

The complexity of NISQ

Received: 19 February 2023

Accepted: 25 August 2023

Published online: 26 September 2023

 Check for updatesSitan Chen^{1,6}✉, Jordan Cotler²✉, Hsin-Yuan Huang^{3,4}✉ & Jerry Li⁵✉

The recent proliferation of NISQ devices has made it imperative to understand their power. In this work, we define and study the complexity class NISQ, which encapsulates problems that can be efficiently solved by a classical computer with access to noisy quantum circuits. We establish super-polynomial separations in the complexity among classical computation, NISQ, and fault-tolerant quantum computation to solve some problems based on modifications of Simon's problems. We then consider the power of NISQ for three well-studied problems. For unstructured search, we prove that NISQ cannot achieve a Grover-like quadratic speedup over classical computers. For the Bernstein-Vazirani problem, we show that NISQ only needs a number of queries logarithmic in what is required for classical computers. Finally, for a quantum state learning problem, we prove that NISQ is exponentially weaker than classical computers with access to noiseless constant-depth quantum circuits.

Fault-tolerant quantum computing promises to offer speed-ups to various computational problems, including simulating quantum systems^{1–5}, factoring large numbers^{6,7}, performing optimization^{8–11}, and solving linear systems of equations¹². While a fault-tolerant quantum computer has not yet been built, recent technological advancements have resulted in the development of new quantum devices that can outperform the best classical supercomputers for certain artificial computational tasks¹³. The opportunities offered by these devices and the challenges they engender have invigorated theorists and experimentalists alike, ushering in a new age of research into quantum computation and information often referred to as the NISQ (noisy intermediate-scale quantum) era¹⁴.

Computation in the NISQ era is modeled by hybrid computation consisting of a classical computer and a noisy quantum device. Most currently available noisy quantum devices, such as those built from superconducting qubits^{13,15–17}, trapped ions^{18–21}, nuclear spins in silicon^{22–24}, or other solid-state systems^{25–27}, are restricted to preparing a noisy initial state, performing noisy quantum gates, and executing a noisy measurement on all qubits to generate a random bitstring. In most circumstances these noisy quantum devices are too weak to perform useful computation on their own. Generally, useful computation is facilitated by a classical computer that repeatedly runs the noisy quantum device with various gate sequences to obtain different

classical output bit strings, and then performs classical post-processing on those strings. Algorithms within this computational model are referred to as hybrid quantum-classical algorithms^{28,29} or NISQ algorithms^{14,30}.

The excitement around NISQ algorithms has led to a plethora of new near-term algorithms targeting different applications, including quantum chemistry^{3,28,31–33}, machine learning^{17,34–39}, combinatorial optimization^{40–43}, linear system solvers^{44–46}, and experimental data analysis^{47–52}. However, to the best of our knowledge, no existing works have rigorously examined the class of all possible NISQ algorithms and studied their inherent computational power. As a result, many important and fundamental questions remain unanswered. In particular, how powerful are NISQ algorithms compared to classical algorithms? Are NISQ algorithms inherently weaker than fault-tolerant quantum algorithms?

In this work, we formalize and study these basic questions through the lens of computational complexity theory. To do so rigorously, we define a complexity class which we believe encapsulates what is possible on the vast majority of existing quantum devices. Our definition is intended to capture the following capabilities of noisy quantum devices, which we alluded to above:

1. Noisy quantum gates: The device can execute noisy two-qubit logic gates. Using quantum logic gates (as opposed to, say,

¹Department of Electrical Engineering and Computer Science, University of California, Berkeley, Berkeley, CA, USA. ²Society of Fellows, Harvard University, Cambridge, MA, USA. ³Institute for Quantum Information and Matter, Caltech, Pasadena, CA, USA. ⁴Department of Computing and Mathematical Sciences, Caltech, Pasadena, CA, USA. ⁵Microsoft Research AI, Redmond, WA, USA. ⁶Present address: John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, USA. ✉e-mail: sitan@seas.harvard.edu; jcotler@fas.harvard.edu; hsinyuan@caltech.edu; jerrli@microsoft.com

more general non-unitary CPTP maps) is standard in existing quantum devices, and it is well-understood that in real-world settings, they will be subject to noise. For concreteness, we consider the standard model of local depolarizing noise per qubit. However, our results extend to more general noise models (see Remarks 2.4, 4.19 and 6.5) in the Supplementary Information.

2. Noisy state preparation at the start: The quantum devices have a fixed number of qubits and as such cannot bring in fresh qubits during the computation. This means that the device must prepare all qubits at the start. Notably, since we assume all quantum gates are subject to noise, this means all qubits will accrue entropy throughout the computation.
3. Noisy measurement at the end: The quantum devices are limited to perform noisy measurements only at the end of the computation, which means the measurement is performed on all qubits simultaneously. From a physical perspective, this constraint arises due to the difficulty of isolating subsets of qubits and measuring them without decohering the residual qubits.

Finally, we consider a classical computer that can repeatedly run the noisy quantum device and analyze the output from the noisy quantum device.

These constraints are chosen to encapsulate the gap between the physical limitations of what we can achieve with existing quantum computers and general quantum computation. We note these considerations preclude the implementation of all known general fault-tolerant quantum computation schemes^{53–59}, but that removing any one of these constraints would already allow for some form of non-trivial quantum fault tolerance^{53,57,58}. The obstruction to fault tolerance can be understood intuitively. The noisy quantum gates cause all qubits to accrue entropy, which cannot be pumped out until the measurement at the end. Since too much entropy would destroy all useful quantum correlations, it is not possible for the noisy quantum devices under the above constraints to perform an arbitrarily long quantum computation. Note that in contrast, if we consider more benign noise which does not increase entropy, such as dephasing or amplitude-damping noise, then there are already schemes for achieving fault-tolerant quantum computation⁵⁷ within the above three constraints.

Motivated by the above considerations, in Section “Definition of NISQ” we formally define the NISQ complexity class to be the set of all problems that can be efficiently solved by a classical computer with access to a noisy quantum device that can (i) prepare a noisy poly(n)-qubit all-zero state, (ii) execute noisy quantum gates, and (iii) perform a noisy measurement on all of the poly(n) qubits. We subsequently show how the NISQ complexity class is situated relative to classical computers, and full-fledged quantum computers with quantum error correction.

Results

In Section “Definition of NISQ” we give an overview of the definition of NISQ. Then, in Section “Super-polynomial separations”, we give two modifications of Simon’s problem which respectively yield a super-polynomial separation between BPP and NISQ, and an exponential separation between NISQ and BQP. In Section “NISQ in three well-studied problems”, we study the NISQ complexity of three well-known problems: unstructured search, Bernstein-Vazirani problem, and shadow tomography. We defer all technical details to the Supplementary Information.

Definition of NISQ

We begin with the definition of the NISQ complexity class, which is visually depicted in Fig. 1. The formal mathematical definition is given in Supplementary Note 2.A. In Supplementary Note 2.B, we recall the standard definition of oracle access in classical and quantum computation – when the oracle is classical, we extend the definition to NISQ by giving oracle access to both the classical computer and the noisy quantum computer.

Definition 2.1. (NISQ complexity class, informal) NISQ contains all problems that can be solved by a polynomial-time probabilistic classical algorithm with access to a noisy quantum device. To solve a problem of size n , the classical algorithm can access a noisy quantum device that can:

1. Prepare a noisy poly(n)-qubit all-zero state;
2. Apply arbitrarily many layers of noisy two-qubit gates;
3. Perform noisy computational basis measurements on all the qubits simultaneously.

All quantum operations are subject to a constant amount of depolarizing noise per qubit.

The definition of a noisy quantum device given above forbids the implementation of all known fault-tolerant quantum computation schemes^{53–56,59}. Hence, it is plausible that there are problems that could be solved efficiently by a fault-tolerant quantum algorithm but not by a NISQ algorithm, i.e. that $\text{NISQ} \subsetneq \text{BQP}$.

The definition of NISQ immediately gives us that

$$\text{BPP} \subseteq \text{NISQ} \subseteq \text{BQP}. \tag{1}$$

The inclusion $\text{BPP} \subseteq \text{NISQ}$ follows from the fact that a NISQ algorithm is a hybrid quantum-classical algorithm that can also run any classical computation. The latter inclusion $\text{NISQ} \subseteq \text{BQP}$ holds because a quantum computer can simulate any noisy quantum device. However, it remains an open question whether $\text{BPP} \subsetneq \text{NISQ}$ and also if $\text{NISQ} \subsetneq \text{BQP}$. The strict inclusion $\text{BPP} \subsetneq \text{NISQ}$ would imply that NISQ algorithms have a super-polynomial speedup over classical algorithms. The second strict inclusion $\text{NISQ} \subsetneq \text{BQP}$ would imply that NISQ

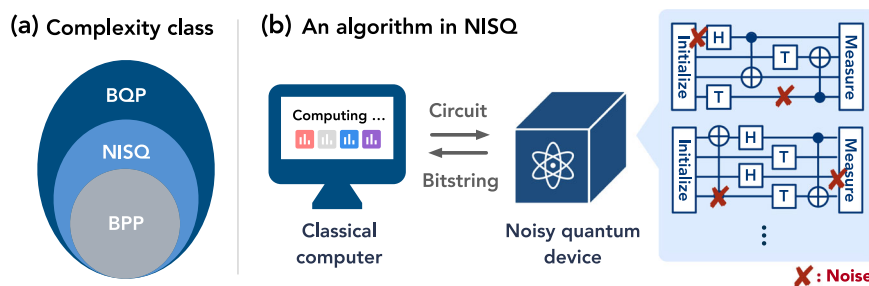


Fig. 1 | Illustration of the NISQ complexity class. **a** Complexity classes: NISQ contains problems that can be solved by classical computation (BPP), and some problems that can be solved by quantum computation (BQP). **b** NISQ algorithm: An algorithm in the complexity class NISQ is modeled by a hybrid quantum-classical

algorithm, where a classical computer can specify the circuit to run on a noisy quantum device and the device would run a noisy version of the circuit and return a random classical bitstring obtained from noisy measurement.

algorithms are not as powerful as fault-tolerant quantum algorithms. Establishing either one of the strict inclusions would imply that $\text{BPP} \subsetneq \text{BQP}$, which is a long-standing open problem in computational complexity theory. In this work, we follow the well-established approach in quantum complexity theory of proving oracle separations. That is, we consider various oracles, which are black boxes implementing certain functions or preparing some states, and show that NISQ algorithms accessing these oracles with noisy queries are strictly stronger or weaker than BPP or BQP algorithms that can access these boxes noiselessly. We consider widely studied oracles like the one for Grover search or for Bernstein-Vazirani problem, as well as variants of oracles like the one for Simon's problem.

Super-polynomial separations

We begin with a set of results involving two modified versions of Simon's problem⁶⁰. Simon's problem is one of the earliest examples of a computational problem demonstrating exponential quantum advantage in query complexity: given a function $f: \{0, 1\}^n \rightarrow \{0, 1\}^n$, the goal is to decide if f is 2-to-1 (with a promise that $f(x) = f(x \oplus s)$ for a secret string $s \in \{0, 1\}^n$) or 1-to-1. In the language of computational complexity theory, it exhibits a classical oracle O for which there is a relativized complexity separation between BPP and BQP, denoted by $\text{BPP}^O \subsetneq \text{BQP}^O$.

In this work, we give two modifications of Simon's problem that yield relativized separations among BPP, NISQ, and BQP. For NISQ, we consider noisy oracle access with local depolarizing noise occurring both before and after the oracle O . We first construct a modification of Simon's problem that requires at least a super-polynomial number of oracle queries for BPP and only a linear number for NISQ. We then give another modification of Simon's problem that requires at least an exponential number of queries for NISQ and only a linear number of queries for BQP. These two results can be summarized as follows:

Theorem 2.2. There is a classical oracle O_1 such that $\text{BPP}^{O_1} \subsetneq \text{NISQ}^{O_1}$.

Theorem 2.3. There is a classical oracle O_2 such that $\text{NISQ}^{O_2} \subsetneq \text{BQP}^{O_2}$.

Theorem 2.2 is established by constructing a robust version of the classical function $f: \{0, 1\}^n \rightarrow \{0, 1\}^n$ in Simon's problem. We denote the robust version as $\tilde{f}: \{0, 1\}^{n'} \rightarrow \{0, 1\}^{n'}$ with more input bits: $n' \gg n$. Each input $x \in \{0, 1\}^{n'}$ corresponds to a large set of inputs $A_x \subseteq \{0, 1\}^{n'}$, such that every $z \in A_x$ produces the same output $\tilde{f}(z) = f(x)$. The new function \tilde{f} is robust to noise, which allows a NISQ algorithm to achieve a super-polynomial speed-up.

The proof of Theorem 2.3 is essentially the opposite of that of Theorem 2.2. We construct a highly non-robust version \tilde{f} of the classical function f in Simon's problem. Any noisy access to \tilde{f} provides exponentially little information, hence any NISQ algorithm would require exponentially more queries than a noiseless quantum algorithm. In fact, the separation we show is even stronger: not only is NISQ exponentially weaker than BQP relative to this oracle, but it is in fact even exponentially weaker than $\text{BPP}^{\text{QNC}^0}$, that is, classical computation assisted by noiseless bounded-depth quantum computation (see Supplementary Note 2.C for definitions).

Our findings give evidence that the computational power of NISQ lies somewhere strictly between BPP and BQP. The detailed proofs are given in Supplementary Note 4. Of course, separations relative to oracles come with the usual provisos inherent in relativized separations, like the fundamental complexity-theoretic barriers to deducing unconditional separations from relativized separations^{61,62}. Another challenge is to instantiate the oracle in practice. While it is impossible to implement the relevant oracles perfectly on NISQ devices, if the oracle can be instantiated with cryptographic functions based on shallow circuits⁶³, then the techniques for proving Theorem 2.2 allow one to implement a robust version of the oracle in NISQ. Despite these

limitations, we view our results as promising evidence that NISQ may be truly intermediate between BPP and BQP.

NISQ in three well-studied problems

Since the modified Simon's problems in Theorems 2.2 and 2.3 are tailored for proving super-polynomial separations, we would also like to study NISQ for more natural problems. We explore three well-known problems in quantum computing: unstructured search, the Bernstein-Vazirani problem, and shadow tomography.

For unstructured search, it is well-known that Grover's algorithm⁶⁴ can achieve a quadratic quantum speedup over any classical algorithm. To find a single marked element among N elements, Grover's algorithm only requires $\mathcal{O}(\sqrt{N})$ queries, whereas any classical algorithm requires at least $\Omega(N)$ queries in the worst case. A natural open question is whether a NISQ algorithm can also achieve such a quadratic speedup. We resolve this open question in the negative by proving the following theorem. The proof is given in Supplementary Note 5.

Theorem 2.4. (Unstructured search) To find a single marked element among N elements, any NISQ algorithm with access to $\text{poly}(N)$ qubits requires at least $\tilde{\Omega}(N)$ queries.

The $\tilde{\Omega}(\cdot)$ neglects logarithmic factors. We stress that the above theorem implies not only that noisy implementation of Grover's algorithm fails to achieve a quadratic speedup but, in fact, that any NISQ algorithm will fail to do so. The intuition behind the proof is that noisy quantum devices can only run for so long before noise overwhelms the system, so it would suffice to prove the lower bound for hybrid quantum-classical algorithms with access to a noiseless quantum device that can run any bounded-depth circuit. While a single run of any noiseless bounded-depth circuit cannot achieve Grover speedup⁶⁵, hybrid quantum-classical algorithms can perform many runs of noiseless bounded-depth circuits that depend adaptively on previous measurement outcomes. We prove that Grover speedup is impossible with hybrid quantum-classical algorithms using tools developed recently in the context of lower bounds for learning quantum states and processes using adaptive single-copy measurements^{47–50,66–68}.

For the second task, the Bernstein-Vazirani problem, we find that a large quantum advantage still remains for NISQ algorithms. Given an unknown n bit string $s \in \{0, 1\}^n$, the Bernstein-Vazirani problem asks how many queries to a function $f(x) = x \cdot s$ are required to learn s . Any classical algorithm requires at least $\Omega(n)$ queries to learn s . However, the Bernstein-Vazirani quantum algorithm can learn the unknown bit string s with just 1 query. We show that the query complexity of this problem in NISQ remains much smaller than the classical query complexity:

Theorem 2.5. (Bernstein-Vazirani) There is a NISQ algorithm that solves the Bernstein-Vazirani problem over n bits in at most $\mathcal{O}(\log n)$ queries.

Perhaps surprisingly, our analysis shows that a simple modification of the original Bernstein-Vazirani algorithm is already sufficiently noise-robust. This result provides optimism that there may be other natural problems for which quantum advantage can be obtained in the NISQ era. The detailed proof of Theorem 2.5 is presented in Supplementary Note 6.

Finally, we consider the problem of predicting many properties in an unknown quantum system, also known as shadow tomography^{69–72}. In particular, we restrict to predicting absolute values of Pauli observables $\{I, X, Y, Z\}^{\otimes n}$: Given many copies of an unknown n -qubit state ρ , the goal is to learn $|\text{tr}(P\rho)|$ for all $P \in \{I, X, Y, Z\}^{\otimes n}$ up to a constant error by processing the quantum state copies. This task has received considerable attention in recent works^{37,48,50,71} which have demonstrated, both theoretically and experimentally, that very simple BQP algorithms can solve this task using only $\mathcal{O}(n)$ copies, while any classical

algorithm that can obtain classical data by performing measurements on individual copies of ρ requires $\Omega(2^n)$ copies. Here we show that this large quantum advantage is damped by the presence of noise. Specifically, we show the following exponential separation between NISQ and BQP. The proof of the following theorem is given in Supplementary Note 7.

Theorem 2.6. (Shadow tomography) Any NISQ algorithm with noise rate λ per qubit requires at least $\Omega((1-\lambda)^{-n})$ copies of ρ to learn $|\text{tr}(P\rho)|$, for all $P \in \{I, X, Y, Z\}^{\otimes n}$ up to a constant error. In contrast, at most $\mathcal{O}(n)$ copies are needed for BQP.

Theorem 2.6 demonstrates that a fault-tolerant quantum algorithm can be exponentially more powerful than any NISQ algorithm in learning quantum systems. On the other hand, for small noise rate λ , the exponential scaling in the lower bound for NISQ algorithms has a base which is close to one. In⁴⁹ it was shown that NISQ algorithms can achieve $(1-\lambda)^{-\mathcal{O}(n)}$ even for more general noise models. This suggests that NISQ algorithms can still perform well for learning quantum systems with a few hundred qubits⁵⁰.

Discussion

By abstracting hybrid quantum-classical computation in the NISQ era into a computational complexity class, our work offers a mathematical framework for reasoning about the potential for noisy quantum advantage. We used tools from quantum query complexity to characterize how NISQ lies between BPP and BQP. On the one hand, the fact that NISQ can be more powerful than BPP provides optimism for the NISQ era and the computational advances it may precipitate. On the other hand, NISQ being less powerful than BQP portends that we will have to wait until the advent of fault-tolerant devices to harness the richest features of quantum computation. Our results for the NISQ complexity of three well-known problems in quantum computing punctuate our outlook by demonstrating specific promises and pitfalls of computation in the NISQ era.

There are many future directions to pursue with the NISQ complexity class. A main open problem is to understand if we could show a separation between BPP, NISQ, and BQP under a standard complexity-theoretic assumption. It would be desirable to have an exponential oracle separation between BPP and NISQ, as opposed to one that is merely super-polynomial. Moreover, one could ask if a similar oracle separation exists between BPP and NISQ with the additional restriction that our quantum gates are spatially local, e.g., the gates are restricted to a two-dimensional geometry. Perhaps under this additional restriction of geometric locality on NISQ, it is possible to better understand the computational complexity of NISQ without relying on oracle separations. Additionally, it would be valuable to analyze the NISQ complexity of other natural quantum algorithms, beyond the ones we studied. Natural targets include the original Simon's problem (as opposed to our variations thereof, see Supplementary Note 10 for one approach in this direction), Forrelation⁷³, Shor's algorithm^{6,7}, linear system solving¹², the recent random oracle result of⁷⁴, and topological data analysis⁷⁵, among many others.

Taking a broader view, our work suggests a roadmap for investigating future quantum devices which may go beyond the NISQ era, but fall short of fault-tolerance. The approach is to formulate a computational complexity class which encapsulates the salient features of whichever quantum devices are contemporary, and then to study that complexity class to make statements about the capabilities of those devices with great generality. In such a future, there should be complexity classes beyond NISQ, but still intermediate to BQP. Whichever generalizations of NISQ prove fruitful in the future, they will have much to tell us about what computation is possible in that future world.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary.

References

- Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996).
- Georgescu, I. M., Ashhab, S. & Nori, F. Quantum simulation. *Rev. Mod. Phys.* **86**, 153 (2014).
- O'Malley, P. J. et al. Scalable quantum simulation of molecular energies. *Phys. Rev. X* **6**, 031007 (2016).
- Babbush, R. et al. Low-depth quantum simulation of materials. *Phys. Rev. X* **8**, 011044 (2018).
- Low, G. H. & Chuang, I. L. Hamiltonian simulation by qubitization. *Quantum* **3**, 163 (2019).
- Shor, P. W. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 124–134 (Ieee, 1994).
- Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Rev.* **41**, 303–332 (1999).
- Crosson, E. & Harrow, A. W. Simulated quantum annealing can be exponentially faster than classical simulated annealing. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, 714–723 (IEEE, 2016).
- Farhi, E. & Harrow, A. W. Quantum supremacy through the quantum approximate optimization algorithm. *arXiv:1602.07674* (2016).
- Brandao, F. G. & Svore, K. M. Quantum speed-ups for solving semidefinite programs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, 415–426 (IEEE, 2017).
- Brandão, F. G. et al. Quantum sdp solvers: Large speed-ups, optimality, and applications to quantum learning. *arXiv:1710.02581* (2017).
- Harrow, A. W., Hassidim, A. & Lloyd, S. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.* **103**, 150502 (2009).
- Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
- Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
- Clarke, J. & Wilhelm, F. K. Superconducting quantum bits. *Nature* **453**, 1031–1042 (2008).
- Reagor, M. et al. Demonstration of universal parametric entangling gates on a multi-qubit lattice. *Sci. Adv.* **4**, eaao3603 (2018).
- Havlíček, V. et al. Supervised learning with quantum-enhanced feature spaces. *Nature* **567**, 209–212 (2019).
- Debnath, S. et al. Demonstration of a small programmable quantum computer with atomic qubits. *Nature* **536**, 63–66 (2016).
- Zhang, J. et al. Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator. *Nature* **551**, 601–604 (2017).
- Friis, N. et al. Observation of entangled states of a fully controlled 20-qubit system. *Phys. Rev. X* **8**, 021012 (2018).
- Wright, K. et al. Benchmarking an 11-qubit quantum computer. *Nat. Commun.* **10**, 1–6 (2019).
- Noiri, A. et al. Fast universal quantum gate above the fault-tolerance threshold in silicon. *Nature* **601**, 338–342 (2022).
- Xue, X. et al. Quantum logic with spin qubits crossing the surface code threshold. *Nature* **601**, 343–347 (2022).
- Madzik, M. T. et al. Precision tomography of a three-qubit donor quantum processor in silicon. *Nature* **601**, 348–353 (2022).
- Weber, J. et al. Quantum computing with defects. *Proc. Natl Acad. Sci.* **107**, 8513–8518 (2010).
- Doherty, M. W. et al. The nitrogen-vacancy colour centre in diamond. *Phys. Rep.* **528**, 1–45 (2013).

27. Choi, S. et al. Observation of discrete time-crystalline order in a disordered dipolar many-body system. *Nature* **543**, 221–225 (2017).
28. McClean, J. R., Romero, J., Babbush, R. & Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *N. J. Phys.* **18**, 023023 (2016).
29. Endo, S., Cai, Z., Benjamin, S. C. & Yuan, X. Hybrid quantum-classical algorithms and quantum error mitigation. *J. Phys. Soc. Jpn.* **90**, 032001 (2021).
30. Bharti, K. et al. Noisy intermediate-scale quantum algorithms. *Rev. Mod. Phys.* **94**, 015004 (2022).
31. Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 1–7 (2014).
32. Kandala, A. et al. Hardware-efficient variational quantum eigen-solver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017).
33. Huggins, W. J. et al. Unbiasing fermionic quantum monte carlo with a quantum computer. *Nature* **603**, 416–420 (2022).
34. Farhi, E. & Neven, H. Classification with quantum neural networks on near term processors. *arXiv:1802.06002* (2018).
35. Schuld, M. & Killoran, N. Quantum machine learning in feature hilbert spaces. *Phys. Rev. Lett.* **122**, 040504 (2019).
36. Huang, H.-Y. et al. Power of data in quantum machine learning. *Nat. Commun.* **12**, 1–9 (2021).
37. Huang, H.-Y., Kueng, R. & Preskill, J. Information-theoretic bounds on quantum advantage in machine learning. *Phys. Rev. Lett.* **126**, 190505 (2021).
38. Caro, M. C. et al. Generalization in quantum machine learning from few training data. *Nat. Commun.* **13**, 1–11 (2022).
39. Cerezo, M., Verdon, G., Huang, H.-Y., Cincio, L. & Coles, P. J. Challenges and opportunities in quantum machine learning. *Nat. Computational Sci.* **2**, 567–576 (2022).
40. Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. *arXiv:1411.4028* (2014).
41. Hadfield, S. et al. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms* **12**, 34 (2019).
42. Zhou, L., Wang, S.-T., Choi, S., Pichler, H. & Lukin, M. D. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Phys. Rev. X* **10**, 021067 (2020).
43. Basso, J., Gamarnik, D., Mei, S. & Zhou, L. Performance and limitations of the qaoa at constant levels on large sparse hypergraphs and spin glass models. *arXiv:2204.10306* (2022).
44. Bravo-Prieto, C. et al. Variational quantum linear solver. *arXiv:1909.05820* (2019).
45. Xu, X. et al. Variational algorithms for linear algebra. *Sci. Bull.* **66**, 2181–2188 (2021).
46. Huang, H.-Y., Bharti, K. & Reberstrost, P. Near-term quantum algorithms for linear systems of equations. *arXiv:1909.07344* (2019).
47. Aharonov, D., Cotler, J. & Qi, X.-L. Quantum algorithmic measurement. *Nat. Commun.* **13**, 1–9 (2022).
48. Chen, S., Cotler, J., Huang, H.-Y. & Li, J. Exponential separations between learning with and without quantum memory. *To appear in FOCS* *arXiv:2111.05881* (2021).
49. Huang, H.-Y., Flammia, S. T. & Preskill, J. Foundations for learning from noisy quantum experiments. *arXiv:2204.13691* (2022).
50. Huang, H.-Y. et al. Quantum advantage in learning from experiments. *Science* **376**, 1182–1186 (2022).
51. Schuster, T. et al. Learning quantum systems via out-of-time-order correlators. *arXiv:2208.02254* (2022).
52. Cotler, J., Schuster, T. & Mohseni, M. Information-theoretic hardness of out-of-time-order correlators. *arXiv:2208.02256* (2022).
53. Aharonov, D. & Ben-Or, M. Fault-tolerant quantum computation with constant error. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, 176–188 (1997).
54. Preskill, J. Fault-tolerant quantum computation. In *Introduction to quantum computation and information*, 213–269 (World Scientific, 1998).
55. Bacon, D., Kempe, J., Lidar, D. A. & Whaley, K. B. Universal fault-tolerant quantum computation on decoherence-free subspaces. *Phys. Rev. Lett.* **85**, 1758 (2000).
56. Raussendorf, R. & Harrington, J. Fault-tolerant quantum computation with high threshold in two dimensions. *Phys. Rev. Lett.* **98**, 190504 (2007).
57. Ben-Or, M., Gottesman, D. & Hassidim, A. Quantum refrigerator. *arXiv preprint arXiv:1301.1995* (2013).
58. Fawzi, O., Grosse, A. & Leverrier, A. Constant overhead quantum fault tolerance with quantum expander codes. *Commun. ACM* **64**, 106–114 (2020).
59. DiVincenzo, D. P. & Aliferis, P. Effective fault-tolerant quantum computation with slow measurements. *Phys. Rev. Lett.* **98**, 020501 (2007).
60. Simon, D. R. On the power of quantum computation. *SIAM J. Comput.* **26**, 1474–1483 (1997).
61. Baker, T., Gill, J. & Solovay, R. Relativizations of the p=?np question. *SIAM J. Comput.* **4**, 431–442 (1975).
62. Aaronson, S. Oracles are subtle but not malicious. In *21st Annual IEEE Conference on Computational Complexity (CCC'06)*, 15–pp (IEEE, 2006).
63. Applebaum, B., Ishai, Y. & Kushilevitz, E. Cryptography in nc⁰. *SIAM J. Comput.* **36**, 845–888 (2006).
64. Grover, L. K. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, 212–219 (1996).
65. Zalka, C. Grover’s quantum searching algorithm is optimal. *Phys. Rev. A* **60**, 2746 (1999).
66. Chen, S., Cotler, J., Huang, H.-Y. & Li, J. A hierarchy for replica quantum advantage. *arXiv:2111.05874* (2021).
67. Bubeck, S., Chen, S. & Li, J. Entanglement is necessary for optimal quantum property testing. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, 692–703 (IEEE, 2020).
68. Chen, S., Huang, B., Li, J. & Liu, A. Tight bounds for quantum state certification with incoherent measurements. *arXiv:2204.07155* (2022).
69. Aaronson, S. Shadow tomography of quantum states. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, 325–338 (2018).
70. Aaronson, S. & Rothblum, G. N. Gentle measurement of quantum states and differential privacy. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 322–333 (2019).
71. Huang, H.-Y., Kueng, R. & Preskill, J. Predicting many properties of a quantum system from very few measurements. *Nat. Phys.* **16**, 1050–1057 (2020).
72. Bădescu, C. & O’Donnell, R. Improved quantum data analysis. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 1398–1411 (2021).
73. Aaronson, S. & Ambainis, A. Correlation: A problem that optimally separates quantum from classical computing. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*, 307–316 (2015).
74. Yamakawa, T. & Zhandry, M. Verifiable quantum advantage without structure. *arXiv:2204.02063* (2022).
75. Akhalwaya, I. Y. et al. Exponential advantage on noisy quantum computers. *arXiv:2209.09371* (2022).
76. Regev, O. & Schiff, L. Impossibility of a quantum speed-up with a faulty oracle. In *International Colloquium on Automata, Languages, and Programming*, 773–781 (Springer, 2008).

Acknowledgements

We would like to thank Isaac Chuang for valuable conversations on complexity classes for learning theory, John Preskill for inspiring discussions on NISQ and its complexity class formulation, and John Wright for bringing⁷⁶ and related works to our attention. S.C. is supported by NSF Award 2103300. J.C. is supported by a Junior Fellowship from the Harvard Society of Fellows, as well as in part by the Department of Energy under grant DE-SC0007870. H.H. is supported by a Google Ph.D. fellowship.

Author contributions

S.C., J.C., H.H., and J.L. contributed equally to this work. All authors developed the theoretical aspects of this work and wrote the manuscript together.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-023-41217-6>.

Correspondence and requests for materials should be addressed to Sitan Chen, Jordan Cotler, Hsin-Yuan Huang or Jerry Li.

Peer review information *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023, corrected publication 2024