E}} \left\| \tilde{\mathcal{E}}(\sigma) - \mathcal{E}(\sigma) \right\|_{1} \leq 0.25.$$
 (F5)

Because $||X||_1 = \max_{O:||O||_{\infty} \le 1} |\operatorname{tr}(OX)|$, the above implies that the algorithm can predict $\operatorname{tr}(O\rho)$ up to an error of 0.25. From Theorem 6, we conclude that the learning algorithm must use at least $\Omega(2^n)$ copies of ρ , which corresponds to $\Omega(2^n)$ accesses to \mathcal{E} .

F.3. Proof of polynomial upper bound in Theorem 11

Theorem 11 establishes an upper bound on the number of times we must access the unknown process \mathcal{E} in the quantum-enhanced setting to construct an approximate model of \mathcal{E} . Note that the theorem only concerns the number of times we run the process \mathcal{E} ; it does not address the computational complexity of the learning procedure. Our strategy for proving the theorem is as follows. First we find an upper bound on the number of elements of an ϵ' -covering net for the set of all quantum processes that can be constructed using up to p two-qubit quantum gates, with distance defined by the diamond norm. Next we explain how to use a quantum hypothesis testing algorithm to find a process $\tilde{\mathcal{E}}$ in the covering net that approximates \mathcal{E} as specified in (F2), if ϵ' is appropriately chosen. This quantum hypothesis testing method can be carried out in the quantum-enhanced setting, but not in the conventional setting. The number of times we must access \mathcal{E} depends on the size of the covering net, and can be shown to scale polynomially with the number of gates p, proving the theorem. An illustration is given in Supp. Fig. 13.

F.3.a. Covering net

First, we construct the covering net for the set S of quantum processes with a fixed n, m, p. An ϵ -covering net of a set S is a subset $\mathcal{N}_{\epsilon} \subseteq S$ such that for every point $x \in S$, there exists a point $y \in \mathcal{N}$ with $||x - y|| \le \epsilon$ in an

appropriate norm.

Recall that a unitary U corresponds to a unitary channel \mathcal{U} defined as

$$\mathcal{U}(\rho) = U\rho U^{\dagger}. \tag{F6}$$

Because two-qubit unitary channels form a bounded set in a finite-dimensional space, the $\tilde{\epsilon}$ -covering net for two-qubit unitary channels has a size of at most

$$\left(\frac{c_1}{\tilde{\epsilon}}\right)^{c_2},$$
 (F7)

where c_1, c_2 are two constants (see, e.g., Section 4.2 in [52]). Here, we consider the norm to be the diamond norm $\|\cdot\|_{\diamond}$ (see Section 3.3 in [53]). The bound in (F7) only pertains to the covering net size when the unitary acts on a fixed set of two qubits. Let us now consider two-qubit unitary channels that can act on any two of the n+m qubits. Because there are $\binom{n+m}{2}$ pairs of qubits that the unitary could act on, the size of the $\tilde{\epsilon}$ -covering net $\mathcal{N}_{\tilde{\epsilon},n+m}$ of all two-qubit gates on an (n+m)-qubit system is upper bounded as follows,

$$|\mathcal{N}_{\tilde{\epsilon},n+m}| \le \binom{n+m}{2} \left(\frac{c_1}{\tilde{\epsilon}}\right)^{c_2}. \tag{F8}$$

To construct an ϵ -covering net for the composed quantum process \mathcal{E} , we need to consider $\tilde{\epsilon} = \epsilon'/p$ in $\mathcal{N}_{\tilde{\epsilon},n+m}$. Consider any sequence of two-qubit unitary channels $\mathcal{U}_1, \ldots, \mathcal{U}_p$ on an (n+m)-qubit system. For each U_i in the sequence, we find the closest unitary channel $\tilde{\mathcal{U}}_i$ in $\mathcal{N}_{\tilde{\epsilon},n+m}$, hence $\left\|\mathcal{U}_i - \tilde{\mathcal{U}}_i\right\|_{\diamond} \leq \tilde{\epsilon}$. Then we can use a telescoping sum and the triangle inequality to see that

$$\left\| \operatorname{tr}_{n+1,\dots,n+m} \left(\mathcal{U}_p \dots \mathcal{U}_1(\rho \otimes |0^m \rangle \langle 0^m |) \right) - \operatorname{tr}_{n+1,\dots,n+m} \left(\tilde{\mathcal{U}}_p \dots \tilde{\mathcal{U}}_1(\rho \otimes |0^m \rangle \langle 0^m |) \right) \right\|_1$$
 (F9)

$$\leq \left\| \mathcal{U}_p \dots \mathcal{U}_1(\rho \otimes |0^m\rangle \langle 0^m|) - \tilde{\mathcal{U}}_p \dots \tilde{\mathcal{U}}_1(\rho \otimes |0^m\rangle \langle 0^m|) \right\|_1 \tag{F10}$$

$$\leq \left\| \mathcal{U}_p \dots \mathcal{U}_1(\rho \otimes |0^m\rangle\langle 0^m|) - \mathcal{U}_p \tilde{\mathcal{U}}_{p-1} \dots \tilde{\mathcal{U}}_1(\rho \otimes |0^m\rangle\langle 0^m|) \right\|$$

$$+ \left\| \mathcal{U}_{p} \tilde{\mathcal{U}}_{p-1} \dots \mathcal{U}_{1}(\rho \otimes |0^{m}\rangle\langle 0^{m}|) - \tilde{\mathcal{U}}_{p} \tilde{\mathcal{U}}_{p-1} \dots \tilde{\mathcal{U}}_{1}(\rho \otimes |0^{m}\rangle\langle 0^{m}|) \right\|_{1}$$
(F11)

$$\leq \left\| \mathcal{U}_{p-1} \dots \mathcal{U}_{1}(\rho \otimes |0^{m}\rangle\langle 0^{m}|) - \tilde{\mathcal{U}}_{p-1} \dots \tilde{\mathcal{U}}_{1}(\rho \otimes |0^{m}\rangle\langle 0^{m}|) \right\|_{1} + \left\| \mathcal{U}_{p} - \tilde{\mathcal{U}}_{p} \right\|_{2}$$
(F12)

$$\leq \left\| \mathcal{U}_{p-2} \dots \mathcal{U}_{1}(\rho \otimes |0^{m}\rangle\langle 0^{m}|) - \tilde{\mathcal{U}}_{p-2} \dots \tilde{\mathcal{U}}_{1}(\rho \otimes |0^{m}\rangle\langle 0^{m}|) \right\|_{1} + \left\| \mathcal{U}_{p-1} - \tilde{\mathcal{U}}_{p-1} \right\|_{2} + \left\| \mathcal{U}_{p} - \tilde{\mathcal{U}}_{p} \right\|_{2}$$
(F13)

$$\leq \dots$$
 (F14)

$$\leq \sum_{i=1}^{p} \left\| \mathcal{U}_i - \tilde{\mathcal{U}}_i \right\|_{\diamond} \tag{F15}$$

$$\leq p\tilde{\epsilon} = \epsilon'.$$
 (F16)

The first inequality uses the fact that taking partial trace does not increase the trace norm. The second inequality uses $||A - B|| \le ||A - C|| + ||C - B||$. The third inequality uses $||\mathcal{E}(X)||_1 \le ||X||_1$ for any CPTP map \mathcal{E} , and $||\mathcal{E}(\rho)||_1 \le ||\mathcal{E}||_{\diamond} ||\rho||_1 = ||\mathcal{E}||_{\diamond}$. The fourth inequality considers the same steps taken in the second and third inequality. Then, using induction, we obtain the formula given in the second-to-last line. The last line uses the fact that $||\mathcal{U}_i - \tilde{\mathcal{U}}_i||_{\diamond} \le \tilde{\epsilon}$ for all i and $\tilde{\epsilon} = \epsilon'/p$.

From the above analysis, we can see that we can find an ϵ' -covering net $\mathcal{N}_{\epsilon',n,m,p}$ for the space of \mathcal{E} with an n-qubit input state, m ancillary qubits, and p two-qubit gates that satisfies

$$|\mathcal{N}_{\epsilon',n,m,p}| \le \left[\binom{n+m}{2} \left(\frac{pc_1}{\epsilon'} \right)^{c_2} \right]^p. \tag{F17}$$

For any \mathcal{E} in the space, we can find an $\tilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$ such that for all n-qubit input states ρ we have

$$\left\| \mathcal{E}(\rho) - \tilde{\mathcal{E}}(\rho) \right\|_1 \le \epsilon'.$$
 (F18)

We will then utilize the ϵ -covering net $\mathcal{N}_{\epsilon',n,m,p}$ in the subsequent proof. An ϵ -covering net of quantum processes have also been used in [8] to establish an information-theoretic bound on quantum advantage in [54] to analyze generalization performance of quantum neural networks.

We will sample N_{in} input states $\rho_1, \ldots, \rho_{N_{\text{in}}}$ from the distribution \mathcal{D} . For each $i \in [N_{\text{in}}]$ and every $\widetilde{\mathcal{E}}_k \in \mathcal{N}_{\epsilon',n,m,p}$ (for ϵ' to be tuned later), we will access the true process \mathcal{E} a number of times N_{out} using ρ_i as the input state, obtaining N_{out} copies of $\mathcal{E}(\rho_i)$. We will store these $N_{\text{in}} \cdot N_{\text{out}}$ states in the quantum memory and run a known algorithm for quantum hypothesis selection [47] to determine for which k the product state $\bigotimes_{i=1}^{N_{\text{in}}} \widetilde{\mathcal{E}}_k(\rho_i)$ is approximately closest to $\bigotimes_{i=1}^{N_{\text{in}}} \mathcal{E}(\rho_i)$. We will argue that if ϵ' is sufficiently small and N_{in} sufficiently large, then the index k that we find will satisfy

$$\mathbb{E}_{\rho \sim \mathcal{D}} \left[\left\| \mathcal{E}(\rho) - \widetilde{\mathcal{E}}_k(\rho) \right\|_1 \right] \le \epsilon \,, \tag{F19}$$

as desired.

We now proceed to the analysis of this protocol. We begin with an estimate for the distance between product states whose components are pairwise far from each other.

Lemma 4. If ρ_1, \ldots, ρ_N and ρ'_1, \ldots, ρ'_N satisfy $\frac{1}{N} \sum_{i=1}^N \|\rho_i - \rho'_i\|_1 \ge \epsilon$, then

$$\left\| \bigotimes_{i} \rho_{i} - \bigotimes_{i} \rho'_{i} \right\|_{1} \ge 2 \left(1 - \left(1 - \epsilon^{2} / 4 \right)^{N/2} \right). \tag{F20}$$

Proof. For convenience, denote $\epsilon_i := \|\rho_i - \rho_i'\|_1$. We have that $F(\rho_i, \rho_i') \leq 1 - \epsilon_i^2/4$, so

$$\left\| \bigotimes_{i} \rho_{i} - \bigotimes_{i} \rho'_{i} \right\|_{1} \ge 2 \left(1 - \sqrt{F\left(\bigotimes_{i} \rho_{i}, \bigotimes_{i} \rho'_{i}\right)} \right) \tag{F21}$$

$$\geq 2\left(1 - \sqrt{\prod_{i=1}^{N} (1 - \epsilon_i^2/4)}\right) \tag{F22}$$

$$\geq 2\left(1 - \left(\frac{1}{N}\sum_{i}(1 - \epsilon_i^2/4)\right)^{N/2}\right) \tag{F23}$$

$$= 2\left(1 - (1 - \epsilon^2/4)^{N/2}\right) \tag{F24}$$

where the first step follows by the standard inequality $\|\rho - \rho'\|_1 \ge 2\sqrt{1 - F(\rho, \rho')}$, the second step follows by tensorization of fidelity, the third step follows by AM-GM and the fact that $\epsilon_i = \|\rho_i - \rho_i'\|_1 \le 2$, and the last step follows from the assumption.

Next, we elaborate on how to select N_{in} . Consider any $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$. By Hoeffding's inequality, because $\rho_1, \ldots, \rho_{N_{\text{in}}}$ are sampled independently and identically distributed from the distribution \mathcal{D} , and $\|\rho - \rho'\| \leq 2$ for all density matrices ρ, ρ' we have

$$\left| \frac{1}{N_{\text{in}}} \sum_{i=1}^{N_{\text{in}}} \left\| \widetilde{\mathcal{E}}(\rho_i) - \mathcal{E}(\rho_i) \right\|_1 - \underset{\rho \sim \mathcal{D}}{\mathbb{E}} \left\| \widetilde{\mathcal{E}}(\rho) - \mathcal{E}(\rho) \right\|_1 \right| \le \epsilon/2$$
 (F25)

with probability at least $1 - \delta'$ provided $N_{\text{in}} = \Omega(\log(1/\delta')/\epsilon^2)$. By a union bound over $\mathcal{N}_{\epsilon',n,m,p}$, the above bound holds simultaneously for all $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$ with probability at least $1 - |\mathcal{N}_{\epsilon',n,m,p}|\delta'$. Henceforth, we condition on this event holding.

The following shows that it suffices to find $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$ for which the product state $\bigotimes_{i=1}^{N_{\text{in}}} \widetilde{\mathcal{E}}(\rho_i)$ is sufficiently close to $\bigotimes_{i=1}^{N_{\text{in}}} \mathcal{E}(\rho_i)$.

Lemma 5. If (F25) holds for all $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$, then if

$$\left\| \bigotimes_{i=1}^{N_{\text{in}}} \widetilde{\mathcal{E}}(\rho_i) - \bigotimes_{i=1}^{N_{\text{in}}} \mathcal{E}(\rho_i) \right\|_1 \le 2 \left(1 - (1 - \epsilon^2 / 16)^{N_{\text{in}}/2} \right)$$
 (F26)

for some $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$, we have that $\mathbb{E}_{\rho \sim \mathcal{D}} \left\| \widetilde{\mathcal{E}}(\rho) - \mathcal{E}(\rho) \right\|_1 \leq \epsilon$.

Proof. We prove the contrapositive. Suppose $\mathbb{E}_{\rho \sim \mathcal{D}} \| \widetilde{\mathcal{E}}(\rho) - \mathcal{E}(\rho) \|_{1} > \epsilon$. Then by (F25), we have

$$\frac{1}{N_{\text{in}}} \sum_{i=1}^{N_{\text{in}}} \left\| \widetilde{\mathcal{E}}(\rho_i) - \mathcal{E}(\rho_i) \right\|_1 \ge \epsilon/2.$$
 (F27)

The lemma follows from Lemma 4.

As we show in the next lemma, there exists an $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$, namely the process in the covering net which is closest to \mathcal{E} , for which (F26) holds, provided ϵ' is sufficiently small.

Lemma 6. For any $\epsilon' > 0$, there exists an $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$ for which

$$\left\| \bigotimes_{i=1}^{N_{\text{in}}} \widetilde{\mathcal{E}}(\rho_i) - \bigotimes_{i=1}^{N_{\text{in}}} \mathcal{E}(\rho_i) \right\|_{1} \le N_{\text{in}} \epsilon'.$$
 (F28)

Proof. Take $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$ satisfying $\|\widetilde{\mathcal{E}}(\rho) - \mathcal{E}(\rho)\|_1 \le \epsilon'$ for all ρ . For convenience, let $\sigma_i := \mathcal{E}(\rho_i)$, $\sigma_i' := \widetilde{\mathcal{E}}(\rho_i)$, and $\delta_i := \sigma_i' - \sigma_i$ for all $i \in \{1, \ldots, N_{\text{in}}\}$. Then

$$\bigotimes_{i=1}^{N_{\text{in}}} \sigma_i' - \bigotimes_{i=1}^{N_{\text{in}}} \sigma_i = \sum_{i=1}^{N_{\text{in}}} \left[\left(\bigotimes_{j=1}^{i-1} \sigma_j' \right) \otimes (\sigma_i' - \sigma_i) \otimes \left(\bigotimes_{j=i+1}^{N_{\text{in}}} \sigma_j \right) \right], \tag{F29}$$

so by the triangle inequality we conclude that $\|\bigotimes_i \sigma'_i - \bigotimes_i \sigma_i\|_1 \leq N_{\text{in}} \epsilon'$ as claimed.

Lemma 5 and Lemma 6 guarantee the existence of a process in $\mathcal{N}_{\epsilon',n,m,p}$ satisfying the desired bound of (F19). To complete the proof of Theorem 11, we will use the following special case of a result from [47] to find a process in the covering net which performs comparably.

Theorem 12 ([47], Theorem 1.5). Suppose we are given m fixed hypothesis states $\sigma_1, \ldots, \sigma_M \in \mathbb{C}^{d \times d}$, parameters $0 < \epsilon, \delta < 1/2$, and access to copies of a state $\rho \in \mathbb{C}^{d \times d}$. Then there is an algorithm that uses

$$N = \mathcal{O}\left(\frac{1}{\epsilon^2} \left(\log^3 M + \alpha \log M\right) \cdot \alpha\right)$$
 (F30)

copies of ρ for $\alpha := \log(\log(1/\eta)/\delta)$ and $\eta := \min_i \|\rho - \sigma_i\|_1$ such that with probability at least $1 - \delta$ the algorithm outputs a $k \in \{1, \ldots, M\}$ for which $\|\rho - \sigma_k\|_1 \le 4\eta$.

We can now put together the ingredients assembled in this section to complete the proof of Theorem 11.

Proof of Theorem 11. We first prove the theorem for constant δ . Take $\epsilon' = c/N_{\rm in}$ for a sufficiently small constant c > 0, $\delta' = 1/(10|\mathcal{N}_{\epsilon',n,m,p}|)$, and $N_{\rm in} = \widetilde{\mathcal{O}}(p/\epsilon^2)$, where $\widetilde{\mathcal{O}}(\cdot)$ hides factors of $\log n, \log m, \log p, \log 1/\epsilon$, so that $N_{\rm in} \geq \Omega(\log(1/\delta')/\epsilon^2)$ and (F25) holds for all $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$ with probability at least 4/5. Note that for some absolute constant C > 0,

$$1 - (1 - \epsilon^2 / 16)^{N_{\rm in}/2} \ge 1 - e^{-N_{\rm in}\epsilon^2 / 32} = 1 - \delta'^C \ge \Omega(1).$$
 (F31)

In contrast, by Lemma 6 and our choice of ϵ' , there exists some $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$ for which

$$\left\| \bigotimes_{i=1}^{N_{\text{in}}} \widetilde{\mathcal{E}}(\rho_i) - \bigotimes_{i=1}^{N_{\text{in}}} \mathcal{E}(\rho_i) \right\|_{1} \le c.$$
 (F32)

Applying Theorem 12 to $\sigma_k = \bigotimes_i \widetilde{\mathcal{E}}_k(\rho_i)$ where $\widetilde{\mathcal{E}}_k$ is the k-th element of $\mathcal{N}_{\epsilon',n,m,p}$, using N_{out} copies of $\rho = \bigotimes_i \mathcal{E}(\rho_i)$ where

$$N_{\text{out}} = \mathcal{O}\left(\frac{1}{\epsilon^2} \left(\log^3 |\mathcal{N}_{\epsilon',n,m,p}| + \alpha \log |\mathcal{N}_{\epsilon',n,m,p}|\right) \cdot \alpha\right), \qquad \alpha := \mathcal{O}(\log \log(1/c)), \tag{F33}$$

we can output a k for which $\|\sigma_k - \rho\|_1 \le 4c$ with probability 4/5. By taking the constant c sufficiently small, we can leverage (F31) and Lemma 5 to conclude that

$$\mathbb{E}_{\rho \sim \mathcal{D}} \left[\left\| \mathcal{E}(\rho) - \widetilde{\mathcal{E}}_k(\rho) \right\|_1 \right] \le \epsilon. \tag{F34}$$

Note that N_{out} in (F33) is dominated by the $\epsilon^{-2} \log |\mathcal{N}_{\epsilon',n,m,p}|$ term since $\alpha = \mathcal{O}(1)$, and so recalling (F17) and our choice of $\epsilon' = c/N_{\text{in}}$ we obtain

$$N_{\text{out}} = \mathcal{O}\left(\frac{p^3}{\epsilon^2}\log^3\left((n+m)pN_{\text{in}}\right)\right) = \mathcal{O}\left(\frac{p^3}{\epsilon^2}\log^3\left((n+m)pN_{\text{in}}\right)\right) = \widetilde{\mathcal{O}}\left(\frac{p^3}{\epsilon^2}\right),\tag{F35}$$

where again $\widetilde{\mathcal{O}}(\cdot)$ hides logarithmic factors in $n, m, p, 1/\epsilon$.

As we require N_{out} copies of $\bigotimes_{i=1}^{N_{\text{in}}} \mathcal{E}(\rho_i)$, we must make $N_{\text{out}} \cdot N_{\text{in}} = \widetilde{\mathcal{O}}(p^4/\epsilon^4)$ accesses to \mathcal{E} . By union bounding over (F25) holding for all $\widetilde{\mathcal{E}} \in \mathcal{N}_{\epsilon',n,m,p}$ and over the success of the algorithm in Theorem 12, we obtain Theorem 11 for $\delta = 2/5$ from the assumption that p = poly(n).

We now describe how to extend this result to general δ by a standard clustering argument. We can run $r:=\Theta(\log(1/\delta))$ independent copies of the above protocol, resulting in indices k_1,\ldots,k_r into $\mathcal{N}_{\epsilon',n,m,p}$ such that for any fixed $i\in[r]$, $\|\sigma_{k_i}-\rho\|_1\leq 4c$ with probability at least 3/5. If $S\subset[r]$ denotes the set of $i\in[r]$ for which $\|\sigma_{k_i}-\rho\|_1\leq 4c$, then by a Chernoff bound, $|S|\geq r/2$ with probability at least $1-\delta$ provided the constant factor in the definition of r is sufficiently large. Condition on the event that $|S|\geq r/2$.

Let k be an index into $\mathcal{N}_{\epsilon',n,m,p}$ for which there are at least r/2 indices $i \in [r]$ for which $\|\sigma_k - \sigma_{k_i}\|_1 \leq 8c$, and output the channel $\widetilde{\mathcal{E}}_k$. Such a k certainly exists: take any $i \in S$ and note that by triangle inequality, for any other $j \in S$ we have

$$\|\sigma_{k_i} - \sigma_{k_j}\|_1 \le \|\sigma_{k_i} - \rho\|_1 + \|\sigma_{k_j} - \rho\|_1 \le 8c.$$
 (F36)

Now observe that regardless of which k we choose that meets the criterion that at least r/2 indices $i \in [r]$ satisfy $\|\sigma_k - \sigma_{k_i}\|_1 \leq 8c$, we must have

$$\|\sigma_k - \rho\|_1 \le 12c. \tag{F37}$$

Indeed, suppose to the contrary. Then for any $i \in S$,

$$\|\sigma_k - \sigma_{k_i}\|_1 \ge \|\sigma_k - \rho\|_1 - \|\sigma_{k_i} - \rho\|_1 > 12c - 4c = 8c, \tag{F38}$$

where the second step is by the definition of S and the assumption that (F37) does not hold. As $|S| \ge r/2$, this yields a contradiction of the fact that there are at least r/2 indices $i \in [r]$ for which $\|\sigma_k - \sigma_{k_i}\|_1 \le 8c$.

If we take the constant c sufficiently small, then (F37) together with (F31) and Lemma 5 allow us to conclude that $\mathbb{E}_{\rho \sim \mathcal{D}}\left[\left\|\mathcal{E}(\rho) - \widetilde{\mathcal{E}}_k(\rho)\right\|_1\right] \leq \epsilon$. As we ran $r := \Theta(\log(1/\delta))$ independent copies of the protocol that we used for constant failure probability δ , we merely incur an additional $\Theta(\log(1/\delta))$ multiplicative overhead in the number of access to \mathcal{E} we must make for general failure probability δ .

G. Predicting highly-incompatible observables with bounded quantum memory

Here we will substantively generalize our results for predicting highly-incompatible observables, given in Appendix D. We show an exponential lower bound on the number of experiments when the size of the additional quantum memory is not large enough.

G.1. Background and statement of results

Let us first recapitulate the setting of our previous results, so as to draw a contrast with our generalization. We have so far considered an experimentalist who is given sequential access to copies of an unknown state ρ . In each measurement round, the experimentalist receives a copy of ρ , and can measure with a POVM. The residual post-measurement state is then discarded, and only the classical data of the POVM outcome is kept. This classical

data can be used to inform the choice of POVM measurement employed in subsequent rounds, i.e. the protocol can be adaptive. This kind of protocol is emblematic of most contemporary and historical experiments in physics.

Note that the information maintained and processed from round to round in the protocols described above is solely classical. With the advent of quantum computers and more flexible quantum memory architectures, a new possibility emerges. Suppose that the unknown state ρ in question is an n-qubit state. Moreover, suppose we have n+k qubit registers under our control. Then we can use those registers however we please, including performing arbitrary quantum information processing. Our only constraint is that each time we receive a new copy of ρ , we must necessarily use n qubits of our registers to hold it. It is thus appropriate to say that we have k qubits of quantum memory, since even when we receive a new state ρ we can still maintain k qubits worth of quantum data. We further allow ourselves to maintain and process an arbitrary amount of classical data, thought of as being stored in a classical device external to our quantum system.

A new question immediately presents itself: are there experimental tasks which are exponentially hard with only k qubits of quantum memory, but easy with k' > k qubits of quantum memory? In [12] this question was answered in the affirmative, but only in the sense of query complexity. That is, it was shown that there is an experimental task which requires $\Omega(2^{(n-k)/3})$ copies of ρ if there are only k qubits of quantum memory; however, the gate complexity of achieving the task is always exponential regardless of the size of k. Let us unpack this result. Note that if k = n, the bound $\Omega(2^{(n-k)/3})$ becomes trivial; indeed, it can be shown that one only requires a modest number of copies of ρ to achieve the specified task if n = k. But even in this case, the total number of quantum operations required is exponentially large in n. Nonetheless, the result is an interesting one: it means that unless k goes as $n - c \log(n)$ (i.e. unless k is logarithmically close to n), the task is exponentially hard. When $k \sim n - c \log(n)$, the task is only polynomially hard in the sense of query complexity.

While the aforementioned result is theoretically interesting, it does not correspond to a quantum memory advantage that could be realized by a quantum device on account of the exponential gate complexity required to achieve the task for a quantum memory of *any* size. Here we ameliorate this issue, and present the first example of a quantum memory advantage in the sense of both query *and* gate complexity, and as such it can be realized on a quantum device. Moreover, we have realized this quantum advantage in our experiments on the Sycamore quantum computer.

Our experimental task has the form of a partially-revealed many-versus-one distinguishing task, closely related to the one in Appendix D. A statement of the new task is as follows:

Task 3 (Expectation value with bounded quantum memory). There is an unknown state ρ which is either

- 1. A maximally mixed state $I/2^n$ on n qubits, or
- 2. The state $(I+P)/2^n$ where $P \in \{I, X, Y, Z\}^{\otimes n} \setminus \{I^{\otimes n}\}$ is a random but fixed Pauli string.

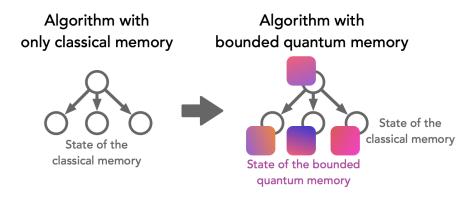
The choice of whether case 1 or case 2 is instantiated is made with equal probability at the outset, and is not revealed. The experimentalist is given access to T copies of the unknown state ρ for a T decided by the experimentalist, and after this an observable O is revealed. The task of the experimentalist is to determine the value of $|\text{tr}(O\rho)|$ using the final state of the n+k qubit registers, along with any classical information that has been stored or processed along the way. In case 1 the operator O is chosen uniformly at random from the non-identity Pauli strings; in case 2 the O is chosen to be the Pauli operator P if the state ρ is $(I+P)/2^n$.

Note that if k = n so that the total number of registers is 2n, then the task can be readily solved using the algorithm given in Appendix D.2. This algorithm is both query and gate efficient: we only require a constant number of copies of ρ (i.e. the number of copies is independent of n) and O(n) gate complexity.

What is difficult is to show that if k < n, then the number of copies of ρ we require to determine $|\text{tr}(O\rho)|$ as per the task above is $\Theta(2^{(n-k)/3})$. We will establish this in the subsections which follow below.

G.2. Review of learning tree framework for bounded quantum memories

Here we provide an exposition of the learning tree framework for bounded quantum memories in [12]. As explained above, suppose we have n+k qubit registers, where we designate k as the quantum memory. Suppose for the moment that k < n. At each round in the protocol, we receive an n-qubit state ρ , which we must hold on the n non-memory registers. (Note that we cannot receive more than one copy of ρ , since we do not have enough registers to hold additional copies on account of k < n.) Then upon receiving and holding ρ , the state of all of our registers can be written as $\rho \otimes \Sigma$, where Σ is the density matrix of the k-qubit quantum memory. The most general operation we can perform on the joint system is a quantum process, followed by a POVM measurement; we



Supplementary Figure 14: Illustration of learning tree representation for algorithms with bounded quantum memory. We consider algorithms with unlimited classical memory and a bounded quantum memory consisting of k qubits. To each node in the tree (corresponding to the state of the classical memory), we associate the k-qubit state of the bounded quantum memory.

can then apply another quantum process followed by another POVM measurement, and so on. However, we can rewrite an alternating sequence of quantum processes and POVM measurements as a single POVM measurement, which we denote by $\{F_i\}_{i=1}^{N-1}$. Writing $F_i = M_i^{\dagger} M_i$, suppose our measurement outputs the *i*th POVM element. Defining

$$A_{M_i}^{\rho}(\mathbf{\Sigma}) := \operatorname{tr}_{1,\dots,n} \left(M_i(\rho \otimes \mathbf{\Sigma}) M_i^{\dagger} \right) ,$$
 (G1)

the reduced density matrix of our quantum memory is

$$\frac{A_{M_i}^{\rho}(\mathbf{\Sigma})}{\operatorname{tr}(A_{M_i}^{\rho}(\mathbf{\Sigma}))} \tag{G2}$$

with probability $\operatorname{tr}(A_{M_i}^{\rho}(\Sigma))$. In the next round, we can leverage our measurement output i to adaptively inform our choice of POVM on

$$\rho \otimes \frac{A_{M_i}^{\rho}(\Sigma)}{\operatorname{tr}(A_{M_i}^{\rho}(\Sigma))}. \tag{G3}$$

Indeed, we can use the information of *all* of our previous POVM outcomes to inform the choice of our next POVM. An illustration is given in Supp. Fig. 14.

The above description is ripe for being cast in the learning tree framework, which we presently articulate. The definition below is the same as Definition 6.1 of [12], albeit with slightly different notation.

Definition 6.1 of [12] (Tree representation of learning states with bounded quantum memory). Let ρ be a fixed, unknown quantum density matrix on n qubits. Suppose we have access to n+k qubit registers. A learning algorithm with a quantum memory of size k can be expressed as a rooted tree \mathcal{T} of depth T, where each node encodes the current state of the quantum memory in addition to the transcript of measurement outcomes the algorithm has seen so far. In particular, the tree satisfies the following properties:

- 1. Each note u is associated with a k-qubit unnormalized density matrix $\Sigma^{\rho}(u)$ corresponding to the current state of the quantum memory.
- 2. For the root r of the tree, $\Sigma^{\rho}(r)$ is an initial state denoted by Σ_0 .
- 3. At each note u, we apply a POVM measurement $\{F_s^u\}_s$ on $\rho \otimes \Sigma^{\rho}(u)$ to obtain a classical outcome s. Each child node v of u is connected through the edge $e_{u,s}$.
- 4. If v is the child note of u connected through the edge $e_{u,s}$, then letting $F_s^u = M_s^{u\dagger} M_s^u$ we have

$$\Sigma^{\rho}(v) := A_{M^u}(\Sigma^{\rho}(u)). \tag{G4}$$

5. Note that for any node u we have that $p^{\rho}(u) := \operatorname{tr}(\mathbf{\Sigma}^{\rho}(u))$ is the probability that the transcript of measurement outcomes observed by the learning algorithm is described by u. Moreover, $\mathbf{\Sigma}^{\rho}(u)/p^{\rho}(u)$ is the (normalized) state of the k-qubit memory at the node u.

Let us unpack the ingredients of this definition. The initial state of our quantum memory is Σ_0 , and we apply some initial choice of POVM $\{F_s^r\}_s$ (where r denotes the 'root' of the tree). If we measure the sth POVM outcome, then the quantum memory is in the unnormalized state $A_{M_s^r}(\Sigma_0)$ with probability $\operatorname{tr}(A_{M_s^r}(\Sigma_0))$. Each outcome s of the POVM corresponds to a child note of the root; thus at the next level of the tree, each node is labelled by the POVM outcome s and the corresponding state of the quantum memory $A_{M_s^r}(\Sigma_0) := \Sigma^{\rho}(s)$. For the next measurement, we can leverage our knowledge of the previous POVM to craft a new POVM to be applied to the present state of the quantum memory. This type of procedure is repeated for many rounds.

To be explicit, suppose that the present state of the quantum memory is $\Sigma^{\rho}(u)$, where the node ρ reflects a sequence or transcript of POVM outcomes which brought us to the present state by an adaptive protocol. We can use this transcript of previous outcomes to choose a new POVM $\{F_s^u\}_s$ that we use to measure $\rho \otimes \Sigma^u(\rho)$, which will result in the output $A_{M_s^u}(\Sigma^{\rho}(u))$ with probability $\operatorname{tr}(A_{M_s^u}(\Sigma^{\rho}(u)))$. The nodes v in the next layer encode the data of the previous measurement outcomes and the latest outcome (i.e., determined by the location of v in the tree), as well as the new (conveniently unnormalized) state of the quantum memory $A_{M_s^u}(\Sigma^{\rho}(u)) := \Sigma^{\rho}(v)$, where here v is connected to v by an edge v, (designating that v is the consequence of the v-th measurement outcome starting from the configuration in v).

G.3. Hardness result for small quantum memories

We will prove the following result using the learning tree framework:

Theorem 13 (Shadow tomography with partial reveal using a bounded quantum memory). Consider Task 3 for learning an expectation value with a bounded quantum memory. Any learning algorithm with n + k qubit registers needs $T \ge \Omega(2^{(n-k)/3})$ copies of ρ to determine $|\operatorname{tr}(O\rho)|$ with probability at least 2/3.

On account of (C15), it suffices to upper bound $\mathbb{E}_P[\text{TV}(p_{I/2^n}, p_{(I+P)/2^n})]$. To do so, we will leverage a key technical result coming from Theorem 1.4 of [12]. But in order to state this technical result, we first need to introduce the notion of good Paulis and bad Paulis. While details are provided in Definition 6.4 of [12], here we explain the essential intuition and key properties.

In the learning protocol, we are trying to distinguish between the maximally mixed state $I/2^n$ and states of the form $(I+P)/2^n$. The intuition is that if the size of our quantum memory k is small relative to n, then it is hard to tell the two kinds of states apart. If this was the case, then any round of the protocol should only reveal a very small amount of distinguishing information. In particular, suppose we are at node u in the learning tree, and so the state of the quantum memory at that node is either $\Sigma_{I/2^n}(u)$ or $\Sigma_{(I+P)/2^n}(u)$ for some Pauli string P. If we consider a POVM $\{F_s^u\}_s$ where $F_s^u = M_s^{u\dagger} M_s^u$, then if we measure some fixed outcome s the new state of the quantum memory will be either $A_{M_s^u}^{I/2^n}(\Sigma_{I/2^n}(u))$ or $A_{M_s^u}^{(I+P)/2^n}(\Sigma_{(I+P)/2^n}(u))$. We would like for $\|A_{M_s^u}^{I/2^n}(\Sigma_{I/2^n}(u)) - A_{M_s^u}^{(I+P)/2^n}(\Sigma_{(I+P)/2^n}(u))\|_1$ to be exponentially small in n-k, in particular relative to some distinguishing measure between $\Sigma_{I/2^n}(u)$ and $\Sigma_{(I+P)/2^n}(u)$. This would mean that starting from node u, the next POVM measurement will not significantly change our ability to distinguish the two possibilities for the resulting memory registers. While we cannot guarantee that such a property holds for all states $(I+P)/2^n$, such a property will hold for some P's. Given a node u, the set of good Paulis P[u] is the set of all Pauli operators satisfying a particular version of the above property for all edges from the root of the tree to u (see Definition 6.4 of [12] for details). The residual Paulis are called the set of bad Paulis. In other words, the good Paulis P[u] designate the states $(I+P)/2^n$ which are hard to distinguish from $I/2^n$ for a particular instantiation of the learning tree, specifically for the sequence of POVMs that get us from the root of said learning tree to the node u. By contrast, the bad Paulis reveal too much information.

We have the following useful Lemma about bad Paulis, which we will soon leverage in the proof of Theorem 13:

Lemma 7 (Fact 6.5 of [12]). For any edge $e_{u,s}$, there are at most $2^{-(n-k)/3} \cdot (4^n - 1)$ bad Paulis P. In particular, along any root-to-leaf path of the learning tree, there are at most $T \cdot 2^{-(n-k)/3} \cdot (4^n - 1)$ Paulis which are bad for some edge along the path.

Equipped with our discussion of good and bad Paulis, we can now state the following technical result from [12]:

Lemma 8 (Following from the proof of Theorem 1.4 of [12]). We have the inequality

$$\frac{1}{4^{n}-1} \sum_{\ell \in \text{leaf}(\mathcal{T})} \sum_{P \in P[\ell]} \left\| \mathbf{\Sigma}_{I/2^{n}}(\ell) - \mathbf{\Sigma}_{(I+P)/2^{n}}(\ell) \right\|_{1} \le T \cdot 2^{-(n-k)/3} \cdot \sqrt{\frac{2^{2n}}{2^{2n}-1}}.$$
 (G5)

We are finally ready to prove Theorem 13.

Proof of Theorem 13. Let us upper bound $\mathbb{E}_P[\text{TV}(p_{I/2^n}, p_{(I+P)/2^n})]$. We have the inequalities

$$\mathbb{E}_{P}\left[\text{TV}(p_{I/2^{n}}, p_{(I+P)/2^{n}})\right] \leq \mathbb{E}_{P}\left[\sum_{\ell} \max(0, p_{I/2^{n}}(\ell) - p_{(I+P)/2^{n}}(\ell))\right]$$
(G6)

$$\leq \mathbb{E}_{P} \left[\sum_{\ell} \min \left(p_{I/2^{n}}(\ell), |p_{I/2^{n}}(\ell) - p_{(I+P)/2^{n}}(\ell)| \right) \right]$$
 (G7)

$$\leq \mathbb{E}_{P} \left[\sum_{\ell} \min(p_{I/2^{n}}(\ell), \|\mathbf{\Sigma}_{I/2^{n}}(\ell) - \mathbf{\Sigma}_{(I+P)/2^{n}}(\ell)\|_{1}) \right]$$
 (G8)

$$\leq \sum_{\ell} \Pr[P \notin P[\ell]] p_{I/2^n}(\ell) + \frac{1}{4^n - 1} \sum_{P \in P[\ell]} \| \mathbf{\Sigma}_{I/2^n}(\ell) - \mathbf{\Sigma}_{(I+P)/2^n}(\ell) \|_1 . \tag{G9}$$

In the first line, we have used that $\mathrm{TV}(p,q) = \frac{1}{2} \sum_i |p_i - q_i| = \sum_{i:p_i \geq q_i} (p_i - q_i) = \sum_i \max(0,p_i - q_i)$. To go from (G6) to (G7) we used $\max(0,a-b) \leq \min(a,|a-b|)$. In going from (G7) to (G8) we leveraged that $|p_{I/2^n}(\ell) - p_{(I+P)/2^n}(\ell)| = |\mathrm{tr}(\mathbf{\Sigma}_{I/2^n}(\ell) - \mathbf{\Sigma}_{(I+P)/2^n}(\ell))| \leq \|\mathbf{\Sigma}_{I/2^n}(\ell) - \mathbf{\Sigma}_{(I+P)/2^n}(\ell)\|_1$. Finally, to go from (G8) to (G9) we used $\sum_{i \in S} \min(a_i,b_i) \leq \sum_{i \in S \setminus R} a_i + \sum_{i \in R} b_i$ for any $R \subseteq S$. By Lemma 7 and the fact that $\sum_{\ell} p_{I/2^n}(\ell) = 1$, we have the simple bound

$$\sum_{\ell} \Pr[P \notin P[\ell]] \, p_{I/2^n}(\ell) \le T \cdot 2^{-(n-k)/3} \tag{G10}$$

and Lemma 8 gives us

$$\frac{1}{4^n - 1} \sum_{P \in P[\ell]} \left\| \mathbf{\Sigma}_{I/2^n}(\ell) - \mathbf{\Sigma}_{(I+P)/2^n}(\ell) \right\|_1 \le T \cdot 2^{-(n-k)/3} \cdot \sqrt{\frac{2^{2n}}{2^{2n} - 1}}.$$
 (G11)

Then in total, we have

$$\mathbb{E}_{P}\left[\text{TV}(p_{I/2^{n}}, p_{(I+P)/2^{n}})\right] \le T \cdot 2^{-(n-k)/3} \left(1 + \sqrt{\frac{2^{2n}}{2^{2n} - 1}}\right). \tag{G12}$$

If the left-hand side is $\Omega(1)$, then we must thus have $T \geq \Omega(2^{(n-k)/3})$, as claimed.