Unsupervised machine learning of quantum phase transitions using diffusion maps

Alexander Lidiak[1,*] and Zhexuan Gong[1,2,†]

[1]*Department of Physics, Colorado School of Mines, Golden, Colorado 80401, USA*
[2]*National Institute of Standard and Technology, Boulder, Colorado 80305, USA*

(Dated: March 18, 2020)

Experimental quantum simulators have become large and complex enough that discovering new physics from the huge amount of measurement data can be quite challenging, especially when little theoretical understanding of the simulated model is available. Unsupervised machine learning methods are particularly promising in overcoming this challenge. For the specific task of learning quantum phase transitions, unsupervised machine learning methods have primarily been developed for phase transitions characterized by simple order parameters, typically linear in the measured observables. However, such methods often fail for more complicated phase transitions, such as those involving incommensurate phases, valence-bond solids, topological order, and many-body localization. We show that the diffusion map method, which performs nonlinear dimensionality reduction and spectral clustering of the measurement data, has significant potential for learning such complex phase transitions unsupervised. This method works for measurements of local observables in a single basis and is thus readily applicable to many experimental quantum simulators as a versatile tool for learning various quantum phases and phase transitions.

With the recent demonstration of quantum supremacy [1], the need for understanding well-controlled experimental quantum systems that cannot be simulated efficiently on a classical computer is growing rapidly. However, experimental data generated by measurements on post quantum-supremacy devices can be too large and complex for traditional data analysis tools to extract useful features from. This in particular poses a major challenge for using quantum simulators to make new discoveries at the frontier of quantum many-body physics, where existing theoretical understanding is often lacking [2]. A promising method to address this challenge is the unsupervised machine learning, which can extract important features from data with little to no a priori understanding of the data [3–11].

While machine learning has become a standard toolbox for data analysis in many areas of physics, including high-energy, astrophysics [5], and condensed-matter physics [12], the use of machine learning in experimental quantum simulators has so far been lacking. For example, the standard approach to demonstrate a quantum phase transition in quantum simulation is to extract some feature from the measurement data related to an "order parameter" [13–15]. However, for the discovery of a new quantum phase or phase transition, it is often unclear what feature or order parameter one should extract from the data. For a simple symmetry breaking phase transition, one can usually find an order parameter that is linear or quadratic in the measured observables. In such a scenario, a common unsupervised machine learning method known as principal component analysis (PCA) can be applied [9, 16, 17]. PCA performs a linear projection of the sample data onto a lower dimensional, principle component subspace which retains as much of the variance of the data as possible. For example, using measurement samples of the spins in a classical Ising model at various temperatures, PCA can automatically identify the average magnetization as the order parameter since it is the first principle component of the data which tracks the most abrupt change of the data with temperature. A follow-up k-means clustering algorithm can then be used to locate the phase transition point [9].

However, to facilitate the discovery of novel quantum phases and phase transitions, it is necessary to develop unsupervised learning methods for quantum systems with phases whose order parameters are often complex, nonlinear functions of local observables [5], cases where PCA fails. Example systems include valence-bond solids [18], quantum spin liquids [19], topologically ordered matter [20], and many-body localized (MBL) systems [21]. A number of nonlinear dimensionality reduction methods in machine learning can be applied for these systems, such as kernel PCA [22], autoencoders [16], and diffusion maps [23]. The success so far is however limited, with the exception that the diffusion map method has been recently used to identify certain topological phases unsupervised [23], which has long been regarded as challenging. In this work, we show that the diffusion map is in fact a rather versatile method that can identify a variety of complex quantum phases. It is also computationally efficient and works for data easily obtained by quantum simulation experiments, such as the measurement of all spins in a single direction [24]. This is in contrast to many machine learning approaches [3, 7, 25–28] that require the entanglement spectrum of quantum states difficult to obtain experimentally.

The main idea of the diffusion map is to reveal the structure of the measurement samples in the configuration space and perform automatic clustering of the samples with a tunable cluster radius [29]. Major changes in the configurations of quantum states can be revealed as a result, suggesting the onset of a phase transition. As examples, we will show how diffusion maps can correctly identify incommensurate phases, valence-bond solid phases, and many-body localized phases, all of which are too complex to be learned using PCA or k-means clustering [30]. Note that our primary goal is to learn phases of quantum matter from experimental data with minimal theoretical understanding. For precisely locating phase transition points, supervised learning methods can be applied

afterwards [3, 7, 25, 26].

General picture.— We will first introduce the diffusion map method and describe a general picture of how it can identify different quantum phases. In the following, the diffusion map is always applied to a collection of measurement samples. Without loss of generality, we assume each sample is from a measurement of N quantum spins in some direction. Each measurement sample thus contains N numbers, denoted by an N-dimensional vector $\boldsymbol{X}_i$ $(i = 1, 2, \cdots, M)$. We obtain M such samples by preparing and measuring the same state M times, a routine practice in quantum simulation experiments [24]. One can also obtain these measurement samples computationally using either direct sampling if exact diagonalization is used, or Monte Carlo sampling if a variational ansatz (such as a matrix product state [31], tensor network state [32], or quantum neural network state [33, 34]) is used.

The diffusion map sets a fictitious diffusion process among the samples based on the distances between samples. First, a distance metric needs to be defined. Here, we use the normalized Euclidean distance between two samples i and j, defined as $d_{ij}^2 \equiv \frac{1}{N}\sum_{k=1}^{N}(X_{ik} - X_{jk})^2$, where X_{ik} is the k^{th} element of the sample vector $\boldsymbol{X}_i$, and $\mathcal{N} \sim N$ is a normalization constant that ensures $d_{ij} \in [0, 1]$. Next, a kernel function is used to associate a transition probability between samples based on their distances. A common choice is the Gaussian kernel: $K_{ij} = e^{-d_{ij}^2/(2\epsilon)}$, where the hyperparameter ϵ controls how fast the transition probability decays with distance. Finally we introduce $P_{ij} = K_{ij}/(\sum_k K_{ik})$ as the normalized probability of the diffusion process from samples i to j.

Since the transition probability between any two samples is nonzero for a finite ϵ, the above-mentioned diffusion process is ergodic in the long time limit. This means the largest eigenvalue of the P matrix is always exactly 1 [23, 29]. If there are clusters of samples in which the samples have at most r spins in different configurations, then up to a time scale $\tau = e^{r/\epsilon}$, the diffusion process will be largely restricted within each cluster. The number of such clusters will correspond to eigenvalues of P that are larger than $1 - \delta$ where $\delta \sim 1/\tau$ [29]. Thus by choosing $\epsilon \sim r/(N \ln \tau)$, we can find the number of clusters at a particular size. For the examples we shall discuss, we keep the diffusion time τ fixed by choosing $\delta = 10^{-2.5}$ (the exact value does not matter as long as it's small).

Schematically shown in Fig. 1, when $\epsilon \lesssim 1/(N \ln \tau)$, each different sample will be identified as one cluster, and the number of clusters in this regime is simply given by the number of unique measurement samples. For a generic quantum many-body state that contains non-negligible weights of exponentially many basis states, almost all measurement samples are different from each other. However, for a many-body localized quantum system, the number of unique samples can be significantly smaller than the number of samples. Thus a diffusion map with such a small ϵ should be able to distinguish an ergodic many-body state from a localized many-body state. On the other hand, when $\epsilon \sim 1/\ln \tau$, only samples with $\sim N$ spins in different configurations will belong to different clusters. If multiple clusters of samples are identified in this regime, one can expect a discrete spontaneous symmetry breaking in the measurement direction of the spins, with the order of the symmetry group equal to the number of clusters. Between the small and large ϵ regimes, we can tune ϵ to reveal the number of clusters with variable sizes. Note that changing ϵ in the diffusion map requires no additional experimental data. As shown below, the intermediate ϵ regime $(1/(N \ln \tau) \lesssim \epsilon \lesssim 1/\ln \tau)$ is the key to distinguish a variety of complex quantum phases that cannot be identified using linear dimensionality reduction methods such as PCA. We emphasize that finding eigenvalues of the diffusion map is computationally efficient for the number of samples $(10^2 - 10^3)$ in a typical quantum simulation experiment [24], and we can compute diffusion maps with different system parameters or hyper-parameters in parallel.

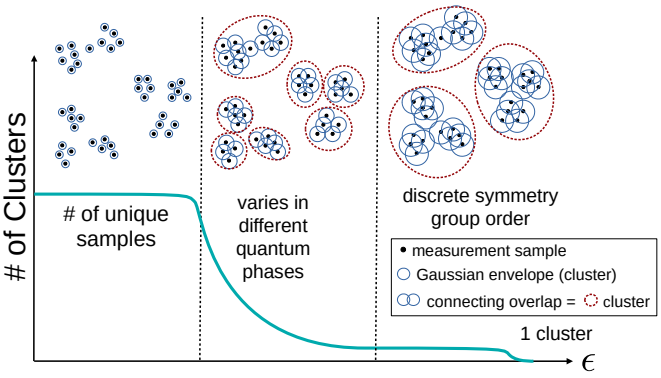

Figure 1. Schematic of how the diffusion map reveals the number of clusters formed by the measurement samples. The hyperparameter ϵ determines the size of the clusters by controlling the width of the Gaussian envelope in the kernel function. The qualitative pictures for the small $[\epsilon \lesssim 1/(N \ln \tau)]$, intermediate, and large $[\epsilon \sim 1/\ln \tau]$ ϵ regimes are shown. For sufficiently large ϵ, the number of clusters always becomes one.

Learning incommensurate phases.— To demonstrate the power of diffusion maps with a tunable ϵ in learning complex quantum phases, we start with a $\mathbb{Z}_n$ transverse-field Ising model (TFIM) (also known as the chiral clock model [35]) with $n = 3$. This model can be simulated using Rydberg atoms experimentally [36, 37] and has a nontrivial incommensurate phase between the usual ferromagnetic and paramagnetic phases of the TFIM. The Hamiltonian of the model reads $H_1 = -f\sum_{j=1}^{N}\tau_j e^{i\theta} - (1-f)\sum_{j=1}^{N-1}\sigma_j\sigma_{j+1}^{\dagger}e^{i\theta} + h.c.$, where $\sigma = \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i2\pi/3} & 0 \\ 0 & 0 & e^{-i2\pi/3} \end{pmatrix}$ and $\tau = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$ are the $\mathbb{Z}_3$ spin operators. Without chirality ($\theta = 0$), the ground state of H_1 undergoes a simple ferromagnetic (FM) to paramagnetic (PM) phase transition when increasing f from 0 to 1. For $\theta > 0$, an incommensurate (IC) phase appears for intermediate values of f (see Fig. 2(a) for a phase diagram), where spin correlations $\langle\sigma_i\sigma_j\rangle$ decay as a power law. The ferromagnetic phase can be easily identified using the average of $\{\langle\sigma_j\rangle\}$, which is nonzero in the FM phase. However, such an order parameter cannot tell the PM phase from the IC phase, as both

phases have vanishing FM order. As a result, PCA (as well as a simple auto-encoder) cannot identify the IC phase and its phase boundary, even with the help of the k-means clustering algorithm [30].

The incommensurate phase and its phase boundary can be numerically identified using entanglement entropy [35, 38], a quantity that is hard to measure in a large experimental system [39]. Here we show that using just the measurement samples of the spin operators $\{\sigma_j\}$, the diffusion map is able to produce a phase diagram in a unsupervised manner that well matches the one obtained using entanglement entropy. Using an intermediate value of ϵ, the number of clusters identified by the diffusion map identifies all three phases of the $\mathbb{Z}_3$ TFIM, as shown in Fig. 2(a). This result can be understood as follows: (1) Deep in the PM phase the measurement samples are approximately drawn from a uniform probability distribution of every possible spin configuration. As a result, two different samples will have on average $N/2$ spins in different configurations. Since in practice the number of samples is often much smaller than the number of spin configurations, each sample will be treated as a separate cluster if $\epsilon \ll 1/\ln\tau$. The number of clusters will thus be close to the number of samples. (2) Deep in the FM phase the ground state has spontaneous symmetry breaking, resulting in one of the 3 FM states ordered in different directions. The samples drawn from each of the 3 FM ordered states should have small distances between each other while the samples drawn from different FM ordered states have close to the maximum possible distance between each other. With $1/(N\ln\tau) \ll \epsilon \ll 1/\ln\tau$, the number of clusters identified will be close to 3. (3) The samples drawn from the IC phase should have a rather diverse set of distances and the number of clusters identified by the diffusion map should vary strongly depending the parameters of the Hamiltonian. We have also obtained a very similar phase diagram for a much larger system size ($N = 100$) of the model using diffusion map [30], with samples generated using matrix product state methods [40]. Note that in both cases no fine tuning of ϵ is needed.

The diffusion map also allows unsupervised learning of the order of the discrete symmetry group underlying a symmetry breaking phase transition. For a large enough ϵ ($\epsilon \sim 1/\ln\tau$), samples in the PM and IC phases will be identified as a single cluster while samples in the FM phase are cleanly sorted into 3 clusters as a result of the Z_3 spontaneous symmetry breaking [see Fig. 2(b)]. Contrast this to the k-means clustering algorithm used frequently in unsupervised machine learning [41], where the number of clusters has to be guessed or predicted using a priori knowledge of the data.

Learning valence-bond solid phase transitions.— Valence-bond solids (VBS) are important in condensed matter physics and quantum information as they are closely related to quantum spin liquids [19], symmetry protected topological order [42, 43], tensor network states [32], and cluster states for quantum computing [44]. Because VBS cannot be identified using an order parameter linear in spin operators, this is another scenario where unsupervised learning methods

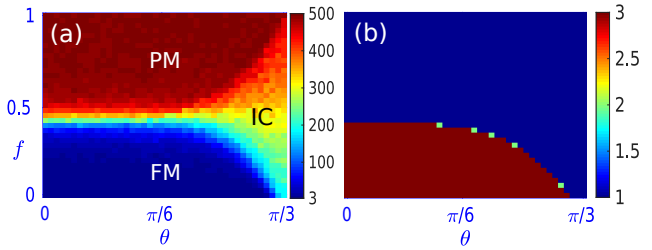

Figure 2. The ground-state phase diagram of the chiral $\mathbb{Z}_3$ TFIM (see H_1) found by the diffusion map method. The measurement samples are obtained from the ground-state found using exact diagonalization with $N = 12$ spins. 500 samples are used for each value of f and θ. (a) $\epsilon = 0.015$, an intermediate value that reveals all three phases and their boundaries. (b) $\epsilon = 0.075$, a large value that causes both the samples in the PM and IC phases to be grouped into one cluster, while in the FM phase three clusters are always identified.

such as PCA will fail while the diffusion maps are useful. As a specific example, we consider a spin-1/2 chain with nearest-neighbor and next-nearest-neighbor antiferromagnetic Heisenberg interactions [45], commonly known as the J_1-J_2 model, with the Hamiltonian $H_2 = \sum_{j=1}^{N}(J_1\boldsymbol{S}_j \cdot \boldsymbol{S}_{j+1} + J_2\boldsymbol{S}_j \cdot \boldsymbol{S}_{j+2})$. For simplicity we set $J_1 = 1$ below. This is a paradigmatic model exhibiting VBS order. In particular, the ground state at $J_2 = 0.5$ is exactly solvable and made of two degenerate VBS (dimer) states, corresponding to two different ways of pairing neighboring spins into spin-1/2 singlets (one with $\boldsymbol{S}_i + \boldsymbol{S}_{i+1} = 0$ for odd i and the other for even i). At around $J_2 = 0.3$, the system is believed to undergo a phase transition from the Luttinger liquid to the VBS/dimer phase [46].

As expected, we find no signatures of the VBS phase transition and no special behavior at $J_2 = 0.5$ using PCA and k-means clustering [30]. With diffusion maps, as shown in Fig. 3, we not only see signatures of the VBS phase transition, but also clearly identify the exactly solvable point of $J_2 = 0.5$, where the ground state has spontaneous translational symmetry breaking even for a finite system size. We thus expect that diffusion maps can be used to detect a generic VBS phase transition at which spontaneous translational symmetry breaking should take place in the thermodynamic limit. The reason is because at the VBS phase transition, diffusion maps can reveal a significant drop in the number of clusters with an intermediate ϵ value (Fig. 3). This can be explained as follows using the J_1-J_2 model example: First, we note that away from the critical point, samples from both dimer states co-exist, while at the critical point, samples only come from either dimer state. The probabilities of finding two samples with a small distance is very low if the two samples are from different dimer states, thus at the critical point, samples are much more likely to be close to each other than when J_2 is away from the critical point. In fact, in the thermodynamic limit, this holds for any samples except those with exactly half of the spins in the opposite direction [30]. As a result, the number of clusters at $J_2 = 0.5$ drops much faster (at roughly

twice the rate [30]) than for J_2 away from 0.5. For a large range of intermediate ϵ values, we can easily identify this critical point.

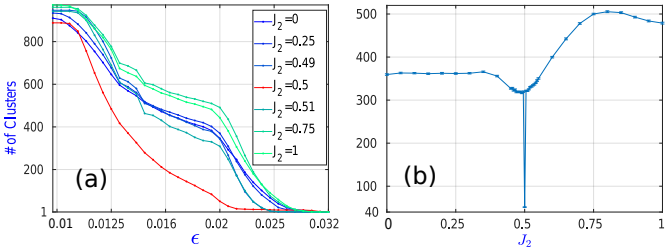

Figure 3. (a) The number of clusters found by diffusion maps when applied to the measurement samples of the J_1-J_2 model (see H_2) as a function of ϵ. 1000 samples are used for each value of J_2 ($J_1 = 1$), obtained from exact diagonalization of H_2 with $N = 24$ spins. (b) An intermediate ϵ ($\epsilon = 0.02$) is chosen to show that the number of clusters drops at the onset of the valence-bond phase transition, with a sharp drop at $J_2 = 0.5$ where the ground state has a spontaneous translational symmetry breaking even for a finite system size. The error bar shows the standard error of the mean from 90 repeated sampling processes.

Learning many-body localization.— A 1D quantum system with tunable disorder can exhibit a dynamical phase transition from a thermal, ergodic phase to a many-body localized (MBL) phase [21]. There is no simple order parameter to detect the MBL phase transition. Theoretically, one can use the inverse participation ratio, level statistics, or entanglement entropy to detect a MBL phase transition [47]. However, these quantities are difficult to obtain experimentally. A more practical way to detect MBL is to use quench dynamics. For example, one can measure local observables after a long-time evolution from some initial, simple-to-prepare product state. Here we show that the diffusion map method [48] can learn the thermal-to-MBL phase transition using the measurement samples obtained in quench dynamics experiments unsupervised. This is different from existing machine learning studies of MBL that require supervised learning [25, 26].

As an example, we study a paradigmatic model exhibiting the thermal-MBL phase transition, i.e. the spin-1/2 Heisenberg model with a random field [25, 47], with the Hamiltonian $H_3 = J\sum_{i=1}^{N} \boldsymbol{S}_i \cdot \boldsymbol{S}_{i+1} + h_i S_i^z$. Here $h_i \in (-hJ, hJ)$ is a random number drawn from uniform distribution and h denotes the disorder strength. It has been numerically found that the thermal-MBL phase transition takes place at the critical disorder strength $h_c \approxeq 3.5 \pm 1.0$. We perform quench dynamics using an initial state with $\langle S_i^z \rangle = \frac{(-1)^i}{2}$ and measure all S_i^z ($i = 1, 2, \cdots, N$) after a long time ($t = 10^4/J$). As mentioned before, since the number of unique samples decreases with increasing disorder strength, we see that the diffusion map with a small ϵ can already indicate the onset of MBL [Fig. 4(a)]. But we can learn more about where the thermal-MBL phase transition occurs by using an intermediate value of ϵ. As shown in Fig. 4, the number of clusters identified by the diffusion map is actually peaked near h_c for a range

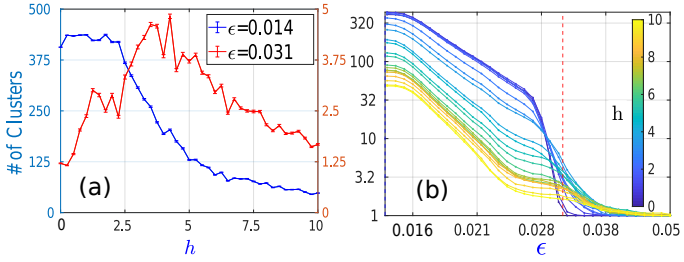

Figure 4. Number of clusters learned by diffusion maps on the measurement samples of the long-time dynamical state of H_3, averaged over 50 disorder realizations. 500 samples are obtained for $N = 16$ spins using exact diagonalization. (a) A small ϵ value (blue curve) leads to the number of samples decreasing rapidly with increased disorder, while the peak of the number of clusters with an intermediate ϵ value (red curve) reveals the approximate location of the thermal-to-MBL phase transition. The error bars show the standard error of the mean calculated over 50 different disorder realizations and 50 repeated sampling processes. (b) There exists a range of intermediate ϵ values where the number of clusters peaks near the critical disorder strength.

of intermediate values of ϵ [48]. Intuitively, this is because deep in the thermal phase, the samples are scattered across the configuration space with similar distances between each other. They are unlikely to form many small clusters. On the other hand, deep in the localized phase the samples are already clustered around the initial state, and the number of clusters should also be small. However, near thermal-MBL phase transitions, the samples form a number of scattered, small clusters in the configuration space due to the competition of disorder and ergodicity. As a result, we expect to see the most clusters near the phase transition. We have also observed this behavior for other spin models exhibiting MBL [30], suggesting that diffusion maps are widely applicable in learning MBL.

Conclusion and Outlook.— We have shown that the diffusion map is a general and versatile unsupervised machine learning method for learning various quantum phases and phase transitions beyond the scope of simple unsupervised learning methods such as PCA and k-means clustering. Compared to traditional data analysis methods, diffusion maps use the full statistics of the measurement samples which contain information of spin correlations at all orders, and thus offers more knowledge of the measured quantum states without demanding more data. The limitations of diffusion maps in learning quantum phases, however, are far from clear. For example, can we learn symmetry protected topological orders [43] using diffusion maps? Do we need to use a different distance metric or kernel function in diffusion maps if it fails to identify certain quantum phases? If we use samples obtained by measurements in more than one local basis, will diffusion maps be able to distinguish classical and quantum phase transitions? And finally, how will other nonlinear dimensionality reduction methods widely used in machine learning, such as t-SNE [49] and DBSCAN [50] compare to the performance of diffusion maps when applied to quantum phase detection?

Acknowledgments.— We thank Lincoln Carr, Eliot Kapit, Cecilia Diniz Behn, and Bo Wu for enlightening discussions related to this work, and the HPC center at Colorado School of Mines for providing computational resources needed in carrying out this work. AL and ZXG acknowledge funding support from the NSF RAISE-TAQS program under Grant No. CCF-1839232.

* alidiak@mines.edu
† gong@mines.edu

[1] F. Arute, K. Arya, R. Babbush, *et al.*, Nature **574**, 505 (2019).
[2] T. H. Johnson, S. R. Clark, and D. Jaksch, EPJ Quantum Technology **1**, 10 (2014).
[3] E. P. Van Nieuwenburg, Y. H. Liu, and S. D. Huber, Nature Physics **13**, 435 (2017).
[4] G. Torlai and R. G. Melko, Physical Review B **94**, 1 (2016).
[5] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Reviews of Modern Physics **91**, 045002 (2019).
[6] P. Broecker, J. Carrasquilla, R. G. Melko, and S. Trebst, Scientific Reports **7** (2017).
[7] J. Carrasquilla and R. G. Melko, Nature Physics **13**, 431 (2017).
[8] P. Broecker, F. F. Assaad, and S. Trebst, arXiv:1707.00663 (2017).
[9] L. Wang, Physical Review B **94** (2016).
[10] K. Ch'ng, J. Carrasquilla, R. G. Melko, and E. Khatami, Physical Review X **7** (2017).
[11] R. T. D'Agnolo and A. Wulzer, Physical Review D **99**, 015014 (2019).
[12] Y. Zhang, A. Mesaros, K. Fujita, S. D. Edkins, M. H. Hamidian, K. Ch'ng, H. Eisaki, S. Uchida, J. C. S. Davis, E. Khatami, and et al., Nature **570**, 484 (2019).
[13] I. Bloch, J. Dalibard, and S. Nascimbène, Nature Physics **8**, 267 (2012).
[14] R. Islam, E. Edwards, K. Kim, S. Korenblit, C. Noh, H. Carmichael, G.-D. Lin, L.-M. Duan, C.-C. Joseph Wang, J. Freericks, and C. Monroe, Nature Communications **2**, 377 (2011).
[15] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, Nature **551**, 601 (2017).
[16] W. Hu, R. R. Singh, and R. T. Scalettar, Physical Review E **95**, 1 (2017).
[17] C. Wang and H. Zhai, Physical Review B **96**, 144432 (2017).
[18] M. E. Zhitomirsky and K. Ueda, Physical Review B **54**, 9007 (1996).
[19] L. Balents, Nature **464**, 199 (2010).
[20] X.-G. Wen, Reviews of Modern Physics **89**, 041004 (2017).
[21] R. Nandkishore and D. A. Huse, Annual Review of Condensed Matter Physics **6**, 15 (2015).
[22] C. Wang and H. Zhai, Frontiers of Physics **13**, 130507 (2018).
[23] J. F. Rodriguez-Nieva and M. S. Scheurer, Nature Physics **15**, 790 (2019).
[24] C. Monroe, W. C. Campbell, L. M. Duan, Z. X. Gong, A. V. Gorshkov, P. Hess, R. Islam, K. Kim, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao, arXiv:1912.07845 (2019).
[25] F. Schindler, N. Regnault, and T. Neupert, Physical Review B **95** (2017).
[26] Y.-T. Hsu, X. Li, D.-L. Deng, and S. Das Sarma, Physical Review Letters **121**, 245701 (2018).
[27] J. Venderley, V. Khemani, and E.-A. Kim, Physical Review Letters **120**, 257204 (2018).
[28] M. Matty, Y. Zhang, Z. Papić, and E.-A. Kim, Physical Review B **100** (2019).
[29] R. R. Coifman and S. Lafon, Applied and Computational Harmonic Analysis **21**, 5 (2006).
[30] See the Supplemental Material for details.
[31] U. Schollwöck, Annals of Physics **326**, 96 (2011).
[32] R. Orús, Annals of Physics **349**, 117 (2014).
[33] G. Carleo and M. Troyer, Science **355**, 602 (2017).
[34] Dong-Ling Deng, Xiaopeng Li, and S. Das Sarma, Physical Review X **7** (2017).
[35] P. Fendley, Journal of Statistical Mechanics: Theory and Experiment **2012**, 0 (2012).
[36] 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, Nature **551**, 579 (2017).
[37] R. Samajdar, S. Choi, H. Pichler, M. D. Lukin, and S. Sachdev, Physical Review A **98**, 023614 (2018).
[38] P. Calabrese and J. Cardy, Journal of Statistical Mechanics: Theory and Experiment **2004**, P06002 (2004).
[39] R. Islam, R. Ma, P. M. Preiss, M. Eric Tai, A. Lukin, M. Rispoli, and M. Greiner, Nature **528**, 77 (2015).
[40] D. Jaschke, M. L. Wall, and L. D. Carr, Computer Physics Communications **225**, 59 (2018).
[41] S. Lloyd, IEEE Transactions on Information Theory **28**, 129 (1982).
[42] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, Physical Review Letters **59**, 799 (1987).
[43] X. Chen, Z.-C. Gu, Z.-X. Liu, and X.-G. Wen, Physical Review B **87** (2013).
[44] F. Verstraete and J. I. Cirac, Physical Review A **70** (2004).
[45] C. K. Majumdar and D. K. Ghosh, Journal of Mathematical Physics **10**, 1388 (1969).
[46] T. Tonegawa and I. Harada, Journal of the Physical Society of Japan **56**, 2153 (1987).
[47] A. Pal and D. A. Huse, Physical Review B **82**, 174411 (2010).
[48] Here we need to use a different kernel in the diffusion map to accommodate the large disparity in the density of measurement samples in the configuration space across the thermal-MBL phase transition, i.e. we will use $K_{ij}/(\sum_k K_{ik} \sum_k K_{kj})$ in the place of K_{ij}. This is a standard practice in applying diffusion maps to data with large variation in density [29, 30].
[49] L. van der Maaten and G. Hinton, Journal of Machine Learning Research **9**, 2579 (2008).
[50] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96 (AAAI Press, 1996) p. 226–231.

Supplementary Material for "Unsupervised machine learning of quantum phase transitions using diffusion maps"

Alexander Lidiak[1,*] and Zhexuan Gong[1,2,†]

[1]*Department of Physics, Colorado School of Mines, Golden, Colorado 80401, USA*
[2]*National Institute of Standard and Technology, Boulder, Colorado 80305, USA*

(Dated: March 18, 2020)

In this supplementary material, we will show additional supporting results for how diffusion maps can learn incommensurate phases (section I), valence-bond solid phases (section II), and many-body localization (section III).

I. LEARNING INCOMMENSURATE PHASES

In this section, we first show that principle component analysis (PCA) together with k-means clustering is unable to learn the incommensurate phase of the $\mathbb{Z}_3$ chiral transverse-field Ising model (see H_1 in the main text) and its boundary. We perform PCA on the same collection of measurement samples used for the diffusion map in Fig 2 of the main text and extract the projection of the sample set onto the first two principle components. We then apply a k-means clustering algorithm to associate each sample $\boldsymbol{X}_i$ with an index $L_i = 1, 2, \cdots, k$. In an attempt to identify all three phases within the samples, we manually set the number of clusters to $k = 3$ (note that the diffusion map does not require such a priori knowledge of the number of distinct phases in the data). We then average L_i for samples belonging to a particular set of Hamiltonian parameters (f, θ), and obtain a phase diagram using this averaged index [see Fig. 1(a)]. While we can identify the ferromagnetic phase and its boundary, there is no clear identification of the paramagnetic or incommensurate phases.

We have also used an auto-encoder included in MATLAB to perform nonlinear dimensionality reduction of the measurement data in substitution of PCA. The auto-encoder trains an artificial neural network to retain as much information of the sample data as possible with two latent variables onto which we then encode each measurement sample (similar to projecting onto the two principle components obtained via PCA). Applying the same k-means clustering algorithm with $k = 3$ leads to a phase diagram shown in Fig. 1(b), which again is unable to identify the incommensurate phase. This is because the incommensurate phase cannot be identified using a single linear or nonlinear function of the measured observables, as the spin configurations in this phase vary strongly with the system parameters (θ and f). The diffusion map method, however, does not try to reduce the dimensionality of the data directly. Instead, it detects the change in the distribution of the measurement samples in configuration space which is often linked to a phase transition.

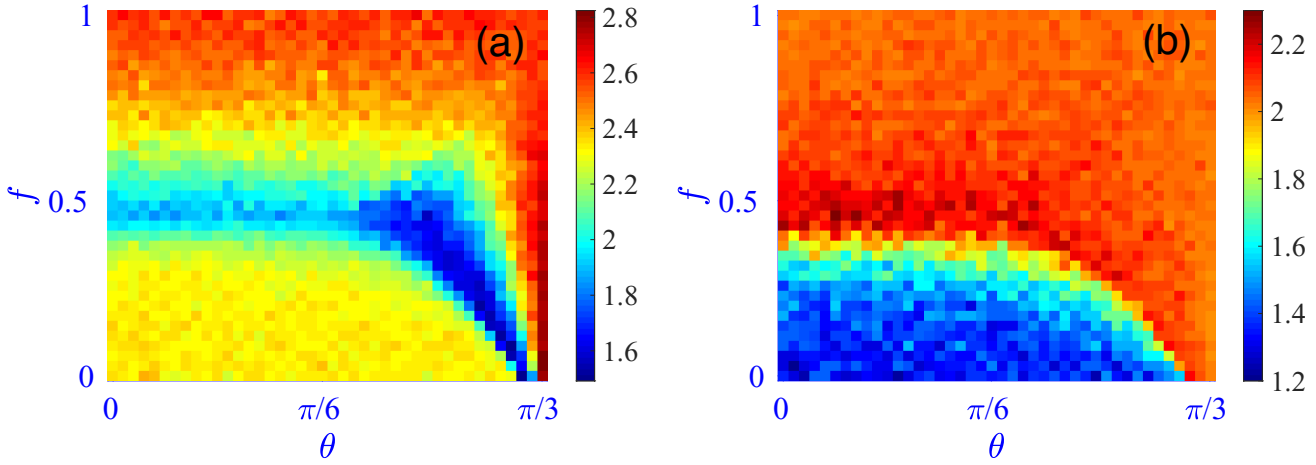

FIG. 1. Phase diagrams generated by PCA (a) and auto-encoder (b) applied on the same measurement samples used in Fig. 2 of the main text. The color represents the index attached by a $k = 3$ k-means clustering algorithm on the projected/compressed measurement sample data, averaged over all measurement samples of the ground state of a particular Hamiltonian.

* alidiak@mines.edu
† gong@mines.edu

In addition, we show the phase diagram obtained using diffusion maps for a much larger system size than that used in Fig. 2 of the main text. We use the OpenMPS library [1] to variationally find the ground state of H_1 (see main text) for $N = 100$ spins using a bond dimension of 200. We then use the Metropolis-Hastings Monte Carlo sampling method [2] to efficiently generate measurement samples of σ_j ($j = 1, 2, \cdots, N$) using the matrix product state (MPS) ansatz that approximates the ground state. We then perform diffusion maps with an intermediate value of ϵ. As shown in Fig. 2(a), the phase diagram obtained using the number of clusters identified by the diffusion map matches well with Fig. 2(a) in the main text and that obtained using the half-system entanglement entropy shown in Fig. 2(b). Note that the noises in the ferromagnetic phases in Fig. 2(b) are due to very small energy gaps in the ground state manifold for $N = 100$ spins such that the variational MPS algorithm may also be subject to spontaneous symmetry breaking.

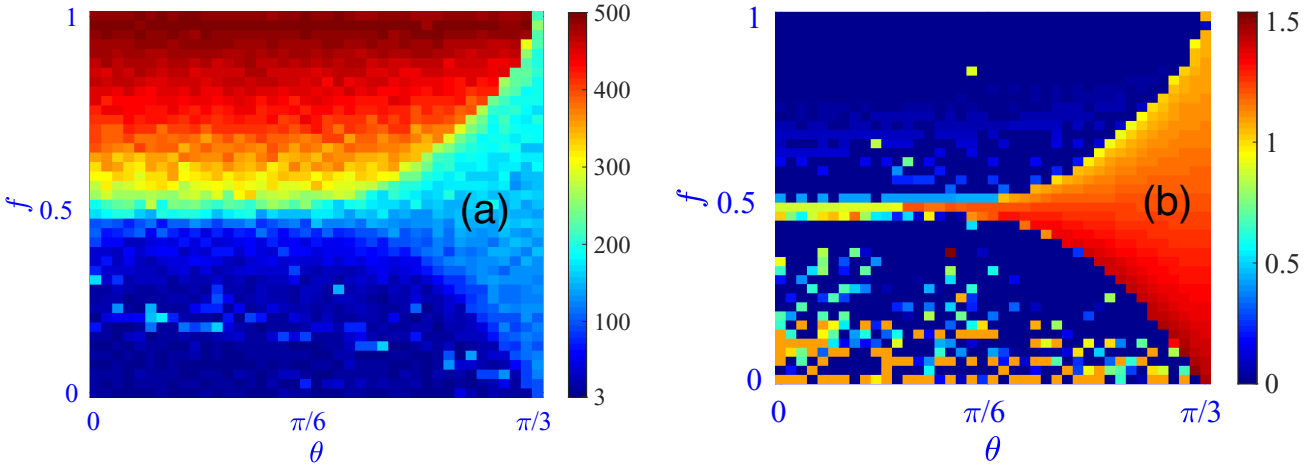

FIG. 2. (a) Ground-state phase diagram of H_1 (see main text) obtained by (a) performing diffusion maps with $\epsilon = 0.004$, $\delta = 10^{-2.5}$ on 500 measurement samples generated for each Hamiltonian using Monte Carlo sampling. (b) calculating the half-system entanglement entropy using MPS methods.

II. LEARNING VALENCE-BOND SOLID PHASES

In this section, we will first show that for the J_1-J_2 model discussed in the main text (H_2), PCA and k-means clustering cannot detect the formation of valence-bond solids (VBS) or the spontaneous symmetry breaking at $J_2 = 0.5$. We first perform PCA on the same measurement samples used in Fig. 3 of the main text. As shown in Fig. 3(a), the efficacy of dimensionality reduction is poor in this case, with many principle components contributing significantly to the variance of the data. Moreover, the first principle component only identifies an anti-ferromagnetic order, which is irrelevant for the VBS phase transition. Keeping the first two principle components, we apply a k-means clustering algorithm with $k = 2$ and plot the average index as a function of J_2 [Fig. 3(b)]. There is no clear signature of a phase transition happening at $J_2 \approx 0.3$ and no indication of the spontaneous translational symmetry breaking at $J_2 = 0.5$. We have also used auto-encoder in place of PCA and find it performs no better.

Next, we will explain in details the sharp drop in the number of clusters at $J_2 = 0.5$ identified by the diffusion map with an intermediate value of ϵ. First, we point out that at $J_2 = 0.5$, we only obtain one of the two degenerate dimer states (one with $\boldsymbol{S}_i + \boldsymbol{S}_{i+1} = 0$ for all odd is and the other for all even is) as the ground state numerically, which is also expected in an actual experiment due to spontaneous symmetry breaking. A small deviation from $J_2 = 0.5$ will lift the degeneracy of the two dimer states and the ground state becomes approximately a superposition of the two dimer states, which we will call the 'combined dimer state' below. Thus to see why the number of clusters identified by diffusion maps suddenly drops at $J_2 = 0.5$, we can compare the results of diffusion maps on a single dimer state with that of the combined dimer state. This comparison is shown in Fig. 4, which is very similar to how the $J_2 = 0.5$ curve compares to the J_2 close to 0.5 curves in Fig. 3(a) of the main text.

To understand why the number of clusters for the single dimer state decreases much more rapidly than the combined dimer state, let us start from the following analysis: If we get two random measurement samples from the combined dimer state, then there is $1/2$ probability that both samples are drawn from either one of the single dimer states (Case I), and $1/2$ probability that the two samples are drawn from two different single dimer states (Case II). The intuition is that the probability of finding two samples that are close to each other is much smaller in Case II than in Case I. For example, in Case I, for a given first sample, the chance of getting the second sample that has zero distance (i.e. identical) to the first sample is always $1/2^{N/2}$. But in Case II, the chance of getting an identical sample is $2/2^N$ because there are only two samples that can be obtained from both of the

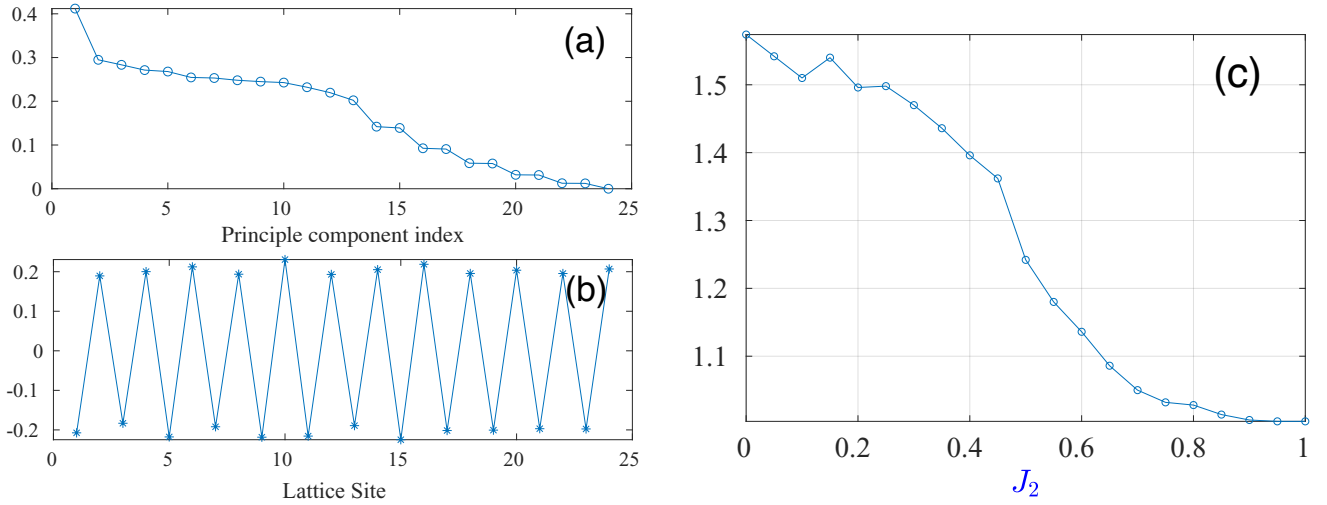

FIG. 3. (a) Principle values of PCA applied to the same measurement data used in Fig. 3 of the main text. (b) Coefficients of the first principle component as a linear combination of each spin's magnetization, which indicates that the first principle component is an anti-ferromagnetic order parameter. (c) The average index assigned by k-means clustering with $k = 2$ to the measurement samples projected onto the first two principle components, as a function of J_2.

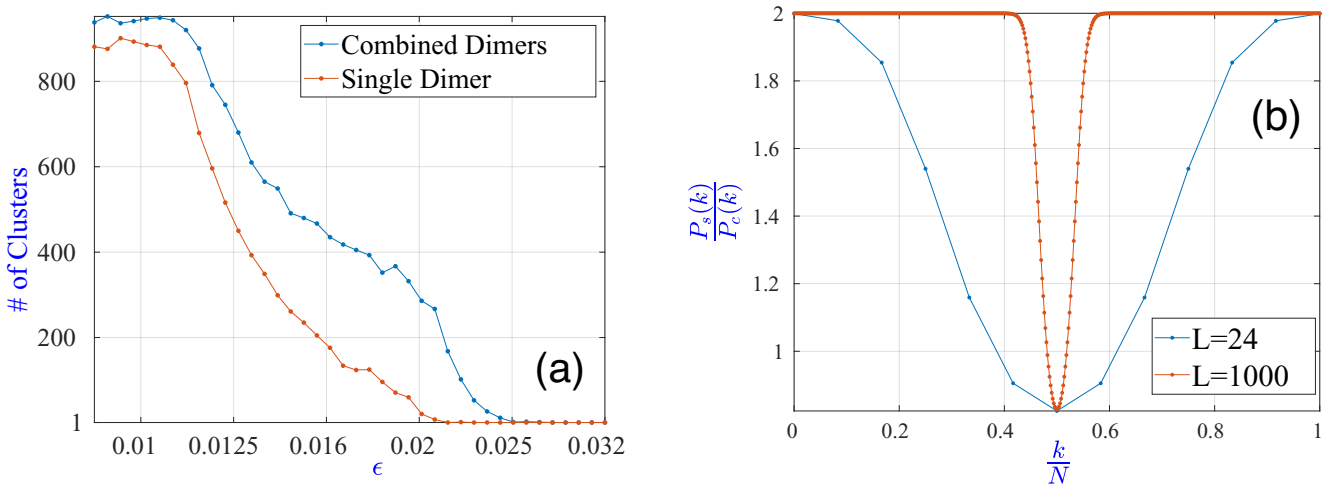

FIG. 4. (a) The number of clusters identified by the diffusion map as a function of ϵ used on the samples drawn from the single dimer state versus the combined dimer state with $N = 24$. (b) The ratio $P_s(k)/P_c(k)$ of the probability of finding two samples with k spins different in the single dimer state to that in the combined dimer state.

dimer states (which are the two perfect antiferromagnetic states), and each only appears with a probability of $1/2^{N/2}$. As a result, we can largely ignore case II in finding two samples close to each other. As we will show below, in the thermodynamic limit ($N \to \infty$), case II can be completely ignored except when we are considering two samples with exactly half of the spins in different directions. After ignoring case II, the probability of finding two samples with k spins different for the combined dimer state, denoted by $P_c(k)$, is only half that of the single dimer state, denoted by $P_s(k)$, in the large N limit.

The above analysis can be made precise mathematically. We find that $P_s(k) = \binom{\frac{N}{2}}{\frac{k}{2}}/2^{\frac{N}{2}}$ and $P_c(k) = [2^{\frac{N}{2}-1}\binom{\frac{N}{2}}{\frac{k}{2}} + \binom{N}{k}]/2^N$. We have plotted the ratio $P_s(k)/P_c(k)$ in Fig. 4(b) for $N = 24$ and $N = 1000$. In the $N \to \infty$ limit, one can show analytically that $P_s(k)/P_c(k) = 2$ for $k \neq N/2$ and $P_s(k)/P_c(k) = 2/(1+\sqrt{2}) \approx 0.83$ for $k = N/2$. Note that both $P_s(k)$ and $P_c(k)$ are symmetric around $k = N/2$.

Because the probability of finding two samples for most distances in the single dimer state is two (or close to two for finite N) times larger than that in the combined dimer state, the number of clusters for the single dimer state will decrease with the cluster radius (proportional to ϵ) at twice the rate of that for the combined dimer state. This twice as fast decay is what we observe in Fig. 4(a) as well as Fig. 3 of the main text which exhibits similar physics.

III. LEARNING MANY-BODY LOCALIZATION

We will first explain why we use a different kernel in the diffusion map for learning many-body localization. In the many-body localized phase, the measurement samples will be clustered around the point of the initial state, while in the thermal phase, the measurement samples will be scattered across the configuration space. This leads to a large variation in the density of samples in the configuration space, and the original Gaussian kernel (see main text) results in the number of clusters always decaying with increasing disorder strength for all values of ϵ. This obscures the physical picture in that when the system goes from the thermal to the MBL phase, samples of the long-time dynamical states will start to form many small clusters due to the interplay of ergodicity and localization. To address this issue, we need a kernel that effectively spreads out samples in the high-density area. Such a kernel is in fact widely used in data science [3], corresponding to an extra normalization of the Gaussian kernel (see the $\alpha = 1$ case in Refs. [3, 4]), i.e. we will use $K'_{ij} = K_{ij}/(\sum_k K_{ik} \sum_k K_{kj})$ as the kernel in the place of K_{ij}. This extra normalization performed on K_{ij} eliminates the density dependence of the samples in extracting the structure of the samples in configuration space [3]. In the context of the thermal-to-MBL phase transition, we see that if the samples are drawn from a disordered state, then most of the samples are close to each other, with only a small amount of samples far away from the rest. Using the kernel given by K'_{ij}, the transition probabilities between samples that are closely packed (far apart) will be reduced (increased), effectively evening the density of samples in the configuration space. We note that for samples without a big variation of density (e.g. those generated from the J_1-J_2 and chiral clock models we studied), using either K_{ij} or K'_{ij} makes little difference.

Finally, we show how diffusion maps can learn the many-body localization phase transition of a different model than the one studied in the main text. This model is a disordered spin-1/2 transverse field Ising chain with next-nearest neighbor interaction, with Hamiltonian:

$$H = -\sum_{i=1}^{N-1} J_i \sigma_i^z \sigma_{i+1}^z + J_2 \sum_{i=1}^{N-2} \sigma_i^z \sigma_{i+2}^z + h \sum_{i=1}^{N} \sigma_i^x. \tag{1}$$

Here the nearest neighbor Ising couplings are disordered as $J_i = J + \delta J_i$ with δJ_i drawn from a uniform random distribution $[-\delta J, \delta J]$. According to Ref. [5], this Hamiltonian undergoes a thermal to MBL phase transition near $\delta J_c = 3.81 \pm 0.04$ (depending also on the energy density of the initial state) when $J = 1$ and $\frac{h}{2} = J_2 = 0.3$. In Fig. 5(a), we show the number of clusters of the diffusion map applied to the measurement samples of $\{\sigma_i^z\}$ drawn from the long-time dynamical state of an initial antiferromagnetic spin state as a function of ϵ. We find that for intermediate values of ϵ, the number of clusters shows a peak around δJ_c [Fig. 5(b)]. This provides further evidence that diffusion maps are able to provide signatures of thermal to MBL phase transitions in general.

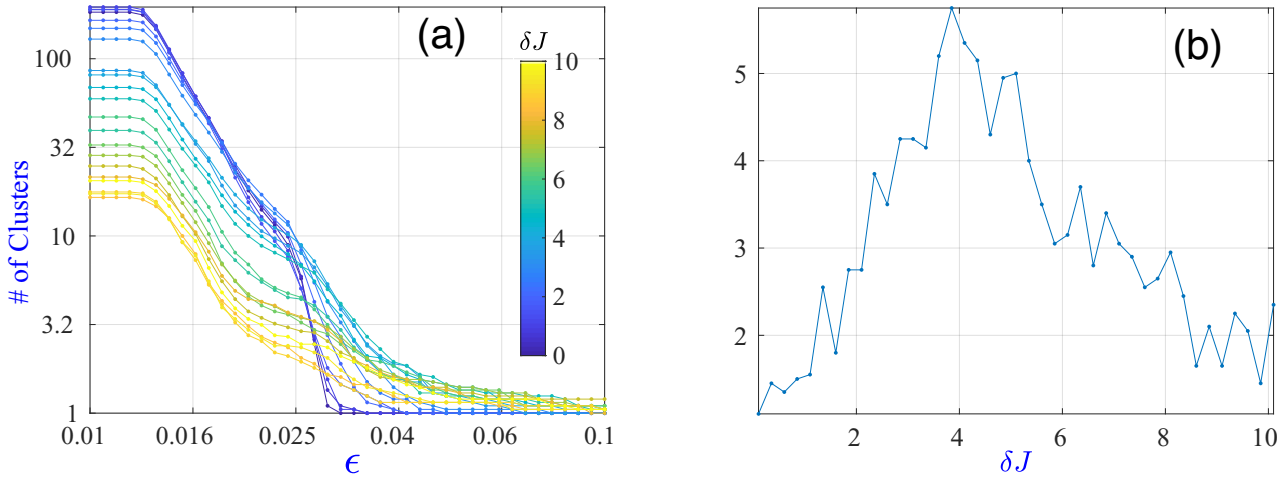

FIG. 5. Number of clusters learned by diffusion maps on the measurement samples of the long-time dynamical state of Eq. (1), averaged over 50 disorder realizations. 500 samples are obtained for $N = 12$ spins using exact diagonalization. (a) The number of clusters as a function of ϵ. (b) For an intermediate value of ϵ ($\epsilon = 0.029$), the number of clusters is peaked around the critical disorder strength $\delta J_c = 3.81$.

[1] D. Jaschke, M. L. Wall, and L. D. Carr, Computer Physics Communications **225**, 59 (2018).

[2] H. G. Katzgraber, "Introduction to monte carlo methods," arXiv:0905.1629 (2009).
[3] R. R. Coifman and S. Lafon, Applied and Computational Harmonic Analysis **21**, 5 (2006).
[4] J. F. Rodriguez-Nieva and M. S. Scheurer, Nature Physics **15**, 790 (2019).
[5] J. A. Kjäll, J. H. Bardarson, and F. Pollmann, Physical Review Letters **113**, 107204 (2014).