RESEARCH ARTICLE

QUANTUM COMPUTING

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], Jarrod R. McClean[3]*

Quantum technology promises to revolutionize how we learn about the physical world. An experiment that processes quantum data with a quantum computer could have substantial advantages over conventional experiments in which quantum states are measured and outcomes are processed with a classical computer. We proved that quantum machines could learn from exponentially fewer experiments than the number required by conventional experiments. This exponential advantage is shown for predicting properties of physical systems, performing quantum principal component analysis, and learning about physical dynamics. Furthermore, the quantum resources needed for achieving an exponential advantage are quite modest in some cases. Conducting experiments with 40 superconducting qubits and 1300 quantum gates, we demonstrated that a substantial quantum advantage is possible with today's quantum processors.

Humans learn about nature through experiments, but until now our ability to acquire knowledge has been hindered by viewing the quantum world through a classical lens. The rapid advancement 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 those physical systems, quantum sensors (1) will eventually be able to transduce (2) quantum information in physical systems directly to a quantum memory (3, 4), in which it can be processed by a quantum computer. Figure 1A 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, 6, 7). In both the conventional and quantum-enhanced settings, multiple copies of the same quantum state are acquired. The crucial distinction is that the copies are measured one at a time in conventional experiments whereas entangling measurements across multiple copies are allowed in quantum-enhanced experiments.

Recent mathematical analyses performed by some of the authors show that there exist properties of an n-qubit system that a quantum machine can learn efficiently whereas the requisite 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 (Fig. 1). We depict conventional and quantum-enhanced experiments for this scenario in Fig. 1B. In conventional experiments, each copy of ρ is measured separately, the measurement data are 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, after which the quantum machine outputs the prediction after processing the quantum data in the quantum memory. We proved that for some tasks, the number of experiments needed to learn a desired property is exponential in n with the conventional strategy, but only polynomial in n using the quantum-enhanced strategy. For suitably defined tasks, we could 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 showed that

quantum-enhanced experiments have a similar exponential advantage in a related scenario shown in Fig. 1C, in which the goal is to learn about a quantum process $\mathcal{E}$ rather than a quantum state ρ. Advantages of entangling measurements over single-copy measurements have been noticed previously (11, 12), but our work goes much further by establishing an advantage that scales exponentially with system size.

Building on previous observations (8, 13), we proved that for a task that entails acquiring information about a large number of noncommuting observables, quantum-enhanced experiments could have an exponential advantage even when the measured quantum state is unentangled. Our work substantially reduces the complexity of the required quantum-enhanced experiments, improving the prospects for near-term implementation. By performing experiments with up to 40 superconducting qubits, we showed that this quantum advantage persisted even when using currently available quantum processors. We also demonstrated quantum advantage in learning the symmetry class of a physical evolution operator, inspired by recent theoretical advances (9, 13). Finally, in a theoretical contribution we rigorously proved that quantum-enhanced experiments have an exponential advantage in learning about the principal component of a noisy state, as previously indicated (14).

In our proof-of-principle experiments, we directly executed 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.

Provable quantum advantage

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 supplementary materials. Our experimental demonstrations are discussed below in the section titled Demonstrations of Quantum Advantage. 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

[1]Institute for Quantum Information and Matter, Caltech, Pasadena, CA, USA. [2]Department of Computing and Mathematical Sciences, Caltech, Pasadena, CA, USA. [3]Google Quantum AI, Venice, CA 90291, USA. [4]Harvard Society of Fellows, Cambridge, MA 02138, USA. [5]Black Hole Initiative, Cambridge, MA 02138, USA. [6]Department of Electrical Engineering and Computer Science, University of California Berkeley, Berkeley, CA, USA. [7]Simons Institute for the Theory of Computing, Berkeley, CA, USA. [8]Microsoft Research AI, Redmond, WA 98052, USA. [9]Institute for Integrated Circuits, Johannes Kepler University Linz, Austria. [10]AWS Center for Quantum Computing, Pasadena, CA 91125, USA.
*Corresponding author. Email: hsinyuan@caltech.edu (H.-Y.H.); jmcclean@google.com (J.R.M.)

Fig. 1. Illustration of quantum-enhanced and conventional experiments.
(**A**) Quantum-enhanced experiments versus conventional experiments. Quantum-enhanced or conventional experiments interface with a quantum or classical machine running a quantum or classical learning algorithm that can store and process quantum or 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 experiences evolution under $\mathcal{E}$. In the conventional setting, the classical machine specifies the input state to $\mathcal{E}$ by using a classical bitstring and obtains classical measurement data (*33*). In the quantum-enhanced setting, the evolution of $\mathcal{E}$ 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.

encodes how the classical memory changes as we obtain more experimental data. We then analyzed 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 (*13*) is given in supplementary materials, section 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 $\{O_1, O_2, \ldots\}$, 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 (*8*, *13*), we required the learner to predict exponentially many observables, which is not possible in practice if the system size is large. To demonstrate the advantage in an actual device, we proved that predicting just the absolute value of one observable requires exponentially many copies in the conventional scenario. By contrast, predicting the entire set of observables can be achieved with a polynomial number of copies in the quantum-enhanced scenario. We thereby established the following constant versus exponential separation. The proof is given in supplementary materials, section 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, whereas 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 $\rho \propto (I + \alpha P)$, in which 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 and 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 supplementary materials section 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 ML task with a quantum advantage is quantum principal component analysis (PCA) (*14*). 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 component of 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 (*15*) and therefore highly informative about the uncorrupted state. To perform quantum PCA, a learning algorithm was introduced in (*14*) on the basis of phase estimation, which requires fault-tolerant quantum computers. One can also obtain information about the principal component of ρ by using more near-term algorithms, such as virtual cooling (*16*), virtual distillation (*17*, *18*), and variational algorithms (*19*, *20*).

Although the quantum PCA algorithm in (*14*) is exponentially faster than known algorithms

based on conventional experiments, this advantage was not proven against all possible algorithms in the conventional scenario. We rigorously established the exponential quantum advantage for performing quantum PCA. The exponential quantum advantage also holds in some of the near-term proposals (*16, 17*). The proofs are provided in supplementary materials section 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, whereas a constant number of experiments will suffice in the quantum-enhanced scenario.

It is worth commenting on recent results in (*21, 22*) 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 high precision requires measuring exponentially many copies of ρ, which takes an exponential number of experiments and exponential time. Hence, the assumptions of (*21, 22*) do not hold here. See (*23*), 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 or classical machine in the quantum-enhanced or conventional setting; see Fig. 1C. We showed 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 $\mathcal{E}$ accurately on average. By 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 supplementary materials, section 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, whereas a polynomial number of experiments will suffice in the quantum-enhanced scenario.

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 simply no way to access enough classical data to perform the specified tasks if the number of experiments is subexponential in n. However, these exponential separations apply in an idealized setting in which quantum states are stored and processed perfectly. This leads us to ask whether access to quantum memory unlocks a substantial quantum advantage under more realistic conditions.

For two different tasks, we have investigated the robustness of the quantum advantage by conducting experiments with a superconducting quantum processor. We consider specialized tasks that maintain exponential quantum advantage and have better noise robustness than the general tasks described in the previous section. 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 $\rho = 2^{-n}(I + \alpha P)$, in which $\alpha = \pm 0.95$, P is a Pauli operator, and both α and 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 then ask the machine to determine which of $|tr(Q_1\rho)|$ and $|tr(Q_2\rho)|$ is larger.

In the conventional scenario in which 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 quantum-enhanced scenario, by contrast, 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. In the quantum-enhanced scenario, we consider two different methods for analyzing the measurement data. The first method uses a specialized formula for estimating $|tr(Q\rho)|$, given in Appendix D2. Figure 2A depicts—as a function of the system size n—the number of experiments needed in the conventional and quantum-enhanced scenarios to achieve 70% prediction accuracy, in which the data from the quantum-enhanced experiments is analyzed by this first method. Also shown is a theoretical lower bound on the number of experiments needed in the conventional scenario, proven in Appendix D4. The first method is explicitly tailored to the structure of this particular learning problem and so cannot be applied readily to other problems. Our second method is more flexible and hence more broadly applicable; we make predictions by feeding the measurement data to a supervised ML model based on a recurrent neural network (*25, 26, 27*), as depicted in Fig. 2B. In contrast to the first method, the ML method does not require prior knowledge about the learning task. We train the neural network with noiseless simulation data for small system sizes ($n < 8$). We then 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 $|tr(Q_1\rho)|$ or $|tr(Q_2\rho)|$ is larger. Figure 2C shows the performance of the ML model as we train the neural network. Despite the noisy storage and processing in the experimental device, we observed a substantial quantum advantage using both the specialized and ML methods. Notably, when using ML, training on smaller systems sufficed for making good predictions on larger systems, a further indication that the measurement data in the quantum-enhanced scenario is so revealing that no special-purpose method is needed to extract a clear signal.

The second task we studied, which pertains to Theorem 3, was inspired by the recent observation that quantum-enhanced experiments can efficiently identify the symmetry class of a quantum evolution operator, whereas conventional experiments cannot (*9, 13*). An unknown n-qubit quantum evolution operator is presented, drawn either from the class of all unitary transformations or 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 on the basis of data obtained from either quantum-enhanced experiments or conventional experiments. An illustration is shown in Fig. 3A.

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. By 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.

Fig. 2. Quantum advantage in learning physical states. (**A**) Quantum advantage in the number of experiments needed to achieve ≥70% accuracy. Here, Q corresponds to results running the best-known strategy for quantum-enhanced experiments, described in Appendix D2, 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 D4. Even running on a noisy quantum processor, quantum-enhanced experiments are seen to vastly outperform the best theoretically achievable conventional results (C, LB). (**B**) Supervised 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) (25, 26). The neurons in the GRU are aggregated to predict an output. (**C**) 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 with noiseless simulation for small system sizes ($n < 8$). We consider the cross entropy (34) as the training loss. Then we use the supervised ML model to make predictions with 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. The purple (Q) and gray (C) dots on the y-axis are the accuracy of the best-known quantum-enhanced and conventional strategy considered in (A). Random guessing yields a prediction accuracy of 0.5.
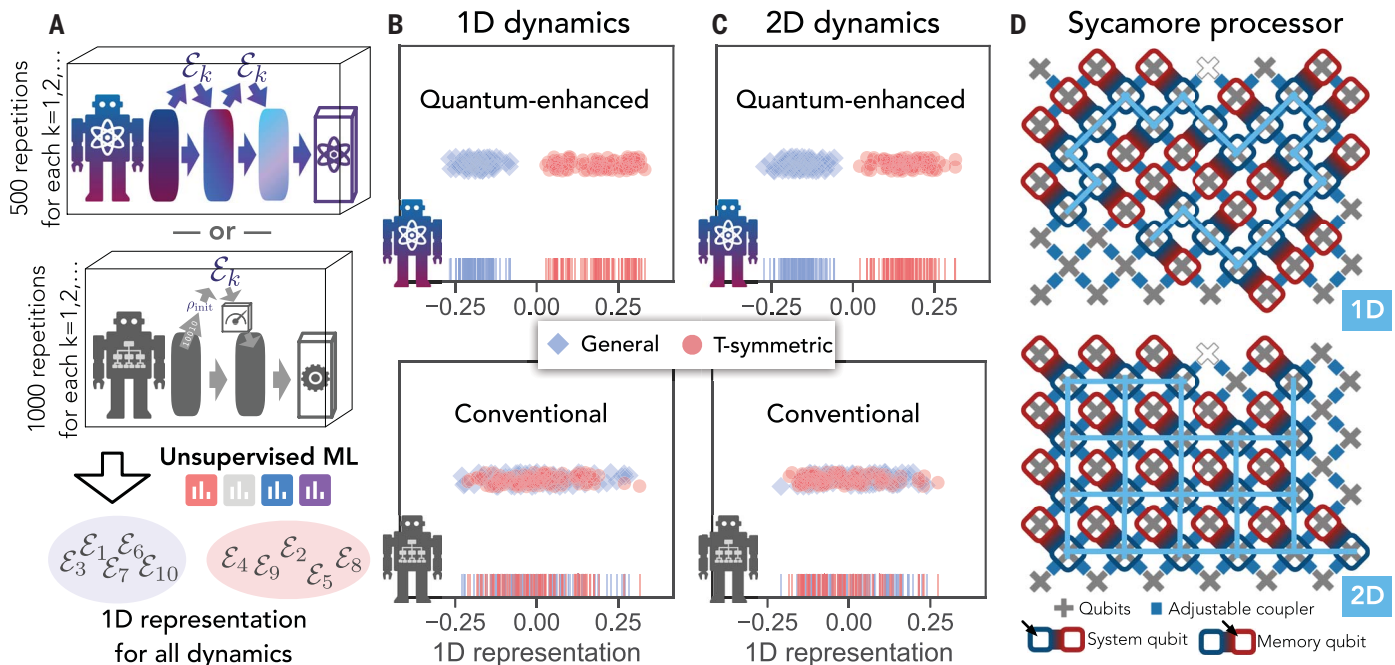
Fig. 3. Quantum advantage in learning physical dynamics. (**A**) Unsupervised ML model. We perform 500 repetitions of quantum-enhanced experiments (each accessing $\mathcal{E}_k$ twice) for every physical process $\mathcal{E}_k$ and feed the data into an unsupervised ML model (Gaussian kernel PCA) (28) to learn a 1D 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 the processes satisfy time-reversal symmetry (blue diamonds) whereas the other half do not (red circles). When fed with data from quantum-enhanced experiments, the ML model accurately discovers the underlying symmetry pattern. By 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.

Each evolution operator is a one-dimensional (1D) or 2D n-qubit quantum circuit as shown in Fig. 3D. After sampling many different evolution operators from both symmetry classes (and obtaining data from each sampled evolution multiple times), we used an unsupervised ML model (kernel PCA) (28) to find a 1D representation of the evolution operators. The representations learned by the unsupervised ML model are shown in Fig. 3, B and C. By using the quantum-enhanced data, the ML model discovers a clean separation between the two symmetry classes, whereas there is no discernable separation into classes when using data from conventional experiments. The signal from the quantum-enhanced experiments was strong enough that the two classes were easily recognized without access to any labeled training data.

In supplementary materials section A4, we analyzed the measurement data using the best-known special-purpose method specifically designed to distinguish general unitary transformations from real orthogonal transformations. We found a quantum advantage similar to that obtained with the ML model. The revelation that unsupervised learning yields results that are competitive with a more customized analysis highlights the potential for discovering previously unknown phenomena with quantum-enhanced measurement strategies. Properties that are blurred beyond recognition by single-copy measurements are brought into sharp relief by two-copy measurements.

Outlook

We have investigated how quantum technology can enhance our ability to discover unknown phenomena occurring in nature. For a variety of tasks, we proved that quantum-enhanced strategies that use 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 when the physical system exhibits only classical correlations. Although many previous studies of quantum advantage have focused on computational tasks with known inputs, our work focused instead on learning tasks in which 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 ML (29,30) and quantum sensing (1).

Our experiments with up to 40 qubits in a superconducting quantum processor showed that a substantial quantum advantage is already evident when using today's noisy intermediate-scale quantum platforms (31).

These experiments demonstrated that supervised and unsupervised ML models (27, 32) employing data obtained from quantum-enhanced experiments could 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 with a quantum computer. Although for now we lack suitably advanced sensors and transducers, we have conducted proof-of-concept experiments in which quantum data were 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.

REFERENCES AND NOTES

1. C. L. Degen, F. Reinhard, P. Cappellaro, *Rev. Mod. Phys.* **89**, 035002 (2017).
2. N. Lauk *et al.*, *Quantum Sci. Technol.* **5**, 020501 (2020).
3. A. I. Lvovsky, B. C. Sanders, W. Tittel, *Nat. Photonics* **3**, 706–714 (2009).
4. E. Dennis, A. Kitaev, A. Landahl, J. Preskill, *J. Math. Phys.* **43**, 4452–4505 (2002).
5. D. Gottesman, T. Jennewein, S. Croke, *Phys. Rev. Lett.* **109**, 070503 (2012).
6. J. Bland-Hawthorn, M. J. Sellars, J. G. Bartholomew, *J. Opt. Soc. Am. B* **38**, A86–A98 (2021).
7. V. Giovannetti, S. Lloyd, L. Maccone, *Nat. Photonics* **5**, 222–229 (2011).
8. H.-Y. Huang, R. Kueng, J. Preskill, *Phys. Rev. Lett.* **126**, 190505 (2021).
9. D. Aharonov, J. Cotler, X.-L. Qi, Quantum algorithmic measurement. arXiv.2004.01372 [quant-ph] (2021)
10. F. Arute *et al.*, *Nature* **574**, 505–510 (2019).
11. A. Peres, W. K. Wootters, *Phys. Rev. Lett.* **66**, 1119–1122 (1991).
12. C. H. Bennett *et al.*, *Phys. Rev. A* **59**, 1070–1091 (1999).
13. S. Chen, J. Cotler, H.-Y. Huang, J. Li, in *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2022), pp. 574–585.
14. S. Lloyd, M. Mohseni, P. Rebentrost, *Nat. Phys.* **10**, 631–633 (2014).
15. B. Koczor, *New J. Phys.* **23**, 123047 (2021).
16. J. Cotler *et al.*, *Phys. Rev. X* **9**, 031013 (2019).
17. W. J. Huggins *et al.*, *Phys. Rev. X* **11**, 041036 (2021).
18. B. Koczor, *Phys. Rev. X* **11**, 031057 (2021).
19. R. LaRose, A. Tikku, É. O'Neel-Judy, L. Cincio, P. J. Coles, *npj Quantum Information* **5**, 1–10 (2019).
20. M. Cerezo, K. Sharma, A. Arrasmith, P. J. Coles, Variational quantum state eigensolver. arXiv.2004.01372 [quant-ph] (2020).
21. E. Tang, *Phys. Rev. Lett.* **127**, 060503 (2021).
22. N.-H. Chia *et al.*, in *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing* (2020), pp. 387–400.
23. J. Cotler, H.-Y. Huang, J. R. McClean, Revisiting dequantization and quantum advantage in learning tasks. arXiv.2112.00811 [quant-ph] (2021).
24. H.-Y. Huang, R. Kueng, J. Preskill, *Nat. Phys.* **16**, 1050–1057 (2020).
25. J. Chung, C. Gulcehre, K. Cho, Y. Bengio, Empirical evaluation of gated recurrent neural networks on sequence modeling, arXiv.1412.3555 [cs.NE] (2014).
26. D. Tang, B. Qin, T. Liu, in *Proceedings of the 2015 conference on empirical methods in natural language processing* (2015), pp. 1422–1432.
27. I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning* (The MIT Press, 2016).
28. B. Schölkopf, A. Smola, K.-R. Müller, *Neural Comput.* **10**, 1299–1319 (1998).
29. J. Biamonte *et al.*, *Nature* **549**, 195–202 (2017).
30. M. Broughton *et al.*, Tensorflow quantum: A software framework for quantum machine learning. arXiv.2003.02989 [quant-ph] (2021).
31. J. Preskill, Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
32. M. Mohri, A. Rostamizadeh, A. Talwalkar, *Foundations of Machine Learning* (The MIT Press, 2018).
33. M. Mohseni, A. T. Rezakhani, D. A. Lidar, *Phys. Rev. A* **77**, 032322 (2008).
34. K. P. Murphy, *Machine Learning: a Probabilistic Perspective* (MIT press, 2012).
35. H.-Y. Huang *et al.*, Code for quantum advantage in learning from experiments, Github (2022); https://github.com/quantumlib/ReCirq/tree/master/recirq/qml_lfe.
36. H.-Y. Huang *et al.*, *Zenodo* (2022); https://zenodo.org/record/6400225#.YoZ23ajMI2w.

SUPPLEMENTARY MATERIALS

science.org/doi/10.1126/science.abn7293
Materials and Methods
Supplementary Text
Figs. S1 to S15
Table S1
Appendices A to G
References (37–68)

Science

Quantum advantage in learning from experiments

Hsin-Yuan HuangMichael BroughtonJordan CotlerSitan ChenJerry LiMasoud MohseniHartmut NevenRyan BabbushRichard KuengJohn PreskillJarrod R. McClean

Learning from quantum experiments

There is considerable interest in extending the recent success of quantum computers in outperforming their conventional classical counterparts (quantum advantage) from some model mathematical problems to more meaningful tasks. Huang *et al.* show how manipulating multiple quantum states can provide an exponential advantage over classical processing of measurements of single-quantum states for certain learning tasks. These include predicting properties of physical systems, performing quantum principal component analysis on noisy states, and learning approximate models of physical dynamics (see the Perspective by Dunjko). In their proof-of-principle experiments using up to 40 qubits on a Google Sycamore quantum processor, the authors achieved almost four orders of magnitude of reduction in the required number of experiments over the best-known classical lower bounds. —YS

View the article online
https://www.science.org/doi/10.1126/science.abn7293
Permissions
https://www.science.org/help/reprints-and-permissions

Use of this article is subject to the Terms of service