





Revealing microcanonical phases and phase transitions of strongly correlated systems via time-averaged classical shadows

Gaurav Gyawali ^{1,2,*} Mabrur Ahmed ^{3,*} Eric W. Aspling ³ Luke Ellert-Beck ^{4,5} and Michael J. Lawler^{3,1,2}

¹*Department of Physics, Cornell University, Ithaca, New York 14853, USA*

²*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*

³*Department of Physics, Applied Physics, and Astronomy, Binghamton University, Binghamton, New York 13902, USA*

⁴*School of Physics and Applied Physics, Southern Illinois University, Carbondale, Illinois 62901, USA*

⁵*Department of Physics, University of Rhode Island, Kingston, Rhode Island 02881, USA*



(Received 20 December 2022; accepted 20 November 2023; published 13 December 2023)

Quantum computers and simulators promise to enable the study of strongly correlated quantum systems. Yet, surprisingly, it is hard for them to compute ground states. They can, however, efficiently compute the dynamics of closed quantum systems. We propose a method to study the quantum thermodynamics of strongly correlated electrons from quantum dynamics. We define time-averaged classical shadows (TACS) and prove it is a classical shadow (CS) of the von Neumann ensemble, the time-averaged density matrix. We then show that the diffusion maps, an unsupervised machine learning algorithm, can efficiently learn the phase diagram and phase transition of the one-dimensional transverse field Ising model both for ground states using CS *and state trajectories* using TACS. It does so from state trajectories by learning features that appear to be susceptibility and entropy from a total of 90 000 shots taken along a path in the microcanonical phase diagram. Our results suggest a low number of shots from quantum simulators can produce quantum thermodynamic data with a quantum advantage.

DOI: [10.1103/PhysRevB.108.235141](https://doi.org/10.1103/PhysRevB.108.235141)

I. INTRODUCTION

Simulation of strongly correlated electrons in the context of quantum chemistry and condensed matter physics is one of the potential areas in which quantum computers will have a significant advantage over their classical counterparts [1–3]. Strong interaction between the ostensibly simple electrons can give rise to novel phases, including high-temperature superconductivity [4,5], strange metallic behavior [6], fractional excitations [7], and quantum spin liquids [8]. Condensed matter physics aims to understand these novel behaviors by studying their phase diagrams and phase-defining features. However, failure of perturbation theory and exponential scaling of the Hilbert space for strongly correlated electrons presents a formidable challenge to classical simulation methods such as exact diagonalization, density matrix renormalization group (DMRG) [9–11], quantum Monte Carlo [12], and dynamical mean-field theory [13]. Whereas, this same challenge provides an exciting opportunity for near-term quantum computers.

Harnessing the power of a quantum computer to simulate quantum systems [14] requires (i) algorithms that can be executed in a reasonable time and (ii) the ability to learn from quantum experiments without exponentially many measurements. Studying the phases via ground state preparation is a QMA-complete problem [15–19], which cannot be carried out in a reasonable time, even with quantum resources. However, performing dynamics on a quantum state is known to be a

BQP-hard problem [20,21], possible within polynomial time. Likewise, it has been shown that shadow tomography [22] methods such as classical shadows (CS) [23–25] are effective at predicting properties using very few measurements. Thus, if we could combine dynamics simulations and classical shadows, we would have an efficient algorithm to simulate condensed matter systems.

We need to prepare low-energy initial states to use dynamics to simulate the dynamics of condensed matter systems and exploit quantum ergodicity [26]. Although preparing ground states of local Hamiltonian on a physical lattice is a challenging problem on a quantum device, it is always possible to prepare some low-energy state with a constant-depth circuit [27,28]. Ergodicity then provides a link between statistical averages and time averages obtained from the dynamics of the low-energy state. It is important that the observables of interest, such as the order parameter, equilibrate before the qubits decohere. Nevertheless, rapid equilibration for most local observables is a feature shared by many interacting quantum systems [29–31]. Thus equilibrium dynamics of low-energy states appear to be a promising route to studying equilibrium quantum phases and phase transitions.

In this manuscript, we present an algorithm for identifying phase diagrams and phase transitions of strongly correlated systems *motivated by how physical quantum systems operate*. It consists of (i) identifying an initial state, (ii) generating state trajectories by evolving this state in time, (iii) using shadow tomography to convert the quantum state to classical data, and (iv) applying unsupervised machine learning methods to discover phases of matter and their phase transitions [32–41]. A schematic overview of our approach is shown in Fig. 1.

*These authors contributed equally to this work.

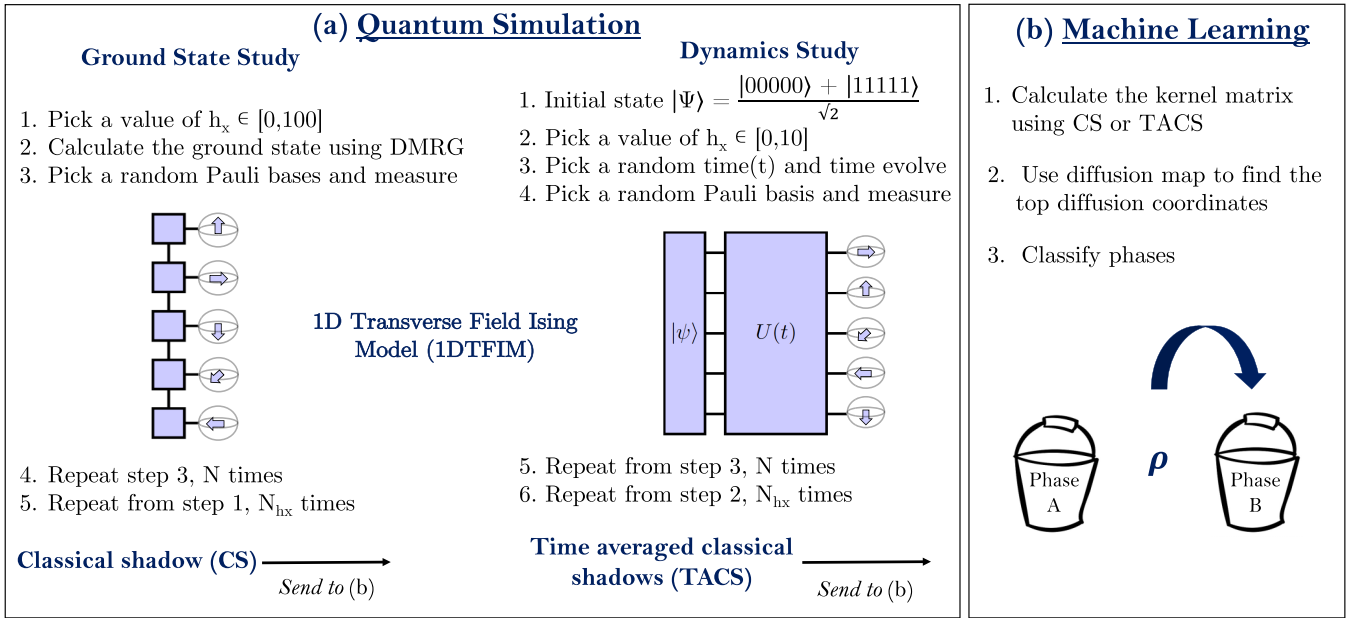


FIG. 1. Schematic overview of our study. (a) Classical shadows (CS) of ground states and time-averaged classical shadows (TACS) from dynamics of a time-reversal invariant GHZ state are generated using quantum simulation. (b) The classical data from quantum simulation is then fed into diffusion maps, an unsupervised machine learning algorithm to learn the phases.

We obtain numerical results using diffusion maps [42,43], an unsupervised machine learning (UL) algorithm [44,45] to learn phase features from unlabeled data. First, we benchmark diffusion maps on CS data from ground states of a 100-qubit one-dimensional transverse field Ising model (1DTFIM) simulation. Diffusion maps identifies the magnetic phase transition, its continuous nature, and the magnetization behavior as a function of the magnetic field (see Sec. II)—another machine-learning-from-CS success story [23–25,46]. Generalizing CS to time-averaged CS (TACS), a shadow-tomographic [22] representation of the time-averaged density matrix [26], we then show, in Sec. III, on a 20 qubit 1DTFIM, diffusion maps also identify the quantum critical region and crossovers along a path in the microcanonical phase diagram from a total of 90 000 shots on state trajectories. Diffusion maps do so efficiently by learning features from TACS that appear to be susceptibility and entropy. Hence, we can efficiently study the phases and phase transitions of strongly correlated electrons by quantum-simulating state trajectories.

II. VON NEUMANN’S MICROCANONICAL ENSEMBLE

A central goal of quantum computing is to build qubits that are completely isolated from their environment. While this is not the case today, the current development of quantum error correction techniques [47,48] promises to have noise-resilient logical qubits in the future. Simulating quantum systems on a quantum computer will therefore take place *within the microcanonical ensemble*. But quantum microcanonical dynamics, the evolution of a closed quantum system under Schrödinger’s equation, does not directly produce the microcanonical ensemble.

Following von Neumann’s 1929 paper [26,49] on the quantum ergodic theorem, it is straightforward to derive a link between time averages and statistical averages using density

matrices. Assuming we start from an initial state $|\psi(0)\rangle$ and evolve under a Hamiltonian H via a quantum circuit algorithm to $|\psi(t)\rangle$, the equilibrium distribution is captured by the von Neumann ensemble, the time average of the density matrix

$$\begin{aligned} \rho_{vN} &= \frac{1}{T} \int_0^T dt |\psi(t)\rangle \langle \psi(t)| \xrightarrow{T \rightarrow \infty} \omega \\ &= \sum_n P_n |\psi(0)\rangle \langle \psi(0)| P_n, \end{aligned} \quad (1)$$

where P_n is a projector onto the n th degenerate subspace of the energy eigenvalues, i.e., $P_n = \sum_{k \in n} |E_k\rangle \langle E_k|$. The $T \rightarrow \infty$ limit, obtained in exponential time [26,31,50,51], results in equilibration of all observables. Existence of ρ_{vN} results in the ergodic principle that time averages of observables should be captured by the statistical averaging with respect to ρ_{vN} . Specifically, in the Schrödinger picture,

$$\langle O \rangle_T = \frac{1}{T} \int_0^T dt \langle \psi(t) | O | \psi(t) \rangle \quad (2)$$

$$= \frac{1}{T} \int_0^T dt \text{Tr}(|\psi(t)\rangle \langle \psi(t)| O) \quad (3)$$

$$= \text{Tr}(\rho_{vN} O) \xrightarrow{T \rightarrow \infty} \text{Tr}(\omega O), \quad (4)$$

where again the $T \rightarrow \infty$ limit produces equilibration in the sense that they are indistinguishable with respect to observable O [31]. Thus the time-averaged density matrix is a link between time averages and statistical averages governed by the von Neumann ensemble ρ_{vN} , a link that holds regardless of whether the system equilibrates.

The connection to Boltzmann’s microcanonical ensemble, obtained by quantizing the classical microcanonical ensemble, is achieved by taking the thermodynamic limit, measuring only coarse-grained observables, requiring nondegenerate energy level spacings/gaps, and considering “typical” initial

states. The coarse-grained observables used by von Neumann were a commuting set of generators of global symmetries, a restriction more recently generalized [52] as part of the development of quantum thermodynamic resource theories [53]. The typical initial states were the first recorded use of typicality arguments [49]. Under these circumstances, von Neumann obtained

$$\omega \sim \rho_{mc} = \frac{1}{\Omega} \sum_{E < E_n < E + \Delta} |E_n\rangle\langle E_n|. \quad (5)$$

Namely, for the purposes of computing $\text{Tr}(O\omega)$, there is no difference between using ω and ρ_{mc} , a maximally mixed state within an energy window $[E, E + \Delta]$ containing Ω states. Von Neumann extended this claim to a second ensemble, $\rho'_{vN} = |\psi(T)\rangle\langle\psi(T)|$, that also satisfies $\rho'_{vN} \sim \rho_{mc}$. The equivalence to ρ_{mc} is also readily proven with the seemingly stronger requirement that each eigenstate satisfies the eigenstate thermalization hypothesis [54–56]. So, in this way, ρ_{vN} reproduces ρ_{mc} .

A central new ingredient in von Neumann's approach to describing the microcanonical ensemble is the initial state, which is never fully forgotten under a unitary evolution. For each initial condition, it's necessary to check whether the time series was run long enough for relevant observables to reach equilibrium. The same observable will equilibrate at different times depending on the initial conditions. It turns out, this time depends on the effective dimension given by $d_{\text{eff}} = 1 / \sum p_k^2$, where $p_k = |\langle\psi||E_k\rangle|^2$ [50,51]. Here, the overlaps p_k measure how many energy eigenstates have a significant weight in $|\psi\rangle$. Remarkably, a large effective dimension results in rapid equilibration. Furthermore, there is a bound on the equilibration time given by the second largest p_k if it is significantly smaller than $1/d_{\text{eff}}$ [50]. The distribution of the p_k 's likely also affects equilibration times [57]. These arguments suggest we choose initial conditions that exhibit a small overlap with most energy levels and or a macroscopic occupation of a single energy level.

Because we will use a machine learning method as part of our study of von Neumann's microcanonical ensemble, we need to compare it to what we already know to validate the approach. In the next section, Sec. III, we will turn to a CS data-driven ground state before continuing to our TACS data-driven thermodynamic study in Sec. V.

III. GROUND STATE DATA

To verify our approach to phase classification and phase-defining feature identification, we first apply it to ground states of the ferromagnetic IDTFIM defined by the Hamiltonian:

$$H_{\text{IDTFIM}} = - \sum_{\langle i,j \rangle} Z_i Z_j + h_x \sum_i X_i, \quad (6)$$

where $\langle \cdot \rangle$ denotes nearest neighbors, Z_i is the Pauli Z operator, and h_x is a parameter proportional to the transverse magnetic field. At zero temperature, this model has a ferromagnetic phase for $|h_x| < 1$, and a paramagnetic phase for $|h_x| > 1$. The ground state study uses diffusion maps to detect the second-order phase transition at $h_x = 1$ for a 100 site IDTFIM. The ground states were generated using density

matrix renormalization group (DMRG) [9–11] with ITENSOR package [58]. Training datasets were generated using two kinds of measurements on these ground states: (i) computational basis measurements to obtain the Z dataset and (ii) measurements on a random Pauli basis to obtain the CS dataset.

A. Classical shadows

Obtaining any meaningful information from a quantum computer requires performing measurements on a quantum state, which is destructive to the quantum information by nature. Since the dimension of the Hilbert space increases exponentially in the number of qubits, a naive strategy to learn the state requires an exponentially large number of copies. Aaronson [22] introduced an alternative method using the notion of shadow tomography, an approximate classical description of the quantum state, in which M properties of a quantum state can be estimated with error ϵ by only $O(\frac{\log^4 M}{\epsilon^2})$ copies of the state. We can think of a shadow as an approximation of a quantum state ρ by summing over measurement outcomes x , obtained by performing measurements on bases b for a quantum state x , i.e.,

$$S[\rho] = \sum_{b,x} P(b) \mathcal{P}_{b,x} \rho \mathcal{P}_{b,x}, \quad (7)$$

where $\mathcal{P}_{b,x}$ is a projector onto the measurement outcome x on basis b , and $P(b)$ is the probability of choosing b .

Based on this notion, Huang *et al.* [23–25] developed an algorithm called classical shadows and showed that it is highly successful at learning the properties of a many-body system. Two kinds of measurement protocols were proposed to construct classical shadows- (i) random Clifford measurements on the entire Hilbert space and (ii) random single-qubit Pauli measurements. Protocol (ii) results in very shallow measurement circuits and thus is more suitable for the NISQ-ERA [59] hardware. After measuring each of the qubits in some random Pauli basis X , Y , or Z with outcomes ± 1 , the post-measurement wavefunction is given by the product state $|s^{(n)}\rangle = \bigotimes_{l=1}^L |s_l^{(n)}\rangle$. Here, $|s_l^{(n)}\rangle \in \{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |i+\rangle, |i-\rangle\}$ is a Pauli basis state to which the l th qubit has collapsed. The classical shadow $S_N[\rho]$ is obtained by summing over N such randomized measurement outcomes as follows [also see Fig. 2(b)]

$$S_N[\rho] = \frac{1}{N} \sum_{n=1}^N |s^{(n)}\rangle \langle s^{(n)}| \quad (8)$$

$$= \frac{1}{N} \sum_{n=1}^N |s_1^{(n)}\rangle \langle s_1^{(n)}| \otimes \cdots \otimes |s_L^{(n)}\rangle \langle s_L^{(n)}|. \quad (9)$$

The underlying quantum state ρ can be approximated by adding the reduced classical shadows [see Fig. 2(a)]. This sum simplifies to the following expression from Ref. [23–25]

$$\rho \approx \sigma_N(\rho) = \frac{1}{N} \sum_{n=1}^N \sigma_1^{(n)} \otimes \cdots \otimes \sigma_L^{(n)}, \quad (10)$$

where

$$\sigma_l^{(n)} = 3|s_l^{(n)}\rangle\langle s_l^{(n)}| - \mathbb{I}. \quad (11)$$

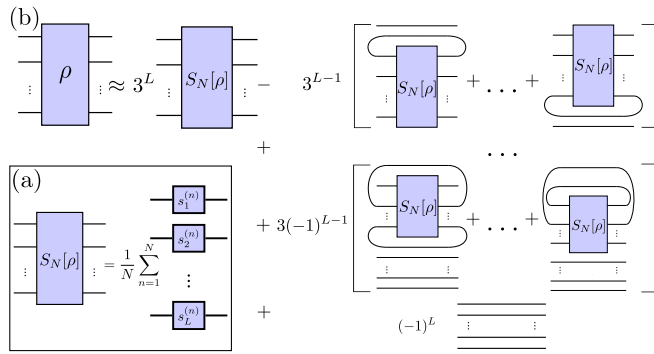


FIG. 2. Diagrammatic description of classical shadows showing a linear relationship between $S_N[\rho]$ and estimator $\sigma_N[\rho] \approx \rho$. (a) The full-density matrix ρ can be approximated by summing over reduced classical shadows with a coefficient that grows exponentially in the number of remaining qubits. (b) A classical shadow $S_N[\rho]$ is obtained by summing over N measurement outcomes on random Pauli bases. For a given N , a reduced density matrix, that involves smaller coefficients in the expansion, can be approximated more accurately compared to the full density matrix.

The definition of $S_N(\rho)$ presented above is different from Refs. [23,24] which defines it to be the dataset of shots itself and not the density matrix obtained from these shots. But both definitions are complete for Fig. 2 (whose derivation from tensor network diagrams is presented in Appendix A) shows the density matrix $S_N(\rho)$ defined above is linearly related to the estimator $\sigma_N(\rho)$ of the quantum state ρ obtained by Refs. [23–25]. Hence, the two definitions are informationally equivalent.

Although estimating the exact density matrix requires $N \rightarrow \infty$, we still desire to predict various linear as well as nonlinear functions of ρ [e.g., $\text{Tr}(O\rho)$ and $\text{Tr}(\rho \ln(\rho))$] respectively. This can be achieved with $N \propto \ln(L)4^k/\epsilon^2$ copies of the state, where k is the locality of operator O [46]. It was shown in Ref. [24] that classical machine learning algorithms can efficiently predict the ground state properties of gapped Hamiltonians in finite spatial dimensions after learning the classical shadows from a training set. An example of interest is classifying the quantum phases of matter. Classifying the symmetry-breaking phases is conceptually simple because it involves calculating $\text{tr}(\rho O)$ for some k -local observable O , such that $\text{tr}(\rho O) \geq 1 \forall \rho \in \text{phase A}$ and $\text{tr}(\rho O) \leq -1 \forall \rho \in \text{phase B}$. Here, phase A includes all the symmetry-preserving states whereas B includes the symmetry-breaking ones.

In contrast to classifying symmetry-breaking phases, capturing continuous phase transitions and classifying topological phases involves nonlinear-in- ρ observables like critical exponents and entropy, which are harder to estimate than linear observables. Learning such nonlinear functions requires an expressive ML model. A central object in kernel-based ML is the kernel function, a local similarity measure in the feature space where the samples live. Ref. [24] proposed a kernel based on mapping from classical shadows to a high-dimensional feature space that includes the polynomial expansion of many-body reduced density matrices. Learning nonlinear functions requires access to k -body reduced density

matrices, where k may be large, but with enough shots, classical shadows can accomplish this. Using such a kernel, Ref. [24] found a rigorous guarantee that a classical ML algorithm can efficiently classify phases of matter, including the topological phases. We will employ this kernel to study the continuous phase transition in the 1DTFIM.

B. Machine learning method: diffusion maps

Let's now turn to the final step in our approach: applying an unsupervised machine learning method called diffusion maps [42,43] to extract features from the shadow tomography data. A diffusion map is a nonlinear dimensionality reduction technique that relies on learning the underlying manifold from which the data points have been generated. Recently, this method was used to identify phases and phase transitions in systems with complex order parameters, which are difficult to learn using linear methods (such as principal component analysis (PCA)). Examples of such phase learning studies include: topological phases and phase transitions [39], incommensurate phases, and many-body localized phases in quantum systems [41], and topological quantum phase transition [40].

In the application of diffusion maps, we imagine a random walk on a dataset $X(x_1, x_2, \dots, x_S)$, where the x_s are estimators $\sigma_N(\rho)$ (see Sec. III A) of density matrices ρ obtained from different points in the phase diagram. These estimators are collections of shots arranged into $N \times L$ matrices; here, L is the number of qubits. The transition probability $P(j|i)$ of jumping from x_i to x_j in a single “time step” is proportional to the kernel function $k(x_i, x_j)$, a non-negative similarity measure between the two data points. Here we use the classical shadow kernel function prescribed in Ref. [24], defined to be for two points x and \tilde{x} in the dataset

$$\begin{aligned} k^{(\text{shadow})}(x, \tilde{x}) &= k^{(\text{shadow})}(\sigma_N(\rho), \sigma_N(\tilde{\rho})) \\ &= \exp\left(\sum_{n, n'=1}^N \frac{\tau}{N^2} \exp\left(\frac{\gamma}{L} \sum_{l=1}^L \text{Tr}[\sigma_l^{(n)} \tilde{\sigma}_l^{(n')}] \right)\right). \end{aligned} \quad (12)$$

γ and τ are hyperparameters. This kernel measures the local similarity between x and \tilde{x} by comparing the trace distance between the CS estimates of all k -reduced density matrices. For the diagonal components ($x = \tilde{x}$) of the kernel matrix, the trace distance between the k -reduced density matrices is the 2nd Renyi entropy. We then construct a transition probability matrix P such that

$$P(j|i) = \frac{k^{(\text{shadow})}(i, j)}{\sum_s k^{(\text{shadow})}(i, s)}. \quad (13)$$

After t time steps of the random walk, the transition probabilities are given by the matrix P^t , where P_{ij}^t gives the probability of going from x_i to x_j in t time steps, it's a sum of the probabilities associated with all of the possible paths to go from x_i to x_j in t time steps. As t increases, the diffusion process unfolds, where data points situated along the overall geometric structure of the dataset become more strongly connected because of the abundance of strongly connected intermediate points along the way.

Given this random walk, we can define a ‘‘diffusion distance’’ to quantify this idea of connectivity between two data points:

$$D_t(x_i, x_j)^2 = \sum_{s=1}^S |P_{is}^t - P_{sj}^t|^2, \quad (14)$$

where the bigger the diffusion distance, the weaker the connection between them. This allows us to map the data points onto a new ‘diffusion space’ so that the diffusion distance in data space is equal to the Euclidean distance in this new space. Following Ref. [42], we will do so with the map:

$$x_i \rightarrow y_i = [\lambda_1^t \psi_1(i), \lambda_2^t \psi_2(i), \dots, \lambda_{S-1}^t \psi_{S-1}(i)] \quad (15)$$

where λ_k and ψ_k are eigenvalues and right eigenvectors of the matrix P^t , $\psi_k(i)$ is the i th element of the k th eigenvector. Then the diffusion distance is,

$$D_t^2(x_i, x_j) = |y_i - y_j|^2 = \sum_{k=1}^{S-1} (\lambda_k^t)^2 [\psi_k(i) - \psi_k(j)]^2. \quad (16)$$

Plotting data in this new space provides an intuitive geometric picture of the data manifold.

The map provides several features we can exploit when interpreting the data. In Eqs. (15) and (16) the $k = 0$ component is ignored because the leading eigenvector, $\psi_0(i) = \frac{1}{\sqrt{S}}$, $\lambda_0 = 1$, is constant for all i by the Perron-Frobenius theorem, but this constant eigenvector impacts the other eigenvectors. We note that adding a constant term $\psi_k(i) \rightarrow \psi_k(i) + C$ to all other eigenvectors preserves the diffusion distance. Additionally, Eq. (16) suggests a dimensionality reduction, as the terms with bigger λ_k will dominate the sum increasingly as $t \rightarrow \infty$. So plotting the data x_i in the truncated space $[\lambda_1^t \psi_1(i), \lambda_2^t \psi_2(i), \lambda_M^t \dots, \psi_M(i)]$ with M determined by keeping only the significant eigenvalues $\lambda_m \gg \lambda_{M+1}$, implies they are accurately separated by distance $D_t(x_i, x_j)$. Lastly, we see that t is arbitrary, and choosing a different t rescales the lengths of each component of the vector y . Hence, the data exists on a hyperplane in Euclidean space up to a certain shift in the origin and a one-parameter rescaling of the axes.

So, using the properties of the diffusion space, we can define *diffusion coordinates* $dc1(i) = A(\psi_1(i) + C)$, $dc2(i) = B(\psi_2(i) + C)$, \dots that map the data x_i onto a Euclidean space that through the choice of constants A, B, C, \dots allow us to interpret the coordinates of each point and visualize the geometry of the data.

IV. PHASE CLASSIFICATION OF GROUND STATES

For the ground state study, we used two datasets, one the computational basis measurements and the other generated via CS tomography. The first (Z dataset) contains qubit measurements only along the Z axis [in Eq. (7), $P(b) = 1$ for $b = Z$, $P(b) = 0$ for all other b]. While the other (CS dataset) has randomized Pauli measurements using the CS method [$P(b) = \frac{1}{3}$ for $b \in \{X, Y, Z\}$]. Both of them contain 200, 100-spin 1DTFIM ground state shots for each state obtained from different h_x values (h_x ranging from 0.1 to 100). Since the Z magnetization is the order parameter for TFIM, the

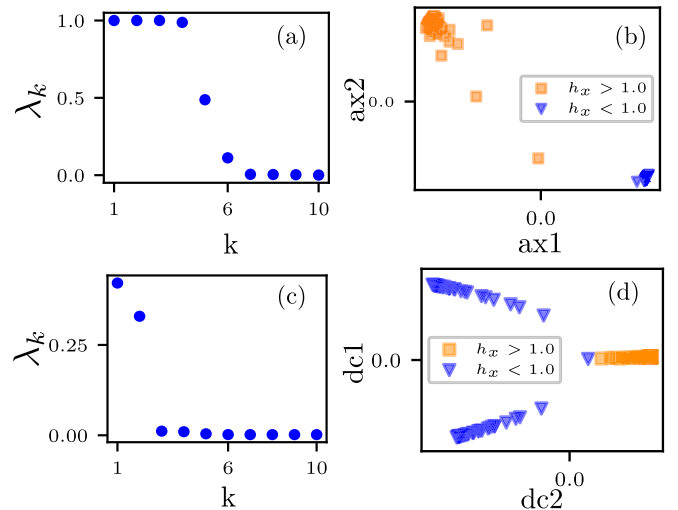


FIG. 3. Learning phases from ground state (CS and Z) data. (a) and (c) show the ten largest eigenvalues of the P matrix (excluding the trivial $k = 0$) for (a) Z data and (c) CS data datasets. (b) Z data points in 5D diffusion space visualized in 2D, using metric MDS. Clustering clearly emerges based on the two phases of the model. (d) CS data points in 2D diffusion space, the figure reveals the symmetry-breaking phase transition. In this case, there is a direct correlation of relevant parameters, the z magnetization and the h_x values, with machine-learned diffusion coordinates $dc1$ and $dc2$, respectively. Notice the CS data show one cluster in (d) with nontrivial geometry associated with the critical point, while the Z data show two clusters in (b) with trivial geometry and no understanding of the critical point.

unsupervised learning algorithm should be able to learn the phases of the model from the Z dataset. Using this knowledge, we compare the Z dataset and the CS data to see if the UL algorithm can successfully identify phases in each case and if so how it does so.

By deploying diffusion maps armed with the shadows kernel function in Eq. (12), utilized for both data sets as our UL model, we are able to identify the phases from both the Z data and the CS data. In both cases, we set $\tau = 1$, $\gamma = 1$. For the Z data, we chose the first five nontrivial eigenvectors as the diffusion space basis vectors because the P -matrix, given in Eq. (13), eigenvalue spectrum shows the first five eigenvalues to be larger than others as demonstrated in Fig. 3(a). Mapping the states from Z data onto this five-dimensional diffusion space, we found that clear clustering emerges based on the phases of the states. We used multidimensional scaling (MDS), a dimensionality reduction method [60] that seeks to preserve point-to-point distances, to project these states onto a 2D plane. We see a clear separation of the two phases even on this 2D reduced space [see Fig. 3(b)], indicating the machine’s success in identifying the two phases.

From the CS data, the unsupervised learning algorithm was also able to learn about the phases and the underlying parameters of the model. We can see in Fig. 3(c) that the P matrix has two nontrivial eigenvalues larger than the rest. The eigenvectors corresponding to these two eigenvalues are the basis vectors of the reduced diffusion space. Figure 3(d) shows all of the ground state classical shadows projected onto

this two-dimensional plane. It shows three groups- the top left and the bottom left are the all-up and all-down states, whereas the group on the center-right are states in the disordered phase. This closely resembles the spontaneous symmetry-breaking phase transition of 1DTFIM [61]. The learned diffusion coordinates (dc1 and dc2) have direct correlations with the magnetization $\langle M_z \rangle$ and the field values h_x respectively, as shown in Fig. 8 and demonstrated in Appendix C.

The above-mentioned clustering is dependent on the number of snapshots N for a given state. However, the dimensionality reduction and the subsequent clustering will settle down after a minimum value of N has been reached (N_c). Increasing the value of N beyond that point does not change the results in any significant way. We find N_c (Z data) $<$ N_c (CS data), so it is easier for the algorithm to learn the phase space structure of the Z data (hence fewer snapshots are required) than the CS data.

A striking feature of the diffusion map results presented in Fig. 3 is the geometry it reveals about the data. In the case of Z data, it finds the data is separated into two distinct clusters [Fig. 3(b)], while in the CS data case, it finds only one cluster but that this cluster has a nontrivial geometry with three curves meeting at the critical point [Fig. 3(d)]. This geometry is directly a consequence of visualizing of the data space through the lens of the kernel function that defines distances between data points via Eqs. (13) and (14). Another kernel function might see the same CS data as separate clusters. Hence, it is a striking feature of the kernel of Eq. (12) that it can capture the full geometry of phase defining features in the TFIM model.

V. MICROCANONICAL DYNAMICS OF THE 1DTFIM

Despite the successful identification of the ground state phases with the CS+ML model, the fact remains that the problem of calculating ground states is a QMA-hard problem [18]. We now turn to an algorithm built using Schrödinger dynamics, a known BQP class algorithm, that aims to reveal the microcanonical phase diagram as a proxy for a ground state study.

Consider now Fig. 4, a sketch of the thermodynamic phase diagram of the 1DTFIM relating internal energy $E = \langle \psi(0) | \hat{H} | \psi(0) \rangle$ to the transverse magnetic field h_x inspired by Ref. [62]. this phase diagram is relevant for a microcanonical dynamics study governed by the entropy $S(E, h_x)$. It exists even for a simulation over a finite time T and with a finite number of spins N and a specific choice of initial conditions but with finite T , finite N , and initial choice-dependent errors that round phase transitions. We present in this figure our expectations for the phase diagram in this context, pointing out phase transitions where the phase diagram will sharpen in the thermodynamic limit. We further highlight the path through the phase diagram carried out by our simulations below, showing that we expect it to cross the quantum critical region, and so be sensitive to the phase diagram at a rounded level even in the long- T , large- N limit.

To reveal the phase diagram expected from microcanonical dynamics presented in Fig. 4, we need an experimentally producible classical representation of the quantum data obtained from a microcanonical dynamics simulation. Noticing that the time-averaged integral amounts to an expectation value of the




FIG. 4. A sketch of the expected 1DTFIM phase diagram at finite T and finite N as a function of internal energy E and transverse magnetic field h_x . This diagram is a modification of the canonical ensemble representation of the phase diagram in Ref. [62], adapted to the microcanonical ensemble. The spectrum is mirror symmetric about $E = 0$ due to the chiral symmetry $\mathcal{C} = ZYZYZY \dots$

pure state density matrix $|\psi(t)\rangle\langle\psi(t)|$ over the probability distribution $P_T(t) = (1/T)(\Theta(T-t) - \Theta(-t))$, where $\Theta(x)$ is the Heaviside step function, we see we can construct *time-averaged classical shadows* (TACS) by the quantum channel

$$\text{TACS}[\rho] = \lim_{T \rightarrow \infty} \int dt \sum_{b,\sigma} P_T(t) P(b) \mathcal{P}_{b,\sigma} |\psi(t)\rangle\langle\psi(t)| \mathcal{P}_{b,\sigma}. \quad (17)$$

Hence by sampling the joint probability distribution $P_T(t)P(b)$ to obtain (t_i, b_i) , $i = 1, \dots, N$, and then measuring one shot σ_i from $|\psi(t_i)\rangle$ in basis b_i we obtain a finite-shot TACS via

$$\text{TACS}_N[\rho] = \sum_{i=1}^N |t_i b_i \sigma_i\rangle\langle t_i b_i \sigma_i|. \quad (18)$$

This approach captures the power of CS tomography and enables an experimental study of microcanonical thermodynamics.

With this shadow tomography method in mind, we ran quantum dynamics simulations of 1DTFIM using the TDVP algorithm [63,64] starting from the GHZ state $|\psi(0)\rangle = \frac{|000\dots\rangle + |111\dots\rangle}{\sqrt{2}}$ to generate TACS data from 500 randomly sampled dimensionless time values between $t = 10.0$ to 20.0 and 187 randomly sampled h_x field values between $h_x = 0.1$ to 10.0 . An example code to generate TACS dataset for 1DTFIM is available in our github repository [65]. These 187 TACS were the data points with which we performed unsupervised learning by constructing the 187×187 kernel matrix using the shadow kernel in Eq. (12) and then using diffusion maps for dimensionality reduction.

A key element needed to obtain reasonable results from the above calculation is an initial state that equilibrates within the chosen time window for observables of interest that are accurately captured by the chosen method of shadow tomography. In the above case, we started from a GHZ state

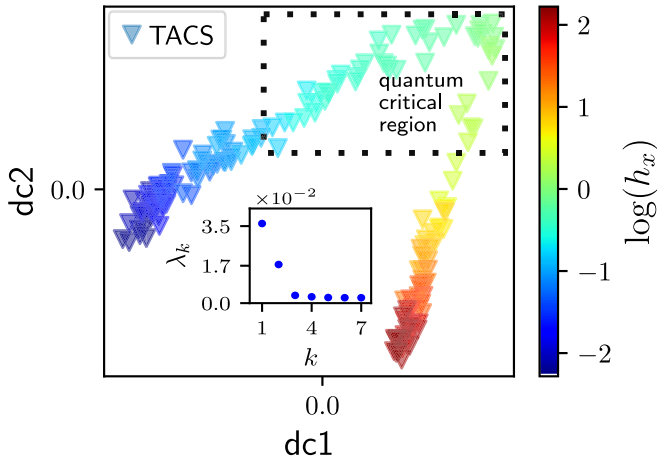


FIG. 5. Phase identification from dynamics data. The eigenvalue spectrum (inset) obtained from diffusion maps shows the two largest eigenvalues corresponding to the two dominant diffusion coordinates dc1 and dc2 (the trivial point $k = 0$, is not shown). The TACS data point largely fall on a two-stranded curve parameterized by h_x in this 2D reduced diffusion space. The quantum critical region [see Appendix D 2] (in green) coincides with the inflection point neighborhood of the curve, with points on the left strand belonging mostly to the ordered phase while points on the right strand belong to the disordered phase.

because it equilibrated efficiently for local observables, as shown in Appendix C. Presumably, this equilibration would occur even faster if we broke the integrability of the 1D TFIM model by adding certain additional terms to the Hamiltonian. Hence, up to possibly finite size effects, thermodynamic observables in our results should behave as expected.

Figure 5 shows our results. The first two eigenvectors were chosen as our diffusion space basis vectors because as Fig. 5 (inset) shows, those are the two dominant nontrivial eigenvalues in the spectrum. The rest of the eigenvectors go to zero as data points increase. Projecting the states onto this two-dimensional diffusion space, we see that the states all fall on a curve in this hyperplane along which the value of h_x increases monotonically, and the inflection point neighborhood of the curve coincides with the quantum critical region (see Appendix D 2). Therefore, as in the ground state study, it is apparent that the unsupervised learning algorithm was able to infer two phases from the data via a single cluster with nontrivial geometry. Presumably, taking paths through the phase diagram closer to the critical point, we would see the inflection point sharpen, leading to a singular point in the data manifold at the critical point. However, unlike the ground state study, it is not obvious what the two diffusion coordinates dc1 and dc2 correspond to. To identify these, we need to study observables capable of capturing the phase-defining features and see which correlate with these learned coordinates.

As a preliminary exploration of phase-defining observables, a straightforward first approach is to check whether the diffusion coordinates obey power laws consistent with the known quantum critical point. In Fig. 6(a), we plot dc2, which diverges as it approaches the critical point with critical exponents $p_- = 0.58 \pm 0.05$ and $p_+ = 0.7 \pm 0.1$. We have shifted the diffusion coordinates by $C = -0.0027$ since the

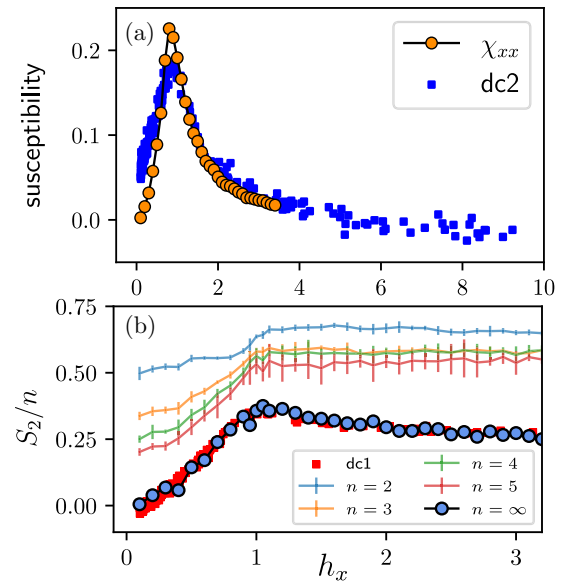


FIG. 6. Interpretation of dc1 and dc2 for microcanonical dynamics of 1DTFIM. The divergent behavior of dc2 qualitatively matches the xx component of the susceptibility, computed using 100k shot TACS data for a ten-site 1DTFIM, denoted by orange circles (b) dc1 matches the Bayesian inference estimate for the second Renyi entropy per site (S_2/n) in the thermodynamic limit ($n = \infty$). Bayesian inference is performed on n -body entropies for $n = 1-5$, also computed using the ten-site dataset.

diffusion distance is invariant under an overall shift of the origin as mentioned in Sec. III B and this renders it positive. Remarkably, this shift *simultaneously renders both dc1 and dc2 positive*. However, our errors in the exponents are hard to estimate. Suppose we view the unknown variable C as a Gaussian distribution. In that case, the corresponding distribution of p_{\pm} from our predictions is highly non-Gaussian (see Appendix D 4). The closest known critical exponent is $\nu = 1$ [66] (see also Wikipedia [67]). However, we found that the observable that qualitatively matches the diverging behavior at $h_x = 1$ is an xx component of the susceptibility, which we define as

$$\chi_{ab} = \frac{1}{L^2} \sum_{ij} \langle \sigma_i^a \sigma_j^b \rangle, \quad (19)$$

where $\sigma^a \in \{X, Y, Z\}$ is a 2×2 Pauli matrix. This is different than the usual definition obtained by summing over the connected correlations. The resemblance between dc2 and susceptibility here is only qualitative, so the critical exponents do not match. Since there are a number of observables, such as χ_{xx} , χ_{yy} , and χ_{zz} , that are equally likely candidates to define the phases, we conjecture that dc2 could be some combination of these.

This leaves the puzzle of determining dc1, which neither diverges nor shows a power law behavior. When inverted, it appears qualitatively similar to the ZZ component of the susceptibility (see Sec. D 3) but we turn to calculate the second Renyi entropy given the diagonal components of the kernel function capture this quantity. Entropy is the key thermodynamic potential of the microcanonical ensemble

that certainly captures phase transitions and from which all important phase-defining features could be extracted. It is also known that the quantum critical region features the interplay of equilibrium and quantum fluctuations leading to the entropy being maximized [62,68]. Renyi entropy is a lower bound for the von Neumann entropy and has the same limits—it vanishes for pure states and reaches N for the maximally mixed state. We do so by first computing the purity using Eq. (10) as follows:

$$\gamma[\rho] = \text{Tr}[\rho^2] \quad (20)$$

$$\approx \frac{1}{N^2} \sum_{n \neq n'=1}^N \text{Tr}[\sigma_1^{(n)} \sigma_1^{(n')}] \times \cdots \times \text{Tr}[\sigma_L^{(n)} \sigma_L^{(n')}], \quad (21)$$

then the Renyi entropy is $S_2 = -\ln_2 \gamma$.

Figure 6(b) presents the Renyi entropy as calculated from TACS shots on a $L = 10$ qubit system using exact diagonalization. We use Bayesian inference, as detailed in Appendix E (see also Ref. [69]) to extract the large- n , large- N predictions with $N = 100\,000$ shots per h_x value and 35 h_x values between $h_x = 0.1$ and 3.3. The range of time sampled for this particular calculation was between $t = 5$ and 25. The results show clear evidence that entropy is maximized in the quantum critical region around $h_x = 1$ as expected from Fig. 2 of Ref. [62]. Error estimates for these values were obtained in Ref. [70], and given by

$$N \geq \frac{4^{n+1} \gamma}{\epsilon^2 \delta} \quad (22)$$

Where $1 - \delta$ is the probability of obtaining a good TACS dataset and ϵ is the additive error. For a $\delta = 0.33$ (67 percentile), and $N = 100\,000$ shots, we find an additive error for the $n = 5$ entropy curve plotted in Fig. 6 at $h_x = 1.0$ of $S_2/n \approx 0.5 \pm 0.24$. This is larger than the observed error shown via error bars in Fig. 6 but only within a factor of order 1. Hence, by using 3 300 000 shots, we have estimated the thermodynamic entropy as a function of h_x that reproduces the expected maxima at the critical point.

Given an estimate of the entropy, we lastly turn to plotting it alongside dc1 to discover that it is highly correlated with this observable. Although dc1 was calculated using only 500 shots per h_x value, orders of magnitude less than the number of shots needed for accurate Bayesian inference estimation, the diffusion process is able to combine information across different h_x values without any supervision.

In summary, the diffusion map was able to learn phase-defining features from TACS and used these features to map the data points as a function of the model parameter h_x onto a curve in the two-dimensional plane with geometry that reveals the quantum phase transition.

VI. OUTLOOK

In this paper, we have identified an approach to studying quantum thermodynamics on a quantum computer in a way that is suitable for studying quantum materials, their phases, and their phase transitions. This approach consists of

(1) preparing a low-depth initial state for which relevant observables are observed to equilibrate within the coherence time of the quantum computer,

(2) time evolving this state using a quantum algorithm to a randomly chosen time point t within some time interval,

(3) extending shadow tomography methods to obtain a physically useful representation of the von Neumann ensemble such as TACS used in this paper, and

(4) employing an unsupervised machine learning method to discover the phase diagram, with kernel methods such as diffusion maps employing well-designed kernels showing promise.

Our approach parallels statistical mechanics calculations on classical Hamiltonians, where a random initial state is prepared, a METROPOLIS Monte Carlo algorithm is run beginning from this state, and data is collected and analyzed using traditional observables and more recently machine learning methods. Our results, demonstrating the existence of a quantum phase transition and the ability to map out regions of the phase diagram by a careful choice of initial conditions, show promise.

There are several resources needed to carry out the microcanonical dynamics simulations. A key resource is a low-energy state that equilibrates within the accessible time scale and can also be prepared easily. For local Hamiltonians on physical lattices, we can always find low-energy states that can be prepared with constant-depth circuits [27]. Another resource to carry out dynamics simulations to a time T in which the relevant local observables equilibrate ($T = 25$ in our simulations). It allows one to exploit advances in variational time evolution algorithms [71–73], which are especially suitable for the NISQ-ERA due to robustness to noise and ability to go beyond the coherence time of quantum computers. Finally, we need the ability to perform time averages by sampling the state $|\psi(t)\rangle\langle\psi(t)|$ at least at the Nyquist rate determined by the bandwidth, which is linear in the system size. We produced a TACS dataset consisting of 500 shots from the equilibrium dynamics starting from a GHZ state at each of 187 points in the phase diagram; these resources were all that were required for diffusion maps to learn the phases and identify the phase transition for a 20-qubit system. Somewhat different resources, a system size of $L = 10$, and $N = 100\,000$ shots at 33 points in the phase diagram were required for Bayesian inference to obtain reasonable estimates for the thermodynamic entropy.

We believe these resources are significantly smaller than those of other proposed methods for studying thermodynamics on quantum computers. For example, the overhead from using ancillas as a heat bath in the existing methods to study thermodynamics on a quantum computer [74–76] is not an issue with our approach. The resources required also open up an exciting possibility of employing a new generation of quantum simulators [77,78] to study quantum thermodynamics as they too can simulate quantum dynamics and are capable of performing randomized measurements. We are currently working on the calculation of sample complexity and establishing rigorous bounds for our predictions. Finally, a possible direction for future research would be to identify and benchmark a strongly interacting system with initial states that will yield a quantum advantage in the near term.

ACKNOWLEDGMENTS

We thank Mikhail Lukin, Katherine Van Kirk, Nishad Maskara, Yanting Teng, Subir Sachdev, Daniel Parker, and Anurag Anshu for useful discussions. This material is based upon work supported by the National Science Foundation under Grants No. OAC-1940243 and OAC-1940260.

APPENDIX A: DIAGRAMMATIC UNDERSTANDING OF CLASSICAL SHADOWS

In this section, we will develop a diagrammatic understanding of the classical shadows. We will work in the superoperator formalism where the indices of the density matrix ρ_{ij} are grouped together to make a vector $|\rho\rangle\rangle$ and the product $A\rho B$ translates to an operator $A \otimes B^T$ acting on the vector $|\rho\rangle\rangle$.

In general, we can view the outcomes obtained from many classical shadows measurements on the same prepared state ρ as defining an ensemble of states $S[\rho]$, where

$$S[\rho] = \sum_{b,x} P(b)\mathcal{P}_{b,x}\rho\mathcal{P}_{b,x} = \left[\sum_{b,x} P(b)\mathcal{P}_{b,x} \otimes \mathcal{P}_{b,x}^T \right] |\rho\rangle\rangle, \tag{A1}$$

and $\mathcal{P}_{b,x}$ is the projector in basis b onto qubit state $|x\rangle$. To break this down into manageable parts, let's start from the one-qubit case and work our way up to N qubits.

One-qubit case. In the one-qubit case, we generate classical shadows samples for bases $b \in \{X, Y, Z\}$ with uniform probabilities, i.e., $P(b) = 1/3$. We can thus express the one-qubit version of (A1) with the following diagram:

$$S[\rho] = \frac{1}{3} \sum_{b,x} \text{---} \boxed{P_{bx}} \text{---} \rho \text{---} \boxed{P_{bx}} \text{---}. \tag{A2}$$

Using cap and cup notation, we can redraw this in the superoperator form:

$$S[\rho] = \frac{1}{3} \sum_{b,x} \text{---} \boxed{P_{bx}} \text{---} \rho \text{---} \boxed{P_{bx}} \text{---} \tag{A3}$$

The highlighted box can be viewed as a superoperator acting on the space of linear operators ρ . Remarkably, this particular superoperator consisting of a product of two projection operators simplifies substantially, i.e.,

$$\sum_{b,x} \text{---} \boxed{P_{bx}} \text{---} \rho \text{---} \boxed{P_{bx}} \text{---} = \left[\text{---} \left| \right| \text{---} + \text{---} \right]. \tag{A4}$$

They amount to the sum of an identity and a cup-cap product. Using this simplification we recognize the one-qubit case as the depolarizing map

$$S[\rho] = \frac{1}{3} \left[\text{---} \boxed{\rho} \text{---} + \text{---} \right], \tag{A5}$$

where we further simplified using

$$\text{---} \boxed{\rho} \text{---} = 1 \tag{A6}$$

Inverting, we can extract the original density matrix (full tomography) via

$$\text{---} \boxed{\rho} \text{---} = 3 \text{---} \boxed{S[\rho]} \text{---} - \text{---} \tag{A7}$$

Going from diagrammatic results to equations gives

$$S[\rho] = \frac{1}{3}(\rho + I) \Rightarrow \rho = 3S[\rho] - I, \tag{A8}$$

which is a single-qubit depolarizing channel.

Two-qubit case. In the two-qubit case, the CS map of Eq. (A1) takes the form

$$\frac{1}{9} \sum_{b,x} \sum_{b',x'} \text{---} \boxed{P_{bx}} \text{---} \rho \text{---} \boxed{P_{bx}} \text{---} \cdot \text{---} \boxed{P_{b'x'}} \text{---} \rho \text{---} \boxed{P_{b'x'}} \text{---} \tag{A9}$$

Again, we can add cups and caps to express it in superoperator form:

$$\text{---} \boxed{P_{bx}} \text{---} \rho \text{---} \boxed{P_{bx}} \text{---} \cdot \text{---} \boxed{P_{b'x'}} \text{---} \rho \text{---} \boxed{P_{b'x'}} \text{---} \tag{A10}$$

This diagram shows that the two-qubit projection operators produce the same structure on each qubit as they did in the one-qubit case. Applying the same simplification as before, we arrive at

$$\frac{1}{9} \left[\text{---} \boxed{\rho} \text{---} + \text{---} \boxed{\rho} \text{---} + \text{---} \boxed{\rho} \text{---} + \text{---} \boxed{\rho} \text{---} \right], \tag{A11}$$

where again we used $\text{Tr}[\rho] = 1$. This expression amounts to a simple sum of all possible reduced density matrices. We can

invert our diagram for the two qubit case by first tracing over one of its qubits to get

$$S[\rho] = \frac{1}{3} \left[\text{Diagram 1} + \text{Diagram 2} \right] \quad (\text{A12})$$

This allows us to rewrite the reduced density matrices in terms of the partially summed $S[\rho]$. Using this relation, we then arrive at the inverse of the two-qubit classical shadow

$$\rho = 9S[\rho] - 3I \otimes \text{Tr}_1 S[\rho] - 3\text{Tr}_2 S[\rho] \otimes I + I \otimes I \quad (\text{A13})$$

As for the one-qubit case, we can convert our diagrams back to algebraic expressions. The results of this two-qubit case amount to the forward expression

$$S[\rho] = \frac{1}{9}(\rho + I \otimes \text{Tr}_1[\rho] + \text{Tr}_2[\rho] \otimes I + I \otimes I) \quad (\text{A14})$$

and the inverse expression

$$\rho = 9S[\rho] - 3I \otimes \text{Tr}_1 S[\rho] - 3\text{Tr}_2 S[\rho] \otimes I + I \otimes I \quad (\text{A15})$$

N-qubit case. The formulas we have derived for the one- and two-qubit cases readily extend to N -qubits. They are

$$S[\rho] = \frac{1}{3^L} \left(\rho + \sum_l \text{Tr}_l \rho + \sum_{l \neq l'} \text{Tr}_{ll'} \rho + \dots + I \right) \quad (\text{A16})$$

for the forward map and for the inverse map

$$\rho = 3^L S[\rho] - 3^{L-1} \sum_l \text{Tr}_l S[\rho] + 3^{L-2} \sum_{l \neq l'} \text{Tr}_{ll'} S[\rho] + \dots + (-1)^L I, \quad (\text{A17})$$

where we have suppressed the presence of identity operators that replace traced out regions for ease of notation. These expressions satisfy $\text{Tr}[\rho] = 1$ and $S[\rho] \geq 0$. The inverse map is not non-negative in general, but should be for the density matrices resulting from this map. Equation (A17) may also be written as

$$\rho = \frac{1}{N} \sum_{n=1}^N \bigotimes_{l=0}^L (3 |b_l^{(n)}\rangle \langle b_l^{(n)}| \sigma_l^{(n)} - 1), \quad (\text{A18})$$

where l denotes the site-index and n denotes the shot index.

APPENDIX B: SYMMETRIES OF 1DTFIM

1. \mathbb{Z}_2 symmetry

The 1DTFIM is invariant under global flipping of the z component of the spin, the \mathbb{Z}_2 symmetry. This unitary symmetry can be expressed as

$$\mathcal{S} = \prod_i X_i. \quad (\text{B1})$$

We can check that the symmetry operator \mathcal{S} commutes with the TFIM Hamiltonian i.e. $[\mathcal{S}, H_{\text{1DTFIM}}] = 0$. This allows us to write H_{1DTFIM} in block diagonal form with each block corresponding to eigenvalues 1 (even parity) and -1 (odd parity) of \mathcal{S} . States in these sectors evolve independently of each other. If we start in GHZ state

$$|\text{GHZ}\rangle = \frac{|00\dots 0\rangle + |11\dots 1\rangle}{\sqrt{2}}, \quad (\text{B2})$$

a time reversal even state, we remain in the even sector under time evolution. Since the magnetization of this state is 0, the magnetization will stay at this value forever. Thus equilibration of the order parameter is not an issue. Since the 1DTFIM Hamiltonian is purely real, it is symmetric under complex conjugation \mathcal{K} , and consequently, the eigenvalues are also real. Hence, it is also symmetric under $\mathcal{T} = \mathcal{S}\mathcal{K}$, i.e., time reversal symmetry.

2. Chiral symmetry

The 1DTFIM is also symmetric under the following chiral operator:

$$\mathcal{C} = ZYZY\dots ZY. \quad (\text{B3})$$

We can check that \mathcal{C} anticommutes with H_{1DTFIM} , i.e., $\{\mathcal{C}, H_{\text{1DTFIM}}\} = 0$, so for every energy eigenstate E , there exists an eigenstate with $-E$, which makes the spectrum of 1DTFIM mirror symmetric about zero energy.

APPENDIX C: EQUILIBRATION FROM INITIAL STATES

An important resource for studying microcanonical phases using quantum dynamics is an initial state that equilibrates within the time scale T accessible to a quantum device. The initial state sets the energy and the symmetry sector of the microcanonical ensemble resulting from time-averaging over $[0, T]$.

In Fig. 7, we present a numerical assessment of equilibration for some observables in the 1DTFIM. Specifically, we plot the evolution of expectation values for operators $\langle Z \rangle$, $\langle X \rangle$ and $\langle ZZ \rangle$ corresponding to two initial states: the ferromagnetic ($|00\dots\rangle$) and GHZ ($\frac{1}{\sqrt{2}}(|00\dots\rangle + |11\dots\rangle)$) states, respectively. We observe that the \mathcal{T} -odd operators such as Z do not equilibrate for the all-up state within the sampling window whereas they are forced to be 0 for the GHZ state by the \mathcal{T} symmetry. Likewise, we find that \mathcal{T} -even operators such as X , ZZ equilibrate and are identical for both initial states, also due to the \mathcal{T} symmetry (the ferromagnetic state is a superposition of \mathcal{T} -even and \mathcal{T} -odd states, and the expectation




FIG. 7. Time evolution of $\langle Z \rangle$, $\langle X \rangle$ and $\langle ZZ \rangle$ for ferromagnetic (all up) and GHZ (all up plus all down) states for ten sites and $h_x = 0.8$. $t \in [5, 25]$ (highlighted in orange) is the sampling window used for generating the TACS data. The order parameter $\langle Z \rangle$ for the ferromagnetic state doesn't equilibrate in this window.

value of \mathcal{T} -even operator for a \mathcal{T} -odd state yields 0). Finally, we observe all equilibrating observables equilibrate within a time scale of 5. For this reason, all of the dynamics results presented in the main manuscript used this numerical evidence time interval to $[5, 25]$ (shown in the highlighted region in Fig. 7 for time averages). Hence, we find numerical evidence for equilibration of local observables, evidence that formed an important basis upon which we carried out our dynamics simulations.

APPENDIX D: INTERPRETING DIFFUSION MAPS

1. Learning physical parameters from 1DTFIM ground state data

The unsupervised learning method we implemented was able to unveil the symmetry-breaking phase transition of 1DTFIM from ground state CS data [Fig. 3(d)]. It did so by generating diffusion coordinates that are related to relevant parameters, the order parameter M_z , and the model parameter h_x of 1DTFIM. Figure 8 shows the correlation between the diffusion coordinates and these parameters.

2. Quantum criticality in TACS kernel matrix

The microcanonical phase diagram being studied here has three characteristically distinct regions—namely, the ordered phase, the quantum critical region, and the disordered phase. The quantum critical region, although does not include the




FIG. 8. Learning 1DTFIM model parameter and order parameter from ground state CS data. Correlation between (a) dc1 and order parameter M_z and (b) dc2 and model parameter h_x .

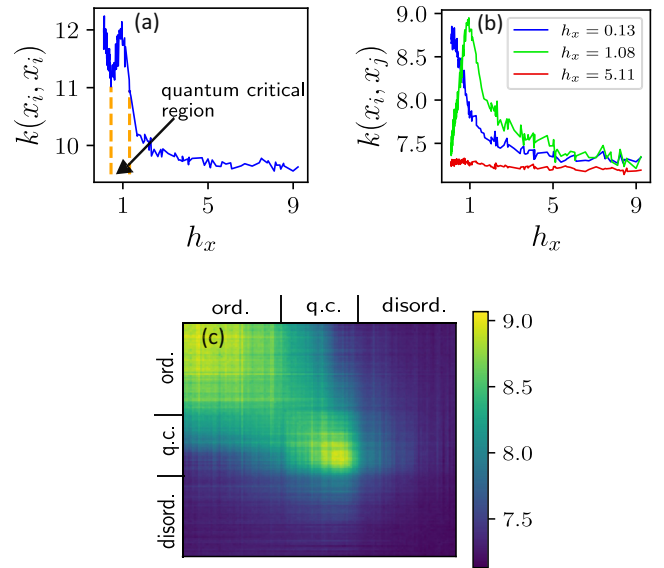


FIG. 9. The three distinct regions of the microcanonical phase diagram depicted via the kernel matrix K . (a) shows the diagonal elements of K against their respective h_x values. The quantum critical region is inferred from this plot. (b) shows three different rows of K , each belonging to a different region in the microcanonical phase diagram. (c) portrays the full matrix K itself, where the distinct regions (ordered, quantum critical, and disordered) can be seen. [Diagonal elements have been ignored in (b) and (c) for better visualization.]

phase transition point, exhibits critical behavior characterized by a singularity in the order parameter and the response functions. This behavior manifests itself in the feature space accompanying the shadows kernel function [Eq. (12)] as it contains the polynomial expansion of the reduced density matrices [24] and can be analyzed from the kernel matrix.

Figure 9 displays how the microcanonical phases reveal themselves in the kernel matrix. Figure 9(a) shows a maximum in the diagonal elements of the kernel matrix in the quantum critical region due to increased correlation length. In the figure, we delineate the quantum critical region in the neighborhood of this peak. Likewise, Figs. 9(b) and 9(c) demonstrate how the kernel function between states behaves in different regions. The ordered states have low entropy; hence, greater “similarity” among themselves makes the kernel function take a higher value than the other regions, dropping sharply as we go out of that region. The disordered states have roughly uniform values for the kernel function with all other states due to their high entropy, and a peak in the critical region as discussed above. These character traits of each of these regions help us identify them from the kernel matrix. However, we don't use the kernel matrix for phase classification. We let the probabilities diffuse and use the diffusion matrix and resulting diffusion coordinates.

3. Qualitative similarity between dc's and susceptibility

Susceptibility is an important quantity of interest to us because it diverges at the critical point. Although our microcanonical dynamics take place at an energy above-the-ground state, we expect the signature of this divergence to be present




FIG. 10. xx , yy , and zz components of the magnetic susceptibility for a ten-site 1DTFIM, computed using 10 000 shot TACS dataset.

in the quantum critical region. With the experience that diffusion coordinates correspond to phase-defining observables in the case of ground states, we plotted the xx , yy , and zz components of the susceptibility in Fig. 10 computed using the 100k shot TACS dataset for the 1DTFIM to compare against $dc1$ and $dc2$. Although χ_{zz} , a natural candidate for $dc1$, looks qualitatively similar to $dc1$ when inverted, we find that the second Renyi entropy is a better fit. Similarly, χ_{xx} behaves qualitatively similar to $dc2$ in the sense that both are sharply peaked at the critical value of $h_x = 1$. However, they do not share the same critical exponents, hence, we cannot make as strong of a claim as for $dc1$.

4. Estimating the critical exponent from the TACS diffusion coordinates

In Fig. 11(a) below, we see that the second diffusion coordinate $dc2$ in TACS diffusion maps approximates a power law in the quantum critical region. In order to estimate the critical exponent, we modeled $dc2$ as

$$dc2(h_x; a, p) = a|h_x - 1|^{-p} + C, \quad (D1)$$

where a and p are fitting parameters and ν is the critical exponent.

It is evident that our estimate of p depends on our choice of C . Figure 11 shows that dependence on either side of the critical point ($h_x = 1$). We can obtain a probability distribution $P(C)$ on C by modeling the Bayesian estimate of the second Renyi entropy S_2/n as a function of $dc1$:

$$S_2/n(h_x; \alpha, C) = \alpha(dc1(h_x) - C). \quad (D2)$$

Here, α and C are fitting parameters. The ordinary least squares fit gives us the optimum value for C (C_{opt}) with the least square error (ϵ). We then model $P(C)$ as a normal distribution, $P(C) = \mathcal{N}(C_{opt}, \sigma = \epsilon^2)$ and plot it together with the dependence of p on the shift C in Fig. 11 to visualize how an error in C translates to an error in ν .

APPENDIX E: BAYESIAN INFERENCE EXTRAPOLATION OF ENTROPY DATA

To infer the entropy in the limit of large N , the number of shots, and large n , the number of qubits, we need to extrapolate the estimates we obtain from CS data. At first glance, this would seem hard to do because the error in our

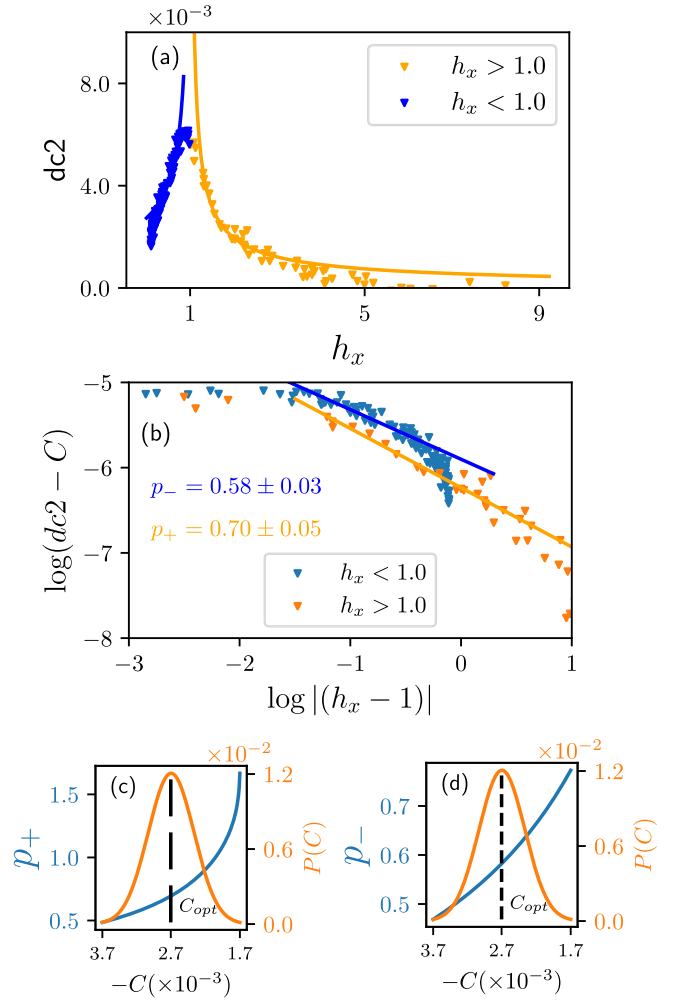


FIG. 11. Modeling the power-law behavior of $dc2$ in the critical region. (a) shows the $dc2$ value of each TACS plotted against their respective h_x values, as well as the power law fit on either side of the critical point. (b) log-log plot of $|h_x - 1|$ vs $(dc2 - C)$, along with straight-line fits, the slopes give us the p values. Here $C = -0.0027$. (c) and (d) show the dependence of the (c) p_+ and (d) p_- , estimated power law exponents on the shift in the diffusion coordinates C discussed in Sec. III B (blue) [dashed line shows the optimum C -value, which was chosen in (a) and (b)], along with the modeled normal distribution on C obtained from error estimates on the shift needed to render $dc1$ positive near $h_x = 0$ (orange). Chosen range of C is $[C_{opt} - 3\sigma, C_{opt} + 3\sigma]$.

estimates grows exponentially with the locality of the observables and entropy is not a local observable. However, the dynamics data we have obtained represents a mixed state with volume-law entanglement and, due to the finite energy—the microcanonical stand-in for temperature—typically has exponentially decaying correlations beyond a correlation length. Hence, we expect the entropy of the reduced density matrix of a region A with n_A qubits will obey $S \propto n_A$ even for small n_A . Our approach to extrapolate the entropy is therefore to build a probability model $p(X|\theta)$ with parameters θ that captures our estimated entropy value data X .

To simplify the calculation of entropy, we will compute the second Renyi entropy of a subregion A : $S_A = -\ln_2 \gamma_A$,

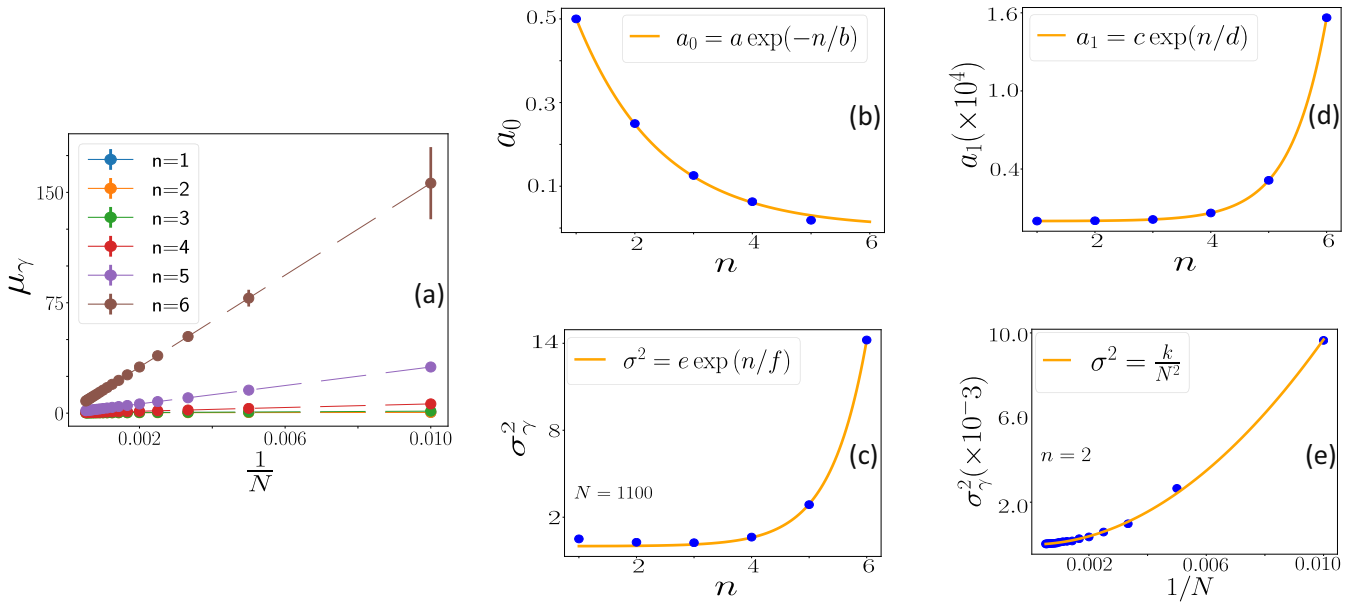


FIG. 12. Developing a model for purity $\gamma = \mathcal{N}(\mu_\gamma, \sigma_\gamma)$ from CS data on the maximally mixed state. (a) μ_γ is linear in $\frac{1}{N}$, so it takes the form: $\mu_\gamma = a_0(n) + \frac{a_1(n)}{N}$. (b) The intercept $a_0(n)$ can be modeled as an inverse exponential function of n , $a_0(n) = a e^{-n/b}$ (c) The slope, $a_1(n)$, can be modeled as an exponential in n , $a_1(n) = c e^{dn}$. (d) and (e) show that the variance σ_γ^2 can be closely approximated by an exponential function in n and linear in $\frac{1}{N^2}$, $\sigma_\gamma^2 = e e^{fn}/N^2$.

$\gamma = \text{Tr} \rho_A^2$. To obtain a model of this entropy as a function of the number of CS shots N and qubits n_A , a data set of values Y given by the entropy S_A and dependent variables X given by (N, n_A) , we generated data from the maximally mixed state $\rho = (1/d)I$, $d = 2^n$ the size of the Hilbert space. The results fit a model of the form (see Fig. 12):

$$\mu_\gamma(n, N) = a e^{-bn} + c e^{dn}/N, \quad \sigma_\gamma = e e^{fn} \quad (\text{E1})$$

with positive parameters a, b, c, d, e, f . Namely, we found the purity γ is linear in $1/N$ but exponential in n .

Given the mean and variance as modeled above, we can then model the probability distribution from which a given data point $(\vec{x}, y) \in (X, Y)$ is a sample as a Gaussian:

$$P(y|\theta, \vec{x}) = \frac{1}{\sqrt{2\pi\sigma^2(\vec{x};\theta)}} e^{-(y-\mu(\vec{x};\theta))^2/2\sigma^2(\vec{x};\theta)}. \quad (\text{E2})$$

Then by Bayes Law, we can learn the posterior

$$p(\theta|Y, X) = \frac{\prod_{(\vec{x}, y) \in (X, Y)} P(y|\theta, \vec{x}) P(\theta)}{P(Y|X)}, \quad (\text{E3})$$

where $P(X, Y) = \int d\theta \prod_{(\vec{x}, y)} P(\vec{x}, y|\theta) P(\theta)$ is called the evidence that provides a sense of how well the model is performing.

The probability of observing a new data point (\vec{x}', y') is then given by the posterior predictive

$$P(\vec{x}', y'|X, Y) = \int d\theta p(\vec{x}', y'|\theta) P(\theta|X, Y). \quad (\text{E4})$$

An estimate of which is obtainable from a set of samples Θ drawn from $P(\theta|X, Y)$

$$P(\vec{x}', y'|X, Y) = \frac{1}{|\Theta|} \sum_{\theta \in \Theta} p(\vec{x}', y'|\theta). \quad (\text{E5})$$

We are then specifically interested in the mean and standard deviation of $P((\infty, \infty), y|X, Y)$. Knowing this, we solve the problem of extrapolating the entropy from a finite number of shots and qubits for the entropy is the mean and our uncertainty in obtaining it is the standard deviation.

It remains then to obtain samples from the posterior $P(\theta|X, Y)$. We could do so using a straightforward Monte Carlo algorithm. For example, starting with an initial choice for the parameters θ_0 , we pick a random direction in parameter space and move an amount δ in that direction to obtain θ_{trial} . We then compute

$$\begin{aligned} \ln(r) &= \ln \frac{P(\theta_{\text{trial}}|X, Y)}{P(\theta_0|X, Y)} \\ &= \sum_{(\vec{x}, y) \in (X, Y)} (\ln P(\vec{x}, y|\theta_{\text{trial}}) - \ln P(\vec{x}, y|\theta_0)) \\ &\quad + \ln P(\theta_{\text{trial}}) - \ln P(\theta_0), \end{aligned} \quad (\text{E6})$$

which simplifies if we choose a uniform distribution for $P(\theta)$. We keep the trial, setting $\theta_1 = \theta_{\text{trial}}$ if a random number q between 0 and 1 satisfies $q < r$ and reject otherwise. Either way, we repeat the process generating ultimately a list Θ of correlated samples θ_i from which we can estimate the entropy and uncertainty from $P((\infty, \infty), y|X, Y)$.

However, a better approach than the METROPOLIS algorithm is to use the NUTS algorithm available in PyMC instead. This algorithm automatically chooses parameters in Hamiltonian Monte Carlo (HMC) and is more efficient than METROPOLIS for Bayesian inference. See Ref. [79].

- [1] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, Quantum chemistry in the age of quantum computing, *Chem. Rev.* **119**, 10856 (2019).
- [2] D. Wecker, M. B. Hastings, N. Wiebe, B. K. Clark, C. Nayak, and M. Troyer, Solving strongly correlated electron models on a quantum computer, *Phys. Rev. A* **92**, 062318 (2015).
- [3] H. Ma, M. Govoni, and G. Galli, Quantum simulations of materials on near-term quantum computers, *npj Comput. Mater.* **6**, 85 (2020).
- [4] N. Nagaosa, Superconductivity and antiferromagnetism in high-*T_c* cuprates, *Science* **275**, 1078 (1997).
- [5] X. Zhou, W.-S. Lee, M. Imada, N. Trivedi, P. Phillips, H.-Y. Kee, P. Törmä, and M. Eremets, High-temperature superconductivity, *Nat. Rev. Phys.* **3**, 462 (2021).
- [6] J. McGreevy, In pursuit of a nameless metal, *Physics* **3**, 83 (2010).
- [7] R. B. Laughlin, Nobel lecture: Fractional quantization, *Rev. Mod. Phys.* **71**, 863 (1999).
- [8] L. Savary and L. Balents, Quantum spin liquids: A review, *Rep. Prog. Phys.* **80**, 016502 (2017).
- [9] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [10] S. R. White, Density-matrix algorithms for quantum renormalization groups, *Phys. Rev. B* **48**, 10345 (1993).
- [11] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011), January 2011 Special Issue.
- [12] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, Cambridge, United Kingdom, 2017).
- [13] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions, *Rev. Mod. Phys.* **68**, 13 (1996).
- [14] R. P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [15] A. Kitaev, A. Shen, and M. Vyalıy, *Classical and Quantum Computation*, Graduate Studies in Mathematics (American Mathematical Society, Providence, Rhode Island, USA, 2002).
- [16] B. O’Gorman, S. Irani, J. Whitfield, and B. Fefferman, Intractability of electronic structure in a fixed basis, *PRX Quantum* **3**, 020322 (2022).
- [17] S. Gharibian and J. Yirka, The complexity of simulating local measurements on quantum systems, *Quantum* **3**, 189 (2019).
- [18] A. D. Bookatz, QMA-complete problems, *Quantum Information and Computation* **14**, 5 (2014).
- [19] J. Kempe, A. Kitaev, and O. Regev, The complexity of the local Hamiltonian problem, *SIAM J. Comput.* **35**, 1070 (2006).
- [20] M. L. Baez, M. Goihl, J. Haferkamp, J. Bermejo-Vega, M. Gluza, and J. Eisert, Dynamical structure factors of dynamical quantum simulators, *Proc. Natl. Acad. Sci.* **117**, 26123 (2020).
- [21] A. Rudi, L. Wossnig, C. Ciliberto, A. Rocchetto, M. Pontil, and S. Severini, Approximating Hamiltonian dynamics with the Nyström method, *Quantum* **4**, 234 (2020).
- [22] S. Aaronson, Shadow tomography of quantum states, in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2018), pp. 325–338.
- [23] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [24] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, *Science* **377**, eabk3333 (2022).
- [25] H.-Y. Huang, S. Chen, and J. Preskill, Learning to predict arbitrary quantum processes, *PRX Quantum* **4**, 040337 (2023).
- [26] J. von Neumann, Proof of the ergodic theorem and the h-theorem in quantum mechanics, *Eur. Phys. J. H* **35**, 201 (2010).
- [27] D. Aharonov, I. Arad, and T. Vidick, The quantum PCP conjecture, *arXiv:1309.7495*.
- [28] A. Anshu, N. P. Breuckmann, and C. Nirkhe, NLTS Hamiltonians from good quantum codes, in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023 (Association for Computing Machinery, New York, NY, USA, 2023), p. 1090-1096.
- [29] H. Wilming, M. Goihl, C. Krumnow, and J. Eisert, Towards local equilibration in closed interacting quantum many-body systems, *arXiv:1704.06291* (2017).
- [30] A. S. L. Malabarba, L. P. García-Pintos, N. Linden, T. C. Farrelly, and A. J. Short, Quantum systems equilibrate rapidly for most observables, *Phys. Rev. E* **90**, 012121 (2014).
- [31] D. Hetterich, M. Fuchs, and B. Trauzettel, Equilibration in closed quantum systems: Application to spin qubits, *Phys. Rev. B* **92**, 155314 (2015).
- [32] S. J. Wetzel, Unsupervised learning of phase transitions: From principal component analysis to variational autoencoders, *Phys. Rev. E* **96**, 022140 (2017).
- [33] L. Wang, Discovering phase transitions with unsupervised learning, *Phys. Rev. B* **94**, 195105 (2016).
- [34] W. Hu, R. R. P. Singh, and R. T. Scalettar, Discovering phases, phase transitions, and crossovers through unsupervised machine learning: A critical examination, *Phys. Rev. E* **95**, 062122 (2017).
- [35] C. Wang and H. Zhai, Machine learning of frustrated classical spin models. I. Principal component analysis, *Phys. Rev. B* **96**, 144432 (2017).
- [36] C. Wang and H. Zhai, Machine learning of frustrated classical spin models (II): Kernel principal component analysis, *Front. Phys.* **13**, 130507 (2018).
- [37] P. Broecker, F. F. Assaad, and S. Trebst, Quantum phase recognition via unsupervised machine learning, *arXiv:1707.00663*.
- [38] W. Zhang, J. Liu, and T.-C. Wei, Machine learning of phase transitions in the percolation and XY models, *Phys. Rev. E* **99**, 032142 (2019).
- [39] J. F. Rodriguez-Nieva and M. S. Scheurer, Identifying topological order through unsupervised machine learning, *Nat. Phys.* **15**, 790 (2019).
- [40] Y. Che, C. Gneiting, T. Liu, and F. Nori, Topological quantum phase transitions retrieved through unsupervised machine learning, *Phys. Rev. B* **102**, 134213 (2020).
- [41] A. Lidiak and Z. Gong, Unsupervised machine learning of quantum phase transitions using diffusion maps, *Phys. Rev. Lett.* **125**, 225701 (2020).
- [42] R. R. Coifman and S. Lafon, Diffusion maps, *Appl Comput Harmon Anal.* **21**, 5 (2006), special Issue: Diffusion Maps and Wavelets.

- [43] J. De la Porte, B. Herbst, W. Hereman, and S. Van Der Walt, An introduction to diffusion maps, in *Proceedings of the 19th Symposium of the Pattern Recognition Association of South Africa (PRASA 2008)*, Cape Town, South Africa (PRASA, Cape Town, South Africa, 2008), pp. 15–25.
- [44] E. Alpaydin, *Introduction to Machine Learning*, Adaptive Computation and Machine Learning series (MIT Press, Cambridge, Massachusetts, USA, 2014).
- [45] C. Bishop, *Pattern Recognition and Machine Learning*, Information Science and Statistics (Springer, New York, 2016).
- [46] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, *Nat. Rev. Phys.* **5**, 9 (2023).
- [47] E. Knill, Quantum computing with realistically noisy devices, *Nature (London)* **434**, 39 (2005).
- [48] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, United Kingdom, 2010).
- [49] S. Goldstein, J. L. Lebowitz, C. Mastrodonato, R. Tumulka, and N. Zanghì, Normal typicality and von Neumann’s quantum ergodic theorem, *Proc. R. Soc. A: Math. Phys. Eng. Sci.* **466**, 3203 (2010).
- [50] C. Gogolin and J. Eisert, Equilibration, thermalisation, and the emergence of statistical mechanics in closed quantum systems, *Rep. Prog. Phys.* **79**, 056001 (2016).
- [51] A. J. Short and T. C. Farrelly, Quantum equilibration in finite time, *New J. Phys.* **14**, 013063 (2012).
- [52] N. Yunger Halpern, P. Faist, J. Oppenheim, and A. Winter, Microcanonical and resource-theoretic derivations of the thermal state of a quantum system with noncommuting charges, *Nat. Commun.* **7**, 12051 (2016).
- [53] S. Vinjanampathy and J. Anders, Quantum thermodynamics, *Contemp. Phys.* **57**, 545 (2016).
- [54] J. M. Deutsch, Quantum statistical mechanics in a closed system, *Phys. Rev. A* **43**, 2046 (1991).
- [55] M. Srednicki, Chaos and quantum thermalization, *Phys. Rev. E* **50**, 888 (1994).
- [56] Z. Gong and R. Hamazaki, Bounds in nonequilibrium quantum dynamics, *Int. J. Mod. Phys. B* **36**, 2230007 (2022).
- [57] N. Margolus, Counting distinct states in physical dynamics, [arXiv:2111.00297](https://arxiv.org/abs/2111.00297).
- [58] M. Fishman, S. R. White, and E. M. Stoudenmire, The ITensor software library for tensor network calculations, *SciPost Phys. Codebases* **4** (2022).
- [59] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [60] I. Borg and P. Groenen, *Modern Multidimensional Scaling: Theory and Applications*, Springer Series in Statistics (Springer, New York, 2013).
- [61] S. Sachdev, *Quantum Phase Transitions*, 2nd ed. (Cambridge University Press, Cambridge, United Kingdom, 2011).
- [62] J. Wu, L. Zhu, and Q. Si, Crossovers and critical scaling in the one-dimensional transverse-field Ising model, *Phys. Rev. B* **97**, 245127 (2018).
- [63] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Unifying time evolution and optimization with matrix product states, *Phys. Rev. B* **94**, 165116 (2016).
- [64] <https://github.com/orialb/TimeEvoMPS.jl>.
- [65] <https://github.com/Lawler-Research-Group/TACS>.
- [66] J. Cardy, *Scaling and Renormalization in Statistical Physics*, Cambridge Lecture Notes in Physics (Cambridge University Press, Cambridge, United Kingdom, 1996).
- [67] Wikipedia contributors, Ising critical exponents (2022), [Online; accessed 12-Dec-2022], https://en.wikipedia.org/wiki/Ising_critical_exponents.
- [68] J. Wu, L. Zhu, and Q. Si, Entropy accumulation near quantum critical points: Effects beyond hyperscaling, *J. Phys.: Conf. Ser.* **273**, 012019 (2011).
- [69] J. M. Lukens, K. J. Law, and R. S. Bennink, A Bayesian analysis of classical shadows, *npj Quantum Inf.* **7**, 113 (2021).
- [70] S. H. Sack, R. A. Medina, A. A. Michailidis, R. Kueng, and M. Serbyn, Avoiding barren plateaus using classical shadows, *PRX Quantum* **3**, 020365 (2022).
- [71] Y.-X. Yao, N. Gomes, F. Zhang, C.-Z. Wang, K.-M. Ho, T. Iadecola, and P. P. Orth, Adaptive variational quantum dynamics simulations, *PRX Quantum* **2**, 030307 (2021).
- [72] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, Theory of variational quantum simulation, *Quantum* **3**, 191 (2019).
- [73] C. Cîrstoiu, Z. Holmes, J. Iosue, L. Cincio, P. J. Coles, and A. Sornborger, Variational fast forwarding for quantum simulation beyond the coherence time, *npj Quantum Inf.* **6**, 82 (2020).
- [74] L. Coopmans, Y. Kikuchi, and M. Benedetti, Predicting Gibbs-state expectation values with pure thermal shadows, *PRX Quantum* **4**, 010305 (2023).
- [75] M. Motta, C. Sun, A. T. K. Tan, M. J. O’Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, and G. K.-L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nat. Phys.* **16**, 205 (2020).
- [76] A. Francis, D. Zhu, C. H. Alderete, S. Johri, X. Xiao, J. K. Freericks, C. Monroe, N. M. Linke, and A. F. Kemper, Many-body thermodynamics on quantum computers via partition function zeros, *Sci. Adv.* **7**, eabf2447 (2021).
- [77] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, *Nature (London)* **551**, 579 (2017).
- [78] A. H. Karamlou, J. Braumüller, Y. Yanay, A. Di Paolo, P. M. Harrington, B. Kannan, D. Kim, M. Kjaergaard, A. Melville, S. Muschinske, B. M. Niedzielski, A. Vepsäläinen, R. Winik, J. L. Yoder, M. Schwartz, C. Tahan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, Quantum transport and localization in 1d and 2d tight-binding lattices, *npj Quantum Inf.* **8**, 35 (2022).
- [79] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, Probabilistic programming in python using PyMC3, *PeerJ Computer Science* **2**, e55 (2016).