Entanglement clustering for ground-stateable quantum many-body states

Michael Matty,[1,*] Yi Zhang [ORCID],[1,2] T. Senthil,[3] and Eun-Ah Kim [ORCID][1,†]

[1]*Department of Physics, Cornell University, Ithaca, New York 14853, USA*
[2]*International Center for Quantum Materials, Peking University, Beijing 100871, China*
[3]*Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

Despite their fundamental importance in dictating the quantum-mechanical properties of a system, ground states of many-body local quantum Hamiltonians form a set of measure zero in the many-body Hilbert space. Hence determining whether a given many-body quantum state is ground-stateable is a challenging task. Here we propose an unsupervised machine learning approach, dubbed Entanglement Clustering ("EntanCl"), to separate out ground-stateable wave functions from those that must be excited-state wave functions using entanglement structure information. EntanCl uses snapshots of an ensemble of swap operators as input and projects these high-dimensional data to two dimensions, preserving important topological features of the data associated with distinct entanglement structure using the uniform manifold approximation and projection. The projected data are then clustered using K-means clustering with $k = 2$. By applying EntanCl to two examples, a one-dimensional free fermion model and the two-dimensional toric code, we demonstrate that EntanCl can successfully separate ground states from excited states with high computational efficiency. Being independent of a Hamiltonian and associated energy estimates, EntanCl offers a new paradigm for addressing quantum many-body wave functions in a computationally efficient manner.

I. INTRODUCTION

Quantum many-body wave functions are complex objects that encode a great deal of information. However, interpreting this information is difficult due to the exponential number of parameters in the wave function and the need for a technique to interpret those parameters. In particular, we are interested in separating out wave functions that can be ground states of local Hamiltonians from the exponentially large space of all wave functions. Unfortunately, such "ground-statable" wave functions likely form a set of measure zero in the full many-body Hilbert space [1–4]. Although the typical approach to wave functions is to measure their energies against a particular Hamiltonian of interest, such ranking by energy is subject to change when details of the Hamiltonian change.

As an alternative to resorting to a Hamiltonian, one could turn to entanglement properties. In particular, given a partitioning of a system into two subregions A and B, the scaling of the (Von Neumann) entanglement entropy $S_A = -\mathrm{Tr}\rho_A \ln \rho_A$, where ρ_A is the reduced density matrix of subregion A, can help determine ground-stateability [5]. Ground-stateable wave functions typically exhibit S_A that scales as the codimension-1

*mfm94@cornell.edu
†eun-ah.kim@cornell.edu

boundary of the cut between subregions A and B (area law), while that of non-ground-stateable wave functions typically scales as a codimension-0 boundary (volume law). Such a distinction has indeed previously been used to distinguish ground-stateable and non-ground-stateable wave functions (see, for example, [6–12]). However, at a practical level, an investigation of the entanglement entropy scaling is often prohibitively expensive, and the finite-size effects can make it challenging to declare area or volume law with confidence. Clearly, a computationally efficient approach to separate out ground-stateable wave functions in an unbiased fashion is much desired.

Here we introduce "EntanCl" (Entanglement Clustering), a machine learning approach designed to learn the entanglement structure of many-body quantum states and separate ground-stateable states from the rest of the Hilbert space in a computationally efficient yet unbiased manner. In this initial study of EntanCl, we do not provide an exhaustive proof of this stated aim, but rather we will primarily study applications to eigenstates of a single Hamiltonian or a family of similar Hamiltonians. Also, we would like to emphasize that the purpose of EntanCl is not to calculate and compare the entanglement entropy of wave functions, but to learn structure in noisy snapshots of an entanglement-based quantity. Increasingly, the quantum condensed-matter community is successfully applying machine learning approaches to various tasks such as phase recognition [13–33], hypothesis tests on experimental data [34,35], and compact representation of many-body wave functions [18,24,36–44]. A common feature among these different problems that motivates the use of machine learning approaches is the need to find structure in

voluminous and complex data. However, the vast majority of applications so far use supervised learning, which requires labeled training data, and researchers' bias gets built into the labeling of the training data. Without the preconceived notion of what makes a wave function ground-stateable, we would like to separate out ground-stateable wave functions by learning the entanglement structure inherent in the many-body wave functions. For this, EntanCl uses Monte Carlo snapshots of the swap operator as the subsystem partition scans over the system. Then it employs uniform manifold approximation and projection (UMAP) [45], which is an *unsupervised* machine learning (ML) approach of manifold learning in high-dimensional spaces to project the data down to a two-dimensional space. The final step of EntanCl is to cluster using K-means clustering.

We will demonstrate the effectiveness of EntanCl by applying the method to many-body states associated with two specific models: a one-dimensional free fermion model and Kitaev's toric code [46] in two dimensions. The models are chosen to be representative of cases in which the ground states and excited states are distinguished by entanglement structure, and they are useful benchmarking cases because we know precisely what the ground states are. For any ML approach to data to be successful, it is critical to select relevant features to be fed into the ML algorithm. Motivated by the previously established importance of entanglement properties in determining ground-stateability, we will use an ensemble of swap operators [47] as feature selectors for our wave functions.

The rest of the paper is organized as follows. In Sec. II, we introduce and describe the three steps of EntanCl. In Sec. III, we apply EntanCl to a simple, one-dimensional free fermion model and study the accuracy of our method in classifying wave functions. In Sec. IV, we apply EntanCl to a strongly correlated problem: Kitaev's toric code [46]. In Sec. V, we summarize our conclusions and discuss possible future applications.

II. METHODS

EntanCl consists of three steps. The first step is to construct the input data of swap operator snapshots. In search of the right feature selection approach, we are inspired by the use of the swap operator in calculating Renyi entropies [47]. The action of the swap operator is illustrated in Fig. 1. The expectation value of the swap operator in the state $|\Psi\rangle = \sum_{\alpha,\beta} C_{\alpha\beta}|\alpha\beta\rangle$ is given by

$$\langle \mathbf{swap}_A \rangle = e^{-S_2} = \sum_{\alpha,\beta,\alpha',\beta'} |C_{\alpha\beta}|^2 |C_{\alpha'\beta'}|^2 \frac{C_{\alpha'\beta}C_{\alpha\beta'}}{C_{\alpha\beta}C_{\alpha'\beta'}}, \quad (1)$$

where S_2 denotes the second Renyi entropy, A denotes a subsystem, the quantum numbers α describe subsystem A, and β describe the remainder of the system. We will *not* take the expectation value, however. Instead, we will variationally sample the swap data for $|\Psi\rangle = \sum_{\alpha,\beta} C_{\alpha\beta}|\alpha\beta\rangle$ according to Eq. (1), where $|C_{\alpha\beta}|^2 \times |C_{\alpha'\beta'}|^2$ plays the role of the sampling weights. To acquire more comprehensive data across the system, we will consider many subsystems A_i to form an ensemble of swap operators $\{\mathbf{swap}_{A_i}\}$.
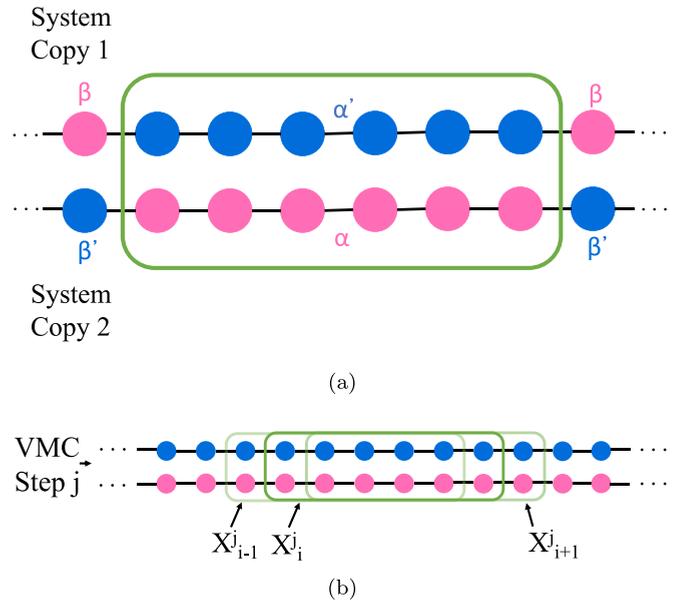

FIG. 1. (a) Schematic depiction of the action of the swap operator on a subsystem **A**. The quantum numbers α describe the subsystem, and β describe the remainder of the system. Since **swap** acts on a doubled Hilbert space, we denote the quantum numbers belonging to one copy by primed variables and those belonging to the other by unprimed variables. The operator **swap**$_A$ switches the primed and unprimed variables within the region **A**. (b) Illustration of our data collection procedure. At each VMC step j, we collect swap data from a collection of subsystems A_i and store each in a vector $\vec{X}^j$ at index i. The collection of $\vec{X}^j$'s forms our complete dataset **X**.

As we sample the swap data with variational Monte Carlo (VMC), we build up a collection of vectors $\mathbf{X} = \{\vec{X}^j\}$ (cf. Fig. 1) where at index i, $\vec{X}^j$ contains the data $C_{\alpha'\beta}C_{\alpha\beta'}/C_{\alpha\beta}C_{\alpha'\beta'}$ sampled from **swap**$_{A_i}$ at VMC step j. The dimensionality of our data is precisely the number of subsystems A_i we choose to consider. This will be of order hundreds of dimensions for the free fermion model and thousands for the toric code. We thus have a high-dimensional data set **X** that contains entanglement information about the wave function $|\Psi\rangle$.

The second step of EntanCl is to project the input data living in the high-dimensional space (typically hundreds or thousands of dimensions) down to two-dimensional space in which clustering can be visualized. Typical applications of unsupervised ML to high-dimensional data sets involve visualizing the data in a low-dimensional space via dimensional reduction. Dimensional reduction algorithms (such as those described in Refs. [48–55]) vary in the way that they approximate the high-dimensional manifold populated by the data and what features of that manifold they try to preserve under projection to the low-dimensional space. We are interested in an algorithm that will allow us to visualize the cluster structure in our swap data set **X**. This is because we expect that those $\vec{X}^j$ obtained from ground-stateable and non-ground-stateable wave functions will appear as two separate clusters due to differing entanglement structure.

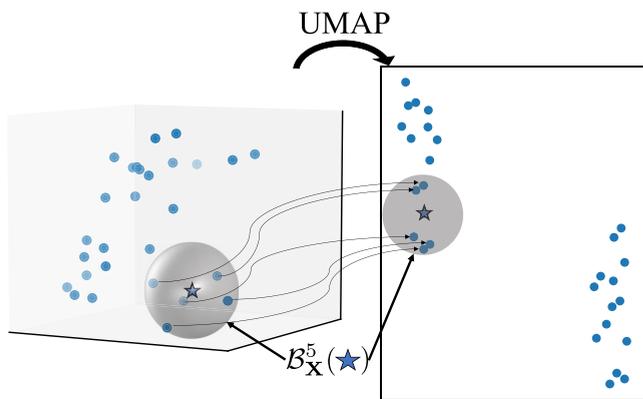We can view clusters from a neighborhood perspective. As an example, in Fig. 2 we consider three-dimensional data

FIG. 2. Schematic illustration of "neighborhood structure" preservation, projecting points in three dimensions to two. The five nearest neighbors of the star are found by application of $\mathcal{B}_{\mathbf{X}}^5$. After projection, we can see that the five nearest neighbors of the point marked with a star remain its five nearest neighbors. Moreover, by preserving local neighborhoods, we have discovered two distinct clusters in the high-dimensional data. For this example, the projection was done by UMAP.

consisting of two clusters: 15 points randomly generated on the upper hemisphere of a unit radius sphere, and 15 generated on the lower hemisphere. Gaussian noise is applied to the coordinates of the points. We then project the points down to two dimensions so as to preserve their local neighborhood structure. In this case, we use UMAP to do the projection. On the right-hand panel of Fig. 2, we can see that in each of the two clusters, the local neighborhoods of each point are entirely contained within the same cluster as the point. To emphasize this, we illustrate a local neighborhood of size five around the point marked by a star. From this we can infer that preserving local neighborhood structure also preserves cluster structure. Formally, define a function $\mathcal{B}_{\mathbf{X}}^m$ such that $\mathcal{B}_{\mathbf{X}}^m(\vec{X}_*) \subseteq \mathbf{X}$ is the set of the m nearest neighbors of $\vec{X}_*$ in $\mathbf{X}$. A cluster is then a subset $\mathbf{C} \subseteq \mathbf{X}$ such that $\mathcal{B}_{\mathbf{X}}^m(\vec{C} \in \mathbf{C}) \subseteq \mathbf{C}$. For visualizing clusters, a natural choice for a dimensional reduction algorithm is then one that preserves neighborhoods after projection.

Algorithms that preserve neighborhood structure [48–52] try to find a mapping $\mathcal{P}$ from the D-dimensional data space to $\mathbb{R}^d$ (again, $\mathbb{R}^2$ for us), such that $\mathcal{P} \circ \mathcal{B}_{\mathbf{X}}^m = \mathcal{B}_{\mathcal{P}(\mathbf{X})}^m \circ \mathcal{P}$, where $\circ$ denotes the usual composition of mappings. Observe that preserving neighborhoods entails not only keeping points within a cluster nearby, but keeping points in separate clusters far away from each other. Common algorithms accomplish this by taking as input a hyperparameter that defines an estimated neighborhood or cluster size, related to the m in our definition of $\mathcal{B}_{\mathbf{X}}^m$. These algorithms treat the effective distance between points outside of a neighborhood as extremely (or sometimes infinitely) far away. One must be sure to choose this hyperparameter large enough (based on the density of the data) that spurious clusters do not appear in the projected data. That is to say, the intersection of the neighborhoods $\mathcal{B}_{\mathbf{X}}^m$ need to contain the entire, true cluster. For our purposes, we use UMAP, which has previously found use in biology [56–64], materials engineering [65], and machine learning [66–68],

but has had limited use in quantum matter [69]. For more details about how UMAP in particular works, see Appendix A. We choose UMAP from the various unsupervised ML algorithms that seek to preserve neighborhood structures for two reasons. First, it led to the clearest projected clustering for our purposes. Second, in contrast to other algorithms like tSNE, UMAP provides us with a generalizable mapping that can be applied immediately to new data without rerunning UMAP.

The final step of EntanCl is to interpret the learned UMAP output using k-means clustering. K-means clustering partitions a set of data points into k clusters by placing k cluster means (centroids) in a way that minimizes the sum of squared distances from each data point to its nearest centroid. A ($k = 2$)-means clustering thus naturally allows us to classify (non-)ground-stateable wave functions in the 2D projected space. For our test cases where we know which cluster corresponds to each type of wave function, we define a metric of accuracy given by assignment to the correct centroid.

III. FREE FERMION MODEL

To establish EntanCl on a simple, known model, we first study a one-dimensional free fermion model. This model is described by the Hamiltonian

$$\mathcal{H} = \sum_i (t_1 b_i^\dagger a_i + t_2 a_{i+1}^\dagger b_i) + \text{H.c.} \qquad (2)$$

This model has two bands with energy gap $\Delta E \sim |t_2 - t_1|$, and we consider the case of half-filling. We report results in terms of the dimensionless, normalized gap $t \equiv |t_2 - t_1|/t_1$. The ground-state Slater determinant wave function of the half-filled system corresponds to completely filling the lower band. The non-ground-stateable eigenstates we consider have some fixed density $n_{\text{ex}} \equiv N_{\text{ex}}/L$ of randomly chosen k-points promoted to the upper band, where L is the system size. This model gives us a testbed to identify ground-state wave functions and non-ground-stateable wave functions in the parameter space of energy gap ΔE and excited k-point density n_{ex}.

The ensemble of swap operators we use in this case is the set of all length six contiguous subsystems of an $L = 100$ chain. Our dataset $\mathbf{X}$ consists of 1000, 100-dimensional swap vectors $\vec{X}^j$ corresponding to the ground state and 1500 corresponding to a non-ground-stateable wave function. We choose an uneven ratio of swap data from the two classes to illustrate that a symmetric amount of data is nonessential to our technique. We project the data to two dimensions via UMAP and assign the projected data points to clusters with k-means. Since we know which swap data points came from (non-)ground-stateable wave functions, we also calculate the accuracy.

Our results are shown in Fig. 3. Figure 3(a) corresponds to a projection with the normalized gap $t = 2$ and excitation density $n_{\text{ex}} = 3\%$. In this case, one can clearly see the success of EntanCl: the data corresponding to the ground-stateable wave function (red) and the non-ground-stateable wave function (green) appear as two well-separated clusters. This case corresponds to an accuracy of 99.12%. In Figs. 3(b) and 3(c) we can see that as both t and excitation density increase, the
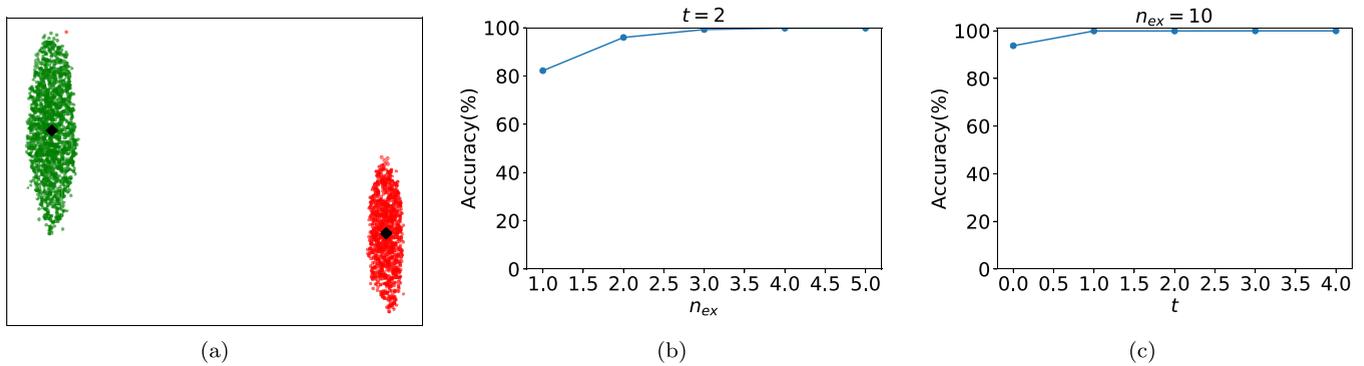
FIG. 3. (a) UMAP projection of swap data obtained from wave functions for the free fermion model. Red dots correspond to swap data from a ground-stateable wave function. Green dots correspond to swap data from a non-ground-stateable wave function with $n_{ex} = 3\%$ and $t = 2$. Black diamonds denote the $(k = 2)$-means clustering centroids. This case has accuracy 96.52%. We also show the accuracy as a function of (b) excitation density n_{ex} at normalized energy gap $t = 2.0$ and (c) t at $n_{ex} = 10\%$. In both cases, accuracy increases as a function of the relevant parameter, and moreover it stays relatively high at the minimum possible value.

accuracy also increases. This makes sense: as both t and n_{ex} increase, the excited state becomes more entangled compared to the ground state as the entanglement entropy scaling transitions from area law to volume law. Moreover, the accuracy stays high even at the lowest possible n_{ex} (80.00% for $t = 2$) and for a gapless system (90.03% for $n_{ex} = 10\%$). This demonstrates that EntanCl is a viable method of identifying the differing entanglement structure in ground-stateable and non-ground-stateable wave functions.

The learned UMAP projection is generalizable. In Fig. 4 we illustrate the results of using the UMAP projection optimized with swap data obtained from the ground-stateable wave function and a single non-ground-stateable wave function (i.e., a single choice of excited k-points) with $t = 2$ and $n_{ex} = 2\%$ to four more non-ground-stateable wave functions with the same t and n_{ex}. We collect 1000 MC samples for

the ground-stateable wave function and 1500 for each non-ground-stateable wave function. The projection map clusters all the data from non-ground-stateable wave functions together, away from the data from the ground-stateable wave function. The accuracy in this case is 84.4%, lower than the 96.6% in Fig. 3(b) for two wave functions. This is because most of the error is non-ground-stateable data being misclassified as ground-stateable. Increasing the amount of data collected from the ground-stateable wave function would increase the accuracy. These results show that the structure that UMAP is learning generalizes well.

We additionally consider an example using swap data from ground-stateable wave functions ($n_{ex} = 0\%$) at normalized gap $t = 0$, 1, and 2. We can see in Fig. 5 that that although the data from the gapped states appear as one cluster, the data from the gapless state appear as a separate cluster. This is presumably due to the logarithmic correction to the entanglement entropy in the gapless system. If we consider as well a non-ground-stateable wave function with $n_{ex} = 10\%$ and $t = 0$ and allow for three clusters instead of two due to the separation of the ground states, we find that the output clusters

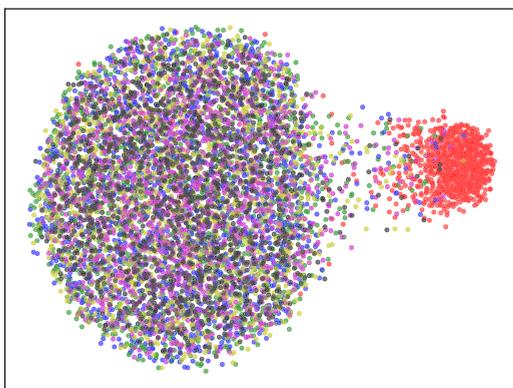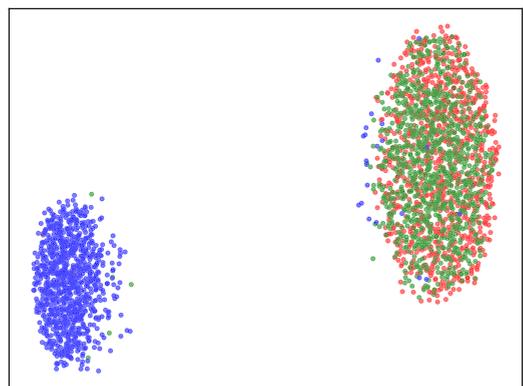FIG. 4. UMAP projection of swap data from free fermion model wave functions at gap $t = 2$ and excitation density $n_{ex} = 2\%$. The UMAP projection was optimized using the ground state and a single excited-state configuration (i.e., single choice of excited k-points). We then use the same mapping on four more excited-state configurations and display the results simultaneously. The ground-state data are shown in red; the other colors correspond to various excited-state configurations. Clearly, subsequent excited states cluster together with each other, and more importantly all cluster separately from the ground state.

FIG. 5. UMAP projection of swap data obtained from the free fermion model for several ground-stateable wave functions ($n_{ex} = 0\%$) at normalized gaps $t = 0$, 1, 2. The data from the gapped states (red and green dots) cluster together, and they do so separately from the data from the gapless state (blue dots).
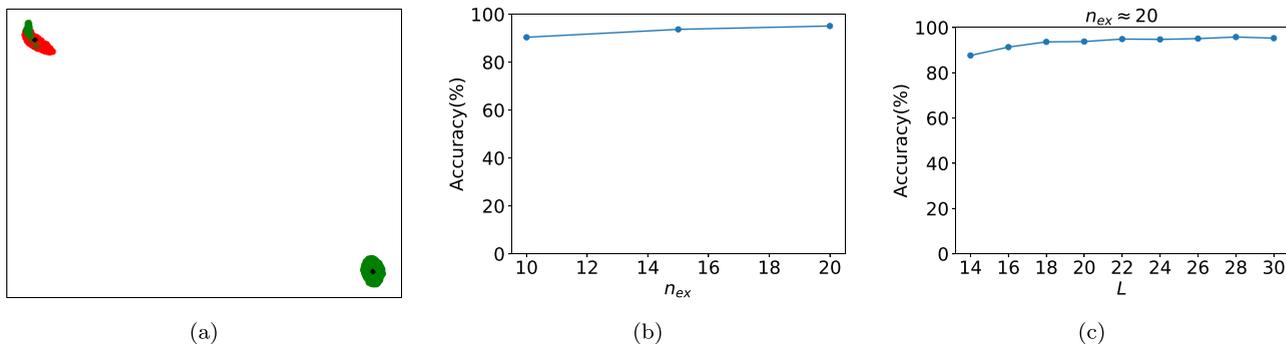
FIG. 6. (a) UMAP projection of swap data obtained from wave functions for the toric code. Red dots again correspond to swap data from ground-stateable wave functions. Green dots correspond to swap data from a non-ground-stateable wave function from a lattice with linear dimension $L = 25$ with spinon density $n_{ex} = 20\%$. Black diamonds denote the ($k = 2$)-means clustering centroids, and this case corresponds to accuracy 95.91%. (b) The accuracy at a fixed lattice size grows with n_{ex}, as expected. (c) Classification accuracy for UMAP projection of toric code wave functions as a function of lattice linear dimension. Data shown are at spinon density $\sim 20\%$. Accuracy increases with system size and plateaus around 95%. A slight nonmonotonicity near the plateau is expected because n_{ex} must be an even integer and is therefore not exactly 20% for all lattice sizes.

with 95.82% accuracy. However, if we include additional excited states with $n_{ex} = 10\%$ and $t = 1, 2$, the accuracy drops to 54%.

IV. TORIC CODE

We now turn to a two-dimensional example: Kitaev's toric code [46]. This is a strongly interacting system whose ground state has topological order, and because it is exactly solvable, we will be able to assess the accuracy of EntanCl. This model is defined on a square lattice with spin-1/2 variables living on the edges. The wave functions that we will consider in this case are eigenstates of the Hamiltonian

$$\mathcal{H} = -\sum_{\square} A_{\square} - \sum_{v} B_{v}, \qquad (3)$$

where the operators

$$A_{\square} = \prod_{i\in\square} \sigma_i^x, \ B_v = \prod_{i\in\partial v} \sigma_i^z \qquad (4)$$

are defined as the product of Pauli σ^x operators around a plaquette and σ^z operators on the edges incident on a vertex v, respectively. Note that we will be working in the σ^z basis.

The ground-state wave function we will consider is the equal amplitude superposition of all lattice configurations of closed loops in the trivial homology class [70]. The non-ground-stateable wave functions we will consider are equal amplitude superpositions of all states with a fixed spinon density (also allowing closed loops) where a spinon is a vertex v with $B_v = -1$. Note that this does not correspond to fixed spinon *locations*, as such wave functions could be made ground states by simply flipping the sign of the B_v's corresponding to the spinon locations. With this model, we will classify wave functions at different values of our control parameter: the spinon density n_{ex}.

We collect swap data at 1000 uncorrelated VMC time steps for each wave function we consider. The ensemble of swap operators we use in this case consists of all rectangular subregions of the lattice, which grows with the linear dimension of the lattice L as L^4. Due to the massively increased

dimensionality of the swap data in this case, we add a preprocessing step to compress the data volume for RAM storage, especially for larger system sizes. We average the swap data for a fixed subsystem width and height over all basepoints for the subsystem. This reduces the dimensionality of the data to L^2, which is sufficiently tractable for our purposes. With this addition to our analysis, we can project the swap data to two dimensions via UMAP [71].

Our results for the toric code are shown in Fig. 6. We find that we can achieve 95.91% accuracy for $n_{ex} = 20\%$ for a lattice with linear dimension $L = 25$ as shown in Fig. 6(a). For a lattice with linear dimension $L = 35$, we get accuracy 99.1% even at $n_{ex} = 5\%$. Once again, for this high accuracy case, the success of the clustering is remarkably clear. In Fig. 6(b), we can see that the accuracy also increases with n_{ex}, as we would expect. Moreover, we do not need such a large system to achieve good accuracy. We can see in Fig. 6(c) that for $n_{ex} = 20\%$, the accuracy of the projection is over 90% already at $L = 16$.

We now turn to a benchmarking exercise for generalizability. Due to topological degeneracy, we have access to four ground-stateable wave functions from the toric code. In Fig. 7 we show the results of optimizing the UMAP projection mapping for an $L = 20$ lattice using the ground-stateable wave function containing only homologically trivial loops and the non-ground-stateable wave function with $n_{ex} = 20\%$. We then generalize the projection map to swap data obtained from the other three ground-stateable wave functions (those with an odd parity of noncontractible loops around one or both cycles of the torus). This is not a surprise given that all states were drawn from the same probability distribution. Nevertheless, it confirms that EntanCl works as expected. In Fig. 7, we can see that the data from the non-ground-stateable wave function (purple dots) cluster separately from the ground-stateable data (other colors), which indeed all cluster together. The accuracy of the collective projection is 98.04%, compared to 95.1% from the initial data used to optimize the projection map. This makes sense because the only errors are non-ground-stateable data being classified as ground-stateable, so adding more ground-stateable data reduces the error. This shows that
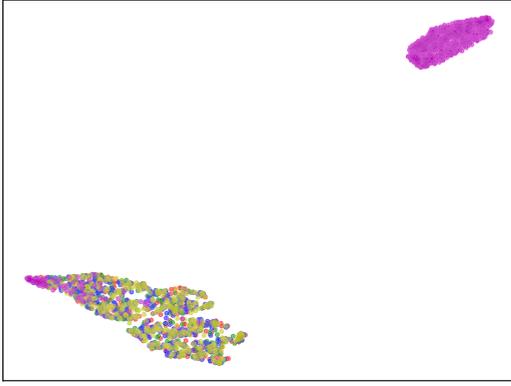
FIG. 7. UMAP projection of swap data obtained from the toric code on a 20×20 lattice for all four topologically degenerate ground states and an excited state at excitation density $n_{ex} = 20\%$. The projection was optimized using only data from the ground state consisting of only homologically trivial loops and the excited state. Then we subsequently apply the projection to the other three ground states. The purple dots are data from the excited state; the other colors are from the ground states. The overall accuracy is 98.04%. All of the ground-stateable data cluster together, and more importantly they cluster separately from the excited state except for the small fraction of errors.
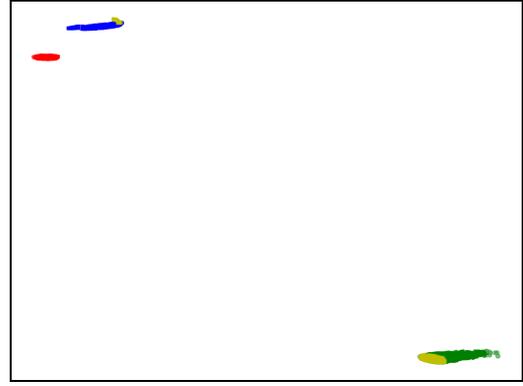

FIG. 8. UMAP projection of swap data obtained from both the free fermion model and the toric code. Red dots correspond to a ground-stateable free fermion wave function with ($n_{ex} = 0\%$, $t = 2$), blue dots to a ground-stateable toric code wave function with no nontrivial loops, green dots to a non-ground-stateable free fermion wave function with ($n_{ex} = 10\%$, $t = 2$), and yellow dots to a non-ground-stateable toric code wave function with $n_{ex} = 10\%$. Using $k = 2$-means clustering, the ground states cluster separately from the excited states with an accuracy of 97.75%. Note, however, that the toric code and free fermion ground states cluster form disjoint (but nearby) clusters.

the learned UMAP projection optimized on one ground state generalizes to other ground states in the presence of topological degeneracy.

Another interesting feature of the clustering in this case is that misclassifications are always excited states being incorrectly classified as ground states. The distinction between the ground state and the excited state is the presence of spinons and the string operators connecting them. To detect the excited nature of the wave function, a swap operator must swap a subsystem in a way that cuts a string operator. We therefore conjecture that misclassifications of MC samples from excited states as ground states is due to VMC configurations in which the string operators connecting spinons are sufficiently short such that very few subsystems pick up the excited character of the wave function.

Additionally, we study the effects of considering wave functions from the toric code and the free fermion model simultaneously. To do so, we extend the dimensionality of the swap data vectors collected from the free fermion wave functions to match that of the toric code data by simply appending zeros. We take data from one ground-stateable wave function from the free fermion model ($n_{ex} = 0\%$, $t = 2$) and one from the toric code (the one with no nontrivial loops). We also take one non-ground-stateable wave function from each at $n_{ex} = 10\%$. In Fig. 8 we can see that with two clusters, both ground-stateable wave functions would be assigned to one cluster and both non-ground-stateable wave functions to the other with accuracy 97.45%. The non-ground-stateable wave functions all cluster together, which is surprising.

The two systems differ in their dimensionality and the geometry of the subsystems used to collect swap data. We contend that this demonstrates that EntanCl has learned something meaningful about the data that goes beyond the specifics of individual datasets. However, at the moment we are unable

to say what that was, and we leave further exploration of this result as a subject of future work. The ground-stateable toric code wave function forms a separate cluster from the ground-stateable free fermion wave function (albeit relatively nearby). Again, the entanglement structure of these wave functions differs: that of the toric code has an additional constant contribution due to topological order. Also note that adding more non-ground-stateable wave functions does not decrease the accuracy. The data from the ground-stateable wave function corresponding to the gapless limit of the free fermion model, however, cluster with the data from the non-ground-stateable wave functions in this case.

In this example and the one presented in Fig. 5, we have seen that although two wave functions may both be ground-stateable, differences in entanglement structure (e.g., log corrections and topological terms) may result in them forming separate clusters in the EntanCl output. The non-ground-stateable wave functions also formed a single cluster in all the cases we studied. This suggests that in more general situations with different kinds of potential ground-stateable wave functions, it may be more useful to assess ground-stateability by looking for multiple clusters in EntanCl output separate from a larger cluster of non-ground-stateable wave functions and not necessarily as a single cluster. We also saw that the gapless ground state is poorly separated from gapped excited states. Thus the entanglement structure of the ground state can cause EntanCl to fail when including both ground and excited states from wave functions with multiple types of entanglement structure.

V. CONCLUSION

In summary we introduced EntanCl, an unsupervised machine learning method to separate out the ground-stateable

wave functions from the exponentially large Hilbert space of many-body wave functions with high computational efficiency. EntanCl consists of three steps: (i) preparation of input data, (ii) projection of the data down to two-dimensional space using UMAP, and (iii) K-means clustering of the projected data. The input data of our choice are matrix elements of an ensemble of swap operators collected as snapshots of individual uncorrelated variational Monte Carlo steps. By using the noisy snapshots as opposed to demanding convergence of the swap operator expectation value, EntanCl gains computational efficiency. For example, for the case of the free fermion with normalized gap $t = 3$ and excitation density $n_{ex} = 0\%$ ($n_{ex}10\%$), it takes $O(10^6–10^7)$ $[O(10^5–10^6)]$ (depending on subsystem length) uncorrelated MC steps and swap evaluations for $\langle swap_A \rangle$ to converge to five digits for a particular subsystem A. We collect data from only 10^3 uncorrelated MC steps and perform 10^5 swap evaluations, an order of magnitude fewer than it takes to converge the average to five digits. Note also that we have all the data we need from those 10^5 swap samples, while in a typical diagnosis of the entropy scaling law, the average must be converged multiple times for finite-size scaling. In this example, doing finite-size scaling for 10 subsystem sizes would take at least an order of magnitude more uncorrelated swap evaluations compared to EntanCl. Moreover, we would like to reemphasize that converging $\langle \mathbf{swap}_A \rangle$ must be done for each wave function.

We applied EntanCl to a simple one-dimensional free fermion model and Kitaev's toric code to find accurate clustering results. Moreover, we established that the learned UMAP projection is generalizable to an expansion of the data set. The clustering errors are found to occur asymmetrically: an excited state may get misplaced into the ground-state cluster but not vice versa. Hence the cluster assignment into excited states will be a reliable way of ruling out ground-stateability of the quantum many-body state. As with any VMC sampling, the quality of the results can depend on the sampling basis due to the basis dependence in the spread of the noise. As we demonstrate in Appendix B, as long as the spread of the noise remains comparable under a basis transformation, EntanCl will work independent of the basis choice.

Using EntanCl on larger collections of wave functions certainly warrants further study. Moreover, we have thus far only considered wave functions obtained from integrable systems and wave functions that are eigenstates of the models we considered. A subject of future work could be to explore applications of EntanCl to both nonintegrable systems and non-eigenstate wave functions. Challenges in these cases may arise in finding efficient ways to sample swap data.

In the same vein of addressing wave functions, a more ambitious approach would be to attempt to reconstruct the Hamiltonian that takes a given wave function as its ground state. There has been recent progress in this direction with concrete proposals [72–75]. However, the Hamiltonian reconstruction is computationally costly as it requires precise measurements of many correlation functions. EntanCl can be a swift first pass that can weed out non-ground-stateable many-body states without reference to Hamiltonians. Furthermore, as a method that can efficiently sort the swap data associated with different quantum many-body states based on the their entanglement structure, we anticipate that

EntanCl will find applications beyond separating out ground-stateable wave functions. For instance, EntanCl will be ideal for studying quantum phase transitions involving a change of entanglement structure due to spontaneous symmetry breaking or topological order [76].

APPENDIX A: OVERVIEW OF UMAP PROCEDURE

The purpose of the uniform manifold approximation and projection (UMAP) algorithm is to create a low-dimensional projection of high-dimensional data such that the nearest neighbors of a data point in high dimensions remain its nearest neighbors in the low-dimensional projection. How many nearest neighbors we try to keep is an input parameter to the algorithm. This is useful for us because data that belong to distinct clusters in the high-dimensional space will not share nearest neighbors between clusters. Thus, in the low-dimensional space, these data should still show up as distinct clusters. Here we give an overview of how this algorithm works.

(i) Let $X = \{X_1, \ldots, X_N\}$ denote our set of input data, where each X_i is an n-dimensional vector. Let $Y = \{Y_1, \ldots, Y_N\}$ denote the output projected data points, where Y_i corresponds to the projection of X_i, and each Y_i is a d-dimensional vector with $d \leqslant n$.

(ii) We would like the data to be uniformly distributed on the underlying manifold because then the collection of local neighborhoods of our data points will provide a good picture of the underlying manifold. UMAP *forces* our data to be uniformly distributed by normalizing the distance from each point to the furthest neighbor we would like to consider. We are also going to assume that there are no isolated points on the underlying manifold, which we will enforce by fixing the distance to the *nearest* neighbor. To do this, we define a local metric d_i for each input data point X_i,

$$d_i(X_j, X_k) = \begin{cases} \frac{1}{r_i} d_{\mathbb{R}^n}(X_j, X_k) - \rho_i & \text{if } i = j \text{ or } i = k, \\ \infty & \text{otherwise,} \end{cases}$$

where $d_{\mathbb{R}^n}$ is the Euclidean metric on $\mathbb{R}^n$, ρ_i fixes the distance to the nearest neighbor to be zero, and r_i fixes the distance to the furthest neighbor we would like to consider. Note that we choose r_i's so that for each d_i, the distance from X_i to its furthest relevant neighbor is the same. For the projected output, we will define local metrics as well. The difference in the projected space is that we know what the underlying

manifold is ($\mathbb{R}^d$), so we know what the true metric is. UMAP still enforces an assumption of local connectivity. Our local metrics for the encoded output Y_i's are then

$$d_i(Y_j, Y_k) = \begin{cases} d_{\mathbb{R}^d}(Y_j, Y_k) - \rho_i & \text{if } i = j \text{ or } i = k, \\ \infty & \text{otherwise.} \end{cases}$$

(iii) Comparisons of distance between our different local metrics are meaningless, which seems to give us no way to assess the quality of a projection. To circumvent this, UMAP considers a new representation of the data: a *neighborhood graph*. To build the graph, UMAP draws an edge between each data point and each of its neighbors up to the furthest one we would like to consider. The edges are weighted, where for an edge from X_i to X_j, the weight of the edge is $\exp(-d_i(X_i, X_j))$. UMAP performs the same procedure for the projected data Y. Note that $d_i(X_i, X_j)$ is not necessarily equal to $d_J(X_j, X_i)$. Thus, the edges drawn between X_i and X_j by d_i and d_j may not have the same weight.

(iv) Next UMAP combines edges so that there is at most one edge between any two points. The edges are combined pairwise, where for a pair of edges with weights α, β, UMAP forms a combined edge with weight $f(\alpha, \beta) = \alpha + \beta - \alpha \cdot \beta$. This process occurs for both the input data X and the projected data Y. The function f is not the unique way to combine edge weights, but it is a choice made by UMAP.

(v) Now we have a neighborhood graph for X and Y with an unambiguous definition of *the* edge between two points. Because the neighborhood graphs for X and Y have the same number of vertices and each vertex is the same degree, we can define an isomorphism between them. We do this by associating projected points with data points, being careful to ensure that if there is an edge between X_i and X_j, the points Y_i and Y_j that we associate with them are also connected by an edge. Thus we can speak unambiguously about a single edge set E. To measure the "similarity" of the two neighborhood graphs, we will use the cross entropy

$$C(E; \mu_\cup, \nu_\cup) \equiv \sum_{e \in E} \mu_\cup(e) \ln \left(\frac{\mu_\cup(e)}{\nu_\cup(e)} \right)$$
$$+ [1 - \mu_\cup(e)] \ln \left(\frac{1 - \mu_\cup(e)}{1 - \nu_\cup(e)} \right),$$

where E is the set of edges, $\mu_\cup(e)$ is the combined weight [as in step (iv)] of an edge in Y, and $\nu_\cup(e)$ is the combined weight of an edge in X. We can minimize the cross entropy using stochastic gradient descent. For each step of the optimization, we move the positions of the encoded points, changing the distance, and therefore the edge weights, between them.

APPENDIX B: EXAMPLE OF BASIS DEPENDENCE

A basis transformation can affect the spread in the VMC data obtained during step (i) of EntanCl by changing the relative magnitudes of the coefficients $C_{\alpha\beta}$ in the wave function
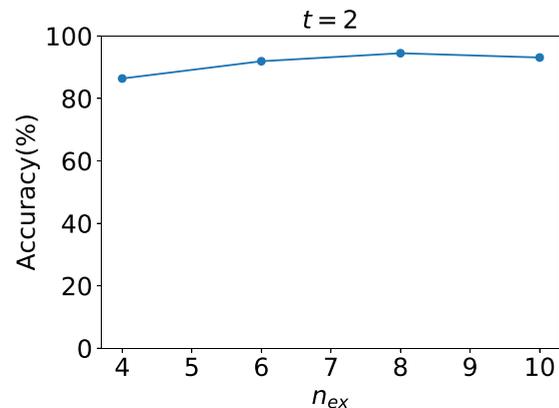

FIG. 9. Here we show the clustering accuracy for swap data obtained from the ground-state wave function of $\mathcal{H}'_k$ [cf. Eq. (B3)] and non-ground-stateable wave functions with normalized energy gap $t = 2$ and varying excitation density n_{ex}. Although the accuracy at similar n_{ex} is lower for the model in this basis than the original [cf. Fig. 3(b)], the accuracy is still high (peaking over 90%) and stays above 80% even at low n_{ex} values.

[cf. Eq. (1)]. This change in the spread of the data can affect the accuracy of the resultant clustering if the neighborhoods of MC samples from ground-stateable wave functions intersect those of non-ground-stateable wave functions in the high-dimensional space. Here we discuss an example of the basis dependence of our results by reexamining the free fermion model of Sec. III under a basis transformation. The k-space Hamiltonian for the original free fermion model is given by

$$\mathcal{H}_k = [t_1 + t_2 \cos(k)]\sigma^x - t_2 \sin(k)\sigma^y, \quad \text{(B1)}$$

where the σ^i's are Pauli matrices. We now consider a new model that differs from the original by an SU(2) unitary transformation with Hamiltonian

$$\mathcal{H}' = \sum_i t_1(a_i^\dagger a_i - b_i^\dagger b_i)$$
$$+ \frac{t_2}{2}(a_{i+1}^\dagger a_i - b_{i+1}^\dagger b_i + b_{i+1}^\dagger a_i - b_{i-1}^\dagger a_i + \text{H.c.}), \quad \text{(B2)}$$

$$\mathcal{H}'_k = [t_1 + t_2 \cos(k)]\sigma^z + t_2 \sin(k)\sigma^y. \quad \text{(B3)}$$

This new model $\mathcal{H}'_k$ describes the same physics as $\mathcal{H}_k$, but it differs by a basis transformation. We show the clustering accuracy results of scaling the excitation density n_{ex} at fixed normalized gap $t = 2$ in Fig. 9. We can see that, as was the case in Fig. 3, the accuracy is high and remains high even at low n_{ex} values. However, the accuracy in this basis is not as high as in the original basis at the same n_{ex} values. This illustrates that noise in the VMC data does indeed carry a basis dependence, but that sampling data in a new basis does not necessarily destroy the separability of the swap data from ground-stateable and non-ground-stateable wave functions.

[1] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, Rev. Mod. Phys. **82**, 277 (2010).

[2] D. N. Page, Average Entropy of a Subsystem, Phys. Rev. Lett. **71**, 1291 (1993).

[3] S. K. Foong and S. Kanno, Proof of Page's Conjecture on the Average Entropy of a Subsystem, Phys. Rev. Lett. **72**, 1148 (1994).

[4] S. Sen, Average Entropy of a Quantum Subsystem, Phys. Rev. Lett. **77**, 1 (1996).

[5] M. Srednicki, Entropy and Area, Phys. Rev. Lett. **71**, 666 (1993).

[6] L. Vidmar, L. Hackl, E. Bianchi, and M. Rigol, Volume Law and Quantum Criticality in the Entanglement Entropy of Excited Eigenstates of the Quantum Ising Model, Phys. Rev. Lett. **121**, 220602 (2018).

[7] L. Vidmar, L. Hackl, E. Bianchi, and M. Rigol, Entanglement Entropy of Eigenstates of Quadratic Fermionic Hamiltonians, Phys. Rev. Lett. **119**, 020601 (2017).

[8] J. P. Keating, N. Linden, and H. J. Wells, Spectra and eigenstates of spin chain hamiltonians, Commun. Math. Phys. **338**, 81 (2015).

[9] M. Storms and R. R. P. Singh, Entanglement in ground and excited states of gapped free-fermion systems and their relationship with Fermi surface and thermodynamic equilibrium properties, Phys. Rev. E **89**, 012125 (2014).

[10] F. Ares, J. G. Esteve, F. Falceto, and E. Sánchez-Burillo, Excited state entanglement in homogeneous fermionic chains, J. Phys. A **47**, 245301 (2014).

[11] V. Alba, M. Fagotti, and P. Calabrese, Entanglement entropy of excited states, J. Stat. Mech.: Theor. Expt. (2009) P10020.

[12] Q. Miao and T. Barthel, Eigenstate entanglement: Crossover from the ground state to volume laws, arXiv:1905.07760.

[13] P. Broecker, J. Carrasquilla, R. G. Melko, and S. Trebst, Machine learning quantum phases of matter beyond the fermion sign problem, Sci. Rep. **7**, 8823 (2017).

[14] P. Broecker, F. F. Assaad, and S. Trebst, Quantum phase recognition via unsupervised machine learning, arXiv:1707.00663.

[15] Y. Zhang and E.-A. Kim, Quantum Loop Topography for Machine Learning, Phys. Rev. Lett. **118**, 216401 (2017).

[16] Y. Zhang, R. G. Melko, and E.-A. Kim, Machine learning $\mathbb{z}_2$ quantum spin liquids with quasiparticle statistics, Phys. Rev. B **96**, 245119 (2017).

[17] L. Wang, Discovering phase transitions with unsupervised learning, Phys. Rev. B **94**, 195105 (2016).

[18] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, Science **355**, 602 (2017).

[19] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, Nat. Phys. **13**, 431 (2017).

[20] E. P. L. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Learning phase transitions by confusion, Nat. Phys. **13**, 435 (2017).

[21] M. J. S. Beach, A. Golubeva, and R. G. Melko, Machine learning vortices at the Kosterlitz-Thouless transition, Phys. Rev. B **97**, 045207 (2018).

[22] K. Ch'ng, J. Carrasquilla, R. G. Melko, and E. Khatami, Machine Learning Phases of Strongly Correlated Fermions, Phys. Rev. X **7**, 031038 (2017).

[23] K. Ch'ng, N. Vazquez, and E. Khatami, Unsupervised machine learning account of magnetic transitions in the Hubbard model, Phys. Rev. E **97**, 013306 (2018).

[24] D.-L. Deng, X. Li, and S. Das Sarma, Machine learning topological states, Phys. Rev. B **96**, 195145 (2017).

[25] Y.-H. Liu and E. P. L. van Nieuwenburg, Discriminative Cooperative Networks for Detecting Phase Transitions, Phys. Rev. Lett. **120**, 176401 (2018).

[26] E. van Nieuwenburg, E. Bairey, and G. Refael, Learning phase transitions from dynamics, Phys. Rev. B **98**, 060301(R) (2018).

[27] T. Ohtsuki and T. Ohtsuki, Deep learning the quantum phase transitions in random two-dimensional electron systems, J. Phys. Soc. Jpn. **85**, 123706 (2016).

[28] F. Schindler, N. Regnault, and T. Neupert, Probing many-body localization with neural networks, Phys. Rev. B **95**, 245134 (2017).

[29] S. J. Wetzel and M. Scherzer, Machine learning of explicit order parameters: From the Ising model to su(2) lattice gauge theory, Phys. Rev. B **96**, 184410 (2017).

[30] S. J. Wetzel, Unsupervised learning of phase transitions: From principal component analysis to variational autoencoders, Phys. Rev. E **96**, 022140 (2017).

[31] N. Yoshioka, Y. Akagi, and H. Katsura, Learning disordered topological phases by statistical recovery of symmetry, Phys. Rev. B **97**, 205110 (2018).

[32] J. Venderley, V. Khemani, and E.-A. Kim, Machine Learning Out-of-Equilibrium Phases of Matter, Phys. Rev. Lett. **120**, 257204 (2018).

[33] M. Matty, Y. Zhang, Z. Papic, and E.-A. Kim, Multi-faceted machine learning of competing orders in disordered interacting systems, Phys. Rev. B **100**, 155141 (2019).

[34] S. Ghosh, M. Matty, R. Baumbach, E. D. Bauer, K. A. Modic, A. Shekhter, J. A. Mydosh, E.-A. Kim, and B. J. Ramshaw, Single-component order parameter in URu_2Si_2 uncovered by resonant ultrasound spectroscopy and machine learning, Sci. Adv. **6**, eaaz4074 (2020).

[35] Y. Zhang, A. Mesaros, K. Fujita, S. D. Edkins, M. H. Hamidian, K. Ch'ng, H. Eisaki, S. Uchida, J. C. Davis, E. Khatami *et al.*, Using machine learning for scientific discovery in electronic quantum matter visualization experiments, arXiv:1808.00479.

[36] Z. Cai and J. Liu, Approximating quantum many-body wave functions using artificial neural networks, Phys. Rev. B **97**, 035116 (2018).

[37] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, Equivalence of restricted Boltzmann machines and tensor network states, Phys. Rev. B **97**, 085104 (2018).

[38] D.-L. Deng, X. Li, and S. Das Sarma, Quantum Entanglement in Neural Network States, Phys. Rev. X **7**, 021021 (2017).

[39] X. Gao and L.-M. Duan, Efficient representation of quantum many-body states with deep neural networks, Nat. Commun. **8**, 662 (2017).

[40] Y. Huang and J. E. Moore, Neural network representation of tensor network and chiral states, arXiv:1701.06246.

[41] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, Self-learning Monte Carlo method, Phys. Rev. B **95**, 041101(R) (2017).

[42] Y. Nomura, A. S. Darmawan, Y. Yamaji, and M. Imada, Restricted Boltzmann machine learning for solving strongly correlated quantum systems, Phys. Rev. B **96**, 205152 (2017).

[43] M. Schmitt and M. Heyl, Quantum dynamics in transverse-field Ising models from classical networks, SciPost Phys. **4**, 013 (2018).

[44] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, Nat. Phys. **14**, 447 (2018).

[45] L. McInnes, J. Healy, and J. Melville, Umap: Uniform manifold approximation and projection for dimension reduction, arXiv:1802.03426.

[46] A. Kitaev, Anyons in an exactly solved model and beyond, Ann. Phys. **321**, 2 (2006), January Special Issue.

[47] M. B. Hastings, I. González, A. B. Kallin, and R. G. Melko, Measuring Renyi Entanglement Entropy in Quantum Monte Carlo Simulations, Phys. Rev. Lett. **104**, 157201 (2010).

[48] J. Tang, J. Liu, M. Zhang, and Q. Mei, Visualizing large-scale and high-dimensional data, in *Proceedings of the 25th International Conference on World Wide Web* (International World Wide Web Conferences Steering Committee, 2016), pp. 287–297.

[49] L. van der Maaten and G. Hinton, Visualizing data using t-SNE, J. Mach. Learn. Res. **9**, 2579 (2008).

[50] R. R. Coifman and S. Lafon, Diffusion maps, Appl. Comput. Harmonic Anal. **21**, 5 (2006), special Issue: Diffusion Maps and Wavelets.

[51] M. Belkin and P. Niyogi, Laplacian eigenmaps and spectral techniques for embedding and clustering, in *Advances in Neural Information Processing Systems* (2002), pp. 585–591.

[52] J. B. Tenenbaum, V. de Silva, and J. C. Langford, A global geometric framework for nonlinear dimensionality reduction, Science **290**, 2319 (2000).

[53] J. W. Sammon, A nonlinear mapping for data structure analysis, IEEE Trans. Comput. **100**, 401 (1969).

[54] J. B. Kruskal, Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis, Psychometrika **29**, 1 (1964).

[55] H. Hotelling, Analysis of a complex of statistical variables into principal components. J. Ed. Psych. **24**, 417 (1933).

[56] E. Becht, L. McInnes, J. Healy *et al.*, Dimensionality reduction for visualizing single-cell data using UMAP, Nat. Biotechnol. **37**, 38 (2019).

[57] A. Diaz-Papkovich, L. Anderson-Trocmé, C. Ben-Eghan, and S. Gravel, UMAP reveals cryptic population structure and phenotype heterogeneity in large genomic cohorts, PLoS Genet. **15**, e1008432 (2019).

[58] K. Polański, M. D. Young, Z. Miao, K. B. Meyer, S. A. Teichmann, and J.-E. Park, BBKNN: Fast batch alignment of single cell transcriptomes, Bioinformatics **36**, 964 (2020).

[59] K. A. Oetjen, K. E. Lindblad, M. Goswami, G. Gui, P. K. Dagur, C. Lai, L. W. Dillon, J. P. McCoy, and C. S. Hourigan, Human bone marrow assessment by single-cell RNA sequencing, mass cytometry, and flow cytometry, JCI Insight **3**, e124928 (2018).

[60] F. O. Bagger, S. Kinalis, and N. Rapin, Bloodspot: A database of healthy and malignant haematopoiesis updated with purified and single cell mRNA sequencing profiles, Nucl. Acids Res. **47**, D881 (2018).

[61] B. S. Clark, G. L. Stein-O'Brien, F. Shiau, G. H. Cannon, E. Davis, T. Sherman, F. Rajaii, R. E. James-Esposito, R. M. Gronostajski, E. J. Fertig, L. A. Goff, and S. Blackshaw, Comprehensive analysis of retinal development at single

cell resolution identifies NFI factors as essential for mitotic exit and specification of late-born cells, bioRxiv 378950 (2018).

[62] A. Kulkarni, A. G. Anderson, D. P. Merullo, and G. Konopka, Beyond bulk: A review of single cell transcriptomics methodologies and applications, Curr. Opin. Biotechnol. **58**, 129 (2019), systems Biology Nanobiotechnology.

[63] G. La Manno, R. Soldatov, A. Zeisel, E. Braun, H. Hochgerner, V. Petukhov, K. Lidschreiber, M. E. Kastriti, P. Lönnerberg, A. Furlan, J. Fan, L. E. Borm, Z. Liu, D. van Bruggen, J. Guo, X. He, R. Barker, E. Sundström, G. Castelo-Branco, P. Cramer *et al.*, Rna velocity of single cells, Nature (London) **560**, 494 (2018).

[64] W. Wolf, 1. die thematisierung von migration, arbeitsmarkt und nationalstaat, in *Entgrenzungsprozesse in Arbeitsmärkten durch transnationale Arbeitsmigration: World Polity und Nationalstaat im 19. Jahrhundert und heute* (Nomos Verlagsgesellschaft mbH & Co. KG, Baden-Baden, 2018), pp. 15–32.

[65] L. Fuhrimann, V. Moosavi, P. O. Ohlbrock, and P. D'acunto, Data-driven design: Exploring new structural forms using machine learning and graphic statics, in *Proceedings of IASS Annual Symposia* (International Association for Shell and Spatial Structures, 2018), Vol. 2018, pp. 1–8.

[66] K. Blomqvist, S. Kaski, and M. Heinonen, Deep convolutional gaussian processes, arXiv:1810.03052.

[67] B. Gaujac, I. Feige, and D. Barber, Gaussian mixture models with wasserstein distance, arXiv:1806.04465.

[68] C. Escolano, M. R. Costa-jussà, and J. A. R. Fonollosa, (self-attentive) autoencoder-based universal language representation for machine translation, arXiv:1810.06351.

[69] X. Li, O. E. Dyck, M. P. Oxley, A. R. Lupini, L. McInnes, J. Healy, S. Jesse, and S. V. Kalinin, Manifold learning of four-dimensional scanning transmission electron microscopy, npj Comput. Mater. **5**, 5 (2019).

[70] A loop is a closed, connected path of edges with the same σ_z eigenvalue, where at least one vertex that intersects the path has two edges of each σ_z eigenvalue incident on it.

[71] Note that in this case all swap matrix elements are either 0 or 1, leading to occasionally redundant $\vec{X}_k$. We take only unique $\vec{X}_k$ here to avoid artificial clusters in the UMAP, but account for the multiplicity in error calculations.

[72] X.-L. Qi and D. Ranard, Determining a local Hamiltonian from a single eigenstate, Quantum **3**, 159 (2019).

[73] E. Bairey, I. Arad, and N. H. Lindner, Learning a Local Hamiltonian from Local Measurements, Phys. Rev. Lett. **122**, 020504 (2019).

[74] E. Chertkov and B. K. Clark, Computational Inverse Method for Constructing Spaces of Quantum Models from Wave Functions, Phys. Rev. X **8**, 031029 (2018).

[75] J. R. Garrison and T. Grover, Does a Single Eigenstate Encode the Full Hamiltonian?, Phys. Rev. X **8**, 021026 (2018).

[76] M. A. Metlitski and T. Grover, Entanglement entropy of systems with spontaneously broken continuous symmetry, arXiv:1112.5166.