nature reviews physics

https://doi.org/10.1038/s42254-022-00535-2

Review article

Check for updates

The randomized measurement toolbox

Andreas Elben[1,2,3,4], Steven T. Flammia[1,5], Hsin-Yuan Huang[1,6], Richard Kueng[7], John Preskill[1,2,5,6], Benoît Vermersch[3,4,8] & Peter Zoller[3,4]

Abstract

Programmable quantum simulators and quantum computers are opening unprecedented opportunities for exploring and exploiting the properties of highly entangled complex quantum systems. The complexity of large quantum systems is the source of computational power but also makes them difficult to control precisely or characterize accurately using measured classical data. We review protocols for probing the properties of complex many-qubit systems using measurement schemes that are practical using today's quantum platforms. In these protocols, a quantum state is repeatedly prepared and measured in a randomly chosen basis; then a classical computer processes the measurement outcomes to estimate the desired property. The randomization of the measurement procedure has distinct advantages. For example, a single data set can be used multiple times to pursue a variety of applications, and imperfections in the measurements are mapped to a simplified noise model that can more easily be mitigated. We discuss a range of cases that have already been realized in quantum devices, including Hamiltonian simulation tasks, probes of quantum chaos, measurements of non-local order parameters, and comparison of quantum states produced in distantly separated laboratories. By providing a workable method for translating a complex quantum state into a succinct classical representation that preserves a rich variety of relevant physical properties, the randomized measurement toolbox strengthens our ability to grasp and control the quantum world.

Sections

[1]Institute for Quantum Information and Matter, Caltech, Pasadena, CA, USA. [2]Walter Burke Institute for Theoretical Physics, Caltech, Pasadena, CA, USA. [3]Institute for Theoretical Physics, University of Innsbruck, Innsbruck, Austria. [4]Institute for Quantum Optics and Quantum Information of the Austrian Academy of Sciences, Innsbruck, Austria. [5]AWS Center for Quantum Computing, Pasadena, CA, USA. [6]Department of Computing and Mathematical Sciences, Caltech, Pasadena, CA, USA. [7]Institute for Integrated Circuits, Johannes Kepler University Linz, Linz, Austria. [8]Université Grenoble Alpes, CNRS, LPMMC, Grenoble, France. e-mail: peter.zoller@uibk.ac.at

Key points

- Increasingly sophisticated quantum simulators and quantum computers are becoming available, but they are difficult to characterize accurately using classically measured data.
- Randomized measurements provide a feasible procedure for converting a many-qubit quantum state to succinct classical data that can later be processed to estimate many properties of interest with rigorous guarantees.
- Randomized measurements are readily implemented in noisy intermediate-scale quantum devices by repeatedly preparing and measuring a quantum state in a randomly selected basis.
- Many applications of randomized measurements have been conceived and experimentally demonstrated, including Hamiltonian simulation tasks, probes of quantum chaos, measurements of non-local order parameters, and comparison of quantum states produced in distantly separated laboratories.
- Experimental imperfections in performing randomized measurements can often be easily mitigated; a wide range of different physical platforms realizing qubits, bosonic and fermionic quantum many-body systems is accessible.
- Viewed as a powerful quantum-to-classical converter, randomized measurements enable the use of classical algorithms to learn and predict properties of quantum systems that may never have been realized before.

Introduction

As far as we know, it is not possible to use classical data to fully and succinctly characterize generic quantum systems of many strongly interacting particles. This observation is both a curse and a blessing. On the one hand, it limits the ability of classical beings like us to grasp the behaviour of complex highly entangled quantum systems. On the other hand, it invites us to build and operate large-scale quantum devices.

The emergence of increasingly powerful quantum technologies has transformed the challenge of characterizing complex quantum systems from a theoretical conundrum to a laboratory imperative. Someday there will be large-scale error-corrected quantum computers to advance the frontiers of science and run useful practical applications. Although such machines may lie far in the future, even today programmable quantum platforms[1,2] can create and control complex states comprising many atoms[3–8], spins[9,10], photons[11,12] or superconducting circuit elements[13], opening unprecedented opportunities for scientific discovery[14].

Experimentalists and theorists working together must develop, perfect and use suitable tools to investigate and exploit the features of many-qubit quantum states that are created in the laboratory. This typically involves preparing and measuring the same quantum state over and over again. With enough repetitions, it is possible to completely characterize an N-qubit state by means of full-state tomography, but this task is hopelessly inefficient, requiring a number of experiments exponential in N (refs.[15–17]) and an amount of classical post-processing of the experimental results that is also exponentially large. Fortunately, a far less complete description of the state is adequate for many purposes[18–20], so that the number of experiments and the amount of classical processing needed can be greatly reduced. Here, we review recent theoretical ideas about how to improve the efficiency of characterizing complex quantum states, and some of the experimental results that flow from these ideas.

The concepts and examples that we discuss in this Review share a common theme. Rather than tailoring the measurements performed in the laboratory to the particular properties one wishes to study, one can instead repeatedly perform measurements that are randomly sampled from a fixed ensemble, and then adapt the classical post-processing of the measurement outcomes to the particular task at hand[21–28]. This randomized measurement (RM) strategy can be surprisingly powerful even when the measurements are simple enough to be performed with adequate precision using today's noisy quantum platforms. A particularly simple procedure is to measure each qubit in a randomly chosen basis. By repeating this procedure of order $\log(L)$ times, and using only efficient classical post-processing, one can accurately estimate the expectation values of any L local operators — the number of experiments needed does not depend at all on the total number of qubits[24]. Randomized single-qubit measurements also enable one to estimate the properties of larger subsystems[22]; in this case, the cost rises exponentially with the size of the subsystem but is still far lower than the cost of complete tomography of the subsystem[22–24]. Alternatively, the global properties of the quantum state can be estimated by using a modest number of measurement repetitions if the measurements are preceded by relatively efficient information scrambling unitary operations executed with a quantum computer or programmable quantum simulator[24,26]. A further advantage of the RM approach is that randomization simplifies the effects of noise, so that imperfections in measurement outcomes can be more easily mitigated by suitably modifying the classical post-processing of the outcomes[29,30].

Many applications of this RM toolbox have already been conceived and executed in experiments using quantum devices. For example, one can estimate the overlap of two quantum states produced in separate laboratories far apart from one another[31,32]. It is possible to probe chaotic quantum dynamics by measuring out-of-time-order correlation functions[33,34], without reversing time evolution or introducing ancilla systems. One can quantify quantum entanglement by measuring entropy[35] and other entanglement measures[36]. It is possible to compute order parameters that characterize topological order or symmetry-protected topological order[37] and estimate the expectation value and variance of a local Hamiltonian[24].

These and other applications have a notable feature. First, repeated RMs map a multiqubit quantum state to succinct classical data. Later, these classical data are processed to investigate properties of interest. Conveniently, the properties to be investigated need not be known when the measurements are performed using the quantum device; rather, one can:

Measure first, ask questions later.

Indeed, some of the applications reviewed here were carried out by reanalysing data that had originally been taken with a different purpose in mind.

An RM protocol may be viewed as a feasible scheme for translating the extravagant quantum information residing in a many-qubit state into a succinct classical representation of the state. This quantum-to-classical conversion process unavoidably discards a vast amount of

information about the state, but the applications mentioned above illustrate that many physically relevant features of the state can survive. Thus, scientists assisted by their powerful classical computers, by pondering and manipulating classical data, can grasp crucial properties of the quantum world that might otherwise remain concealed.

Experimental recipe and post-processing of the measurements

We start by reviewing a particular RM scheme using repeated single-qubit measurements to construct a 'classical shadow' of a quantum state, state a rigorous guarantee on the accuracy of estimated operator expectation values based on classical shadows, and describe an application of randomized single-qubit measurements to estimating the purity of a many-qubit state.

Data acquisition protocol

Consider a quantum system consisting of N qubits with associated Hilbert space $(\mathbb{C}^2)^{\otimes N}$. An RM (see Fig. 1) consists of the following steps. First, the quantum many-body state ρ of interest is prepared in the device. Second, a unitary operation U, selected at random from a suitable ensemble of unitary operations, is applied to ρ. For concreteness, we consider here local random operations $U = \bigotimes_{n=1}^{N} U_n$, applied to each qubit independently. The individual single-qubit rotations U_n ($n = 1,\ldots, N$) are sampled from ensembles of single-qubit unitary operations that evenly cover the Bloch sphere of each qubit. Examples of such unitary designs[38,39] include the single-qubit Clifford group and the full unitary group $U(2)$ encompassing all single-qubit transformations. We note that more general choices of multiqubit random operations are possible and are discussed in the final section. Third, a projective measurement in the computational basis $\{|\mathbf{s}\rangle\}$ is performed, with outcome bit string $\mathbf{s} = (s_1,\ldots, s_N)$ and $s_n \in \{0, 1\}$ for $n = 1,\ldots, N$. These three steps are then repeated K times with a fixed unitary U. Subsequently, the entire procedure is repeated with M independently sampled unitaries U such that in total MK experimental runs are performed.

In summary, MK experimental runs are executed, each of which is characterized by N single-qubit unitaries $U_1^{(m)}, \ldots, U_N^{(m)}$ that only depend on m, and an N-bit outcome $\mathbf{s}^{(m,k)} = (s_1^{(m,k)}, \ldots, s_N^{(m,k)})$ that depends on both $m = 1,\ldots, M$ and $k = 1,\ldots, K$. Storing both up to floating point accuracy is comparatively cheap, requiring storage of $\mathcal{O}(MKN)$ floating point numbers in total.

Post-processing protocol

After completing a full experiment, one can use the obtained data to extract information about the underlying many-body system. It is instructive to consider two extreme examples. For the sake of simplicity, we also assume that the single-qubit rotations U_i are sampled from the single-qubit Clifford group, the discrete ensemble that randomly permutes Pauli matrices $\{X, Y, Z\}$. That is, $U_n Z U_n^\dagger = W_n \in \{X, Y, Z\}$ for $1 \le n \le N$.

Setting $M = 1$ means that one repeats the same RM over and over ($K > 1$ times). This is equivalent to measuring a random string of Pauli observables, namely $U_1^\dagger Z U_1 \otimes \cdots \otimes U_n^\dagger Z U_N = W_1 \otimes \cdots \otimes W_N$, a total of K times. This, in turn, allows one to approximate the expectation value $\mathrm{tr}(W_1 \otimes \cdots \otimes W_N \rho)$ and compatible subsystem marginals, for example, $\mathrm{tr}(W_1 \otimes \cdots \otimes W_l \otimes \mathbb{I}^{\otimes(N-l)}\rho)$ for $1 \le l \le N$ (see Fig. 2). Expectation values of other Pauli observables that do not commute with $W_1 \otimes \cdots \otimes W_N$ are off limits, though.

The other extreme case drives us into more interesting territory. One samples a total of M random Pauli strings $W_1^{(m)} \otimes \cdots \otimes W_N^{(m)}, 1 \le m \le M$ and measures each of them exactly once, $K = 1$. A single measurement outcome does not enable the reliable approximation of any of the original Pauli expectation values. But one can combine samples across different measurements to predict many (subsystem) expectation values. Take $X \otimes Y \otimes Z \otimes \mathbb{I}^{\otimes(N-3)}$ as a concrete example. The outcome from measuring $W_1^{(m)} \otimes \cdots \otimes W_N^{(m)}$ provides useful statistical information if and only if $W_1^{(m)} = X, W_2^{(m)} = Y$ and $W_3^{(m)} = Z$. If one assigns all single-qubit unitaries U_n uniformly at random, these accordances occur with probability $(1/3)^3$. In turn, one can expect that a total of $M \ge 3^3/\varepsilon^2$ randomly selected N-qubit Pauli measurements provides enough statistical data to approximate $\mathrm{tr}(X \otimes Y \otimes Z \otimes \mathbb{I}^{\otimes(n-3)}\rho)$ up to an error ε. Interestingly, this argument only depends on the size w of the subsystem where the Pauli strings of interest act non-trivially ($w = 3$ for our example). The actual qubit locations and Pauli strings of interest ($X \otimes Y \otimes Z \otimes \mathbb{I}^{\otimes(N-3)}$ in our example) do not matter at all. Note that because the Pauli operators form a basis (unnormalized) for the w-qubit subsystem, one can estimate any observable on the w-qubit marginal in this fashion. In fact, as summarized in Theorem 1 below, one can extend this argument to cover (very) many size-w expectation values in one go.

The actual prediction step is also relatively straightforward. We restrict our attention to measurement settings that are compatible with the Pauli expectation value $o = \mathrm{tr}(O\rho)$ in question and use empirical averaging of compatible outcomes to obtain an approximation $\hat{o}$ of o. The following formula succinctly summarizes such an estimation process:

$$\hat{o} = \frac{1}{M}\sum_{m=1}^{M} \mathrm{tr}(O\hat{\rho}^{(m)}) \quad \text{where} \tag{1}$$

$$\hat{\rho}^{(m)} = \frac{1}{K}\sum_{k=1}^{K}\bigotimes_{n=1}^{N}\left(3(U_n^{(m)})^\dagger |s_n^{(m,k)}\rangle\langle s_n^{(m,k)}| U_n^{(m)} - \mathbb{I}\right) \tag{2}$$

combines the mth measurement settings $(U_1^{(m)}, \ldots, U_N^{(m)})$ with the K associated outcomes $(s_1^{(m,k)}, \ldots, s_N^{(m,k)})$ to produce an approximation of the underlying N-qubit quantum state ρ. The collection $\{\hat{\rho}^{(m)}\}_{m=1,\ldots,M}$

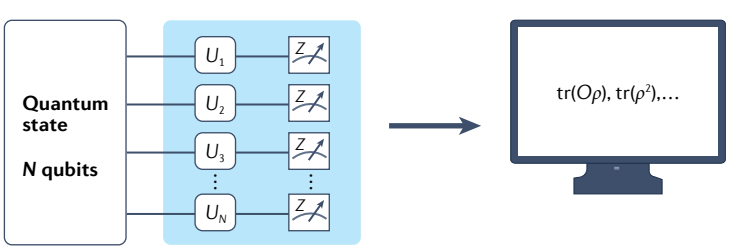

Fig. 1 | Randomized measurements. Randomized measurements are made on a N-qubit quantum state ρ via the application of local random unitaries $U = \bigotimes_{n=1}^{N} U_n$ and subsequent projective measurements performed in the computational Z-basis. Through classical post-processing of the outcomes, many properties of ρ, such as observable expectation values $\mathrm{tr}(O\rho)$ and the purity $\mathrm{tr}(\rho^2)$, can be estimated. Remarkably, this estimation is provably efficient if we restrict attention to properties of (arbitrary) subsystems of constant size.

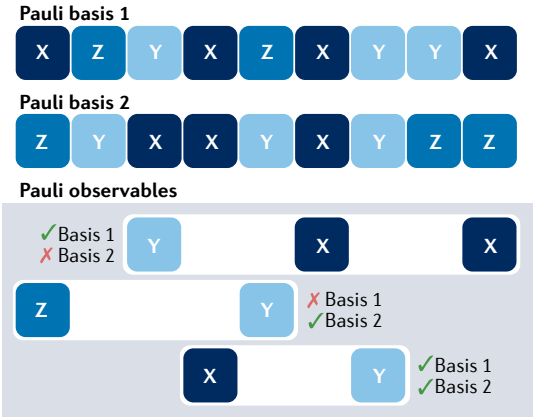

Fig. 2 | Pauli observables. A Pauli string denotes the basis one measures in for each qubit. A Pauli observable O given by a tensor product of $\{\mathbb{I}, X, Y, Z\}$, where $\mathbb{I}$ is the identity and X,Y,Z are the Pauli matrices is compatible with the string if the non-identity part of O matches the string.

is called a classical shadow of ρ (refs.[24,25]). This procedure works for arbitrary target observables O (not just Pauli expectation values) and ensembles of single-qubit random unitaries that cover the Bloch sphere evenly (not just Clifford unitaries).

Remarkably, RMs give access not only to observables but also to polynomial functionals of the density matrix. In fact, RMs were first envisioned to estimate the purity $P_2 = \mathrm{tr}(\rho^2)$ by means of the following formula[21–23,35]

$$\hat{P}_2 = \frac{2^N}{MK(K-1)} \sum_{m=1}^{M} \sum_{\substack{k,k'=1 \\ k\neq k'}}^{K} (-2)^{-D[\mathbf{s}^{(m,k)},\mathbf{s}^{(m,k')}]}. \tag{3}$$

One first calculates a weighted average of Hamming distances $D[\mathbf{s}^{(m,k)}, \mathbf{s}^{(m,k')}]$ of two distinct outcomes $\mathbf{s}^{(m,k)}$ and $\mathbf{s}^{(m,k')}$ that belong to the same measurement setting m. Subsequently, one averages over different measurement settings. Although the precise form of equation (3) follows from derivations elsewhere[23,35], some intuition might be gained as follows: one observes bitstrings $\mathbf{s}^{(m,k)}$ according to their Born probabilities $P_U(\mathbf{s}^{(m,k)}) = |\langle \mathbf{s}^{(m,k)}|(U^{(m)})^\dagger \rho U^{(m)}|\mathbf{s}^{(m,k)}\rangle|^2$. In equation (3), the second-order correlations of (estimations of) these Born probabilities are averaged over local random unitaries. Such an average must correspond to a second-order functional of ρ which is invariant under local random unitary transformations, that is, in this case, the purity.

Alternatively, one can estimate the purity by replacing distinct copies of ρ by distinct classical shadows (equation (2)) and average over all possible choices[24,36]:

$$\hat{P}_2 = \frac{1}{M(M-1)} \sum_{m\neq m'} \mathrm{tr}(\hat{\rho}^{(m)}\hat{\rho}^{(m')}). \tag{4}$$

This estimation procedure also extends to arbitrary polynomials of the density matrix[24,36,40,41].

With both estimators, one can access purities $\mathrm{tr}(\rho_A^2)$ of reduced density matrices $\rho_A = \mathrm{tr}_{A^c}(\rho)$ of arbitrary subsystems A (with complement A^c) by restriction during the post-processing. For a fixed total number of experimental runs MK, equation (3) achieves a more accurate estimate for many repetitions K of a few measurement settings M, whereas equation (4) performs better for many different measurement settings with few repetitions each. In addition, the estimator in equation (3) is expected to be more robust against miscalibration of the random unitaries (gate-independent unitary errors) than equation (4) because the estimator in equation (3) depends solely on the measured bitstrings, and matrix elements of the applied random unitaries do not appear explicitly.

Rigorous theory and history

The post-processing rules introduced in equations (1) and (2) can be equipped with rigorous error bounds. Here, we present an exemplary performance guarantee that is valid for evenly distributed ensembles of single-qubit unitaries, such as the full unitary group $U(2)$ or the Clifford group.

Theorem 1. $M \propto \log(L)4^w/\varepsilon^2$ independent randomized measurements suffice to ε-approximate an entire collection of L subsystem-size-w expectation values with high success probability.

For the special case of Pauli expectation values, an improved scaling of $M \propto \log(L)3^w/\varepsilon^2$ RMs readily follows from the arguments provided above. Historically, this result for Pauli operators pre-dates Theorem 1 and was first proven (Appendix D of ref.[42]) by a slightly different, yet equally simple, argument to the one presented here. This result in turn was influenced by the earlier but quadratically weaker bound in ref.[43]. The general case displayed in Theorem 1 is based on the arguments presented in ref.[24]. The actual error bound implicitly works in the single-shot limit ($K = 1$). But multiple repetitions for each measurement setting ($K > 1$) can only further improve performance.

Theorem 1 contains an interesting tradeoff between subsystem size (which enters exponentially) and the number of observables (which enters logarithmically). For instance, $M \propto \log(N)/\varepsilon^2$ randomized N-qubit Pauli measurements already suffice to ε-approximate all two-body Pauli expectation values in a system with N qubits. And, remarkably, the statement is completely independent of the underlying quantum state ρ. Similarly, for the measurement of the purity of (reduced) density matrices of (sub)systems with size $w \leq N$, for estimators in both equations (3) and (4), the required number of experimental runs MK to obtain a given accuracy ε scales only (exponentially) with subsystem size w (refs.[24,36,44]). Asymptotic scalings $M \propto 2^w/\varepsilon^2$ have been rigorously shown for equation (4)[24,36] and also observed for equation (3)[23,35,44].

Theorem 1 showcases that it can be much easier (and more reliable) to accurately approximate certain properties of an unknown state than to estimate the full state ρ. This related problem, called quantum state tomography, has a long and prominent history[45]. Fundamental lower bounds assert that ε-accurate quantum state tomography of an N-qubit system must require exponentially many samples in general ($MK \geq 4^N/\varepsilon^2$; see refs.[17,46–48]). Substantial improvements are only possible if the state in question is known to have (very) advantageous structure, for example matrix product states with polynomial bond dimension[49] or neural network states[50].

The idea of bypassing quantum state tomography, that is the full reconstruction of the quantum state ρ, and directly predicting (subsystem) expectation values $\mathrm{tr}(O\rho)$ is also known as shadow estimation[18–20]. In its original form, shadow estimation does not have an exponential dependence on subsystem size w, but does require loading multiple state copies into a quantum memory and performing entangling quantum computations on them. The procedure discussed here can be viewed as a more near-term variant of this idea. But it also draws inspiration from refs.[21,51] and resource-efficient approaches to quantum state tomography[52–54].

Purity measurements in a trapped-ion quantum simulator

To give an example of application, let us take ref.[35] as an illustrative example of RMs. In this study, the goal was to measure the purity in a trapped-ion quantum simulator, which is relevant for checking that a quantum device works as intended, meaning that the realized quantum state is pure. Subsystem purities can be used to quantify entanglement within quantum many-body systems in terms of the second Rényi entropy[55] $S_2(\rho_A) = -\log_2(P_2)$, where A denotes the subsystem. Trapped-ion quantum simulators contain an array of ions, $N = 10$ or $N = 20$ in this case, and each encodes a qubit using two long-lived electronic states. These can then be manipulated using focused laser beams. The system was propagated from an initial Néel state $|\psi\rangle = |01\rangle^{\otimes(N/2)}$ to an entangled state $|\psi(t)\rangle = e^{-iH_{XY}t}|\psi\rangle$ using a Hamiltonian $H_{XY} = \sum_{i<j} J_{ij}(\sigma_i^+\sigma_j^- + \text{h.c.})$, with $J_{ij} \approx J/|i-j|^\alpha$, $0 < \alpha < 3$, and σ_i^+ (σ_i^-) spin-1/2 raising (lowering) operators. Owing to preparation errors, dephasing, spontaneous emission and others, the system after the evolution time t is described by a density matrix $\rho(t)$. RMs were implemented by sampling individual single-qubit rotations from the circular unitary ensemble and decomposing them into rotations along the z and x axes. This process is illustrated in Fig. 3a. Importantly, it only requires local Z-rotations, while the ions were rotated along the X-axis via a global beam. The total data acquisition involved $M = 500$ RM settings with $K = 150$ single-shot repetitions each. Post-processing was based on the averaged purity formula in equation (3) which is designed to process many repetitions per measurement setting ($K \gg 1$).

The plot in Fig. 3b highlights that this RM protocol faithfully estimates second Rényi entropies for various different subsystems A and evolution times t. In particular, we observe that the second Rényi entropy of the total system ($i = 10$) remains almost constant over time at a small value ~0.4 (corresponding to a large purity $P_2 \approx 0.8$). This shows that the state is slightly affected by preparation and measurement errors, but the dynamics is almost perfectly unitary. Considering subsystems, the second Rényi entropy increases as a function of time and becomes larger than the entropy of the total system, a conclusive signature of quantum entanglement. Such entanglement growth has also been measured recently with superconducting qubits using RMs[56].

The experimental data of ref.[35] have been recently reanalysed[31,36,40,57] to access other entanglement properties, in particular using the classical shadow framework[36].

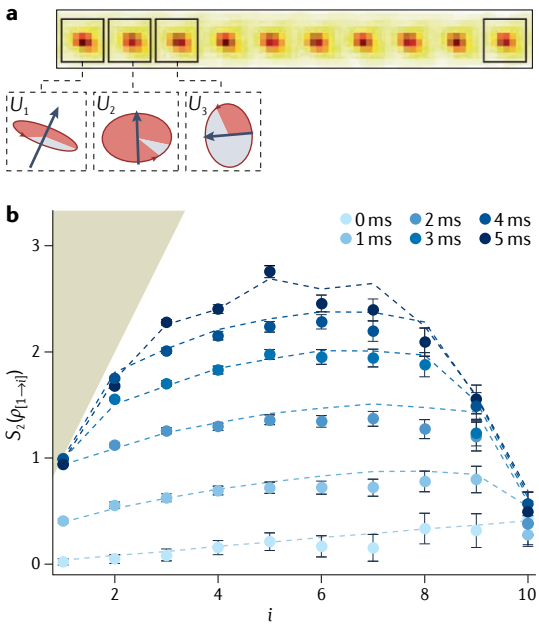

Fig. 3 | Estimation of the second Rényi entropy with randomized measurements. a, Randomized measurements (RMs) were implemented[35] in a 10-qubit trapped-ion system using single-qubit random unitaries sampled uniformly from the full unitary group $U(2)$ followed by a computational basis measurement. In this case, the unitaries $U_n = R_n^X(-\pi/2)R_n^Z(\alpha_n)R_n^X(\pi/2)R_n^Z(\beta_n)$ can be decomposed into uniform X-rotations $R_n^X(\pm\pi/2) = \exp(\mp iX_n\pi/4)$ and local Z-rotations $R_n^Z(\alpha_n) = \exp(-iZ_n\alpha_n/2)$ with rotation angles α_n, β_n (illustrated in the figure for $n = 1,2,3$). **b**, Experimental results[35] for the second Rényi entropy $S_2(\rho_A) = -\log_2(P_2)$ of partitions $A = [1,2,\ldots,i]$ ($i = 1,\ldots,10$) in a system with $N = 10$ ions in total. Coloured dots correspond to different evolution times $t = 0,\ldots,5$ ms, with error bars denoting the standard error of the mean. Dotted lines display results of numerical simulations including decoherence effects. Maximally mixed states with maximal Rényi entropy lie on the boundary of the shaded area. Panels **a**,**b** are adapted with permission from ref.[35], AAAS.

Applications of randomized measurements

We now turn to the numerous applications of RMs. As described in the introduction, these applications span many areas, including probing quantum many-body physics, quantum simulation, noise diagnostics of quantum systems, machine learning (ML) of properties of quantum systems, variational quantum algorithms and quantum computation with noisy intermediate-scale quantum (NISQ) devices, and more. Note that many noise diagnostics models also use randomness to reliably extract figures of merit that capture the execution error of a quantum evolution. Techniques such as randomized benchmarking actually predate the RM toolbox, but are different in scope and purpose. We discuss them in the Supplementary Information.

Characterization of topological order

Topological quantum phases of matter are exotic phases of matter characterized by global correlations[58]. There is increasing interest in realizing topological quantum phases in synthetic quantum devices in the context of quantum simulation, and in topological quantum computing. However, by their very definition, topological phases cannot be detected by local measurements. Thus, their identification and characterization in quantum simulation experiments pose a substantial challenge. RM protocols have been proposed as an experimental tool to address this question, to detect and classify topological quantum phases.

First, RM protocols have been designed to measure many-body topological invariants of symmetry-protected topological (SPT) phases[37]. These invariants are highly non-local and/or nonlinear correlators of the many-body wavefunction that unambiguously identify SPT phases[59]. In the following, we discuss the example of the reflection invariant $\mathcal{Z}_R$ which detects topological phases that emerge from the existence of a spatial reflection symmetry. As depicted schematically in Fig. 4a, $\mathcal{Z}_R$ is given as $\mathcal{Z}_R = \text{tr}(\mathcal{R}_I\rho_I)$ where $\mathcal{R}_I$ is the 'partial' reflection operator acting as $\mathcal{R}_I\ |s_1, \ldots, s_{2n}\rangle = |s_{2n}, s_{2n-1}, \ldots, s_2, s_1\rangle \equiv |\mathcal{R}_I(\mathbf{s})\rangle$ on the subsystem I containing $2n$ spins symmetrically distributed across the central bond. If n is large compared with the correlation length of the system, the quantity $\widetilde{\mathcal{Z}}_R = \mathcal{Z}_R/\sqrt{[P_2(\rho_{I_1}) + P_2(\rho_{I_2})]/2}$ acts as a topological order parameter, taking a quantized value ±1 depending on whether the phase is trivial or topological. (Here $I = I_1 \cup I_2$, where I_1 denotes the n spins just to the left of the central bond, and I_2 denotes the n spins just to the right, and $P_2(\rho_{I_{1,2}}) = \text{tr}(\rho_{I_{1,2}}^2)$ are the purities of the corresponding reduced density matrices).

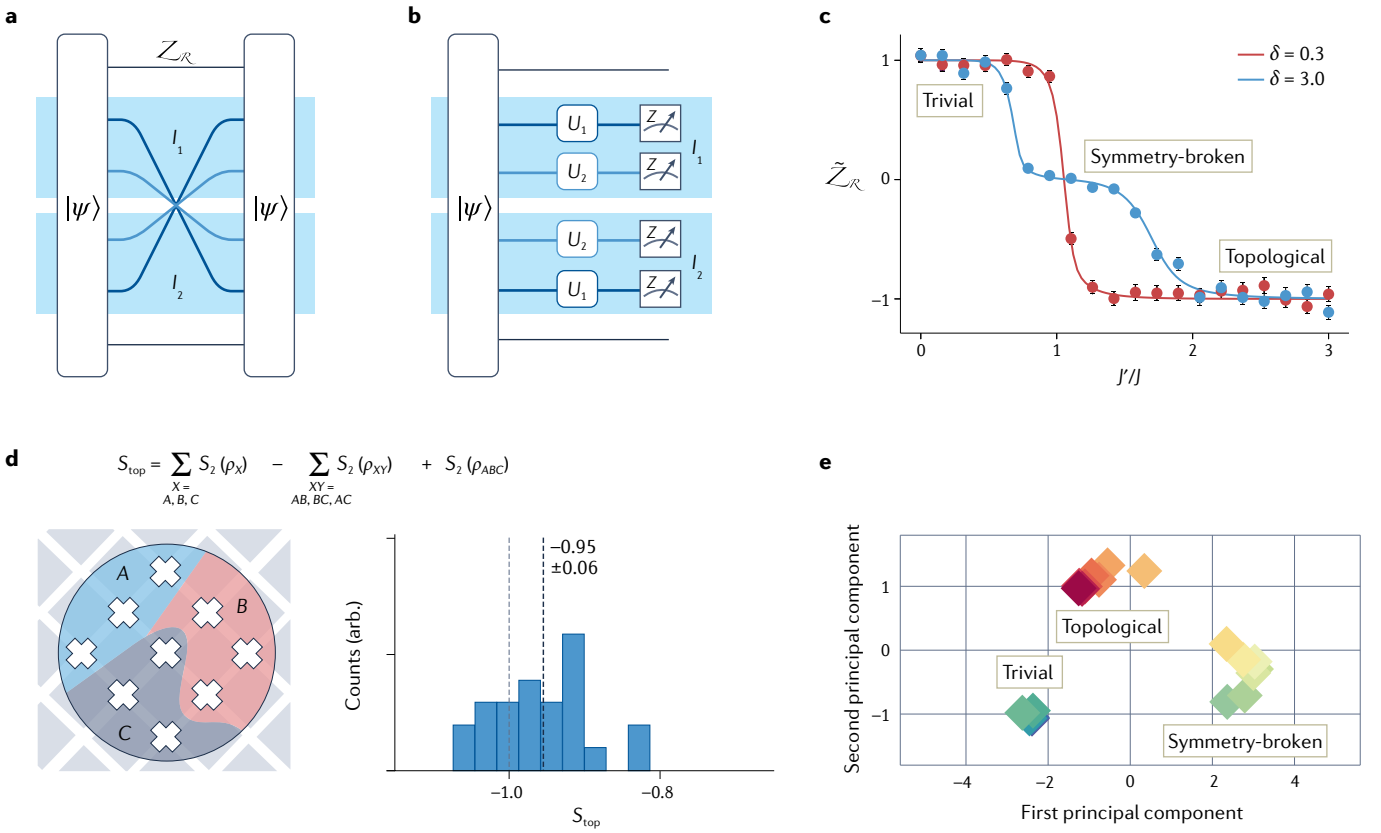

Fig. 4 | Detecting topological order with randomized measurements. a,**b**, The many-body topological invariant $\mathcal{Z}_R$ associated with spatial reflection symmetry can be inferred from statistical correlations of randomized measurements (RMs), implemented with local random unitaries $U_1, U_2, \ldots, U_n$ (here $n = 2$) applied symmetrically around the central bond[37]. **c**, Quantized values of the normalized invariant $\tilde{\mathcal{Z}}_R$ reveal two symmetry-protected topological phases and a symmetry-broken phase in a bond-alternating XXZ model with spin-exchange coefficients J and J' and spin-exchange anisotropy coefficient δ (ref.[37]). Dots represent estimations from a finite number of simulated RMs with error bars indicating statistical errors, lines are obtained numerically in a system with $N = 48$ spins and a subsystem I consisting of $n = 6$ pairs of spins. **d**, The topological entanglement entropy S_{top} is defined from second Rényi entropies of various combinations of three connected partitions A,B,C. The graph shows a histogram of the measured topological entanglement entropies S_{top} for different choices of subsystems A,B,C (consisting of a total of 9 qubits) in a system of 31 superconducting qubits implementing the ground state of the toric code[61]. Left and right dashed lines correspond to the ideal and mean measured value of the topological entanglement entropy, respectively. **e**, Two-dimensional feature space uncovered by an unsupervised machine learning model based on classical shadows[65]. Each coloured diamond corresponds to a different quantum state in one of the three phases (trivial, symmetry-broken, topological) of the bond-alternating XXZ model, with colours corresponding to different values of J/J'. The 2D feature space clusters the three phases perfectly. With a standard clustering algorithm[66], the machine learning model would uncover the three phases with high accuracy. Panels **a**,**b**,**c** are adapted with permission from ref.[37], AAAS; panel **d** is adapted with permission from ref.[61], AAAS; and panel **e** is adapted with permission from ref.[65], AAAS.

Figure 4b illustrates the RM protocol to access $\tilde{\mathcal{Z}}_R$ using random unitaries whose spatial distribution is reflection-symmetric across the centre bond. As shown in ref.[37], the statistics of the bitstrings collected in I map to $\mathcal{Z}_R$, where

$$\hat{\mathcal{Z}}_R = \frac{2^n}{MK} \sum_{m=1}^{M} \sum_{k=1}^{K} (-2)^{D\left[\mathbf{s}^{(m,k)}, \mathcal{R}_I(\mathbf{s}^{(m,k)})\right]} \tag{5}$$

is an unbiased estimator of $\mathcal{Z}_R$.

Figure 4c illustrates the protocol in the context of the bond-alternating XXZ model $H_{\text{XXZ}} = \sum_{i=1}^{N-1} J_i (\sigma_i^+ \sigma_{i+1}^- + \sigma_i^- \sigma_{i+1}^+ + \delta \sigma_i^z \sigma_{i+1}^z)$ with spin-exchange coefficients $J_i = J$ for odd $i = 1,3,\ldots$ and $J_i = J'$ for even $i = 2,4,\ldots$ and spin-exchange anisotropy coefficient δ (ref.[37]). This model hosts SPT phases that are protected by spatial reflection and also time-reversal and internal symmetries. Dots in Fig. 4c denote simulated RMs, from which $\mathcal{Z}_R$ and $P_2(\rho_{I_{1,2}})$ have been estimated using equations (5) and (3), respectively. Symmetry-protected topologically trivial ($J'/J \ll 1$) and non-trivial phases ($J'/J \gg 1$) are clearly identified by estimated $\tilde{\mathcal{Z}}_R \approx 1$ and $\tilde{\mathcal{Z}}_R \approx -1$, respectively. In contrast, a large spin-exchange anisotropy $\delta \gg 1$ leads to spontaneous symmetry breaking with $\tilde{\mathcal{Z}}_R \approx 0$.

RM protocols for other topological invariants associated with time-reversal and internal symmetries have also been developed[37], providing a versatile toolbox to identify SPT phases. RM protocols have also been developed to access the many-body Chern number revealing topological order in certain fractional quantum Hall states[60].

RMs were also used to identify topological order in a 31-qubit quantum computer implementing the toric code model[61]. In this case, RM gave access, via the measurement of the purity in connected partitions A, B and C (see Fig. 4d), to the topological entanglement entropy

S_{top} (refs.[62–64]). The topological entanglement entropy takes a quantized value $S_{top} = -1$ in a topologically ordered phase and thus serves as an order parameter to detect the topological character of the toric code. Figure 4d shows results of RMs for a partition of nine qubits whose average corresponds to the topological entanglement entropy.

Machine learning for quantum many-body problems

Classical data obtained from RMs can be used to construct a succinct classical representation of any quantum system. This succinct representation encodes a wide range of properties of the quantum system, opening new opportunities for addressing quantum many-body problems using classical methods such as ML. Here, we focus on the task of classifying quantum phases of matter. Results on provably efficient classical ML algorithms for predicting properties of quantum ground states are discussed in a recent paper[65].

Suppose that one is presented with a collection of quantum states drawn from various quantum phases of matter, where appropriate order parameters for classifying the phases are not known in advance. The goal is to discover the underlying phase structure, and to assign each state to the appropriate phase with high accuracy. One way to proceed is to perform an RM on each state, converting it to a succinct classical shadow, and then to attempt to classify the shadows using a suitably designed classical ML algorithm.

Numerical experiments have confirmed that classical ML algorithms can classify quantum phases successfully. In one such experiment[65], the states to be classified were ground states of the bond-alternating XXZ model with $N = 300$ spins, for various values of the Hamiltonian parameters. As discussed in the previous subsection, this model has three gapped phases: a trivial phase, a broken-symmetry phase, and a symmetry-protected topological phase. For each quantum state, 500 RMs were simulated to construct a classical shadow. An unsupervised classical ML model mapped the classical shadows to a high-dimensional feature space; a projection of this feature space onto its 2D principal subspace is illustrated in Fig. 4e. As the figure indicates, states belonging to the same phase cluster tightly in the feature space found by the classical ML model; hence standard clustering algorithms such as k-means clustering[66] can group the states into distinct phases with near-perfect accuracy. Similar results were found for other models with multiple phases.

Recent work[65] also provides rigorous theoretical support for these empirical studies. Consider a supervised learning scenario, where the training data consist of classical shadows of quantum states, each accompanied by a label indicating the phase to which that state belongs. The ML algorithm then predicts the phase label for new classical shadows different from those encountered during training. Under a plausible physically motivated assumption, one may prove that the classical ML model can learn to classify phases accurately with amounts of training data and classical processing that scale polynomially with the total system size. The key assumption is that the phases can be distinguished accurately by a nonlinear function of marginal density operators of subsystems of constant size. For phases with an energy gap, this assumption is reasonable because one expects the phase to be revealed in subsystems that are larger than the correlation length but independent of the total system size. Under this assumption, the classical ML model, receiving classical shadows of quantum states as input, not only learns to classify phases accurately, but also constructs an explicit classification function. The numerical results indicate that the ML model classifies phases accurately even in an unsupervised scenario, where the training data are unlabelled.

Quantum chaos diagnostics

Prominent diagnostics of many-body quantum chaos are out-of-time-ordered correlators (OTOCs), which detect the scrambling of quantum information by revealing how local perturbations spread as a function of time[67–69]. In their simplest form, OTOCs at infinite temperature can be written as $O(t) = \mathrm{tr}(\rho_\infty W(t)VW(t)V)$, where W, V are local operators that act on small subsystems, $W(t) = e^{-iHt}We^{iHt}$ is a time-evolved operator in the Heisenberg picture, determined by the Hamiltonian H, and $\rho_\infty \propto \mathbb{1}$ denote the maximally mixed 'infinite temperature' state.

RMs allow one to extract OTOCs of unitary operators V, W at infinite temperature from statistical correlations of two separate experiments[33]. Importantly, no ancilla degrees of freedom and only forward time evolution are required. The idea is to generate in both experiments the same randomized initial state via the application of local random unitaries U to a simple computational basis state $|\psi_0\rangle$. In the first experiment, this quantum state evolves for time t, and we measure an expectation value $\langle W(t)\rangle_1 = \langle\psi_0|U^\dagger W(t)U|\psi_0\rangle$; for example, a Pauli operator on site i is measured. The second experiment is similar, except that the operator V (for example, a Pauli operator at a different site j) is applied before the time evolution. One then obtains a different expectation value $\langle W(t)\rangle_2 = \langle\psi_0|U^\dagger V^\dagger W(t)VU|\psi_0\rangle$. This is repeated for many randomized initial states, and the statistical correlations between the two measurements $\langle W(t)\rangle_1$ and $\langle W(t)\rangle_2$ can be directly mapped to OTOCs.

At initial times, the measurement of W at site i yields the same outcome in both experiments regardless of the application of V at a different site j: that is, $\langle W(t)\rangle_2 = \langle W(t)\rangle_1$. These maximal statistical correlations correspond to the maximal initial value of the OTOC. With time, the information about the application of V on site j spreads ('scrambles') through the entire system, and the measurement $\langle W(t)\rangle_2$ in the second experiment differs in general from the first measurement $\langle W(t)\rangle_1$. This decay of correlation between the measurement outcomes directly corresponds to the decay of OTOCs in scrambling quantum systems. Here, what governs the scalings of statistical errors is not directly the system size N, as in other protocols mentioned above, but the details of the dynamics: at early times, $\langle W(t)\rangle_1 \approx \langle W(t)\rangle_2$ is of order 1, and thus the OTOC can be estimated accurately with a small number of measurements. However, for chaotic systems at long times, $\langle W(t)\rangle_{1,2} \approx (1/2)^{N/2}$, so that the estimation of the OTOC with a small relative error requires approximately 2^N measurements. We note that the protocol can also be extended to access finite-temperature OTOCs. In this case, one needs to sample global random states, which are distributed according to their overlaps with the thermal state of interest[33].

Using the described protocol, infinite temperature OTOCs have been measured experimentally in a trapped-ion quantum simulator[34] to study the scrambling of quantum information in quantum spin models with tunable long-range interactions, and also in an nuclear magnetic resonance experiment[70].

Recently, such families of RM protocols based on propagating random initial states have been extended to access other quantum chaos diagnostics[71–73]. In particular, an RM protocol has been proposed[73] to access the spectral form factor $K(t) = |\mathrm{tr}(\exp(-iHt))|^2$ – a quantum chaos diagnostic that is directly connected to the statistics of eigenlevels of the Hamiltonian H. It can be used to test predictions of random matrix theory and universal aspects of thermalization in many-body quantum systems. The key idea of this protocol is to apply, within a single experimental run, the same local random unitaries before and after the time evolution described by the evolution operator $\exp(-iHt)$. From the statistics of final computational basis measurements one can

infer[73] $K(t)$, that is, (the product of) the traces of $\exp(-iHt)$ and its adjoint. This extends the RM toolbox to access genuine properties of dynamical quantum processes, without reference to an initial state or measured observable (see also ref.[74] for an extension of the classical shadows framework to quantum processes).

Fidelity estimation

Suppose one wishes to prepare a target state ψ, which we assume is pure for simplicity. If instead, owing to experimental limitations, one can only prepare the state ρ, how can one tell how close ρ is to ψ? As discussed above, full tomography could achieve this but is very expensive. The RM toolbox provides another answer using the protocol of direct fidelity estimation (DFE)[75,76]. In DFE, one first chooses an operator basis of convenient observables, for example the N-qubit Pauli operators. The fidelity between ψ and ρ when ψ is pure reduces to $F(\psi,\rho) = \mathrm{tr}(\psi\rho)$. Next, one expands ρ and ψ in the Pauli basis as $\rho = \sum_j a_j W_j/2^{N/2}$ and $\psi = \sum_j b_j W_j/2^{N/2}$ for W_j, the jth N-qubit Pauli operator. Expanding the fidelity, one finds that

$$F(\psi,\rho) = \mathrm{tr}(\psi\rho) = \sum_j a_j b_j = \sum_j \left(\frac{a_j}{b_j}\right) b_j^2,$$

where the latter sum is over the support of b_j. The purpose of rewriting in this latter form is that $\sum_j b_j^2 = 1$ for a pure state ψ, so the fidelity has been reinterpreted as an 'expected value' over a known distribution. Furthermore, the ratio $a_j/b_j = \mathrm{tr}(W_j\rho)/\mathrm{tr}(W_j\psi)$ is an observable quantity since a_j can be estimated empirically and b_j is known from the known target state. By doing Monte Carlo importance sampling of the distribution b_j^2 for about $\mathcal{O}(1/\varepsilon^2)$ samples, one obtains a random collection of observables $\{W_k\}$ that one can then estimate using, for example, shadow estimation, and it yields a randomized estimate of the fidelity F that is accurate to within $F \pm \varepsilon$ with high probability.

Although DFE requires sampling only $\mathcal{O}(1/\varepsilon^2)$ Pauli observables, independent of N, there is still some important scaling with N in this protocol. First, for generic states, the choice of Pauli operators will include samples of very high-weight Pauli strings that must be estimated, and these are generally more difficult to estimate as a single two-outcome measurement or as an inferred observable from many single-qubit measurements. Second, to estimate the fidelity of a generic state requires the ability to resolve a Pauli observable to a precision $\pm(1/2)^{N/2}$. Thus, the worst-case complexity for the total number of measurements (not just observables) is $\mathcal{O}(2^N/\varepsilon^2)$ which is still exponential in N, albeit better than full quantum tomography by another factor of 2^N. In addition, in the most general case, just describing the distribution b_j^2 from the target state requires listing all 4^N probabilities, which quickly becomes expensive with increasing N. However, this worst-case behaviour can significantly overestimate the scaling for many important cases where most of the probability mass b_j^2 in the target state is concentrated on a relatively small number of low-weight Pauli operators, and where one can easily compute (or estimate) this distribution. Important examples of such clustering behaviour are stabilizer states with local stabilizer generators and high-temperature Gibbs states of local Hamiltonians.

The above method generalizes naturally to quantum processes (see refs.[75,76] for details) and was first used to efficiently estimate the process fidelity of a Toffoli gate in a superconducting transmon architecture[77]. Another notable use of the method was to validate the fidelity of a 14-qubit state preparation[78] in an ion trap, as shown in Fig. 5a. A variant of DFE that uses a simpler Pauli measurement

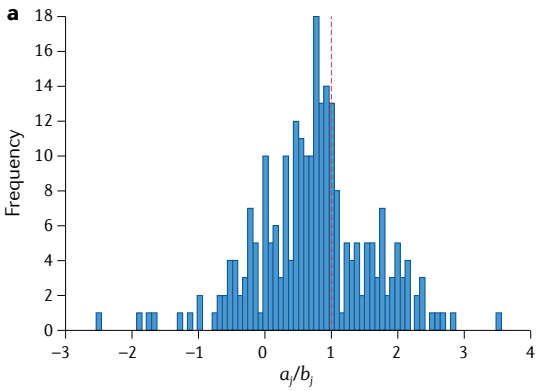

b

	IBM_2	IBM_3	UMD_1	IonQ_2
IBM_3	0.294(9)	1	0.127(7)	0.129(6)
UMD_1	0.157(6)	0.127(7)	1	0.896(3)
IonQ_2	0.156(7)	0.129(6)	0.896(3)	1
Simulation	0.143(6)	0.102(7)	0.956(2)	0.908(2)

Fig. 5 | Fidelity estimation with randomized measurements. a, Direct fidelity estimation allows to extract the fidelity between an experimentally prepared quantum state ρ and and a pure theoretical target states ψ as the mean of the distribution of the random variable (a_j/b_j) (see main text). The histogram shows the experimental data for a 14-qubit quantum state prepared in trapped-ion quantum simulator with estimated mean (fidelity) of 0.75 ± 0.05 (ref.[78]). The red dashed line indicates the ideal case of unit fidelity. **b**, Cross-device fidelities of seven-qubit quantum volume states prepared in various quantum devices based on superconducting qubits (IBM_3) and trapped ions (UMD_1, IonQ_2). Data from ref.[32]. Panel **a** is adapted with permission from ref.[78], Springer Nature Ltd.

scheme and achieves nearly optimal sample complexity was proposed and tested on four-qubit entangled states in a trapped-ion device[79,80].

Although DFE allows the comparison between an ideal state or process and an imperfect experimental implementation, RM also enables direct comparison between two experiments[31]. Here, the goal is to measure a fidelity between two mixed states ρ_1, ρ_2, defined as

$$F_{\max} = \frac{\mathrm{tr}(\rho_1\rho_2)}{\max[\mathrm{tr}(\rho_1^2), \mathrm{tr}(\rho_2^2)]}. \quad (6)$$

The density matrices ρ_1 and ρ_2 represent quantum states realized in different experimental devices, which may be separated by a large distance, not operating at the same time, and not using the same physical systems. Both density matrices can refer to subsystems of constant size of large quantum many-body systems.

Such a comparison between different devices is relevant in the context of verifying and benchmarking quantum computers and simulators[81,82]: one might gain confidence in the result of a quantum computation or simulation by running the computation or simulation on various devices and comparing, through quantitative measures such as $F_{\max}$, the outcome of one with the other[31,82].

The key to determining $F_{\max}$ is the measurement of the overlap $\mathrm{tr}(\rho_1\rho_2)$. With RM, this can be achieved as follows: first, generate a set of random unitaries U, for instance made of random single-qubit rotations and send them via classical communication to the two devices;

second, apply these same unitaries to both ρ_1 and ρ_2, followed by computational basis measurements. The overlap $\mathrm{tr}(\rho_1\rho_2)$ can be then extracted from the statistical correlations between the outcomes obtained in both devices[31]. One can understand this result as follows: if the two states are identical, $\rho_1 = \rho_2 = \rho$, the bitstrings measured in both devices will be picked from the same distribution $\langle \mathbf{s}|U\rho U^\dagger|\mathbf{s}\rangle$: that is, one will observe perfect correlations between the two experiments. If, instead, the two states are different, the outcomes will be typically uncorrelated. Importantly, whereas the required number of measurements is exponential in the (sub)system size, the cost is strongly reduced compared with performing full quantum state tomography in both devices. This allows one to access (sub)system sizes beyond the regime of full quantum state tomography.

A proof-of-principle demonstration of this protocol was presented[31] by reanalysing the data of ref.[35] to compare the experimentally prepared quantum states of up to 10 qubits with a theoretical simulation, or with a different quantum state prepared in the same experiment. Comparison of quantum devices across different qubit technology has been achieved[32] by comparing entangled quantum states consisting of up to 13 qubits prepared on six different quantum devices based on trapped ions or superconducting qubits (see Fig. 5b).

There are also approaches to specifically estimate fidelities of random quantum states generated by random quantum circuits. In this case, one can define and extract the fidelity between prepared quantum states and their theoretical target without needing to add another layer of randomness in the measurement stage. Variations of this idea are known as cross-entropy benchmarking[83–85] and random circuit sampling[86,87]. This approach can be generalized to quantum states generated via ergodic Hamiltonian dynamics using the concept of projected state ensembles[88,89].

Hamiltonian and Liouvillian learning

RM can be used to learn dynamical variables that govern quantum evolution such as Hamiltonians and more generally Lindbladians. There are many approaches to Hamiltonian and Lindbladian learning in the literature, but the quantum part of nearly all of them boils down to estimating the expectation values of low-weight Pauli observables, so each of these algorithms can benefit from the RM toolbox.

Let H be an unknown Hamiltonian, and suppose one is given the ability to prepare the ground state ψ_0. Perhaps surprisingly, when H is sufficiently generic, low-weight Pauli observables with respect to ψ_0 contain enough information to reconstruct H up to an overall scale factor (and an unobservable energy shift)[90–92]. The argument is remarkably simple and works even for a steady state ρ, not just the ground state ψ_0.

Expand H in the Pauli basis as $H = \sum_j c_j W_j$, and suppose that $c_j = 0$ whenever the support of W_j exceeds w qubits for some $w = \mathcal{O}(1)$. Define the matrix $K_{lm} = \mathrm{i}\,\mathrm{tr}(\rho[W_l, W_m])$. If ρ is a steady state then $[H, \rho] = 0$, and furthermore for any observable O we have $\mathrm{tr}(\rho[O, H]) = \mathrm{tr}([H, \rho]O) = 0$. Now consider the action of K on the vector of Hamiltonian couplings c. By linearity,

$$\begin{aligned}(Kc)_l &= \mathrm{i}\sum_m \mathrm{tr}(\rho[W_l, W_m])c_m \\ &= \mathrm{i}\,\mathrm{tr}(\rho[W_l, H]) \\ &= \mathrm{i}\,\mathrm{tr}([H, \rho]W_l) = 0.\end{aligned} \tag{7}$$

Thus, c is in the kernel of K.

This suggests a procedure for estimating c. First, estimate the matrix elements of K by preparing the steady state ρ and measuring the Pauli observables $\mathrm{i}[W_l, W_m]$ where each W_l or W_m is at most w-body. This gives us an estimate $\hat{K} \approx K$, from which we can return an estimate $\hat{c} \approx c$ by finding a right singular vector of $\hat{K}$ with the least singular value as a proxy for estimating the kernel of K. (Since this vector is only specified up to a scalar multiple, one must in general do a Rabi-type experiment to pin down the overall scale factor that completely determines an estimate $\hat{c}$).

The precision of this estimate will depend on several factors. First, one must be able to prepare a steady state ρ. This can be done by time-averaging the results of the experiment[42,92]. Previous work[42] describes a protocol based on Gaussian quadratures that achieves super-exponential convergence to a steady state in the number s of time averages, meaning that the trace distance to a steady state scales $\propto \exp(-\mathcal{O}(s\log(s)))$, so this step is very efficient. Second, the matrix elements of K must be estimated to sufficient precision. Notice, however, that all of the matrix elements K_{lm} are specified by Pauli observables of weight at most $2w - 1$ (since disjoint Paulis commute) when H has w-body couplings. Therefore, the K_{lm} can be estimated using classical shadows specialized to Pauli observables, and it is here that RMs come in. Finally, there must be sufficient signal in the observables so that the 'approximate kernel' (the space spanned by the least singular value) of $\hat{K}$ is 1D. If the approximate kernel is degenerate, or if the gap between the two least singular values is very small, then the problem is ill-conditioned and cannot be solved to useful precision.

This technique has been generalized in several directions. It has been shown[42] that Bayesian priors on the Hamiltonians can be incorporated to speed up convergence, and that well-characterized auxiliary control fields can be used to enhance the precision as well. A further generalization is to replace the steady state ρ by a fixed point of a Lindbladian. It was shown[93] that this generalization admits similar guarantees as in the Hamiltonian case. Hamiltonian learning with this approach can also be done using the dynamics of a quenched quantum system[94].

Another avenue for generalization is provided by learning the 'entanglement Hamiltonian'[95]. An entanglement Hamiltonian is a Hamiltonian that describes the mixed state $\rho_A = \mathrm{e}^{-\mathcal{H}_A}/Z$ obtained when tracing out half of a bipartite pure state ψ_{AB}. Knowing the entanglement Hamiltonian for subsystems A and B and their spectrum is equivalent to knowing the Schmidt decomposition of ψ_{AB}, and thus contains the complete information about the bipartite entanglement across this cut. One study[95] uses RM data from ref.[35] to estimate the parameters of physically motivated ansätze for the entanglement Hamiltonian. Up to 7-qubit subsystems of a 20-qubit trapped-ion quantum simulator are considered in various states following a quantum quench. For each state, the learned entanglement Hamiltonian is then experimentally verified by independently estimating the fidelity of the Gibbs state defined by learned entanglement Hamiltonian and the quantum state realized in the laboratory. Figure 6 shows the estimated fidelities as a function of time after the quantum quench for various ansätze for $\mathcal{H}_A$. At early and late times, $\mathcal{H}_A$ is well approximated by a deformation $\widetilde{H}_A = \sum_{i,j\in A}\widetilde{J}_{ij}(\sigma_i^+\sigma_j^- + \mathrm{h.c.})$ of the system Hamiltonian $H_A = \sum_{i,j\in A}J_{ij}(\sigma_i^+\sigma_j^- + \mathrm{h.c.})$ with $J_{ij} = J/|i-j|^\alpha$ and $\alpha \approx 1.24$ (dark blue points). At intermediate times, multibody corrections $\widetilde{K}_A^{(i)}$ of increasing complexity ($i = 1,2$, light blue, red points) are important. Here, $\widetilde{K}_A^{(1)}$ contains lattice momentum terms of the form $\sim\sigma_i^x\sigma_j^y$ ($i,j \in A$) and $\sim\sigma_i^x\sigma_j^y\sigma_k^z$ ($i,j,k \in A$) which are suggested by predictions of conformal field theory[96]. $\widetilde{K}_A^{(2)}$ includes all further two-body and three-body terms compatible with global magnetization conservation.

A completely different approach to Hamiltonian learning has also been developed that works in a different regime. Several papers[97–99] have considered learning a quantum Hamiltonian H from the Gibbs state $\rho_{\beta H}$

at a small inverse-temperature β. They showed that estimating expectations $\mathrm{tr}(W\rho_{\beta H})$ of few-qubit Pauli operators W enables accurate reconstruction of the Hamiltonian H when H couples only $k = \mathcal{O}(1)$ qubits at a time and each qubit partakes in at most $\ell = \mathcal{O}(1)$ interactions. However, it is assumed that the support of the non-zero interactions is known in advance. This family is a large and natural class, but some familiar systems such as two-body Hamiltonians with power-law interactions lie outside it. Their algorithms are unconditional in the sense that they work for any Hamiltonian in the family as long as one can prepare the Gibbs state. Assuming this state preparation, the estimation of $\mathrm{tr}(W\rho_{\beta H})$ can be accomplished by performing randomized local measurements.

Aside from learning Hamiltonians from steady states or Gibbs states, one can also use short-time dynamics of the time evolution to learn a Hamiltonian. There are many approaches for this, but one that is notable for its use of RM is the sparse Hamiltonian learning method of ref.[100]. The authors of this study apply random Pauli gates to the state after some short-time dynamics, so that the time-evolved state is effectively linear in the Hamiltonian. From measurements of these states, they obtain estimates of the Hamiltonian coefficients. They prove a guarantee that their method converges to the true Hamiltonian whenever it is sparse in the Pauli basis, regardless of the weight of the unknown Pauli operators in the support. Although this approach is technically efficient, the randomized Pauli operators and the short-time requirement mean that the sampling overhead is around $\mathcal{O}(1/\varepsilon^8)$ to achieve precision ε, so further improvements will be needed to make this practical. It is known that a related approach[101] can improve these guarantees when the unknown Hamiltonian terms are assumed to be of the form $\omega_{ij}Z_iZ_j$ between arbitrary pairs of qubits.

Although most studies listed above focus on qubit Hamiltonians, learning Hamiltonians of infinite-dimensional quantum systems is equally relevant. One study[102] performed Hamiltonian tomography in a nearest-neighbour transmon qubit architecture by estimating couplings among up to six qubits. The authors of this study were able to synthesize these estimates across 27 total qubits to provide a comprehensive picture of the couplings in the device.

Perhaps the most general approach to learning Hamiltonians and Lindbladians is given by notions related to gate set tomography (GST)[103,104]. GST grew out of the need for self-consistent estimates of a quantum process together with the noisy measurements and state preparations used as probes[105]. It works by modelling the noise on an entire collection of gates, measurements and preparations, and then fitting all of these models to the measurements done on a large number of circuits of varying length. Smart choices of subcircuits known as 'germs' attempt to amplify small errors to make them easier to estimate. This basic idea has been used in dozens of experiments on one and two qubits[104], and recently it was used to characterize the one- and two-qubit gates in a three-qubit experiment with two donor nuclear spins in silicon and their shared electron spin[106]. Although GST initially focused on estimating gate sets in the form of completely positive maps, this focus broadened in experiments[106,107] to estimating Lindbladian generators using the same or similar data fitting techniques.

Several attempts to improve on these ideas are now being explored both theoretically and experimentally. On the theory side, compressive GST[108] is a formulation of GST as a tensor completion problem in an effort to bring down the (substantial) computational complexity of the method and provide rigorous convergence guarantees. Fast Bayesian tomography (FBT)[109] is a variant that allows the use of prior information or side information (from randomized benchmarking experiments for example) to speed up GST. In one experiment[109], FBT was performed on two spin qubits in silicon so quickly that the post-processing was faster than the data acquisition, demonstrating that FBT can in principle be used as an online algorithm. Lastly, a special case of a Lindbladian, called a Pauli–Lindbladian, was learnt in a study[110] using techniques similar to the averaged circuit eigenvalue sampling framework[111] discussed in the Supplemental Information and elsewhere[112]. Although this class of Lindbladians is substantially narrower than a general Lindbladian in a qubit system, it has the advantage that the learned noise can be error-mitigated very efficiently. This allows the authors[110] to perform probabilistic error cancellation on a superconducting quantum device despite the presence of crosstalk errors.

Variational quantum-classical algorithms

Variational quantum-classical algorithms use NISQ devices as special-purpose quantum co-processors in tandem with a classical computer to solve complicated optimization problems, for instance finding the ground-state energy of a many-body Hamiltonian or variationally

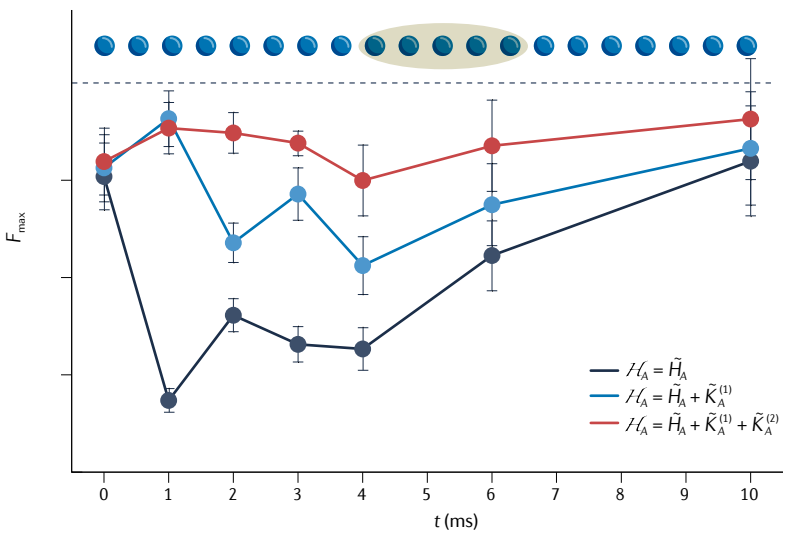

Fig. 6 | Hamiltonian learning with randomized measurements. Entanglement Hamiltonian tomography[95] serves to efficiently learn and independently verify the entanglement Hamiltonian $\mathcal{H}_A$, parameterizing the reduced density matrix $\rho_A \propto \exp(-\mathcal{H}_A)$ of a subsystem A (here five qubits, indicated by the shaded area). This has been experimentally demonstrated[95] using RM data collected[35] in a trapped-ion quantum simulator with a total number of $N = 20$ qubits (shown as blue circles at the top of the panel). Shown is the experimentally estimated fidelity $F_{\max}$ of the learned $\rho_A \propto \exp(-\mathcal{H}_A)$ at various times t after a quantum quench and for different physically motivated ansätze for $\mathcal{H}_A$ constructed as deformation of the system Hamiltonian $\tilde{H}_A$ and including further multibody corrections $\tilde{K}_A^{(i)}$ ($i = 1,2$) (different colours; see main text for details). Figure is adapted with permission from ref.[95], Springer Nature Ltd.

compressing quantum circuits[113]. Such variational hybrid approaches typically require many evaluations of complicated cost-functions through measurements on a variational quantum state. In particular, for quantum chemistry[114–116] and quantum field theory applications[117], where the cost function is represented by the energy of a complicated many-body Hamiltonian, this poses an important practical obstacle, as each cost-function evaluation requires many measurements in multiple settings.

RMs can provide a big advantage here, since they allow for jointly estimating many observables (for example, Hamiltonian terms) based on the same RM data via classical shadows. According to Theorem 1, the required number of measurements only scales logarithmically in the number of observables of interest – an exponential improvement over direct estimation protocols that estimate observables one by one. An example[24] is illustrated in Fig. 7 where the number of RMs $\propto \log(L)$ required to estimate the variance $\mathrm{tr}(H^2\rho) - \mathrm{tr}(H\rho)^2$ of a many-body Hamiltonian H in a given state ρ, where H contains $\propto L^2$ terms, outperforms an optimized 'hand-crafted' scheme[117] for large system sizes L.

However, the required number of measurements scales exponentially with the size w of the observables of interest, for example the locality of the quantum many-body Hamiltonian. This can quickly become an issue for applications to fermionic systems such as those encountered in quantum chemistry, where Jordan–Wigner encodings produce Hamiltonian terms with high weight. Several improvements of the elementary RM protocol are known that address this issue:

- Importance sampling of measurement settings uses knowledge of the observables of interest[118,119] or of the underlying quantum many-body state[44].
- Derandomization[120] replaces the randomized measurements by fixed measurements which, for a specified set of target observables, are guaranteed to predict these observables at least as well as the randomized protocol, and perform much better in some cases; see Fig. 7 and also recent work[121] for an experimental demonstration.
- Fermionic classical shadows[122] use fermionic Gaussian random unitaries to access k-body reduced density matrices of fermionic states encoded in a qubit system with optimal sampling complexity (up to a logarithmic factor).

Finally, we note that other schemes for observable estimation, based on grouping compatible Pauli observables[123], adaptive selection of measurement settings[124] or sampling of high-weight matrix elements[125], have been shown to provide further advantage in certain situations.

Machine learning in quantum-enhanced feature space

To use NISQ devices for general ML problems, a class of supervised learning models using quantum-enhanced feature spaces was proposed[126,127]. These quantum machine learning (QML) models are trained to predict outputs, such as a real number or a discrete label, given an input vector. The input vector, referred to as a feature vector, is transformed into a higher-dimensional quantum-enhanced feature vector (a quantum state) using NISQ devices. The QML models then train a linear function over the quantum-enhanced feature vectors via convex optimization. Because of the convex landscape, the global optimum can always be found efficiently, without encountering any barren plateau problem[128].

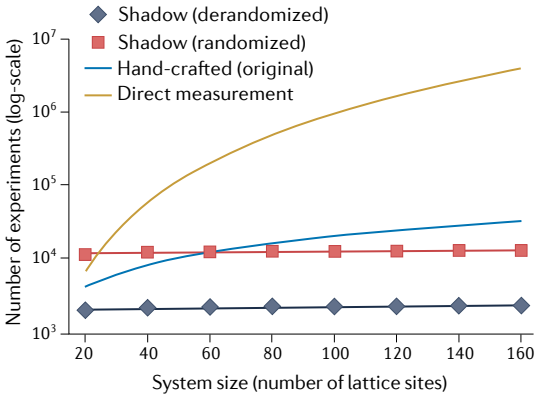

Fig. 7 | Variational quantum algorithm using randomized measurements. The number of randomized measurements (red) to estimate the Hamiltonian variance $\mathrm{tr}(H^2\rho) - \mathrm{tr}(H\rho)^2$ of the Schwinger model, describing one-dimensional quantum electrodynamics, scales only logarithmically with system size, providing, for large systems, an exponential advantage over a direct measurement of Hamiltonian terms (brown) or optimized 'hand-crafted' schemes (light blue)[117]. This can be further improved by derandomization, where each random choice is replaced with the one that yields the best expected performance in the randomized protocol[120] (dark blue). Adapted with permission from ref.[24], Springer Nature Ltd.

Whereas training these QML models is always easy, a recent paper[129] proves that the original proposal[126,127], which constructs the quantum-enhanced feature vector as a quantum state, can have a poor prediction performance. For simple learning tasks, the authors show[129] that the prediction performance can be significantly worse than classical ML models even if the QML model can perfectly fit the training data. The issue stems from the fact that most pairs of quantum states have an exponentially small fidelity overlap, which leads to bad generalization performance on new unseen data through a learning-theoretic analysis[129]. To resolve this issue, the authors[129] propose a different approach for constructing the quantum-enhanced feature vectors using RM. Each quantum-enhanced feature vector corresponds here to a finite set of properties of the quantum state estimated with RM instead of the full quantum state itself. The RM variant yields a simple proof of quantum speed-up over any classical ML model assuming a highly plausible cryptographic conjecture[129] and provides a quadratic speed-up over the original proposal of the quantum-enhanced supervised learning model[126,127] in the number of measurements[130]. This work[129] also shows empirically that the RM variant yields a higher prediction accuracy over conventional classical ML models and the original QML proposal[126,127].

Higher-order polynomial functionals of the density matrix

The use of RM to measure the purity $\mathrm{tr}(\rho^2)$ is illustrated in the example discussed in the Supplementary Information. Expectation values of higher-order k-copy observables of the form $\mathrm{tr}(O\rho^{\otimes k})$ have also been estimated via RM. In particular, the classical shadow formalism enables one to access the expectation value of any such k-copy observable O by cross-correlating k different shadows; this extends equation (4), the special case where $k = 2$ and O is the swap operator, to more general k and O (ref.[24]).

As an important application, RM can be used to access moments $\mathrm{tr}((\rho_{AB}^{T_A})^k)$ of the partial transpose $\rho_{AB}^{T_A}$ of a density matrix ρ_{AB} describing two subsystems A and B (refs.[36,131]). These moments are the expectation values $\mathrm{tr}((\rho_{AB}^{T_A})^k) = \mathrm{tr}((\Pi_A \otimes \Pi_B)\rho_{AB}^{\otimes k})$ of two different cyclic

permutation operations Π_A, Π_B acting on k copies of A and k copies of B, respectively[132,133]. There are inequalities that are satisfied by such moments if $\rho_{AB}^{T_A}$ is a non-negative operator; thus, if these inequalities are found to be violated, then $\rho_{AB}^{T_A}$ must have a negative eigenvalue, and it follows that ρ_{AB} is entangled[36,57,131,134,135]. This property was used to experimentally demonstrate mixed-state entanglement in (sub)systems consisting of up to seven qubits[36]. Similarly, symmetry-resolved entropies[40], the quantum Fisher information[41], and moments of 'realigned' density matrices[136] can be interpreted as expectation values of multicopy observables and inferred from RM. A recent paper[41], in particular, presents closed-form formulas for estimating the numbers of measurements that are required to measure an arbitrary k-copy observable O via classical shadows.

Polynomials of the density matrix can also be used to detect and quantify multipartite entanglement. One way to achieve this is by studying moments of the outcomes of RM, implemented with independent local random unitary operations on each subsystem[27,28,137–141]. Such a procedure does not require a common reference frame shared by the constituents, and is robust against local miscalibration of the measurement basis. It is thus well suited to detect multipartite entanglement in quantum networks with distantly separated parties[28,139,142]. Detecting multipartite entanglement in this way has been experimentally demonstrated with four-photon quantum states[27]. Another (fourth-order) polynomial function of the density matrix that can be estimated using RMs is the stabilizer Rényi entropy, which quantifies the non-Clifford resources needed to prepare a quantum state[143,144].

Finally, we note that RM recently inspired numerical sampling techniques for tensor networks that allow one to estimate polynomial functionals of tensor network states without potentially costly contraction of multiple replicas[145].

Challenges and perspectives

We list below some of the current challenges and open questions for the RM toolbox, which we have grouped into four main topics: assessment and mitigation of errors occurring during the measurement process, use of different random unitaries ensembles as measurement primitives, implementation of RM in quantum systems that are not qubits, and applications of RMs for learning tasks.

Robustness and error mitigation

In the era of NISQ devices, quantum operations are necessarily altered by noise and decoherence. This applies in particular to the measurement process itself. Any practical procedure for learning properties of a quantum system must thus be equipped with sufficient robustness against noise – that is, the ability to make accurate predictions even in the presence of (a certain level of) noise.

Prediction procedures based on RM involve an average over an ensemble of random unitary rotations. As such, the influence of noise that alters the application of the random unitary and the projective measurement can often be reduced to an averaged noise channel. This averaged noise channel can be efficiently learned from calibration experiments, in order to provide robust estimations.

First, estimation formulas presented in refs.[22,31,146] for purity (see equation (3)) and fidelity estimation involve only the measured bitstrings $\mathbf{s}$. No information on the applied random unitaries U, other than the assertion that they are picked from an ensemble that covers the unitary group evenly (a unitary 2-design), is required. Thus, the estimation procedure is insensitive to gate-independent unitary errors, for example random over- and under-rotations[22,31,35,56,61,146]. Second, in the presence of well-characterized depolarization or simple qubit readout errors, the estimation procedure can be corrected based on calibration experiments to provide unbiased estimations[31,56,61,146].

In the classical shadows formalism, robust estimations can also be performed effectively in the presence of an unknown noise channel (satisfying certain standard assumptions such as gate-independence and time-independence). This is achieved via a calibration step that uses a state that can be prepared easily, say the state $|0\rangle^{\otimes N}$ (refs.[29,30,147]). By applying random unitaries that twirl the unknown noise, one can put the noise into the form of a purely stochastic Pauli channel. The calibration step characterizes this Pauli channel; thus the noise can be compensated in the classical post-processing of the RM results when expectation values of observables are estimated. If the noise is not too strong, and is Markovian, gate-independent and time-independent, this procedure effectively eliminates the bias, reducing the effect of the noise on estimated observables until it becomes comparable to the noise floor arising from the sampling error in the characterization of the Pauli channel. In this case, the procedure requires approximately the same number of measurements as the standard shadow estimation scheme[29,30].

Local versus global random unitaries

So far we have discussed in detail implementations of RMs with local random unitaries corresponding to single-qubit rotations, but the ideas in the RM toolbox extend well beyond this. This includes, in particular, global random unitaries which scramble information across the entire system, implemented as quantum circuits or, in an approximate way, as random quenches.

Quantum information aspects. This idea of using global random unitaries goes back to early work, even before their local counterparts[21,22,51]. Global random unitaries come with analytic expressions that are well suited for estimating certain global state properties. However, in contrast to local random unitaries, the type of post-processing matters a lot. Special-purpose formulas, such as equation (3) for the purity, have global counterparts that can be computed efficiently[21,22]. This, however, may not be the case for global variants of general-purpose formulas, such as equation (2), because these require explicit knowledge of the random unitaries in question. This issue can quickly become challenging, owing to the curse of dimensionality (a general unitary acting on N qubits has roughly 4^N degrees of freedom). A notable exception is random multiqubit Clifford circuits[24,26]. These are global scrambling unitaries that nonetheless come with efficient classical post-processing (via the Gottesman–Knill theorem[148]).

In the quantum circuit framework, one can also consider interpolations between local and global random unitaries. In recent work[149], the authors use tensor network techniques to study random unitaries based on shallow-depth circuits. In such a setup, the circuit depth provides a convenient tuning knob: local properties are best estimated with very shallow circuits (realizing 'quasilocal random unitaries'). In contrast, the number of experimental runs to estimate certain global properties such as the global state fidelity decreases exponentially with the circuit depth. Generalizing this approach, other authors[150] consider ensembles of random unitaries whose distribution is invariant under multiplication with Pauli strings. This comprises local and global random unitaries as well as unitaries generated by shallow-depth random circuits. A drawback of both approaches is the lack of a simple analytical expression for classical shadows and succinct property prediction based on the obtained RM data (in contrast to RMs implemented with local random unitaries: see equations (1) and (2)).

Experimental aspects. The choice of implementing local versus global random unitaries also depends on the level of experimental control that is available in a given quantum hardware setting. This choice can be made based on how the process of creating of these unitaries is affected by decoherence. In a Hamiltonian spin system, RMs with global random unitaries can be implemented approximately, for example at the level of an approximate unitary 2-design[22,51,146,151]. The idea is to use random quenches built from engineered time-dependent disorder potentials, requiring coherent interactions during a time window that increases with system size. With a quantum computer, one can also implement random unitaries with Clifford circuits, with a size scaling quadratically with the number of qubits[152]. From this perspective of decoherence, local random unitaries seem thus to have an advantage over global random unitaries, as the required coherence time does not increase with system size. However, for implementing local random unitaries with high fidelity in programmable quantum simulators of trapped ions[35] or Rydberg atoms[153], for example, an important assumption is that any residual interaction between qubits can be turned off or made negligible. Global random unitaries can be created instead in presence of an interaction 'background', and may be thus seen as more appealing for certain experimental setups. Finally, between these two extreme cases of local–global random unitaries, the shallow-depth random circuits of refs.[149,150] can also be seen as a good compromise to realize random unitaries in an interacting system, while requiring a limited coherence time (related to the depth of the circuit). In addition, robust estimation schemes[29,30] to further mitigate the effect of noise and decoherence might be transferred to this setting[150].

Beyond qubits

The RM toolbox can be extended to other quantum systems beyond qubits, in particular to systems consisting of qudits (d-level systems with $d \geq 3$) and to fermionic and bosonic systems using global random unitaries.

Qudits. Although the RM toolbox is traditionally discussed for qubits or spin-1/2 particles in quantum computing and quantum simulation, atomic platforms in particular offer naturally high-dimensional internal state spaces that can serve as qudits in hardware-efficient universal quantum computing[154], and can also represent spins $S > 1/2$ in quantum simulation. We emphasize that the present protocols and applications, discussed for qubits, generalize to these cases (see ref.[23]), and random unitary operations and state-resolved measurements on single qudits are readily implementable on existing platforms of trapped ions or Rydberg tweezer arrays[5,8], for instance. In these settings, classical shadows can furthermore be constructed from general positive operator-valued measurements[155,156]. Single-qubit symmetric informationally complete measurements[157,158] seem an ideal candidate, because they may alleviate the need for choosing measurement settings at random while still allowing for fast state reconstruction and fast estimation of arbitrary polynomial functionals of the quantum state. Recently, this approach has been implemented in both superconducting circuit architectures[159] and trapped-ion quantum computers[160].

Fermions and bosons. Beyond qubits and qudits, programmable quantum many-body systems can be engineered with bosonic or fermionic particles as basic constituents. A seminal example is provided by ultracold fermionic atoms in optical lattices described by a 2D Fermi–Hubbard model with repulsive interactions, where state of the art experiments achieve single-site control and site-resolved single-shot readout via a quantum gas microscope[161]. These setups provide the toolbox to prepare and study strongly correlated equilibrium phases and non-equilibrium dynamics. We note that local random unitaries as discussed above are typically not available in these experiments, as physical Hamiltonians generating unitaries are constrained by conservation laws, such as atom number conservation in closed systems. Instead, an RM toolbox can be developed based on global random unitaries, which have a block structure inherited from the conservation of particles[22,51,146,162]. Finally, we note that the RM toolbox has also been transferred to photonic quantum devices implementing randomized global Clifford measurements using spatial light modulators and single-photon detection[163].

Learning about the quantum world using classical machines

For a classical machine to learn, store and manipulate any object of interest, one must construct a classical representation of the object. RMs provide a powerful set of tools for converting quantum systems into efficient classical representations that capture many aspects of the original quantum object. These tools bridge the gap between the quantum and classical worlds. Any algorithm originally designed for the purpose of learning in a classical world can now be used to learn about a quantum-mechanical world by using RMs as a quantum-to-classical converter.

Classical algorithms for learning in a classical world are capable of predicting what would happen in scenarios never encountered before[164–167]. Some well-known examples include outperforming the best human players in games, answering questions after reading an article, and identifying potential illnesses in the human body. By combining these classical algorithms with quantum-to-classical converters, we envisage that classical machines may one day achieve a powerful ability to predict the behaviour of the quantum world as well. Potential applications range from predicting the properties of exotic quantum systems that have not previously been realized in the laboratory[65] to designing better quantum computers and discovering new physical phenomena. Thus we anticipate that quantum-to-classical conversion enabled by the RM toolbox will have a vital role in unravelling some of nature's deepest secrets.

Published online: 02 December 2022

References

1. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
2. Altman, E. et al. Quantum simulators: architectures and opportunities. *PRX Quant.* **2**, 017003 (2021).
3. Gross, C. & Bloch, I. Quantum simulations with ultracold atoms in optical lattices. *Science* **357**, 995–1001 (2017).
4. Schäfer, F., Fukuhara, T., Sugawa, S., Takasu, Y. & Takahashi, Y. Tools for quantum simulation with ultracold atoms in optical lattices. *Nat. Rev. Phys.* **2**, 411–425 (2020).
5. Browaeys, A. & Lahaye, T. Many-body physics with individually controlled Rydberg atoms. *Nat. Phys.* **16**, 132–142 (2020).
6. Morgado, M. & Whitlock, S. Quantum simulation and computing with Rydberg-interacting qubits. *AVS Quant. Sci.* **3**, 023501 (2021).
7. Blatt, R. & Roos, C. F. Quantum simulations with trapped ions. *Nat. Phys.* **8**, 277–284 (2012).
8. Monroe, C. et al. Programmable quantum simulations of spin systems with trapped ions. *Rev. Mod. Phys.* **93**, 025001 (2021).
9. Kloeffel, C. & Loss, D. Prospects for spin-based quantum computing in quantum dots. *Annu. Rev. Condens. Matter Phys.* **4**, 51–81 (2013).
10. Burkard, G., Ladd, T. D., Nichol, J. M., Pan, A. & Petta, J. R. Semiconductor spin qubits. *Rev. Mod. Phys.* (in the press); preprint available at https://arxiv.org/abs/2112.08863.
11. Slussarenko, S. & Pryde, G. J. Photonic quantum information processing: a concise review. *Appl. Phys. Rev.* **6**, 041303 (2019).
12. Pelucchi, E. et al. The potential and global outlook of integrated photonics for quantum technologies. *Nat. Rev. Phys.* **4**, 194–208 (2021).
13. Kjaergaard, M. et al. Superconducting qubits: current state of play. *Annu. Rev. Condens. Matter Phys.* **11**, 369–395 (2020).

14. Alexeev, Y. et al. Quantum computer systems for scientific discovery. *PRX Quant.* **2**, 017001 (2021).
15. Flammia, S. T., Gross, D., Liu, Y.-K. & Eisert, J. Quantum tomography via compressed sensing: error bounds, sample complexity and efficient estimators. *New J. Phys.* **14**, 095022 (2012).
16. Haah, J., Harrow, A. W., Ji, Z., Wu, X. & Yu, N. Sample-optimal tomography of quantum states. In *STOC'16 — Proc. 48th Annual ACM SIGACT Symposium on Theory of Computing*, 913–925 (ACM, 2016).
17. O'Donnell, R. & Wright, J. Efficient quantum tomography. In *STOC'16 — Proc. 48th Annual ACM SIGACT Symposium on Theory of Computing*, 899–912 (ACM, 2016).
18. Aaronson, S.Shadow tomography of quantum states. In *STOC'18 — Proc. 50th Annual ACM SIGACT Symposium on Theory of Computing*, 325–338 (ACM, 2018).
19. Aaronson, S. & Rothblum, G. N. Gentle measurement of quantum states and differential privacy. In *STOC'19 — Proc. 51st Annual ACM SIGACT Symposium on Theory of Computing*, 322–333 (ACM, 2019).
20. Bădescu, C. & O'Donnell, R. Improved quantum data analysis. In *STOC '21 — Proc. 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 1398–1411 (ACM, 2021).
21. van Enk, S. J. & Beenakker, C. W. J. Measuring $\mathrm{Tr}\rho^n$ on single copies of ρ using random measurements. *Phys. Rev. Lett.* **108**, 110503 (2012).
22. Elben, A., Vermersch, B., Dalmonte, M., Cirac, J. I. & Zoller, P. Rényi entropies from random quenches in atomic Hubbard and spin models. *Phys. Rev. Lett.* **120**, 50406 (2018).
23. Elben, A., Vermersch, B., Roos, C. F. & Zoller, P. Statistical correlations between locally randomized measurements: a toolbox for probing entanglement in many-body quantum states. *Phys. Rev. A* **99**, 1–12 (2019).
24. Huang, H.-Y., Kueng, R. & Preskill, J. Predicting many properties of a quantum system from very few measurements. *Nat. Phys.* **16**, 1050–1057 (2020).
25. Paini, M. & Kalev, A. An approximate description of quantum states. Preprint at https://arxiv.org/abs/1910.10543 (2019).
26. Morris, J. & Dakić, B. Selective quantum state tomography. Preprint at https://arxiv.org/abs/1909.05880 (2019).
27. Knips, L. et al. Multipartite entanglement analysis from random correlations. *npj Quant. Inf.* **6**, 51 (2020).
28. Ketterer, A., Wyderka, N. & Gühne, O. Characterizing multipartite entanglement with moments of random correlations. *Phys. Rev. Lett.* **122**, 120505 (2019).
29. Chen, S., Yu, W., Zeng, P. & Flammia, S. T. Robust shadow estimation. *PRX Quantum* **2**, 030348 (2021).
30. Koh, D. E. & Grewal, S. Classical shadows with noise. *Quantum* **6**, 776 (2022).
31. Elben, A. et al. Cross-platform verification of intermediate scale quantum devices. *Phys. Rev. Lett.* **124**, 10504 (2020).
32. Zhu, D. et al. Cross-platform comparison of arbitrary quantum computations. Preprint at https://arxiv.org/abs/2107.11387 (2021).
33. Vermersch, B., Elben, A., Sieberer, L. M., Yao, N. Y. & Zoller, P. Probing scrambling using statistical correlations between randomized measurements. *Phys. Rev. X* **9**, 21061 (2019).
34. Joshi, M. K. et al. Quantum information scrambling in a trapped-ion quantum simulator with tunable range interactions. *Phys. Rev. Lett.* **124**, 240505 (2020).
35. Brydges, T. et al. Probing Rényi entanglement entropy via randomized measurements. *Science* **364**, 260–263 (2019).
36. Elben, A. et al. Mixed-state entanglement from local randomized measurements. *Phys. Rev. Lett.* **125**, 200501 (2020).
37. Elben, A. et al. Many-body topological invariants from randomized measurements in synthetic quantum matter. *Sci. Adv.* **6**, eaaz3666 (2020).
38. Gross, D., Audenaert, K. & Eisert, J. Evenly distributed unitaries: on the structure of unitary designs. *J. Math. Phys.* **48**, 052104 (2007).
39. Dankert, C., Cleve, R., Emerson, J. & Livine, E. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Phys. Rev. A* **80**, 012304 (2009).
40. Vitale, V. et al. Symmetry-resolved dynamical purification in synthetic quantum matter. *SciPost Phys.* **12**, 106 (2022).
41. Rath, A., Branciard, C., Minguzzi, A. & Vermersch, B. Quantum fisher information from randomized measurements. *Phys. Rev. Lett.* **127**, 260501 (2021).
42. Evans, T. J., Harper, R. & Flammia, S. T. Scalable Bayesian Hamiltonian learning. Preprint at https://arxiv.org/abs/1912.07636 (2019).
43. Cotler, J. & Wilczek, F. Quantum overlapping tomography. *Phys. Rev. Lett.* **124**, 100401 (2020).
44. Rath, A., van Bijnen, R., Elben, A., Zoller, P. & Vermersch, B. Importance sampling of randomized measurements for probing entanglement. *Phys. Rev. Lett.* **127**, 200503 (2021).
45. Banaszek, K., Cramer, M. & Gross, D. Focus on quantum tomography. *New J. Phys.* **15**, 125020 (2013).
46. Haah, J., Harrow, A. W., Ji, Z., Wu, X. & Yu, N. Sample-optimal tomography of quantum states. *IEEE Trans. Inf. Theory* **63**, 5628–5641 (2017).
47. O'Donnell, R. & Wright, J. Efficient quantum tomography II. In *STOC'17 — Proc. 49th Annual ACM SIGACT Symposium on Theory of Computing*, 962–974 (ACM, 2017).
48. Chen, S., Huang, B., Li, J., Liu, A. & Sellke, M. Tight bounds for state tomography with incoherent measurements. Preprint at https://arxiv.org/abs/2206.05265 (2022).
49. Cramer, M. et al. Efficient quantum state tomography. *Nat. Commun.* **1**, 149 (2010).
50. Torlai, G. et al. Neural-network quantum state tomography. *Nat. Phys.* **14**, 447–450 (2018).
51. Ohliger, M., Nesme, V. & Eisert, J. Efficient and feasible state tomography of quantum many-body systems. *New J. Phys.* **15**, 015024 (2013).
52. Sugiyama, T., Turner, P. S. & Murao, M. Precision-guaranteed quantum tomography. *Phys. Rev. Lett.* **111**, 160406 (2013).
53. Kueng, R., Rauhut, H. & Terstiege, U. Low rank matrix recovery from rank one measurements. *Appl. Comput. Harmon. Anal.* **42**, 88–116 (2017).
54. Guta, M., Kahn, J., Kueng, R. & Tropp, J. A. Fast state tomography with optimal error bounds. *J. Phys. A* **53**, 204001 (2020).
55. Eisert, J., Cramer, M. & Plenio, M. B. Colloquium: Area laws for the entanglement entropy. *Rev. Mod. Phys.* **82**, 277–306 (2010).
56. Vovrosh, J. & Knolle, J. Confinement and entanglement dynamics on a digital quantum computer. *Sci. Rep.* **11**, 11577 (2021).
57. Neven, A. et al. Symmetry-resolved entanglement detection using partial transpose moments. *npj Quantum Inf.* **7**, 152 (2021).
58. Zeng, B., Chen, X., Zhou, D.-L. & Wen, X.-G. *Quantum Information Meets Quantum Matter. From Quantum Entanglement to Topological Phases of Many-body Systems* (Springer, 2019).
59. Pollmann, F. & Turner, A. M. Detection of symmetry-protected topological phases in one dimension. *Phys. Rev. B* **86**, 125441 (2012).
60. Cian, Z.-P. et al. Many-body Chern number from statistical correlations of randomized measurements. *Phys. Rev. Lett.* **126**, 050501 (2021).
61. Satzinger, K. J. et al. Realizing topologically ordered states on a quantum processor. *Science* **374**, 1237–1241 (2021).
62. Kitaev, A. & Preskill, J. Topological entanglement entropy. *Phys. Rev. Lett.* **96**, 110404 (2006).
63. Levin, M. & Wen, X. G. Detecting topological order in a ground state wave function. *Phys. Rev. Lett.* **96**, 110405 (2006).
64. Flammia, S. T., Hamma, A., Hughes, T. L. & Wen, X. G. Topological entanglement Rényi entropy and reduced density matrix structure. *Phys. Rev. Lett.* **103**, 261601 (2009).
65. Huang, H.-Y., Kueng, R., Torlai, G., Albert, V. V. & Preskill, J. Provably efficient machine learning for quantum many-body problems. *Science* **377**, eabk3333 (2022).
66. Hartigan, J. A. & Wong, M. A. Algorithm AS 136: a *K*-means clustering algorithm. *J. R. Stat. Soc. C* **28**, 100–108 (1979).
67. Swingle, B. Unscrambling the physics of out-of-time-order correlators. *Nat. Phys.* **14**, 988–990 (2018).
68. Lewis-Swan, R. J., Safavi-Naini, A., Kaufman, A. M. & Rey, A. M. Dynamics of quantum information. *Nat. Rev. Phys.* **1**, 627–634 (2019).
69. Liu, H. & Sonner, J. Quantum many-body physics from a gravitational lens. *Nat. Rev. Phys.* **2**, 615–633 (2020).
70. Nie, X. et al. Detecting scrambling via statistical correlations between randomized measurements on an NMR quantum simulator. Preprint at https://arxiv.org/abs/1903.12237 (2019).
71. Qi, X.-L., Davis, E. J., Periwal, A. & Schleier-Smith, M. Measuring operator size growth in quantum quench experiments. Preprint at https://arxiv.org/abs/1906.00524 (2019).
72. Garcia, R. J., Zhou, Y. & Jaffe, A. Quantum scrambling with classical shadows. *Phys. Rev. Research* **3**, 033155 (2021).
73. Joshi, L. K. et al. Probing many-body quantum chaos with quantum simulators. *Phys. Rev. X* **12**, 011018 (2022).
74. Levy, R., Luo, D. & Clark, B. K. Classical shadows for quantum process tomography on near-term quantum computers. Preprint at https://arxiv.org/abs/2110.02965 (2021).
75. Flammia, S. T. & Liu, Y.-K. Direct fidelity estimation from few Pauli measurements. *Phys. Rev. Lett.* **106**, 230501 (2011).
76. da Silva, M. P., Landon-Cardinal, O. & Poulin, D. Practical characterization of quantum devices without tomography. *Phys. Rev. Lett.* **107**, 210404 (2011).
77. Fedorov, A., Steffen, L., Baur, M., da Silva, M. P. & Wallraff, A. Implementation of a Toffoli gate with superconducting circuits. *Nature* **481**, 170–172 (2011).
78. Lanyon, B. P. et al. Efficient tomography of a quantum many-body-system. *Nat. Phys.* **13**, 1158–1162 (2017).
79. Seshadri, A., Ringbauer, M., Monz, T. & Becker, S. Theory of versatile fidelity estimation with confidence. Preprint at https://arxiv.org/abs/2112.07947 (2021).
80. Seshadri, A., Ringbauer, M., Blatt, R., Monz, T. & Becker, S. Versatile fidelity estimation with confidence. Preprint at https://arxiv.org/abs/2112.07925 (2021).
81. Kliesch, M. & Roth, I. Theory of quantum system certification. *PRX Quantum* **2**, 010201 (2021).
82. Carrasco, J., Elben, A., Kokail, C., Kraus, B. & Zoller, P. Theoretical and experimental perspectives of quantum verification. *PRX Quantum* **2**, 010102 (2021).
83. Boixo, S. et al. Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **14**, 595–600 (2018).
84. Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
85. Wu, Y. et al. Strong quantum computational advantage using a superconducting quantum processor. *Phys. Rev. Lett.* **127**, 180501 (2021).
86. Bouland, A., Fefferman, B., Nirkhe, C. & Vazirani, U. On the complexity and verification of quantum random circuit sampling. *Nat. Phys.* **15**, 159–163 (2018).
87. Liu, Y., Otten, M., Bassirianjahromi, R., Jiang, L. & Fefferman, B. Benchmarking near-term quantum computers via random circuit sampling. Preprint at https://arxiv.org/abs/2105.05232 (2021).

88. Choi, J. et al. Emergent randomness and benchmarking from many-body quantum chaos. Preprint at https://arxiv.org/abs/2103.03535 (2021).
89. Cotler, J. S. et al. Emergent quantum state designs from individual many-body wavefunctions. Preprint at https://arxiv.org/abs/2103.03536 (2021).
90. Garrison, J. R. & Grover, T. Does a single eigenstate encode the full Hamiltonian? *Phys. Rev. X* **8**, 021026 (2018).
91. Qi, X.-L. & Ranard, D. Determining a local Hamiltonian from a single eigenstate. *Quantum* **3**, 159 (2019).
92. Bairey, E., Arad, I. & Lindner, N. H. Learning a local Hamiltonian from local measurements. *Phys. Rev. Lett.* **122**, 020504 (2019).
93. Bairey, E., Guo, C., Poletti, D., Lindner, N. H. & Arad, I. Learning the dynamics of open quantum systems from their steady states. *New J. Phys.* **22**, 032001 (2020).
94. Li, Z., Zou, L. & Hsieh, T. H. Hamiltonian tomography via quantum quench. *Phys. Rev. Lett.* **124**, 160502 (2020).
95. Kokail, C., van Bijnen, R., Elben, A., Vermersch, B. & Zoller, P. Entanglement Hamiltonian tomography in quantum simulation. *Nat. Phys.* **17**, 936–942 (2021).
96. Calabrese, P. & Cardy, J. Entanglement entropy and conformal field theory. *J. Phys. A* **42**, 504005 (2009).
97. Anshu, A., Arunachalam, S., Kuwahara, T. & Soleimanifar, M. Sample-efficient learning of interacting quantum systems. *Nat. Phys.* **17**, 931–935 (2021).
98. Haah, J., Kothari, R. & Tang, E. Optimal learning of quantum Hamiltonians from high-temperature Gibbs states. Preprint at https://arxiv.org/abs/2108.04842 (2021).
99. Rouzé, C. & França, D. S.Learning quantum many-body systems from a few copies. Preprint at https://arxiv.org/abs/2107.03333 (2021).
100. Yu, W., Sun, J., Han, Z. & Yuan, X. Practical and efficient Hamiltonian learning. Preprint at https://arxiv.org/abs/2201.00190 (2022).
101. Seif, A., Hafezi, M. & Liu, Y.-K. Compressed sensing measurement of long-range correlated noise. Preprint at https://arxiv.org/abs/2105.12589 (2021).
102. Hangleiter, D., Roth, I., Eisert, J. & Roushan, P. Precise Hamiltonian identification of a superconducting quantum processor. Preprint at https://arxiv.org/abs/2108.08319 (2021).
103. Blume-Kohout, R. et al. Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography. *Nat. Commun.* **8**, 14485 (2016).
104. Nielsen, E. et al. Gate set tomography. *Quantum* **5**, 557 (2021).
105. Merkel, S. T. et al. Self-consistent quantum process tomography. *Phys. Rev. A* **87**, 062119 (2013).
106. Madzik, M. T. et al. Precision tomography of a three-qubit donor quantum processor in silicon. *Nature* **601**, 348–353 (2022).
107. Samach, G. O. et al. Lindblad tomography of a superconducting quantum processor. Preprint at https://arxiv.org/abs/2105.02338 (2022).
108. Brieger, R., Roth, I. & Kliesch, M. Compressive gate set tomography. Preprint at https://arxiv.org/abs/2112.05176 (2021).
109. Evans, T. et al. Fast bayesian tomography of a two-qubit gate set in silicon. *Phys. Rev. Appl.* **17**, 024068 (2022).
110. van den Berg, E., Minev, Z. K., Kandala, A. & Temme, K. Probabilistic error cancellation with sparse Pauli–Lindblad models on noisy quantum processors. Preprint at https://arxiv.org/abs/2201.09866 (2022).
111. Flammia, S. T. Averaged circuit eigenvalue sampling. Preprint at https://arxiv.org/abs/2108.05803 (2021).
112. Harper, R., Flammia, S. T. & Wallman, J. J. Efficient learning of quantum noise. *Nat. Phys.* **16**, 1184–1188 (2020).
113. Cerezo, M. et al. Variational quantum algorithms. *Nat. Rev. Phys.* **3**, 625–644 (2021).
114. Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 4213 (2014).
115. O'Malley, P. J. J. et al. Scalable quantum simulation of molecular energies. *Phys. Rev. X* **6**, 031007 (2016).
116. Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017).
117. Kokail, C. et al. Self-verifying variational quantum simulation of lattice models. *Nature* **569**, 355–360 (2019).
118. Hadfield, C., Bravyi, S., Raymond, R. & Mezzacapo, A. Measurements of quantum Hamiltonians with locally-biased classical shadows. *Comm. Math. Phys.* **391**, 951–967 (2022).
119. Hillmich, S., Hadfield, C., Raymond, R., Mezzacapo, A. & Wille, R. Decision diagrams for quantum measurements with shallow circuits. In *International Conference on Quantum Computing and Engineering* 24–34 (IEEE, 2021).
120. Huang, H.-Y., Kueng, R. & Preskill, J. Efficient estimation of Pauli observables by derandomization. *Phys. Rev. Lett.* **127**, 030503 (2021).
121. Zhang, T. et al. Experimental quantum state measurement with classical shadows. *Phys. Rev. Lett.* **127**, 200501 (2021).
122. Zhao, A., Rubin, N. C. & Miyake, A. Fermionic partial tomography via classical shadows. *Phys. Rev. Lett.* **127**, 110504 (2021).
123. Yen, T.-C., Ganeshram, A. & Izmaylov, A. F. Deterministic improvements of quantum measurements with grouping of compatible operators, non-local transformations, and covariance estimates. Preprint at https://arxiv.org/abs/2201.01471 (2022).
124. Shlosberg, A. et al. Adaptive estimation of quantum observables. Preprint at https://arxiv.org/abs/2110.15339 (2021).
125. Kohda, M. et al. Quantum expectation-value estimation by computational basis sampling. *Phys. Rev. Research* **4**, 033173 (2022).
126. Havlíček, V. et al. Supervised learning with quantum-enhanced feature spaces. *Nature* **567**, 209–212 (2019).
127. Schuld, M. & Killoran, N. Quantum machine learning in feature Hilbert spaces. *Phys. Rev. Lett.* **122**, 040504 (2019).
128. McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812 (2018).
129. Huang, H.-Y. et al. Power of data in quantum machine learning. *Nat. Commun.* **12**, 2631 (2021).
130. Haug, T., Self, C. N. & Kim, M. Large-scale quantum machine learning. Preprint at https://arxiv.org/abs/2108.01039 (2021).
131. Zhou, Y., Zeng, P. & Liu, Z. Single-copies estimation of entanglement negativity. *Phys. Rev. Lett.* **125**, 200502 (2020).
132. Horodecki, P. Measuring quantum entanglement without prior state reconstruction. *Phys. Rev. Lett.* **90**, 167901 (2003).
133. Carteret, H. A. Noiseless quantum circuits for the peres separability criterion. *Phys. Rev. Lett.* **94**, 040502 (2005).
134. Carteret, H. A. Estimating the entanglement negativity from low-order moments of the partially transposed density matrix. Preprint at https://arxiv.org/abs/1605.08751 (2016).
135. Yu, X.-D., Imai, S. & Gühne, O. Optimal entanglement certification from moments of the partial transpose. *Phys. Rev. Lett.* **127**, 060504 (2021).
136. Liu, Z. et al. Detecting entanglement in quantum many-body systems via permutation moments. Preprint at https://arxiv.org/abs/2203.08391 (2022).
137. Tran, M. C., Dakic, B., Arnault, F., Laskowski, W. & Paterek, T. Quantum entanglement from random measurements. *Phys. Rev. A* **92**, 050301(R) (2015).
138. Tran, M. C., Dakic, B., Laskowski, W. & Paterek, T. Correlations between outcomes of random measurements. *Phys. Rev. A* **94**, 042302 (2016).
139. Ketterer, A., Wyderka, N. & Gühne, O. Entanglement characterization using quantum designs. *Quantum* **4**, 325 (2020).
140. Imai, S., Wyderka, N., Ketterer, A. & Gühne, O. Bound entanglement from randomized measurements. *Phys. Rev. Lett.* **126**, 150501 (2021).
141. Ketterer, A., Imai, S., Wyderka, N. & Gühne, O. Statistically significant tests of multiparticle quantum correlations based on randomized measurements. Preprint at https://arxiv.org/abs/2012.12176 (2021).
142. Knips, L. A moment for random measurements. *Quantum Views* **4**, 47 (2020).
143. Leone, L., Oliviero, S. F. E. & Hamma, A. Stabilizer Rényi entropy. *Phys. Rev. Lett.* **128**, 050402 (2022).
144. Oliviero, S. F. E., Leone, L., Hamma, A. & Lloyd, S. Measuring magic on a quantum processor. Preprint at https://arxiv.org/abs/2204.00015 (2022).
145. Feldman, N., Kshetrimayum, A., Eisert, J. & Goldstein, M. Entanglement estimation in tensor network states via sampling. *PRX Quantum* **3**, 030312 (2022).
146. Vermersch, B., Elben, A., Dalmonte, M., Cirac, J. I. & Zoller, P. Unitary n -designs via random quenches in atomic Hubbard and spin models: application to the measurement of Rényi entropies. *Phys. Rev. A* **97**, 023604 (2018).
147. van den Berg, E., Minev, Z. K. & Temme, K. Model-free readout-error mitigation for quantum expectation values. *Phys. Rev. A* **105**, 032620 (2022).
148. Gottesman, D. The Heisenberg representation of quantum computers. Preprint at https://arxiv.org/abs/quant-ph/9807006 (1998).
149. Hu, H.-Y., Choi, S. & You, Y.-Z. Classical shadow tomography with locally scrambled quantum dynamics. Preprint at https://arxiv.org/abs/2107.04817 (2021).
150. Bu, K., Koh, D. E., Garcia, R. J. & Jaffe, A. Classical shadows with Pauli-invariant unitary ensembles. Preprint at https://arxiv.org/abs/2202.03272 (2022).
151. Nakata, Y., Hirche, C., Koashi, M. & Winter, A. Efficient quantum pseudorandomness with nearly time-independent Hamiltonian dynamics. *Phys. Rev. X* **7**, 021006 (2017).
152. Koenig, R. & Smolin, J. A. How to efficiently select an arbitrary Clifford group element. *J. Math. Phys.* **55**, 122202 (2014).
153. Notarnicola, S. et al. A randomized measurement toolbox for Rydberg quantum technologies. Preprint at https://arxiv.org/abs/2112.11046 (2021).
154. Ringbauer, M. et al. A universal qudit quantum processor with trapped ions. Preprint at https://arxiv.org/abs/2109.06903 (2021).
155. García-Pérez, G. et al. Learning to measure: adaptive informationally complete generalized measurements for quantum algorithms. *PRX Quant.* **2**, 040342 (2021).
156. Nguyen, H. C., Bönsel, J. L., Steinberg, J. & Gühne, O. Optimising shadow tomography with generalised measurements. Preprint at https://arxiv.org/abs/2205.08990 (2022).
157. Renes, J. M., Blume-Kohout, R., Scott, A. J. & Caves, C. M. Symmetric informationally complete quantum measurements. *J. Math. Phys.* **45**, 2171–2180 (2004).
158. Scott, A. J. Tight informationally complete quantum measurements. *J. Phys. A* **39**, 13507–13530 (2006).
159. Fischer, L. E. et al. Ancilla-free implementation of generalized measurements for qubits embedded in a qudit space. *Phys. Rev. Research* **4**, 033027 (2022).
160. Stricker, R. et al. Experimental single-setting quantum state tomography. Preprint at https://arxiv.org/abs/2206.00019 (2022).
161. National Academies of Sciences, Engineering, and Medicine. *Manipulating Quantum Systems: An Assessment of Atomic, Molecular, and Optical Physics in the United States* (National Academies Press, 2020); https://www.nap.edu/catalog/25613/manipulating-quantum-systems-an-assessment-of-atomic-molecular-and-optical.

162. Naldesi, P. et al. Fermionic correlation functions from randomized measurements in programmable atomic quantum devices. Preprint at https://arxiv.org/abs/2205.00981 (2022).
163. Struchalin, G., Zagorovskii, Y. A., Kovlakov, E., Straupe, S. & Kulik, S. Experimental estimation of quantum state properties from classical shadows. *PRX Quantum* **2**, 010307 (2021).
164. LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998).
165. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning. Adaptive Computation and Machine Learning* (MIT Press, 2016); http://www.deeplearningbook.org/.
166. Silver, D. et al. Mastering the game of go without human knowledge. *Nature* **550**, 354–359 (2017).
167. Jumper, J. et al. Highly accurate protein structure prediction with alphafold. *Nature* **596**, 583–589 (2021).

Acknowledgements

A.E. acknowledges funding by the German National Academy of Sciences Leopoldina under grant no. LPDS 2021-02 and by the Walter Burke Institute for Theoretical Physics at Caltech. J.P. acknowledges funding from the US Department of Energy Office of Science, Office of Advanced Scientific Computing Research (DE-NA0003525, DE-SC0020290), and the National Science Foundation (NSF) (PHY-1733907). The Institute for Quantum Information and Matter is an NSF Physics Frontiers Center. B.V. acknowledges funding from the French National Research Agency (ANR-20-CE47-0005, JCJC project QRand) and from the Austrian Science Foundation (FWF, P 32597 N). P.Z. acknowledges support by the US Air Force Office of Scientific Research (AFOSR) via IOE grant no. FA9550-19-1-7044 LASCEM, by the European Union's Horizon 2020 research and innovation programme under grant agreement no. 817482 (PASQuanS), and by the Simons Collaboration on Ultra-Quantum Matter, which is a grant from the Simons Foundation (651440).

Author contributions

The authors contributed equally to all aspects of the article.

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/s42254-022-00535-2.

Correspondence should be addressed to Peter Zoller.

Peer review information *Nature Reviews Physics* thanks Barbara Terhal and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.