ARTICLES
https://doi.org/10.1038/s41567-020-0932-7

nature physics

Predicting many properties of a quantum system from very few measurements

Hsin-Yuan Huang[1,2] ✉, Richard Kueng[1,2,3] and John Preskill[1,2,4]

Predicting the properties of complex, large-scale quantum systems is essential for developing quantum technologies. We present an efficient method for constructing an approximate classical description of a quantum state using very few measurements of the state. This description, called a 'classical shadow', can be used to predict many different properties; order $\log(M)$ measurements suffice to accurately predict M different functions of the state with high success probability. The number of measurements is independent of the system size and saturates information-theoretic lower bounds. Moreover, target properties to predict can be selected after the measurements are completed. We support our theoretical findings with extensive numerical experiments. We apply classical shadows to predict quantum fidelities, entanglement entropies, two-point correlation functions, expectation values of local observables and the energy variance of many-body local Hamiltonians. The numerical results highlight the advantages of classical shadows relative to previously known methods.

Making predictions based on empirical observations is a central topic in statistical learning theory and is at the heart of many scientific disciplines, including quantum physics. For this latter field, predictive tasks, such as estimating target fidelities, verifying entanglement and measuring correlations, are essential for building, calibrating and controlling quantum systems. Recent advances in the size of quantum platforms[1] have pushed traditional prediction techniques—like quantum state tomography—to the limit of their capabilities. This is mainly due to the curse of dimensionality: the number of parameters needed to describe a quantum system scales exponentially with the number of its constituents. Moreover, these parameters cannot be accessed directly, but must be estimated by measuring the system. An informative quantum-mechanical measurement is both destructive (wavefunction collapse) and yields only probabilistic outcomes (Born's rule). Hence, many identically prepared samples are required to estimate accurately even a single parameter of the underlying quantum state. Furthermore, all of these measurement outcomes must be processed and stored in memory for subsequent prediction of relevant features. In summary, reconstructing a full description of a quantum system with n constituents (for example, qubits) necessitates a number of measurement repetitions exponential in n, as well as an exponential amount of classical memory and computing power.

Several approaches have been proposed to overcome this fundamental scaling problem. These include matrix product state (MPS) tomography[2] and neural network tomography[3,4]. Both require only a polynomial number of samples, provided that the underlying state has suitable properties. However, for general quantum systems, these techniques still require an exponential number of samples. See Supplementary Section 3 for details.

Pioneering a conceptually very different line of research, Aaronson[5] pointed out that demanding full classical descriptions of quantum systems may be excessive for many concrete tasks. Instead it is often sufficient to accurately predict certain properties of the quantum system. In quantum mechanics, interesting properties are often linear functions of the underlying density matrix ρ, such as the expectation values $\{o_i\}$ of a set of observables $\{O_i\}$:

$$o_i(\rho) = \mathrm{trace}(O_i\rho) \quad 1 \le i \le M \tag{1}$$

The fidelity with a pure target state, entanglement witnesses and the probability distribution governing the possible outcomes of a measurement are all examples that fit this framework. A nonlinear function of ρ, such as entanglement entropy, may also be of interest. Aaronson coined the term[5,6] 'shadow tomography' for the task of predicting properties without necessarily fully characterizing the quantum state, and he showed that a polynomial number of state copies already suffice to predict an exponential number of target functions. Although very efficient in terms of samples, Aaronson's procedure is very demanding in terms of quantum hardware; a concrete implementation of the proposed protocol requires exponentially long quantum circuits that act collectively on all the copies of the unknown state stored in a quantum memory.

In this Article, we combine the mindset of shadow tomography[5] (predict target functions, not the full state) with recent insights from quantum state tomography[7] (rigorous statistical convergence guarantees) and the stabilizer formalism[8] (efficient implementation). The result is a highly efficient protocol that learns a minimal classical sketch S_ρ—the classical shadow—of an unknown quantum state ρ that can be used to predict arbitrary linear function values (equation (1)) by a simple median-of-means protocol. A classical shadow is created by repeatedly performing a simple procedure: apply a unitary transformation $\rho \mapsto U\rho U^\dagger$, and then measure all the qubits in the computational basis. The number of times this procedure is repeated is called the 'size' of the classical shadow. The transformation U is randomly selected from an ensemble of unitaries, and different ensembles lead to different versions of the procedure that have characteristic strengths and weaknesses. In a practical scheme, each ensemble unitary should be realizable as an efficient quantum circuit. We consider random n-qubit Clifford circuits and tensor products of random single-qubit Clifford circuits as important

[1]Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, CA, USA. [2]Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA, USA. [3]Institute for Integrated Circuits, Johannes Kepler University Linz, Linz, Austria. [4]Walter Burke Institute for Theoretical Physics, California Institute of Technology, Pasadena, CA, USA. ✉e-mail: hsinyuan@caltech.edu

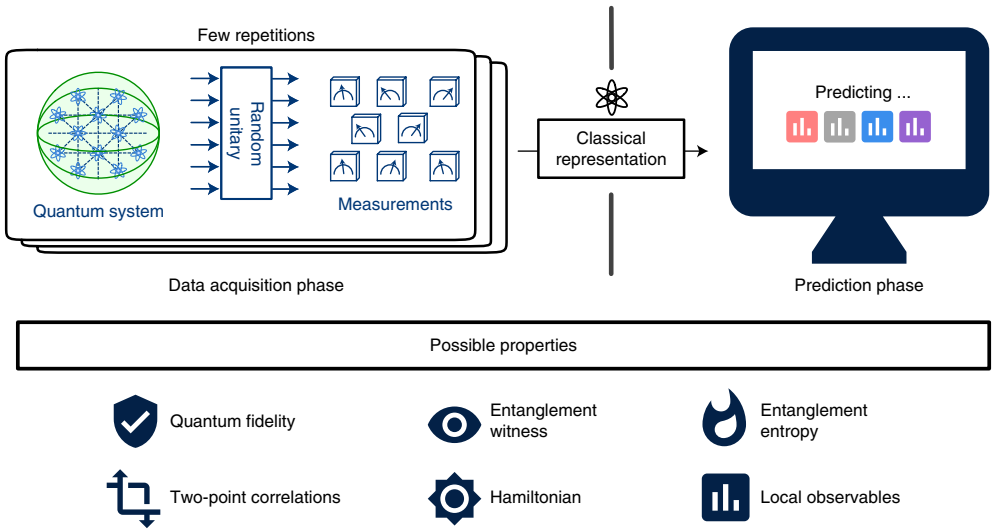

Fig. 1 | An illustration for constructing a classical representation, the classical shadow, of a quantum system from randomized measurements. In the data acquisition phase, we perform a random unitary evolution and measurements on independent copies of an n-qubit system to obtain a classical representation of the quantum system—the classical shadow. Such classical shadows facilitate accurate prediction of a large number of different properties using a simple median-of-means protocol.

special cases. These two procedures turn out to complement each other nicely. Figure 1 provides a visualization and a list of important properties that can be predicted efficiently.

Our main theoretical contribution equips this procedure with rigorous performance guarantees. Classical shadows with size of order $\log(M)$ suffice to predict M target functions in equation (1) simultaneously. Most importantly, the actual system size (number of qubits) does not enter directly. Instead, the number of measurement repetitions N is determined by a (squared) norm $\| O_i \|^2_{\text{shadow}}$. This norm depends on the target functions and the particular measurement procedure used to produce the classical shadow. For example, random n-qubit Clifford circuits lead to the Hilbert–Schmidt norm. On the other hand, random single-qubit Clifford circuits produce a norm that scales exponentially in the locality of target functions, but is independent of system size. The resulting prediction technique is applicable to current laboratory experiments and facilitates the efficient prediction of few-body properties, such as two-point correlation functions, entanglement entropy of small subsystems and expectation values of local observables.

In some cases, this scaling may seem unfavourable. However, we rigorously prove that this is not a flaw of the method, but an unavoidable limitation rooted in quantum information theory. By relating the prediction task to a communication task[9], we establish fundamental lower bounds highlighting that classical shadows are (asymptotically) optimal.

We support our theoretical findings by conducting numerical simulations for predicting various physically relevant properties over a wide range of system sizes. These include quantum fidelity, two-point correlation functions, entanglement entropy and local observables. We confirm that prediction via classical shadows scales favourably and improves on powerful existing techniques—such as machine learning—in a variety of well-motivated test cases. An open-source release for predicting many properties from very few measurements is available at https://github.com/momohuang/predicting-quantum-properties.

Procedure

Throughout this work we restrict attention to n-qubit systems and ρ is a fixed, but unknown, quantum state in $d=2^n$ dimensions. To extract meaningful information, we repeatedly perform a simple measurement procedure: apply a random unitary to rotate the state ($\rho \mapsto U\rho U^\dagger$) and perform a computational-basis measurement. The unitary U is selected randomly from a fixed ensemble. On receiving the n-bit measurement outcome $|\hat{b}\rangle \in \{0,1\}^n$, we store an (efficient) classical description of $U^\dagger|\hat{b}\rangle\langle\hat{b}|U$ in classical memory. It is instructive to view the average (over both the choice of unitary and the outcome distribution) mapping from ρ to its classical snapshot $U^\dagger|\hat{b}\rangle\langle\hat{b}|U$ as a quantum channel:

$$\mathbb{E}\left[U^\dagger|\hat{b}\rangle\langle\hat{b}|U\right] = \mathcal{M}(\rho) \;\Rightarrow\; \rho = \mathbb{E}\left[\mathcal{M}^{-1}\left(U^\dagger|\hat{b}\rangle\langle\hat{b}|U\right)\right] \quad (2)$$

This quantum channel $\mathcal{M}$ depends on the ensemble of (random) unitary transformations. Although the inverted channel $\mathcal{M}^{-1}$ is not physical (it is not completely positive), we can still apply $\mathcal{M}^{-1}$ to the (classically stored) measurement outcome $U^\dagger|\hat{b}\rangle\langle\hat{b}|U$ in a completely classical post-processing step. ($\mathcal{M}$ is invertible if the ensemble of unitary transformations defines a tomographically complete set of measurements; see Supplementary Section 1.) In doing so, we produce a single classical snapshot $\hat{\rho} = \mathcal{M}^{-1}\left(U^\dagger|\hat{b}\rangle\langle\hat{b}|U\right)$ of the unknown state ρ from a single measurement. By construction, this snapshot exactly reproduces the underlying state in expectation (over both unitaries and measurement outcomes): $\mathbb{E}[\hat{\rho}] = \rho$. Repeating this procedure N times results in an array of N independent, classical snapshots of ρ:

$$\begin{aligned} \mathrm{S}(\rho;N) = \Big\{ &\hat{\rho}_1 = \mathcal{M}^{-1}\left(U_1^\dagger|\hat{b}_1\rangle\langle\hat{b}_1|U_1\right), \ldots, \\ &\hat{\rho}_N = \mathcal{M}^{-1}\left(U_N^\dagger|\hat{b}_N\rangle\langle\hat{b}_N|U_N\right) \Big\} \end{aligned} \quad (3)$$

We call this array the classical shadow of ρ. Classical shadows of sufficient size N are expressive enough to predict many properties of the unknown quantum state efficiently. To avoid outlier corruption, we split the classical shadow into equally sized chunks and construct several, independent sample mean estimators. Subsequently, we predict linear function values (1) via median of means estimation[10,11]. This procedure is summarized in Algorithm 1. For many

physically relevant properties O_i and measurement channels $\mathcal{M}$, Algorithm 1 can be carried out very efficiently without explicitly constructing the large matrix $\hat{\rho}_i$.

Median of means prediction with classical shadows can be defined for any distribution of random unitary transformations. Two prominent examples are (1) random n-qubit Clifford circuits and (2) tensor products of random single-qubit Clifford circuits. Example (1) results in a clean and powerful theory, but also practical drawbacks, because $n^2/\log(n)$ entangling gates are needed to sample from n-qubit Clifford unitaries. The corresponding inverted quantum channel is $\mathcal{M}_n^{-1}(X) = (2^n+1)X - \mathbb{1}$. Example (2) is equivalent to measuring each qubit independently in a random Pauli basis. Such measurements can be routinely carried out in many experimental platforms. The corresponding inverted quantum channel is $\mathcal{M}_P^{-1} = \bigotimes_{i=1}^n \mathcal{M}_1^{-1}$. We refer to examples (1)/(2) as random Clifford/Pauli measurements, respectively. In both cases, the resulting classical shadow can be stored efficiently in a classical memory using the stabilizer formalism.

Algorithm 1. Median of means prediction based on a classical shadow $S(\rho, N)$.

1 **function** LinearPredictions($O_1, \ldots, O_M$, $S(\rho; N)$, K)

2 Import $S(\rho; N) = [\hat{\rho}_1, \ldots, \hat{\rho}_N]$ ▷ Load classical shadow

3 Split the shadow into K equally-sized parts and set ▷ Construct K estimators of ρ

$\hat{\rho}_{(k)} = \frac{1}{\lfloor N/K \rfloor} \sum_{l=(k-1)\lfloor N/K \rfloor + 1}^{k\lfloor N/K \rfloor} \hat{\rho}_l$

4 **for** $i = 1$ to M **do**

5 Output $\hat{o}_i(N,K) = \text{median}\left\{\text{tr}\left(O_i\hat{\rho}_{(1)}\right), \ldots, \text{tr}\left(O_i\hat{\rho}_{(K)}\right)\right\}$. ▷ Median of means estimation

Rigorous performance guarantees

Theorem 1 (informal version). *Classical shadows of size N suffice to predict M arbitrary linear target functions* $\text{tr}(O_1\rho), \ldots, \text{tr}(O_M\rho)$ *up to additive error ϵ given that* $N \geq (\text{order}) \log(M) \max_i ||O_i||^2_{\text{shadow}}/\epsilon^2$. *The definition of the norm $||O_i||_{\text{shadow}}$ depends on the ensemble of unitary transformations used to create the classical shadow.*

We refer to Supplementary Section 1 for background, a detailed statement and proofs. Theorem 1 is most powerful when the linear functions have a bounded norm that is independent of system size. In this case, classical shadows allow for predicting a large number of properties from only a logarithmic number of quantum measurements.

The norm $||O_i||_{\text{shadow}}$ in Theorem 1 plays an important role in defining the space of linear functions that can be predicted efficiently. For random Clifford measurements, $||O||^2_{\text{shadow}}$ is closely related to the Hilbert–Schmidt norm $\text{tr}(O^2)$. As a result, a large collection of (global) observables with a bounded Hilbert–Schmidt norm can be predicted efficiently. For random Pauli measurements, the norm scales exponentially in the locality of the observable, not the actual number of qubits. For an observable O_i that acts non-trivially on (at most) k qubits, $||O_i||^2_{\text{shadow}} \leq 4^k ||O_i||^2_\infty$, where $||\cdot||_\infty$ denotes the operator norm. (This scaling can be further improved to 3^k if O_i is a tensor product of k single-qubit observables.) This guarantees the accurate prediction of many local observables from a much smaller number of measurements.

Illustrative example applications

Quantum fidelity estimation. Suppose we wish to certify that an experimental device prepares a desired n-qubit state. Typically, this target state $|\psi\rangle\langle\psi|$ is pure and highly structured, for example, a Greenberger–Horne–Zeilinger (GHZ) state[12] for quantum communication protocols or a toric code ground state[13] for fault-tolerant quantum computation. Theorem 1 asserts that a classical shadow (Clifford measurements) of dimension-independent size suffices to accurately predict the fidelity of any state in the lab with any pure target state. This improves on the best existing result on direct fidelity estimation[14] which requires $O(2^n/\epsilon^4)$ samples in the worst case. Moreover, a classical shadow of polynomial size allows for estimating an exponential number of (pure) target fidelities all at once.

Entanglement verification. Fidelities with pure target states can also serve as (bipartite) entanglement witnesses[15]. For many (but not all[16]) bipartite entangled states ρ, there exists a constant α and an observable $O = |\psi\rangle\langle\psi|$ such that $\text{tr}(O\rho) > \alpha \geq \text{tr}(O\rho_s)$, for all (bipartite) separable states ρ_s. Establishing $\text{tr}(O\rho) > \alpha$ verifies the existence of entanglement in the state ρ. Any $O = |\psi\rangle\langle\psi|$ that satisfies the above condition is known as an entanglement witness for the state ρ. Classical shadows (Clifford measurements) of logarithmic size allow for checking a large number of potential entanglement witnesses simultaneously.

Predicting expectation values of local observables. Many near-term applications of quantum devices rely on repeatedly estimating a large number of local observables. For example, low-energy eigenstates of a many-body Hamiltonian may be prepared and studied using a variational method, in which the Hamiltonian, a sum of local terms, is measured many times. Classical shadows constructed from a logarithmic number of random Pauli measurements can efficiently estimate polynomially many such local observables. Because only single-qubit Pauli measurements suffice, this measurement procedure is highly efficient. Potential applications include quantum chemistry[17] and lattice gauge theory[18].

Predicting expectation values of global observables (non-example). Classical shadows are not without limitations. In our examples, the size of classical shadows must either scale with $\text{tr}(O_i^2)$ (Clifford measurements) or must scale exponentially in the locality of O_i (Pauli measurements). Both quantities can simultaneously become exponentially large for non-local observables with large Hilbert–Schmidt norm. A concrete example is the Pauli expectation value of a spin chain: $\langle P_{i_1} \otimes \cdots \otimes P_{i_n} \rangle_\rho = \text{tr}(O_1\rho)$, where $\text{tr}(O_1^2) = 2^n$ and $k = n$ (non-local observable). In this case, classical shadows of exponential size may be required to accurately predict a single expectation value. In contrast, a direct spin measurement achieves the same accuracy with only of order $1/\epsilon^2$ copies of the state ρ.

Matching information-theoretic lower bounds

The non-example above raises an important question: does the scaling of the required number of measurements with Hilbert–Schmidt norm or with the locality of observables arise from a fundamental limitation, or is it merely an artefact of prediction with classical shadows? A rigorous analysis reveals that this scaling is no mere artefact; rather, it stems from information-theoretic reasons.

Theorem 2 (informal version). *Any procedure based on single-copy measurements, that can predict any M linear functions* $\text{tr}(O_i\rho)$ *up to additive error ϵ, requires at least* (order) $\log(M)\ \max_i ||O_i||^2_{\text{shadow}}/\epsilon^2$ *measurements.*

Here, $||O_i||^2_{\text{shadow}}$ could be taken as the Hilbert–Schmidt norm $\text{tr}(O_i^2)$ or as a function scaling exponentially in the locality of O_i. The proof results from embedding the abstract prediction procedure into a communication protocol. Quantum information theory imposes fundamental restrictions on any quantum communication protocol and allows us to deduce stringent lower bounds. See Supplementary Sections 7 and 8 for details and proofs.

The two main technical results complement each other nicely. Theorem 1 equips classical shadows with a constructive performance

guarantee: an order of $\log(M)\max_i\|O_i\|^2_{\text{shadow}}/\epsilon^2$ single-copy measurements suffice to accurately predict an arbitrary collection of M target functions. Theorem 2 highlights that this number of measurements is unavoidable in general.

Predicting nonlinear functions

The classical shadow $S(\rho;N)=\{\hat{\rho}_1,\ldots,\hat{\rho}_N\}$ of the unknown quantum state ρ may also be used to predict nonlinear functions $f(\rho)$. We illustrate this with a quadratic function $f(\rho)=\mathrm{tr}(O\rho\otimes\rho)$, where O acts on two copies of the state. Because $\hat{\rho}_i$ is equal to the quantum state ρ in expectation, one could predict $\mathrm{tr}(O\rho\otimes\rho)$ using two independent snapshots $\hat{\rho}_i,\hat{\rho}_j$, $i\neq j$. Because of independence, $\mathrm{tr}(O\hat{\rho}_i\otimes\hat{\rho}_j)$ correctly predicts the quadratic function in expectation:

$$\mathbb{E}\,\mathrm{tr}(O\hat{\rho}_i\otimes\hat{\rho}_j)=\mathrm{tr}(O\mathbb{E}\hat{\rho}_i\otimes\mathbb{E}\hat{\rho}_j)=\mathrm{tr}(O\rho\otimes\rho) \tag{4}$$

To reduce the prediction error, we use N independent snapshots and symmetrize over all possible pairs: $\frac{1}{N(N-1)}\sum_{i\neq j}\mathrm{tr}(O\hat{\rho}_i\otimes\hat{\rho}_j)$. We then repeat this procedure several times and form their median to further reduce the likelihood of outlier corruption (similar to median of means). Rigorous performance guarantees are presented in Supplementary Section 1.C. This approach readily generalizes to higher-order polynomials using U-statistics[19].

One particularly interesting nonlinear function is the second-order Rényi entropy, $-\log(\mathrm{tr}(\rho_A^2))$, where A is a subsystem of the n-qubit quantum system. We can rewrite the argument in the log as $\mathrm{tr}(\rho_A^2)=\mathrm{tr}(S_A\rho\otimes\rho)$, where S_A is the local swap operator of two copies of subsystem A, and use classical shadows to obtain very accurate predictions. The required number of measurements scales exponentially in the size of subsystem A, but is independent of total system size. Probing this entanglement entropy is a useful task and a highly efficient specialized approach has been proposed in ref. [20]. We compare this method of Brydges and colleagues to classical shadows in the numerical experiments.

For nonlinear functions, unlike linear ones, we have not derived an information-theoretic lower bound on the number of measurements needed, although it may be possible to do so by generalizing our methods.

Numerical experiments

One of the key features of prediction with classical shadows is scalability. The data acquisition phase is designed to be tractable for state-of-the-art platforms (Pauli measurements) and future quantum computers (Clifford measurements), respectively. The resulting classical shadow can be stored efficiently in classical memory. For many important features—such as local observables or global features with efficient stabilizer decompositions—scalability, moreover, extends to the computational cost associated with median of means prediction.

These design features allowed us to conduct numerical experiments for a wide range of problems and system sizes (up to 160 qubits). The computational bottleneck is not feature prediction with classical shadows, but generating synthetic data, that is, classically generating target states and simulating quantum measurements. Needless to say, this classical bottleneck does not occur in actual experiments. We then use this synthetic data to learn a classical representation of ρ and use this representation to predict various interesting properties.

Machine-learning-based approaches[3,4] are among the most promising alternative methods that have applications in this regime, where the Hilbert space dimension is roughly comparable to the total number of silicon atoms on earth ($2^{160}\simeq 10^{48}$). For example, a recent version of neural network quantum state tomography (NNQST) is a generative model that is based on a deep neural network trained on independent quantum measurement outcomes with local SIC/tetrahedral positive-operator valued measures (POVMs)[21]. In this section, we consider the task of learning a classical representation of an unknown quantum state, and using the representation to predict various properties, addressing the relative merit of classical shadows and alternative methods.

Predicting quantum fidelities via Clifford measurements. Here we focus on classical shadows based on random Clifford measurements, which are designed to predict observables with a bounded Hilbert–Schmidt norm. When the observables have efficient representations—such as efficient stabilizer decompositions—the computational cost for performing median of means prediction can also be efficient. (The runtime of Algorithm 1 is dominated by the cost of computing quadratic functions $\left\langle\hat{b}\middle|UOU^\dagger\middle|\hat{b}\right\rangle$ in 2^n dimensions. If $O=|\psi\rangle\langle\psi|$ is a stabilizer state, the Gottesman–Knill theorem allows for evaluation in $\mathcal{O}(n^2)$-time.) An important example is the quantum fidelity with a target state. In ref. [3], the viability of NNQST is demonstrated by considering GHZ states with a varying number of qubits n. Numerical experiments highlight that the number of measurement repetitions (size of the training data) to learn a neural network model of the GHZ state that achieves a target fidelity of 0.99 scales linearly in n. We have also implemented NNQST for GHZ states and compared it to median of means prediction with classical shadows. Figure 2a confirms the linear scaling of NNQST and the assertion of Theorem 1: classical shadows of constant size suffice to accurately estimate GHZ target fidelities, regardless of the actual system size. In addition, we have also tested the ability of both approaches to detect potential state preparation errors. More precisely, we consider a scenario where the GHZ source introduces a phase error with probability $p\in[0,1]$:

$$\rho_p=(1-p)\left|\psi^+_{\text{GHZ}}(n)\right\rangle\left\langle\psi^+_{\text{GHZ}}(n)\right|+p\left|\psi^-_{\text{GHZ}}(n)\right\rangle\left\langle\psi^-_{\text{GHZ}}(n)\right|,$$
$$\left|\psi^\pm_{\text{GHZ}}(n)\right\rangle=\tfrac{1}{\sqrt{2}}\left(|0\rangle^{\otimes n}\pm|1\rangle^{\otimes n}\right) \tag{5}$$

We learn a classical representation of the GHZ source and subsequently predict the fidelity with the pure GHZ state. Figure 2b highlights that the classical shadow prediction accurately tracks the decrease in target fidelity as the error parameter p increases. NNQST, in contrast, seems to consistently overestimate this target fidelity. In the extreme case ($p=1$), the true underlying state is completely orthogonal to the target state, but NNQST nonetheless reports fidelities close to one. This shortcoming arises because the POVM-based machine-learning approach can only efficiently estimate an upper bound on the true quantum fidelity efficiently. To estimate the actual fidelity, an exceedingly large number of measurements is needed. Similar experiments are described in Supplementary Section 2, where we focus on toric code ground states and entanglement witnesses, respectively.

Predicting two-point correlation and subsystem entanglement entropy (Pauli measurements). Classical shadows based on random Clifford measurements excel at predicting quantum fidelities. However, random Clifford measurements can be challenging to implement in practice, because many entangling gates are needed to implement general Clifford circuits. Next we consider classical shadows based on random local Pauli measurements, which are easier to perform experimentally. The subsystem properties can be predicted efficiently by constructing the reduced density matrix from the classical shadow. Therefore, the computational complexity scales exponentially only in the subsystem size, rather than the size of the entire system. Our numerical experiments confirm that classical shadows obtained using random Pauli measurements excel at predicting few-body properties of a quantum state, such as two-point correlation functions and subsystem entanglement entropy.

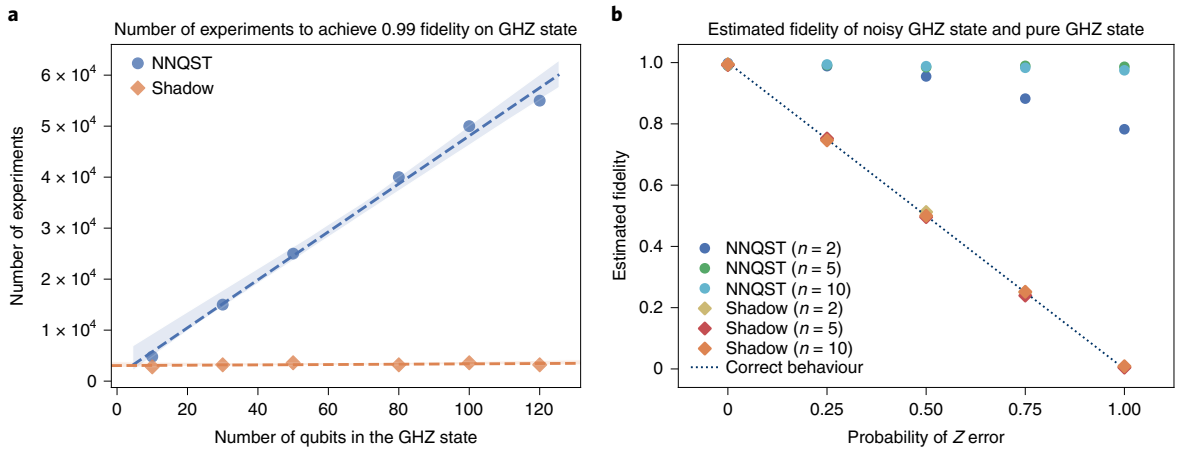

Fig. 2 | Predicting quantum fidelities using classical shadows (Clifford measurements) and NNQST. a, Number of measurements required to identify an *n*-qubit GHZ state with 0.99 fidelity. The shaded regions show the s.d. of the needed number of experiments over 10 independent runs. The dashed lines are the linear regression lines for the number of experiments under different system sizes. **b**, Estimated fidelity between a perfect GHZ target state and a noisy preparation, where *Z* errors can occur with probability $p \in [0, 1]$, under 6×10^4 experiments. The dotted line represents the true fidelity as a function of *p*. NNQST can only estimate an upper bound on quantum fidelity efficiently, so we consider this upper bound for NNQST and use quantum fidelity for the classical shadow.

Two-point correlation functions. NNQST has been shown to predict two-point correlation functions effectively[3]. Here, we compare classical shadows with NNQST for two physically motivated test cases: ground states of the antiferromagnetic transverse field Ising model in one dimension (TFIM) and the antiferromagnetic Heisenberg model in two dimensions. The Hamiltonian for TFIM is $H = J\sum_i \sigma_i^Z \sigma_{i+1}^Z + h\sum_i \sigma_i^X$, where $J > 0$, and we consider a chain of 50 lattice sites. The critical point occurs at $h = J$ and exhibits power-law decay of correlations rather than exponential decay. The Hamiltonian for the two-dimensional (2D) Heisenberg model is $H = J\sum_{\langle i,j\rangle} \vec{\sigma}_i \cdot \vec{\sigma}_j$, where $J > 0$, and we consider an 8×8 triangular lattice. We follow the approach in ref. [3], where the ground state is approximated by a tensor network found using the density matrix renormalization group (DMRG). Random Pauli measurements on the ground state may then be simulated using this tensor network. The two methods are compared in Fig. 3. In Fig. 3a,b, we can see that both the classical shadow (with Pauli measurements) and NNQST perform well at predicting two-point correlations. However, NNQST has a larger error for the 2D Heisenberg model; note that, for larger separations (the lower right corner of the surface plot), NNQST produces some fictitious oscillations that are not visible in the results from DMRG and classical shadows. The two approaches use the same quantum measurement data; the only difference is the classical post-processing. In Fig. 3c we compare the cost of this classical post-processing, finding roughly a 10^4 times speed-up in classical processing time using the classical shadow instead of NNQST.

Subsystem entanglement entropies. An important nonlinear property that can be predicted with classical shadows is subsystem entanglement entropy. The required number of measurements scales exponentially in subsystem size, but is independent of the total number of qubits. Moreover, these measurements can be used to predict many subsystem entanglement entropies at once. This problem has also been studied extensively in ref. [20], where a specialized approach (which we refer to here as the 'Brydges et al. protocol') was designed to efficiently estimate second-order Rényi entanglement entropies using random local measurements. In ref. [20], a random unitary rotation is reused several times. Predictions using classical shadows could also be slightly modified to adapt to this scenario. Results from our numerical experiments are shown in Fig. 4. In Fig. 4a, we predict the entanglement entropy for all subsystems of size ≤ 2 from only 2,500 measurements of the approximate ground state of the disordered Heisenberg model in one dimension. This is a prototypical model for studying many-body localization[22]. The ground state is approximated by a set of singlet states $\{\frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)\}$ found using the strong-disorder renormalization group[23,24]. Both the classical shadow protocol and the Brydges et al. method use random single-qubit rotations and basis measurements to find a classical representation of the quantum state; the only difference between the methods is in the classical post-processing. For these small subsystems, we find that the prediction error of the classical shadow is smaller than the error of the Brydges et al. protocol. In Fig. 4b, we consider predicting the entanglement entropy in a GHZ state for system sizes ranging from $n = 4$ to $n = 10$ qubits. We focus on the entanglement entropy of the subsystem with system size $n/2$ on the left side. Note that this entanglement entropy is equal to one bit for any system size *n*. To achieve an error of 0.05, classical shadows require several times fewer measurements and the discrepancy increases as we require smaller error.

Application to quantum simulation of the lattice Schwinger model (Pauli measurements). Simulations of quantum field theory using quantum computers may someday advance our understanding of fundamental particle physics. Although high-impact discoveries may still be a way off, notable results have already been achieved in studies of 1D lattice gauge theories using quantum platforms.

For example, in ref. [18], a 20-qubit trapped ion analogue quantum simulator was used to prepare low-energy eigenstates of the lattice Schwinger model (1D quantum electrodynamics). The authors prepared a family of quantum states $\{|\psi(\theta)\rangle\}$, where θ is a variational parameter, and computed the variance of the energy $\langle(\hat{H} - \langle\hat{H}\rangle_\theta)^2\rangle_\theta$ for each value of θ. Here, $\hat{H}$ is the Hamiltonian of the model, and $\langle\hat{O}\rangle_\theta = \langle\psi(\theta)|\hat{O}|\psi(\theta)\rangle$ is the expectation value of the operator $\hat{O}$ in the state $|\psi(\theta)\rangle$. Because energy eigenstates, and only energy eigenstates, have vanishing energy dispersion, adjusting θ to minimize the variance of energy prepares an energy eigenstate.

After solving the Gauss law constraint to eliminate the gauge fields, the Hamiltonian $\hat{H}$ of the Schwinger model is 2-local, though not geometrically local in one dimension. Hence the quantity $\langle(\hat{H} - \langle\hat{H}\rangle_\theta)^2\rangle_\theta$ is a sum of expectation values of 4-local observables, which can be measured efficiently using a classical shadow

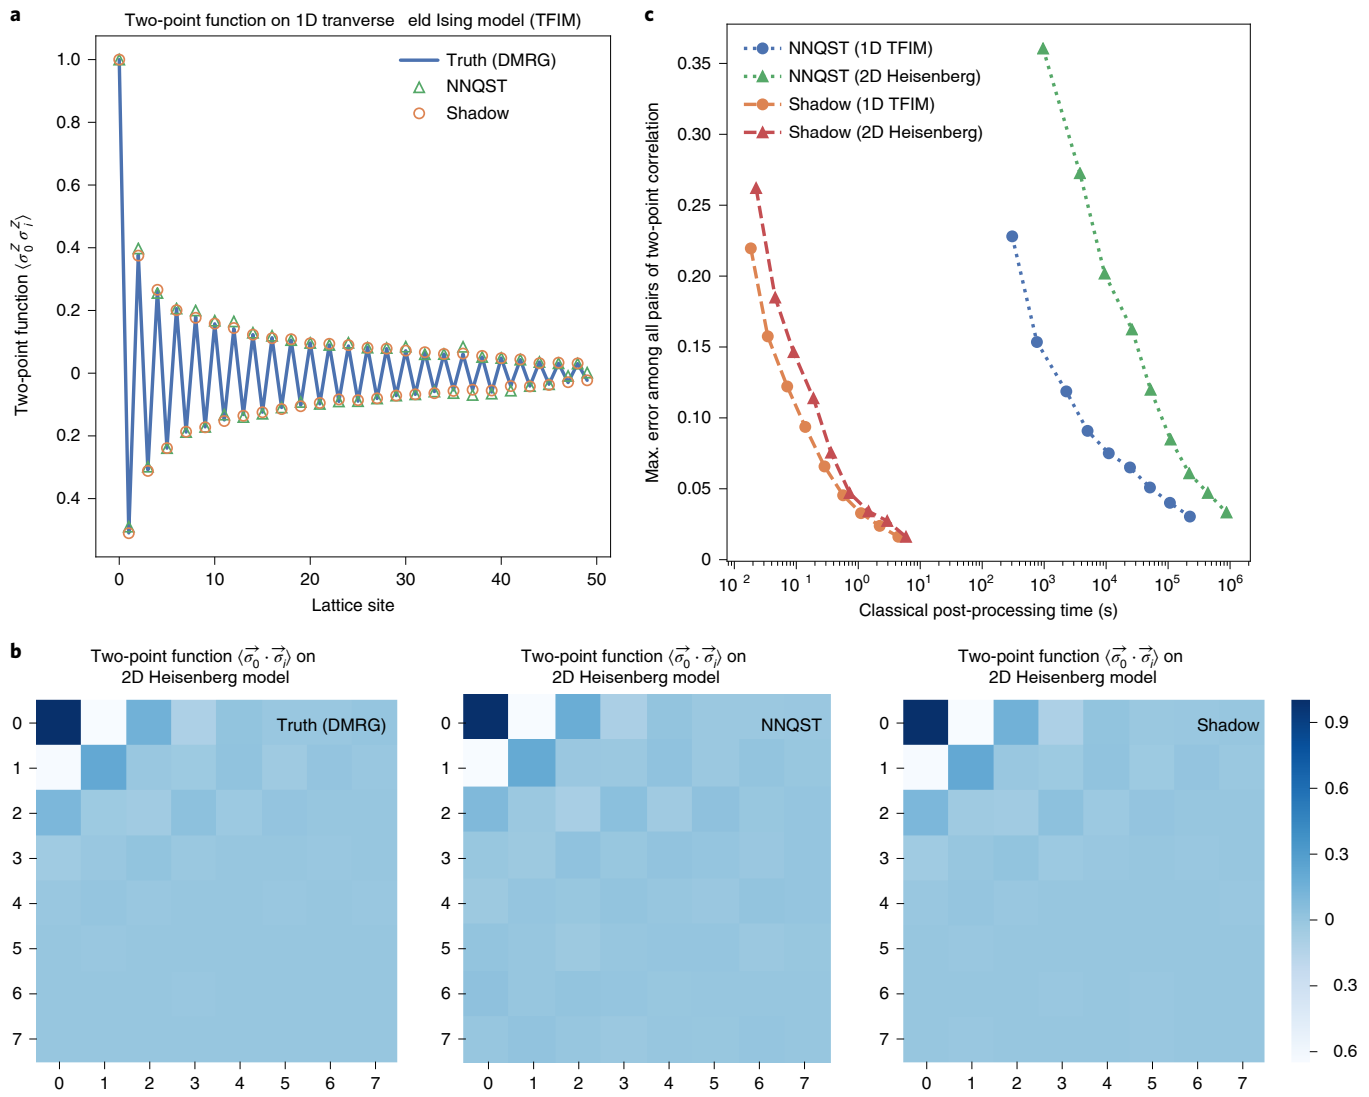

Fig. 3 | Predicting two-point correlation functions using classical shadows (Pauli measurements) and NNQST. a, Predictions of two-point functions $\langle\sigma_0^z\sigma_i^z\rangle$ for ground states of the 1D critical antiferromagnetic TFIM with 50 lattice sites. These are based on 2^{19} random Pauli measurements. **b**, Predictions of two-point functions $\langle\vec{\sigma}_0\cdot\vec{\sigma}_i\rangle$ for the ground state of the 2D antiferromagnetic Heisenberg model with 8×8 lattice sites. The predictions are based on 2^{19} random Pauli measurements. **c**, Classical processing time (CPU time in seconds) and prediction error (the largest among all pairs of two-point correlations) over different numbers of measurements: $\{2^{11}, ..., 2^{19}\}$. The quantum measurement scheme in classical shadows (Pauli) is the same as the POVM-based neural network tomography (NNQST) in ref. [3]. The only difference is the classical post-processing. As the number of measurements increases, the processing time increases, while the prediction error decreases.

derived from random Pauli measurements. This is illustrated in Fig. 5a. Figure 5b compares the performance of classical shadows to the measurement scheme for 4-local observables designed in ref. [18], and also to a recent method[25] for measuring local observables, as well as the standard approach that directly measures all observables independently.

The results show, for the methods we considered, the number of copies of the quantum state needed to measure the expectation value of all 4-local Pauli observables in $\langle(\hat{H}-\langle\hat{H}\rangle_\theta)^2\rangle_\theta$ with an error equivalent to measuring each of these observables at least 100 times. In ref. [18], such a relatively small number of measurements per local observable already yielded results comparable to theoretical predictions based on exact diagonalization. We find that the performance of the classical shadow method is better than the method used in ref. [18] only for system size larger than 50 qubits, and may actually be worse for small system sizes. However, classical shadows provide a good prediction for any set of local observables, while the method of ref. [18] was hand-crafted for the particular task of estimating the variance of the energy in the Schwinger model.

To make a more apt comparison, we constructed a deterministic version of classical shadows, using a fixed set of measurements rather than random Pauli measurements, specifically adapted for the purpose of estimating $\langle(\hat{H}-\langle\hat{H}\rangle_\theta)^2\rangle_\theta$ in the lattice Schwinger model. This deterministic collection of Pauli measurements is obtained by a powerful technique called derandomization[26,27]. This procedure simulates the classical shadow scheme based on randomized measurements and makes use of the rigorous performance bound we developed. When a coin is tossed in the randomized scheme to decide which measurement to perform next, the next measurement in the derandomized version is chosen to have the best possible performance bound for the rest of the protocol. It turns out that this derandomization of the classical shadow method can be carried out very efficiently (full details will appear in upcoming work). Not surprisingly, the derandomized version, also included in Fig. 5,

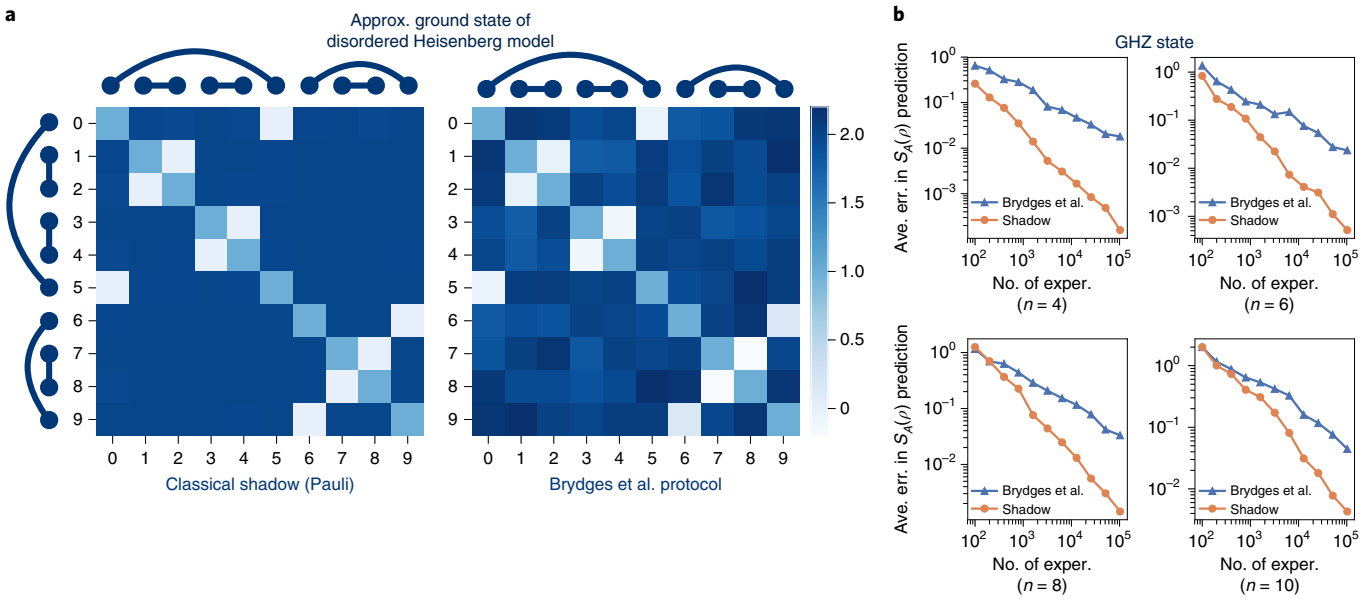

Fig. 4 | Predicting entanglement Rényi entropies using classical shadows (Pauli measurements) and the Brydges et al. protocol. a, Prediction of second-order Rényi entanglement entropy for all subsystems of size at most two in the approximate ground state of a disordered Heisenberg spin chain with 10 sites and open boundary conditions. The classical shadow is constructed from 2,500 quantum measurements. The predicted values using the classical shadow visually match the true values with a maximum prediction error of 0.052. The Brydges et al. protocol[20] results in a maximum prediction error of 0.24. **b**, Comparison of classical shadows and the Brydges et al. protocol[20] for estimating second-order Rényi entanglement entropy in GHZ states. We consider the entanglement entropy of the subsystem with size $n/2$ on the left side.

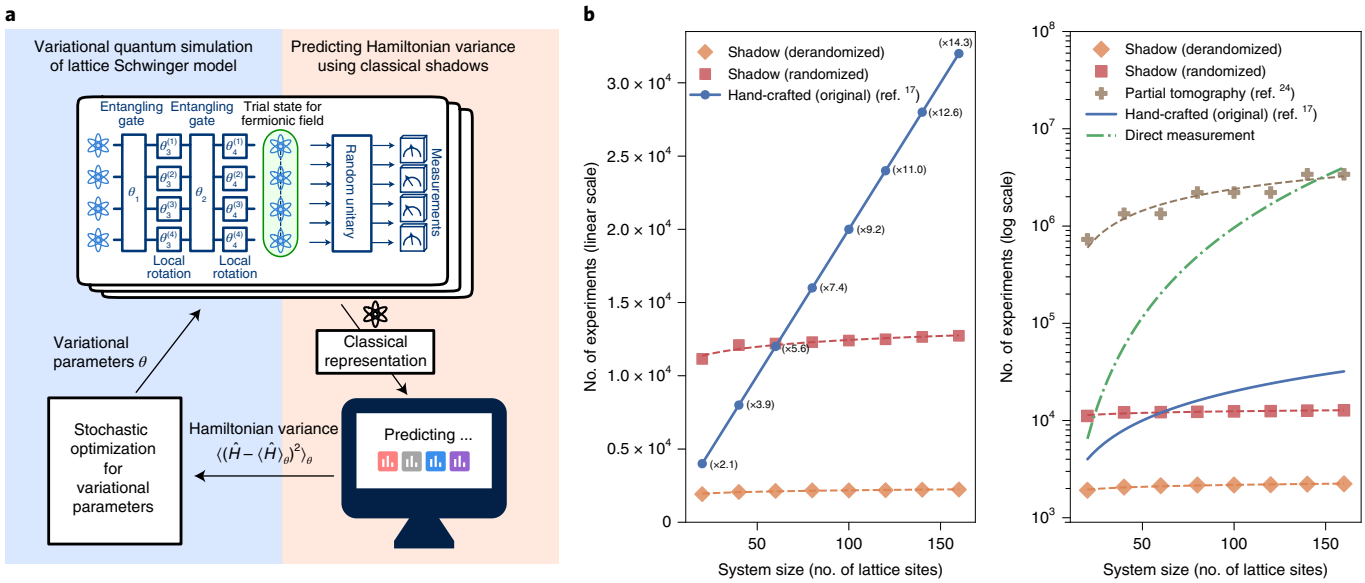

Fig. 5 | Application of classical shadows (Pauli measurements) to variational quantum simulation of the lattice Schwinger model. a, An illustration of variational quantum simulation and the role of classical shadows. **b**, Comparison between different approaches in the number of measurements needed to predict all 4-local Pauli observables in the expansion of $\langle(\hat{H}-\langle\hat{H}\rangle_\theta)^2\rangle_\theta$ with an error equivalent to measuring each Pauli observable at least 100 times. We include a linear-scale plot that compares classical shadows with the original hand-designed measurement scheme in ref. [18] (left) and a log-scale plot that compares with other approaches (right). In the linear-scale plot, ($\times T$) indicates that the original scheme uses T times the number of measurements compared to classical shadows (derandomized).

outperforms the randomized version by a considerable margin. We then find that the derandomized classical shadow method is significantly more efficient than the other methods we considered, including the hand-crafted method from ref. [18]. Finally, we emphasize that the derandomization procedure is fully automated (see https://github.com/momohuang/predicting-quantum-properties for open-source code) and not problem-specific. It could be used for any prespecified set of local observables.

Outlook

A classical shadow is a succinct classical description of a quantum state, which can be extracted by performing reasonably simple

single-copy measurements on a reasonably small number of copies of the state. We have shown that, given its classical shadow, many properties of a quantum state can be accurately and efficiently predicted with a rigorous performance guarantee. In the case of classical shadows based on random Pauli measurements, our methods are feasible using current quantum platforms, and our numerical experiments indicate that many properties can be predicted more efficiently using classical shadows than by using other methods. We therefore anticipate that classical shadows will be useful in near-term experiments characterizing noise in quantum devices and exploring variational quantum algorithms for optimization, materials science and chemistry. Our results also suggest a variety of avenues for further theoretical exploration. Can the classical shadow of a quantum state be updated efficiently as the state undergoes time evolution governed by a local Hamiltonian? Can we use classical shadows to predict properties of quantum channels rather than states? What are the applications of classical shadows based on other ensembles of unitary transformations, for example ensembles of shallow random quantum circuits? More broadly, by mapping many-particle quantum states to succinct classical data, classical shadows open opportunities for applying classical machine-learning methods to numerous challenging problems in quantum many-body physics[4,28,29], such as the classification of quantum phases of matter and the simulation of strongly correlated quantum phenomena.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41567-020-0932-7.

Received: 20 October 2019; Accepted: 6 May 2020;
Published online: 22 June 2020

References

1. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
2. Cramer, M. et al. Efficient quantum state tomography. *Nat. Commun.* **1**, 149 (2010).
3. Carrasquilla, J., Torlai, G., Melko, R. G. & Aolita, L. Reconstructing quantum states with generative models. *Nat. Mach. Intell.* **1**, 155–161 (2019).
4. Torlai, G. et al. Neural-network quantum state tomography. *Nat. Phys.* **14**, 447–450 (2018).
5. Aaronson, S. Shadow tomography of quantum states. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing* (*STOC 2018*) 325–338 (ACM, 2018)
6. 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* (*STOC 2019*) 322–333 (ACM, 2019).
7. Guta, M., Kahn, J., Kueng, R. J. & Tropp, J. A. Fast state tomography with optimal error bounds. *J. Phys. A* **53**, 204001 (2020).
8. Gottesman, D. *Stabilizer Codes and Quantum Error Correction* PhD thesis, Caltech (1997).
9. Fano, R. M. *Transmission of Information*: *A Statistical Theory of Communications* (MIT Press, 1961).
10. Jerrum, M. R., Valiant, L. G. & Vazirani, V. V. Random generation of combinatorial structures from a uniform distribution. *Theoret. Comput. Sci.* **43**, 169–188 (1986).
11. Nemirovsky, A. S. & Yudin, D. B. *Problem Complexity and Method Efficiency in Optimization* (Wiley-Interscience, 1983).
12. Greenberger, D. M., Horne, M. A. & Zeilinger, A. in *Bell's Theorem, Quantum Theory and Conceptions of the Universe. Fundamental Theories of Physics* Vol. 37 (ed. Kafatos, M.) 69–72 (Springer, 1989).
13. Dennis, E., Kitaev, A., Landahl, A. & Preskill, J. Topological quantum memory.*J. Math. Phys.* **43**, 4452–4505 (2002).
14. Flammia, S. T. & Liu, Y.-K. Direct fidelity estimation from few Pauli measurements. *Phys. Rev. Lett.* **106**, 230501 (2011).
15. Gühne, O. & Tóth, G. Entanglement detection. *Phys. Rep.* **474**, 1–75 (2009).
16. Weilenmann, M., Dive, B., Trillo, D., Aguilar, E. A. & Navascués, M. Entanglement detection beyond measuring fidelities. *Phys. Rev. Lett.* **124**, 200502 (2020).
17. Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017).
18. Kokail, C. et al. Self-verifying variational quantum simulation of lattice models. *Nature* **569**, 355–360 (2019).
19. Hoeffding, W. in *Breakthroughs in Statistics* 308–334 (Springer, 1992).
20. Brydges, T. et al. Probing Rényi entanglement entropy via randomized measurements. *Science* **364**, 260–263 (2019).
21. Renes, J. M., Blume-Kohout, R., Scott, A. J. & Caves, C. M. Symmetric informationally complete quantum measurements. *J. Math. Phys.* **45**, 2171–2180 (2004).
22. Nandkishore, R. & Huse, D. A. Many-body localization and thermalization in quantum statistical mechanics. *Annu. Rev. Condens. Matter Phys.* **6**, 15–38 (2015).
23. Dasgupta, C. & Ma, S.-k Low-temperature properties of the random Heisenberg antiferromagnetic chain. *Phys. Rev. B* **22**, 1305 (1980).
24. Ma, S.-k, Dasgupta, C. & Hu, C.-k Random antiferromagnetic chain. *Phys. Rev. Lett.* **43**, 1434 (1979).
25. Bonet-Monroig, X., Babbush, R. & O'Brien, T. E. Nearly optimal measurement scheduling for partial tomography of quantum states. Preprint at https://arxiv.org/pdf/1908.05628.pdf (2019).
26. Raghavan, P. Probabilistic construction of deterministic algorithms: approximating packing integer programs.*J. Comput. Syst. Sci.* **37**, 130–143 (1988).
27. Spencer, J. Ten lectures on the probabilistic method. In *CBMS-NSF Regional Conference Series in Applied Mathematics* 2nd edn, Vol. 64 (SIAM, 1994).
28. Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).
29. Carrasquilla, J. & Melko, R. G. Machine learning phases of matter. *Nat. Phys.* **13**, 431–434 (2017).
30. Paini, M. & Kalev, A. An approximate description of quantum states. Preprint at https://arxiv.org/pdf/1910.10543.pdf (2019).

Data availability

Source data are available for this paper. All other data that support the plots within this paper and other findings of this study are available from the corresponding author upon reasonable request.

Code availability

Source code for an efficient implementation of the proposed procedure is available at https://github.com/momohuang/predicting-quantum-properties.

Acknowledgements

We thank V. Albert, F. Brandão, M. Endres, I. Roth, J. Tropp, T. Vidick, M. Weilenmann and J. Wright for valuable input and inspiring discussions. L. Aolita and G. Carleo provided helpful advice regarding presentation. Our gratitude extends, in particular, to J. Iverson, who helped us in devising a numerical sampling strategy for toric code ground states. We also thank M. Paini and A. Kalev for informing us about their related work[30], where they discussed succinct classical 'snapshots' of quantum states obtained from randomized local measurements. H.-Y.H. is supported by the Kortschak Scholars Program. R.K. acknowledges funding provided by the Office of Naval Research (award no. N00014-17-1-2146) and the Army Research Office (award no. W911NF121054). J.P. acknowledges funding from ARO-LPS, NSF and DOE. The Institute for Quantum Information and Matter is an NSF Physics Frontiers Center.

Author contributions

H.-Y.H. and R.K. developed the theoretical aspects of this work. H.-Y.H. conducted the numerical experiments and wrote the open-source code. J.P. conceived the applications of classical shadows. H.-Y.H., R.K. and J.P. wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at https://doi.org/10.1038/s41567-020-0932-7.

Correspondence and requests for materials should be addressed to H.-Y.H.

Peer review information *Nature Physics* thanks Yi-Kai Liu and other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.