



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Learning phase transitions from dynamics

Evert van Nieuwenburg, Eyal Bairey, and Gil Refael

Phys. Rev. B **98**, 060301 — Published 9 August 2018

DOI: [10.1103/PhysRevB.98.060301](https://doi.org/10.1103/PhysRevB.98.060301)

Learning phase transitions from dynamics

Evert van Nieuwenburg,^{1,*} Eyal Bairey,^{2,*} and Gil Refael¹

¹*Institute for Quantum Information and Matter, Caltech, Pasadena, California 91125, USA*

²*Physics Department, Technion, 3200003, Haifa, Israel*

We propose the use of recurrent neural networks for classifying phases of matter based on the dynamics of experimentally accessible observables. We demonstrate this approach by training recurrent networks on the magnetization traces of two distinct models of one-dimensional disordered and interacting spin chains. The obtained phase diagram for a well-studied model of the many-body localization transition shows excellent agreement with previously known results obtained from time-independent entanglement spectra. For a periodically-driven model featuring an inherently dynamical time-crystalline phase, the phase diagram that our network traces in a previously-unexplored regime coincides with an order parameter for its expected phases.

Introduction - Machine learning is emerging as a novel tool for identifying phases of matter¹⁻¹⁵. At its core, this problem can be cast as a classification problem in which data obtained from physical systems are assigned a class (i.e. a phase) using machine learning methods. This approach has enabled autonomous detection of order parameters^{2,5,6}, phase transitions^{1,3} and entire phase diagrams^{4,7,16,17}. Simultaneous research effort at the interface between machine learning and many-body physics has focussed on the use of neural networks for efficient representations of quantum wavefunctions¹⁸⁻²⁶, drawing a parallel between deep networks and the renormalization group²⁷⁻²⁹. Overall, these studies exemplify the power of machine learning for extracting information from physical data without detailed physical input. In particular, it shows potential for identifying novel phases through automatic processing of large-scale data; possibly identifying features that may have been missed before.

So far, however, these methods have relied only on static properties of the underlying physical systems, such as raw state configurations sampled from Monte Carlo simulations^{1,15} or entanglement spectra obtained using exact diagonalization^{3,11,17}. To our knowledge, the study of phase transitions from dynamics of physical observables has not been addressed.

Here, we suggest a machine learning approach to distinguish between phases based on dynamics of measurable quantities. Specifically, we introduce the use of recurrent neural networks (RNNs), designed for processing sequential data such as time-traces. This approach does not rely on thermal equilibrium, and applies very naturally to time-dependent systems. It is therefore particularly suited for the identification of dynamical as well as Floquet phases³⁰⁻³⁹.

We first test our method on a system with two inherently different dynamical behaviours, namely a 1D system with a many-body localization transition⁴⁰⁻⁴³. Machine learning methods applied on entanglement spectra of eigenstates were used to obtain a phase diagram of the same model¹¹, as well as on a slightly different model featuring two distinct MBL phases¹⁷. Here, we insist on using only experimentally relevant (i.e. measurable) quantities such as the magnetization of individual spins.

We find that the network succeeds at distinguishing between the ergodic and localized phases of this model, recovering phase boundaries similar to those obtained by previous methods.

We then apply our method to a periodically driven model, featuring among its three phases one which is unique to the time-dependent setting, namely a time crystal⁴⁴⁻⁵⁰. Indeed the method distinguishes between the time-crystalline, Floquet-ergodic and Floquet-MBL⁵¹⁻⁵³ phases of this model.

In the following section, we first introduce the essentials of recurrent neural networks. We refer the reader to Ref.⁵⁴ for an extensive introduction to the non-recurrent feed-forward neural network. After we have introduced the network essentials, we outline the procedure we refer to as ‘blanking’ for training the network on a set of physics data. This framework is independent of the underlying model, and serves as the main supervised learning scheme in our work. Next we turn to introducing the models and the results mentioned earlier, and conclude with a critical evaluation of the obtained results.

Recurrent networks - Because we wish to be able to capture non-equal-time correlations in the magnetization traces, we choose to train a recurrent neural network (RNN) to distinguish dynamical regimes. A recurrent neural network is a neural network in which one or multiple outputs are fed back into the network as inputs, as illustrated in Fig. 1. Such a recurrence creates a feedback loop that allows information that was fed into the network to persist in a self-consistent manner. This is ideal for analyzing sequences in which the value at a particular point of that sequence may depend on the previous entries. Consequently, RNNs are well suited for dealing with sequential data or other types of data for which a kind of ‘memory’ or temporal dependence is beneficial.

It is particularly useful to introduce the idea of ‘unrolling’ a recurrent part of a network. In Figure 1 we show a (subsection of) a neural network \mathcal{N} with inputs $x(t)$ and outputs $y(t)$, the latter being fed back into the inputs. We think of t here as a discrete parameter, such that inputs and outputs are computed at timesteps t , $t + 1$, etcetera. The feedback should now be understood such that at timestep t , the network receives both $x(t)$




FIG. 1. **Unrolling a recurrent network.** On the left a (subpart of a) neural network \mathcal{N} is shown with output feeding back into input, making it into a recurrent neural network. On the right the unrolled version of the same network is shown, detailing that the output at step t is fed back as an input for time step $t + 1$. The recurrent connections have their own weights that are optimized during training.

and $y(t - 1)$ as its inputs, and produces $y(t)$ from them via an intermediate step. This is most easily visualized by the unrolled network shown in Fig. 1. Namely, the network keeps track of an internal state $h(t)$, which is updated according to $h(t) = f(h(t - 1), x(t))$. The function f represents the free parameters that we wish to learn by training the network. Given $h(t)$, the output $y(t) = g(h(t))$ is computed via another learnable function. The training of such a network is done in a supervised manner identical to the standard feed-forward networks, except that it can be thought of as done ‘unrolled layer’ by ‘unrolled layer’. There are various choices for the functions f and g introduced above, and we so-called Long Short Term Memory networks (LSTMs) ⁵⁵. We expect that the recurrence allows the network to build a better model governing the dynamics, which helps it in the task of classifying the inputs.

Blanking - Since training the recurrent network requires labeled data (it is a supervised method), we use physics intuition to label the data only in the extremities of the phase space we consider, i.e. in the limits where we are confident about the physics of the system. The network is trained only in these regimes, and hence instead of the network seeing all the data we effectively ‘blank out’ outside of the known limits. This blanking tests the network’s ability to extract the underlying essential model of the data from these limits, and apply it to unseen data as a form of generalization. Care must be taken also that one supplies the network with enough and representative data such that a consistent model can, at least in principle, be extracted. As an important check we have tested that the predictions of the network are insensitive to adding slightly more or slightly less labeled data at the extremities (i.e. by shrinking or enlarging the blanked out region); that the network’s confidence is correlated with its accuracy ⁵⁶, and that the network assigns a confused output to phases it had not encountered during training ⁷⁰.

Additionally, one must check for and prevent the possibility of the network learning examples by heart (i.e.

overfitting). We will employ dropout ⁵⁷ and weight decay (l_2 regularization) to do so. We remark that empirically for models with disorder the many realizations and their variety even for a given disorder strength seem to already build in an inherent robustness against overfitting. The actual training of the network is done by minimizing the cross-entropy using the Adam optimizer ⁵⁸. Additionally, we remark that the usual test-set validation can not be performed in the blanked region, since the network is not trained there.

Given the number n of regions in which we know the physics (i.e. the number of expected phases), our networks are constructed with a softmax output layer with n neurons. Thus, the networks take a sequence of magnetizations and output a probability distribution $\mathbf{p} = (p_1, \dots, p_n)$ over the n phases. This distribution describes the probability that the network assigns for the input sequence to belong to each of the phases $1, \dots, n$. We then measure the confusion (uncertainty) of the network by examining the reduced distribution on the two most likely phases. Namely, assuming the probabilities are ordered by decreasing magnitudes ($p_1 \geq p_2 \geq \dots$), we define the confusion as $C = -\log_2(p_1/(p_1 + p_2))$. The confusion C vanishes when the network confidently predicts a specific phase ($p_1 = 1, p_2 = 0$), and it takes the maximal value of unity whenever the network cannot decide between two or more phases ($p_1 = p_2$). Whenever the network changes its prediction from one phase to another at a certain value of an underlying parameter, the peak in C surrounding this value can hence indicate the corresponding transition region.

MBL transition - We consider the random-field Heisenberg model ⁵⁹:

$$H = \sum_i J S_i \cdot S_{i+1} + w_i S_i^z. \quad (1)$$

The length of the chain is given by L , and the on-site disorders w_i are drawn independently and uniformly from the interval $[-W, W]$. This Hamiltonian exhibits a transition between a delocalized and a many-body localized state at a critical disorder strength that depends on the energy density of the state under consideration ^{59–62}. The dynamics of initial product states of spin polarization differs substantially between the two phases: while spins in the many-body localized phase retain long-term correlation with their initial configuration, in the delocalized phase this correlation is lost on much shorter timescales as expected from an ergodic system ^{63–66}. In what follows we will be considering the dynamics of initial states that evolve in time under the Hamiltonian of Eq. 2, by performing exact time evolution on systems of size $L = 20$.

For the purpose of obtaining a phase diagram, we probe dynamics at various energy densities. Similarly to Ref. ⁶⁰, we measure energy density by a parameter ε interpolating between the minimal and maximal eigenenergies E_0, E_{max} of each disorder realization. For each disorder realization we calculate E_0, E_{max} , and pick the product state in the S_z basis ($|\uparrow, \uparrow, \downarrow, \uparrow, \dots\rangle$ etc.) whose energy ex-




FIG. 2. Detecting the MBL transition in the random-field Heisenberg model [2](#). In the left panel, we show the dependence of the network’s confusion C on the number of LSTM neurons, N , for $\epsilon = 0.5$ and a fixed set of parameters: dropout 0.2, $l_2 = 0.01$, batchsize of 64 and 25 training epochs. The right panel shows the resulting phase diagram (the colorbar represents the confusion C) in the ϵ versus W plane, obtained with $N = 32$ and averaged over 10 re-trainings.

pectation value is closest to $E = E_0 + \epsilon(E_{max} - E_0)$. We numerically evolve this initial state in time and measure $\langle \sigma_i^z \rangle(t)$ for each of the spins.

The input to our networks therefore consists of these magnetization time-traces from $t = 0$ to $t = 500$, which we sample at 50 equally spaced points, and hence is of shape $(L, 50)$ for each disorder realization. At all energy densities considered we assume that the weak disorder regime ($W \leq 0.5$) is ergodic whilst the strong disorder regime ($W \geq 7.5$) is many-body localized. We therefore train the network on magnetization traces from these two extreme regimes. At low disorder these traces are labelled by a label $\mathbf{p} = (1, 0)$, and at high disorder the label assigned is $\mathbf{p} = (0, 1)$. Any data for disorder strengths in the interval $W \in [0.5, 7.5]$ is therefore blanked out.

We fix the network architecture to have a single hidden layer of N LSTM neurons with a dropout rate of 0.2 and l_2 regularization of 0.01, followed by a softmax layer to output a probability of the input being ergodic or non-ergodic. In the results below, we have re-trained the network $k = 10$ times with identical parameters but different initial conditions. The results are averaged over these training cases.

We analyze the dependence of the output on the number N of LSTM units in the left panel of Fig. [2](#), and find that with 32 neurons we are able to converge the results for fixed batchsize 64 and 25 epochs. This training was done on the $\epsilon = 0.5$ data, and uses the confidence enhancement introduced in [11](#). In order to gain a better understanding of what the LSTM neurons are doing, we analyze the case of a single LSTM neuron trained on

a single-spin subsystem in the Supplemental Material [70](#). To obtain the phase diagram, we repeat the training process over the 2-dimensional parameter space of energy density (13 values equally spaced between $\epsilon = 0.2$ and $\epsilon = 0.8$) and disorder strength (64 values equally spaced between $W = 0.125$ and $W = 8$), with 50 disorder realizations for each point. The obtained phase diagram is shown in the right panel of Fig. [2](#), and shows good agreement with the phase diagram obtained from static entanglement spectra in Ref. [11](#).

Time Crystals - Next, we consider the following binary Floquet Hamiltonian acting on a one-dimensional spin-1/2 chain:

$$H = \begin{cases} (g - \epsilon) \sum_i \sigma_i^x, & 0 < t < T_1 \\ \sum_i J_i \sigma_i^z \sigma_{i+1}^z + B_i^z \sigma_i^z, & T_1 < t < T_2. \end{cases} \quad (2)$$

Where J_i , B_i are random variables distributed independently and uniformly in the interval $[0, 0.5]$, g is fixed to $\pi/2$, and $T_1 + T_2 = T$. This is a slight variation of the model studied in [46](#), where we took a different distribution for the bond terms J_i . To the best of our knowledge, our exact model has not been studied before, and its phase diagram has not been mapped out. Moreover, it serves as a case where a phase is inherently dynamical and cannot be studied in a static setting.

We are interested in the effect of the driving parameter ϵ on the resulting phase of the system. A guideline for the phases is provided through the long-time imbalance,

$\mathcal{I}(t)$, defined as:

$$\mathcal{I}(t) = \frac{1}{L} \mathbf{m}(t) \cdot \mathbf{m}(0), \quad (3)$$

where the i -th component of $\mathbf{m}(t)$ is the expectation value of σ_i^z at time t . This definition of the imbalance is the direct generalization of that typically used when the initial state is only taken to be one with a charge-density-wave ordering^{63–66}.

The long-time imbalance shows three distinct behaviours as a function of the driving parameter ϵ (orange line in Fig. 3). If $\epsilon = \pi/2$, the drive term is just an identity operator and the system is governed by a many-body localized Hamiltonian. Subsequently, for ϵ sufficiently close to $\pi/2$, the imbalance retains a value close to its initial one, indicating a trivial Floquet-MBL phase. For intermediate values of ϵ , the long-time imbalance vanishes, indicating a transition to a Floquet-ergodic phase. Interestingly, below a critical value of ϵ , the long-term imbalance retains a value that is close to its initial one in magnitude, but flips sign every driving period. In this regime the system’s response is periodic in $2T$ rather than T , leading to the nomenclature ‘time crystal’.

We proceed with training a RNN on time traces of $\mathbf{m}(t)$ identically to the case of the previously discussed MBL system, apart from having three regions in phase space where we train the network instead of two. Namely, for ϵ close to 0 we assign the time-crystalline label, for $\epsilon \approx 0.7$ we assign the Floquet-ergodic label and for $\epsilon = \pi/2$ we assign the Floquet-MBL label. We again use 32 LSTM units, dropout 0.2 and $l_2 = 0.01$ with Adam optimization. When evaluated on a data-set with many more ϵ available, the resulting 1D phase diagram is shown in Fig. 3 (green and gray lines).

Discussion and Prospects - The main point considered in this work was the study of *dynamics* using machine learning methods, and doing so using *experimentally* available measurements. We employed recurrent neural networks, rather than their non-recurrent variants. There are multiple motivations, apart from the input being sequences, for using such an approach over more common non-recurrent feed-forward networks. First, since the data is fed into the network one time-step at a time, the number of network parameters does not scale with the number of time-steps. This also means that the same recurrent network can be easily trained on various lengths of data. In contrast, a regular feed-forward network would need to be input with all of the data at once, leading to a large initial input layer compatible with a fixed input length. We have studied whether the use of recurrent neurons provides a more direct way of extracting what feature of the data the neurons use to output their guess. By training one or multiple LSTM units on single magnetization curves, it is possible to identify neuron behavior⁷⁰. We speculate that it might be possible to extract from these results a ‘dynamical order parameter’, which takes the full magnetization traces into account.




FIG. 3. A recurrent neural network distinguishes between three dynamical phases of a time-dependent model, after being trained on example curves $\mathbf{m}(t)$ at $\epsilon = 0, 0.7$ and $\pi/2$. The gray curves show the outputs of the three neurons assigned to recognize each of the three phases (time-crystalline, Floquet-ergodic and Floquet-MBL). In green (with dots) the confusion C of the network is shown, indicating two transition points between these phases. In orange we show the long-time imbalance $\mathcal{I}(t)$ measured at an odd driving period, taking a negative value in the time-crystalline phase; a vanishing value in the Floquet-ergodic phase; and a positive value in the Floquet-MBL phase. The phase boundaries extracted by the network are consistent with, and seem more sharp than, the imbalance.

Using the networks, we constructed dynamical phase diagrams for the MBL transition and a driven model featuring a time-crystalline phase, thereby circumventing the need to manually construct a threshold criterion or dynamical order parameter for locating the phase boundary. Rather, such a threshold was automatically determined from the data. By considering the LSTM neuron outputs for the single spin case⁷⁰, we were able to gain some understanding of the behaviour that the network latches onto.

We emphasize that obtaining a phase boundary from data can hence only be as accurate as the available data. The boundary we obtain for the MBL transition is at a slightly lower disorder strength than that of the exact diagonalization results in Ref.⁶⁰, but agrees well with that obtained using the machine learned entanglement spectra of Ref.¹¹. The alternative of finding a non-machine learned proxy to serve as an indicator, such as the imbalance for the MBL transition⁷⁰, can be ambiguous. If sufficient data are available, we expect consistency of that data to be the judge of where the transition happens. It may be possible to use the same criterion in a feed-

back system between a machine learning algorithm and experiments, where measurements performed on the experiment are chosen to improve the phase boundary.

For the MBL transition in particular, we mention that a more detailed investigation should also take into account the possibility of a Griffiths phase, possibly showing up as a region where the network prediction is increasingly uncertain as system size increases. Such a finite-size scaling can indeed be successfully attempted us-

ing machine learned data^{1,15}, and provides a useful and interesting alternative for locating a phase boundary.

Being able to train recurrent neural networks on time-traces of data poses the question of whether such methods can be used to enhance the prediction of dynamics, i.e. in numerical time evolution simulations. Such questions are being actively addressed in order to provide accurate control over e.g. single qubits in decohering environments and noisy measurements⁶⁷.

-
- * Equal contribution
- ¹ J. Carrasquilla and R. G. Melko, Nat. Phys. **13**, 431 (2017), ISSN 1745-2473, 1605.01735, URL <http://dx.doi.org/10.1038/nphys4035>.
 - ² L. Wang, Phys. Rev. B **94**, 195105 (2016), ISSN 2469-9950, 1606.00318, URL <https://link.aps.org/doi/10.1103/PhysRevB.94.195105>.
 - ³ E. P. L. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Nat. Phys. **13**, 435 (2017), ISSN 1745-2473, 1610.02048, URL <http://arxiv.org/abs/1610.02048>; <http://www.nature.com/doi/10.1038/nphys4037>.
 - ⁴ P. Broecker, F. Assaad, and S. Trebst, arXiv:1707.00663 (2017).
 - ⁵ S. J. Wetzel and M. Scherzer, arXiv:1705.05582 (2017), 1705.05582, URL <http://arxiv.org/abs/1705.05582>; <https://arxiv.org/pdf/1705.05582.pdf>.
 - ⁶ S. J. Wetzel, arXiv:1703.02435 (2017), 1703.02435, URL <http://arxiv.org/abs/1703.02435>.
 - ⁷ Y.-H. Liu and E. P. L. van Nieuwenburg (2017), 1706.08111, URL <http://arxiv.org/abs/1706.08111>.
 - ⁸ K. Ch'ng, J. Carrasquilla, R. G. Melko, and E. Khatami, arXiv:1609.02552 (2016), 1609.02552, URL <http://arxiv.org/abs/1609.02552>.
 - ⁹ K. Ch'ng, N. Vazquez, and E. Khatami, Phys. Rev. E **97**, 013306 (2018), URL <https://link.aps.org/doi/10.1103/PhysRevE.97.013306>.
 - ¹⁰ P. Broecker, J. Carrasquilla, R. G. Melko, and S. Trebst, arXiv:1608.07848 (2016), 1608.07848, URL <http://arxiv.org/abs/1608.07848>.
 - ¹¹ F. Schindler, N. Regnault, and T. Neupert, *Physical Review B* **95**, 245134 (2017)
 - ¹² T. Ohtsuki and T. Ohtsuki, J. Phys. Soc. Japan **85**, 123706 (2016), ISSN 0031-9015, 1610.00462, URL <http://journals.jps.jp/doi/10.7566/JPSJ.85.123706>.
 - ¹³ L.-F. Arsenault, A. Lopez-Bezanilla, O. A. von Lilienfeld, and A. J. Millis, Phys. Rev. B **90**, 155136 (2014), ISSN 1098-0121, 1408.1143, URL <https://link.aps.org/doi/10.1103/PhysRevB.90.155136>.
 - ¹⁴ L.-F. Arsenault, O. A. von Lilienfeld, and A. J. Millis, arXiv:1506.08858 (2015), 1506.08858, URL <http://arxiv.org/abs/1506.08858>.
 - ¹⁵ M. J. Beach, A. Golubeva, and R. G. Melko, arXiv:1710.09842 (2017), 1710.09842, URL <https://arxiv.org/abs/1710.09842>.
 - ¹⁶ N. Yoshioka, Y. Akagi, and H. Katsura, arXiv e-prints (2017), 1709.05790, URL <http://arxiv.org/abs/1709.05790>.
 - ¹⁷ J. Venderley, V. Khemani, and E.-A. Kim (2017), arXiv:1711.00020.
 - ¹⁸ G. Carleo and M. Troyer, Science **355**, 602 (2017), ISSN 0036-8075, 1606.02318, URL <http://www.sciencemag.org/lookup/doi/10.1126/science.aag2302>.
 - ¹⁹ M. Schmitt and M. Heyl (2017), 1707.06656, URL <http://arxiv.org/abs/1707.06656>.
 - ²⁰ Z. Cai, arXiv:1704.05148 (2017), 1704.05148, URL <http://arxiv.org/abs/1704.05148>.
 - ²¹ Y. Huang and J. E. Moore (2017), 1701.06246, URL <http://arxiv.org/abs/1701.06246>.
 - ²² D.-L. Deng, X. Li, and S. D. Sarma, arXiv:1609.09060 (2016), 1609.09060, URL <http://arxiv.org/abs/1609.09060>.
 - ²³ Y. Nomura, A. Darmawan, Y. Yamaji, and M. Imada (2017), arXiv:1709.06475v1, URL <https://arxiv.org/pdf/1709.06475.pdf>.
 - ²⁴ D.-L. Deng, X. Li, and S. D. Sarma (2017), 1701.04844, URL <http://arxiv.org/abs/1701.04844>.
 - ²⁵ X. Gao and L.-M. Duan, arXiv:1701.05039 (2017), 1701.05039, URL <http://arxiv.org/abs/1701.05039>.
 - ²⁶ G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, arXiv:1703.05334 (2017), 1703.05334, URL <http://arxiv.org/abs/1703.05334>.
 - ²⁷ P. Mehta and D. J. Schwab, arXiv:1410.3831 (2014), ISSN 01406736, 1410.3831, URL <http://arxiv.org/abs/1410.3831>.
 - ²⁸ M. Koch-Janusz and Z. Ringel, Nature Physics (2018), 1704.06279, URL <https://doi.org/10.1038/s41567-018-0081-4>; <https://doi.org/10.1038/s41567-018-0081-4>.
 - ²⁹ S.-H. Li and L. Wang (2018), 1802.02840, URL <http://arxiv.org/abs/1802.02840>.
 - ³⁰ R. Moessner and S. L. Sondhi, Nature Physics **13**, 424 EP (2017), URL <http://dx.doi.org/10.1038/nphys4106>.
 - ³¹ N. H. Lindner, G. Refael, and V. Galitski, Nature Physics **7**, 490 EP (2011), article, URL <http://dx.doi.org/10.1038/nphys1926>.
 - ³² P. Titum, E. Berg, M. S. Rudner, G. Refael, and N. H. Lindner, Phys. Rev. X **6**, 021013 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevX.6.021013>.
 - ³³ F. Nathan, M. S. Rudner, N. H. Lindner, E. Berg, and G. Refael, Phys. Rev. Lett. **119**, 186801 (2017), URL <https://link.aps.org/doi/10.1103/PhysRevLett.119.186801>.
 - ³⁴ D. V. Else, B. Bauer, and C. Nayak, Phys. Rev. X **7**, 011026 (2017), URL <https://link.aps.org/doi/10.1103/PhysRevX.7.011026>.
 - ³⁵ M. Heyl, *Dynamical quantum phase transitions: a review* (2017), arXiv:1709.07461.
 - ³⁶ C. W. von Keyserlingk and S. L. Sondhi, Phys. Rev. B **93**, 245145 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevB.93.245145>.

- ³⁷ D. V. Else and C. Nayak, Phys. Rev. B **93**, 201103 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevB.93.201103>.
- ³⁸ H. C. Po, L. Fidkowski, T. Morimoto, A. C. Potter, and A. Vishwanath, Phys. Rev. X **6**, 041070 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevX.6.041070>.
- ³⁹ A. C. Potter, T. Morimoto, and A. Vishwanath, Phys. Rev. X **6**, 041001 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevX.6.041001>.
- ⁴⁰ D. Basko, I. Aleiner, and B. Altshuler, Annals of Physics **321**, 1126 (2006), ISSN 0003-4916, URL <http://www.sciencedirect.com/science/article/pii/S0003491605002630>.
- ⁴¹ V. Oganesyan and D. A. Huse, Phys. Rev. B **75**, 155111 (2007), URL <https://link.aps.org/doi/10.1103/PhysRevB.75.155111>.
- ⁴² R. Nandkishore and D. A. Huse, Annual Review of Condensed Matter Physics **6**, 15 (2015), URL <https://doi.org/10.1146/annurev-conmatphys-031214-014726>.
- ⁴³ D. A. Abanin and Z. Papi, Annalen der Physik **529**, 1700169 (2017), ISSN 1521-3889, 1700169, URL <http://dx.doi.org/10.1002/andp.201700169>.
- ⁴⁴ D. V. Else, B. Bauer, and C. Nayak, Phys. Rev. Lett. **117**, 090402 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevLett.117.090402>.
- ⁴⁵ V. Khemani, A. Lazarides, R. Moessner, and S. L. Sondhi, Phys. Rev. Lett. **116**, 250401 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevLett.116.250401>.
- ⁴⁶ N. Y. Yao, A. C. Potter, I.-D. Potirniche, and A. Vishwanath, Phys. Rev. Lett. **118**, 030401 (2017), URL <https://link.aps.org/doi/10.1103/PhysRevLett.118.030401>.
- ⁴⁷ F. Wilczek, Phys. Rev. Lett. **109**, 160401 (2012), URL <https://link.aps.org/doi/10.1103/PhysRevLett.109.160401>.
- ⁴⁸ S. Choi, J. Choi, R. Landig, G. Kucsko, H. Zhou, J. Isoya, F. Jelezko, S. Onoda, H. Sumiya, V. Khemani, et al., Nature **543**, 221 EP (2017), URL <http://dx.doi.org/10.1038/nature21426>.
- ⁴⁹ J. Zhang, P. W. Hess, A. Kyprianidis, P. Becker, A. Lee, J. Smith, G. Pagano, I.-D. Potirniche, A. C. Potter, A. Vishwanath, et al., Nature **543**, 217 EP (2017), URL <http://dx.doi.org/10.1038/nature21413>.
- ⁵⁰ C. W. von Keyserlingk, V. Khemani, and S. L. Sondhi, Phys. Rev. B **94**, 085112 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevB.94.085112>.
- ⁵¹ P. Ponte, Z. Papić, F. m. c. Huvneers, and D. A. Abanin, Phys. Rev. Lett. **114**, 140401 (2015), URL <https://link.aps.org/doi/10.1103/PhysRevLett.114.140401>.
- ⁵² A. Lazarides, A. Das, and R. Moessner, Phys. Rev. Lett. **115**, 030402 (2015), URL <https://link.aps.org/doi/10.1103/PhysRevLett.115.030402>.
- ⁵³ P. Bordia, H. Lüschen, U. Schneider, M. Knap, and I. Bloch, Nature Physics **13**, 460 EP (2017), article, URL <http://dx.doi.org/10.1038/nphys4020>.
- ⁵⁴ M. Nielsen, *Neural Networks and Deep Learning* (Determination Press, 2015).
- ⁵⁵ S. Hochreiter and J. Schmidhuber, Neural computation **9**, 1735 (1997).
- ⁵⁶ C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, ArXiv e-prints (2017), 1706.04599.
- ⁵⁷ N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, Journal of machine learning research **15**, 1929 (2014).
- ⁵⁸ D. P. Kingma and J. L. Ba, pp. 1–15 (2015), ISSN 09252312, 1412.6980v9, URL <http://arxiv.org/abs/1412.6980>.
- ⁵⁹ A. Pal and D. A. Huse, Phys. Rev. B **82**, 174411 (2010).
- ⁶⁰ D. J. Luitz, N. Laflorencie, and F. Alet, Physical Review B **91**, 081103 (2015), URL <https://journals.aps.org/prb/abstract/10.1103/PhysRevB.91.081103>.
- ⁶¹ M. Serbyn, Z. Papić, and D. A. Abanin, Phys. Rev. X **5**, 41047 (2015).
- ⁶² S. Bera, H. Schomerus, F. Heidrich-Meisner, and J. H. Bardarson, Phys. Rev. Lett. **115**, 046603 (2015), URL <https://link.aps.org/doi/10.1103/PhysRevLett.115.046603>.
- ⁶³ S. Iyer, V. Oganesyan, G. Refael, and D. A. Huse, Phys. Rev. B **87**, 134202 (2013), URL <https://link.aps.org/doi/10.1103/PhysRevB.87.134202>.
- ⁶⁴ S. Gopalakrishnan, M. Knap, and E. Demler, Phys. Rev. B **94**, 094201 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevB.94.094201>.
- ⁶⁵ M. Schreiber, S. S. Hodgman, P. Bordia, H. P. Lüschen, M. H. Fischer, R. Vosk, E. Altman, U. Schneider, and I. Bloch, Science **349**, 842 (2015), ISSN 0036-8075, 1501.05661v1, URL <http://www.sciencemag.org/cgi/doi/10.1126/science.aaa7432>.
- ⁶⁶ D. J. Luitz, N. Laflorencie, and F. Alet, Phys. Rev. B **93**, 060201 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevB.93.060201>.
- ⁶⁷ R. Gupta and M. J. Biercuk, ArXiv e-prints (2017), 1712.01291.
- ⁶⁸ M. Abadi et al., arXiv:1603.04467 (2016), 1603.04467, URL <http://arxiv.org/abs/1603.04467>.
- ⁶⁹ F. Chollet et al., *Keras*, <https://github.com/fchollet/keras> (2015).
- ⁷⁰ See Supplemental Material at [URL will be inserted by publisher] for analysis of the interpretability and robustness of the network predictions, comparison with a conventional order parameter for the MBL transition and a detailed explanation of the training procedure.

ACKNOWLEDGMENTS

E.v.N. gratefully acknowledges financial support from the Swiss National Science Foundation through grant P2EZP2-172185. E.v.N. also acknowledges fruitful discussions with Manuel Endres. E.B. is grateful to Netanel Lindner for his support and acknowledges financial support from the European Research Council (ERC) under the European Union Horizon 2020 Research and Innovation Programme (Grant Agreement No. 639172). G.R. is grateful to the the NSF for funding through the grant DMR-1040435 as well as the Packard Foundation. We are grateful for support from the IQIM, an NSF physics frontier center funded in part by the Moore Foundation. The authors used the TensorFlow⁶⁸ backend for Keras⁶⁹. E.v.N. and E.B. contributed equally to this work.