

Discriminative Cooperative Networks for Detecting Phase Transitions

Ye-Hua Liu*

*Institute for Theoretical Physics, ETH Zurich, 8093 Zurich, Switzerland
and Département de Physique and Institut Quantique, Université de Sherbrooke, J1K 2R1 Québec, Canada*

Evert P. L. van Nieuwenburg[†]

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



(Received 16 October 2017; published 26 April 2018)

The classification of states of matter and their corresponding phase transitions is a special kind of machine-learning task, where physical data allow for the analysis of new algorithms, which have not been considered in the general computer-science setting so far. Here we introduce an unsupervised machine-learning scheme for detecting phase transitions with a pair of discriminative cooperative networks (DCNs). In this scheme, a guesser network and a learner network cooperate to detect phase transitions from fully unlabeled data. The new scheme is efficient enough for dealing with phase diagrams in two-dimensional parameter spaces, where we can utilize an active contour model—the snake—from computer vision to host the two networks. The snake, with a DCN “brain,” moves and learns actively in the parameter space, and locates phase boundaries automatically.

DOI: [10.1103/PhysRevLett.120.176401](https://doi.org/10.1103/PhysRevLett.120.176401)

The richness of states of matter, together with the power of machine-learning techniques for recognizing and representing patterns, are revealing new methods for studying *emergent phenomena* in condensed matter physics. Paradigms in machine learning have been nicely mapped to those in physics. For example, the classification techniques in machine learning have been applied in detecting classical and quantum phase transitions [1–14], the artificial-neural-network architecture has inspired a high-quality ansatz for many-body wave functions [15–22], the generative power of energy-based statistical models is utilized to accelerate Monte Carlo simulations [23–30], and regression has aided material-property prediction [31–37]. Moreover, basic notions from both physics and machine learning can mutually inspire new insights, e.g., a relation between deep learning and the renormalization group [38–43].

In physics, the phase (e.g., magnetic vs nonmagnetic phase) is most efficient in summarizing material properties. When changing tuning parameters (e.g., temperature), the material properties may change discontinuously, which is called a phase transition. Machine-learning phase transitions are possible from two angles. In the *supervised* approach, physics knowledge is used to provide answers in limiting cases and the machine learner is asked to extrapolate to the transition point [1]. In the *unsupervised* approach, no such knowledge is assumed and the transition is sought by other means [2,7,11,12].

The confusion scheme proposed previously by us is a *hybrid* method [5], where no knowledge of the limiting cases is needed but the learning is still carried out in a supervised manner. Specifically, one first guesses a transition point and

then tries to train the machine with this guess. When the guess is correct, the machine learner achieves the highest performance. Here we gain the ability to find transitions at the cost of having to repeat the training for many guesses, which is computationally expensive.

In this work, we extend the confusion scheme by training a “guesser” together with the “learner.” This leads to a fully automated scheme—the *discriminative cooperative networks* (DCNs). In addition, phase transitions in two-dimensional (2D) parameter spaces share many common aspects with image-feature detection in computer vision. However, in images the data are the colors, whereas in physics they can be arbitrary results of measurements whose features might not be apparent to the human eye. This inspires us to use an active contour method [44], combined with the DCN scheme, to perform automated searching of phase boundaries in 2D phase diagrams.

We consider data that can be ordered along a tuning parameter λ . At various values of λ the data are described by $\mathbf{d}(\lambda)$, and can be thought of as a vector of real numbers—results of physical measurements at λ . We describe a neural network on an abstract level as a map \mathcal{N} that takes data $\mathbf{d}(\lambda)$ and infers the probability distribution $\mathcal{N}(\mathbf{d}(\lambda)) = (p_A, p_B, \dots)$, where p_i represents the probability of $\mathbf{d}(\lambda)$ belonging to phase i . Since the data are indexed by λ , this can be simplified by considering the probability distribution $\mathbf{L}(\lambda)$ directly on λ . At each λ only a single probability (corresponding to the correct phase) should equal to unity, and the rest zero. With phase transitions, the distribution varies with λ discontinuously, e.g., for a transition at $\lambda = \lambda_c$ between two phases A and B there are two components

$L_A(\lambda) = \Theta(\lambda_c - \lambda)$ and $L_B(\lambda) = \Theta(\lambda - \lambda_c)$, where Θ is the Heaviside step function.

For supervised learning, a large body of $\mathbf{d}(\lambda)$ with the corresponding correct answer $\mathbf{L}(\lambda)$ has to be known beforehand, and the neural network \mathcal{N} is trained with the goal $\mathcal{N}(\mathbf{d}(\lambda)) \rightarrow \mathbf{L}(\lambda)$. To achieve this goal, parameters $\mathbf{W}_{\mathcal{N}}$ that characterize the neural network are adjusted during training to minimize a cost function $C[\mathcal{N}(\mathbf{d}(\lambda)), \mathbf{L}(\lambda)]$, quantifying the mismatch between the network's prediction and the known answer. C depends implicitly on the parameters $\mathbf{W}_{\mathcal{N}}$ through \mathcal{N} and can be minimized using gradient descent methods.

Typical machine-learning data live in high-dimensional feature spaces in an *unordered* fashion. The number of ways to separate them into two classes is 2^N where N is the size of the data set. For phase transitions, however, all data are *ordered* in the parameter space, and for a single transition point, the number of ways is merely $N + 1$. In physics, it is affordable to enumerate all these possibilities to find the most reasonable separation point. This observation led to the confusion scheme [5], where one guesses the transition point $\lambda_g \rightarrow \lambda_c$ and then trains the learner network \mathcal{N} . By monitoring the number of “correctly” classified samples according to this guess—the performance, the true value for λ_g can be deduced. It turns out the true value is the guess for which the performance is optimal, because here the assigned probabilities in $\mathbf{L}(\lambda)$ and the structures in $\mathbf{d}(\lambda)$ are the most consistent, such that the learner network is *least confused* by the training.

In the previous proposal, we searched for the optimal λ_g by a brute-force scan of the parameter space. For phase transitions in higher-dimensional parameter spaces, this approach is inefficient. In this work we introduce the guesser network \mathcal{G} . It performs the map $\lambda \rightarrow \mathcal{G}(\lambda)$, representing the probabilities of λ belonging to each possible phase. That is, now the guesser provides $\mathbf{L}(\lambda)$. The guesser is itself characterized by a set of parameters $\mathbf{W}_{\mathcal{G}}$ on which we wish to perform gradient descent. The overall cost function of the learner \mathcal{N} and guesser \mathcal{G} is now $C[\mathcal{N}(\mathbf{d}(\lambda)), \mathcal{G}(\lambda)]$, see Fig. 1(a). In this way, we have promoted the human input \mathbf{L} to an active agent \mathcal{G} . During training, the learner \mathcal{N} tries to learn the data according to the suggested labels $\mathcal{G}(\lambda)$ obtained from the guesser, and the guesser tries to provide a better set of labels—they cooperatively optimize the cost C .

We first assume one-dimensional (1D) parameter space with two phases, and propose a logistic-regression guesser network with one(two) input(output) neuron(s): $\mathcal{G}_{A,B}(\lambda) = f[s_{A,B}(\lambda - \lambda_g)/\sigma]$, where $f(x) = 1/(1 + e^{-x})$ is the logistic (sigmoid) function, $A(B)$ denotes the first(second) output neuron, and $s_{A,B} = -, +$. The guesser is hence characterized by two parameters λ_g and σ , setting, respectively, the guessed transition point and the sharpness of the transition. Gradient descent can be performed on both λ_g and σ .

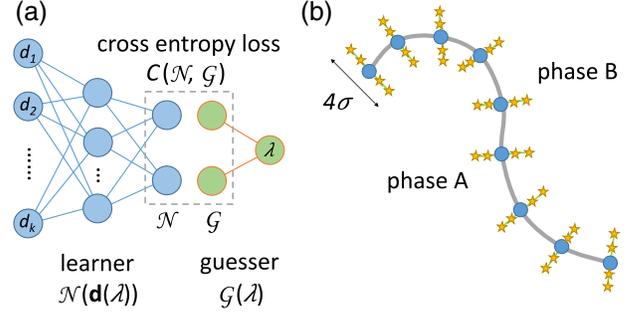


FIG. 1. Schematics for the proposed algorithms. (a) The DCN scheme for learning phase transitions solely, from the data set $\{(\lambda, \mathbf{d}(\lambda))\}$, where λ is the tuning parameter and $\mathbf{d}(\lambda)$ is a vector of measurements at λ . (b) The DCN snake. The blue circles are the snake nodes, the green lines denote normal directions at the nodes. Samples (stars) are generated in the normal direction at each node and are assigned a label according to its distance to the snake. Snakes could be open or closed, and can move to the correct phase boundary (gray line) automatically. The open snake in this figure has 9 nodes and during motion generates minibatches of training data with minibatch size $N_b = 36$ (N_b is much larger in real simulations).

We use the cross entropy cost function $C(\mathcal{N}, \mathcal{G}) = -\log \mathcal{N} \cdot \mathcal{G} - \log(1 - \mathcal{N}) \cdot (1 - \mathcal{G})$, which is suitable for classification problems. The gradient of C on the guesser network is obtained by the following equations:

$$\begin{aligned} \frac{\partial C}{\partial \mathcal{G}} &= -\log \mathcal{N} + \log(1 - \mathcal{N}), \\ \frac{\partial \mathcal{G}_{A,B}}{\partial \lambda_g} &= -\frac{s_{A,B}}{4\sigma \cosh^2[(\lambda - \lambda_g)/2\sigma]}, \\ \frac{\partial \mathcal{G}}{\partial \sigma} &= \frac{\lambda - \lambda_g}{\sigma} \frac{\partial \mathcal{G}}{\partial \lambda_g}. \end{aligned} \quad (1)$$

These equations fully determine the dynamics of the guesser: $\Delta \lambda_g = -\alpha_{\lambda_g} \partial C / \partial \lambda_g$ and $\Delta \sigma = -\alpha_{\sigma} \partial C / \partial \sigma$, where α_{λ_g} and α_{σ} are the learning rates for the two parameters, respectively. The dynamics of the learner follows $\Delta \mathbf{W}_{\mathcal{N}} = -\alpha_{\mathcal{N}} \partial C / \partial \mathbf{W}_{\mathcal{N}}$ with another independent learning rate $\alpha_{\mathcal{N}}$, here the gradient is obtained by the back-propagation algorithm [45].

At this point, one could conceptually regard the guesser and learner together as one compound agent, capable of *self-learning*. We call this scheme *discriminative cooperative networks* (DCNs), with the name inspired by the powerful *generative adversarial networks* (GANs) [46] for generating samples resembling the training data.

The DCN scheme is efficient because there is no need for repetitive training at each guess. This allows us to move to higher-dimensional parameter spaces. Here we focus on 2D since physics studies usually report phase diagrams in 2D parameter spaces. Inspired by the computer vision techniques for finding image features, we use an active contour

model—the snake [44]—for the parametrization of the guesser.

In computer vision, the snake is a discretized curve of linked nodes, $\mathbf{r}(s) = (x(s), y(s))$, parametrized by $s \in [0, 1]$ (for closed snakes) or $s \in [0, 1]$ (for open snakes), see Fig. 1(b). The nodes can move actively under “image forces,” which are the minus gradients of an “external energy,” with respect to the snake nodes. Specifically, the external energy is the total potential energy $E_{\text{external}} = \int_0^1 ds \phi(\mathbf{r}(s))$, with the potential $\phi(\mathbf{r})$ proportional to the local color intensity (gradient of color intensity) for line (edge) detection. To keep the snake smooth, internal forces are also introduced, which are derived from the internal energy $E_{\text{internal}} = \int_0^1 ds (\alpha |(\partial \mathbf{r} / \partial s)|^2 + \beta |(\partial^2 \mathbf{r} / \partial s^2)|^2)$. Increasing α makes for a more “elastic” snake by preventing stretching and β a more “solid” snake by preventing bending. The snake evolves in time to lower its total energy $E_{\text{total}} = E_{\text{external}} + E_{\text{internal}}$, and the equation of motion $\dot{\mathbf{v}} \propto -\delta E_{\text{total}} / \delta \mathbf{r}$ is implemented numerically [44].

In this work, we combine the DCN scheme in artificial intelligence with the snake in computer vision, and the result is an *intelligent snake*. To do this, we replace the conventional image force in computer vision with the machine-learning gradient $\delta E_{\text{external}} / \delta \mathbf{r} \rightarrow \partial C / \partial \lambda_g$. The 1D DCN scheme requires training data from both sides of the guessed transition point. This implies, for the 2D case, a width of the snake. The width, denoted again by σ , is generically different at each node, and enables the snake to sense its surroundings by selecting training samples in its vicinity within this length scale, as shown in Fig. 1(b). Specifically the sample points are drawn at each node perpendicularly to the snake, with distances uniformly picked in $[-2\sigma, 2\sigma]$. The 2D guesser function is then locally the same as in the 1D case, evaluated by each node in its perpendicular direction. For implementation details, see Ref. [47]. We note the probing of data within a window (in searching for distinct phases) is a powerful concept that is also successfully used in Ref. [10].

Ising model.—We test our scheme on 1D parameter space by studying the classical Ising model on the square lattice:

$$H = -J \sum_{\langle i,j \rangle} s_i s_j, \quad Z = \sum_{\{s\}} e^{-H/T}, \quad (2)$$

where $s_i = \{-1, 1\}$ are the Ising spins, $\langle i, j \rangle$ denotes nearest neighbors with coupling J , and the tuning parameter is the temperature T . This model has a thermal phase transition from the ferromagnetic phase (with aligned spins) to the paramagnetic phase (with random spins) when the temperature is increased across $T_c \sim 2.27J$. The training data $\mathbf{d}(\lambda) = \{s\}_T$ are spin configurations drawn from a Monte Carlo simulation on an L by L square lattice. We select 100 temperatures uniformly from $0.1J$ to $5J$ and prepare 100 samples at each temperature. Every

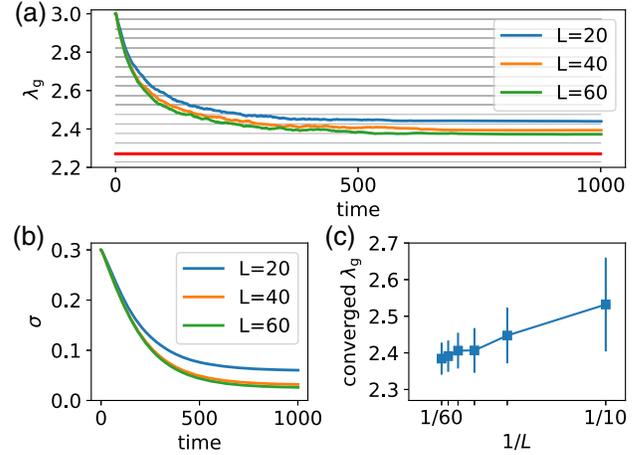


FIG. 2. DCN scheme for the Ising transition. (a) Starting from a higher guess of the transition point, the gradient on the guesser pushes it to move down. The red line marks the exact transition point and the gray lines the temperatures (in the range of the figure) for generating Monte Carlo samples. (b) During training the width σ decreases, meaning the combined self-learner is able to distinguish the two phases sharper. Training on samples from larger lattices is faster and more accurate. (c) Finite-size effect on the converged guess λ_g , where the length of error bars denotes the converged σ . Network architecture: fully connected with L^2 input neurons, L hidden neurons, and 2 output neurons. Hyperparameters for training: minibatch size $N_b = 100$; initial learning rates $\alpha_N = 0.1$, $\alpha_{\lambda_g} = 0.025$, $\alpha_\sigma = 0.001$, decay rate 0.995; dropout keep probability 0.8, ℓ_2 regularization 0.0001. We have set a lower bound for the width $T > 0.01$ and used the minibatch stochastic gradient descent optimizer [49] with batch normalization [50].

minibatch consists of $N_b = 100$ random samples, one from each temperature [48]. Time is measured by the number of learned minibatches. During training, the guesser moves toward the exact transition point $T_c \sim 2.27J$ and decreases the width σ because the discrimination is sharper and sharper (Fig. 2). λ_g does not converge to the exact value when increasing L , because the networks most likely learn the order parameter, which is the simplest, but not the sharpest signal for detecting phase transition. In a future study we investigate the possibility for the networks to learn also the fluctuations of order parameters.

Bose-Hubbard model.—As a first example for applying the DCN scheme in 2D parameter spaces, we choose the Bose-Hubbard model:

$$H = -J \sum_{\langle i,j \rangle} (b_i^\dagger b_j + b_j^\dagger b_i) + \sum_i \left(\frac{U n_i (n_i - 1)}{2} - \mu n_i \right), \quad (3)$$

where b^\dagger/b is the bosonic creation or annihilation operator. Regarding the Hubbard interaction U as the energy unit, for each chemical potential μ , the model

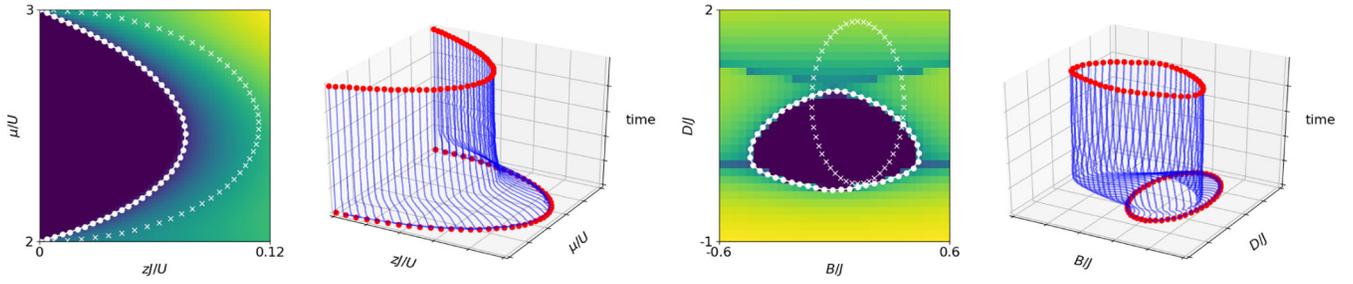


FIG. 3. DCN snake for 2D parameter spaces. (left panels) The Mott insulator to superfluid transition in the Bose Hubbard model. (right panels) The topological transition in the spin-1 antiferromagnetic Heisenberg chain. Color plots (purple to yellow goes from zero to nonzero values) show the average hopping for the Bose Hubbard model, and the difference between the largest two eigenvalues of the reduced density matrix for half of the Heisenberg chain. For the Bose Hubbard model we create a large open snake with the head and tail fixed at integer chemical potentials $\mu/U = 2, 3$. In this case the snake shrinks and stops at the correct phase boundary. For the Heisenberg model, we create a large and offset closed snake, which then moves, rotates, shrinks, and finally stays at the Haldane pocket. Parameters for snakes: number of nodes 50; dynamic width at each node is initialized to $T = 0.06$ (normalized by the ranges of parameters) and clipped to $0.03 < T < 0.08$; regularizations $\alpha = 0.002$, $\beta = 0.4$, $\gamma = 0.25$ (see Ref. [44] for details). Network architecture: fully connected with 80 input neurons, 80 hidden neurons, and 2 output neurons. Hyperparameters for training: minibatch size $N_b = 1500$; initial learning rates $\alpha_{\mathcal{N}} = 0.01$, $\alpha_{\lambda_j} = 0.0008$, $\alpha_{\sigma} = 0.0002$, decay rate 0.9999; dropout keep probability 0.8, ℓ_2 regularization 0.0001. We used the ADAM optimizer [54] because the inputs are sparse for these models [49].

has a quantum phase transition (at zero temperature) from the Mott insulating state to the superfluid state, when the hopping J is increased [51]. A useful indicator of this phase transition is the average hopping $\langle K \rangle$ where $K = \sum_{\langle i,j \rangle} (b_i^\dagger b_j + b_j^\dagger b_i)$. Note the notion of K is *unknown to the initial untrained snake*, otherwise the problem reduces to computer vision where machine learning is not needed. The critical point J_c reaches local maxima when the system is at commensurate fillings, corresponding to half-integers μ/U . A phase diagram of this system results in the series of well-known Mott lobes. We use the mean-field theory developed in Ref. [52] to generate vector data $\mathbf{d}(\lambda_1, \lambda_2) = \mathbf{F}(J, \mu)$, where F_n with $n = 0, 1, \dots, \infty$ denotes the amplitude for having n bosons per site, and a cutoff of $n_{\max} = 79$ is chosen for numerics. We target the third Mott lobe with $2 \leq \mu/U \leq 3$, and the snake successfully captures the phase boundary as shown in Fig. 3. In this case the phase boundary touches the boundary of the parameter space, so we use an open snake with a fixed head and tail at known transition points. The snake's motion is then restricted to shrinking or expanding. It is important to emphasize here that we have used knowledge of only two points along the $J = 0$ axis in the whole phase diagram, and that the training data seen by the snake is not the average hopping as shown in the background, but the vector data $\mathbf{F}(J, \mu)$ mentioned above [53].

Spin-1 Heisenberg chain.—We now move to a quantum phase transition beyond mean-field theory. We choose the spin-1 antiferromagnetic Heisenberg chain with anisotropy and transverse magnetic field:

$$H = J \sum_i \mathbf{S}_i \cdot \mathbf{S}_{i+1} + \sum_i [D(S_i^z)^2 - BS_i^x], \quad (4)$$

where S_i^a are 3 by 3 matrices satisfying $[S_i^a, S_j^b] = i\hbar\delta_{ij}\sum_c \epsilon_{abc}S_i^c$. In the 2D parameter space of magnetic field B/J vs anisotropy D/J , this model has a pocket named the Haldane phase—a topologically nontrivial phase—around zero magnetic field and anisotropy [55]. The transition across the boundary of this pocket can be detected by a change in the degeneracy structure of the entanglement spectrum (eigenvalues of the reduced density matrix for part of the spin chain in the ground state), but again the initial untrained snake is unaware of this. For the training data, we simulate an infinite chain with translational invariance using infinite time-evolving block decimation [56] with bond dimension $m = 80$, and record all m eigenvalues $\{\epsilon_1, \dots, \epsilon_m\}$ of the reduced density matrix when the chain is cut by half at a bond, i.e., $\mathbf{d}(\lambda_1, \lambda_2) = \{\epsilon\}_{B,D}$. The result is shown in Fig. 3. In this model the phase boundary is closed and located near the center of the parameter space. For this reason we use a closed (periodic) snake whose motion now also contains translation and rotation.

In this Letter we have proposed the discriminative cooperative networks, capable of self-consistently finding transition points. The high efficiency of this scheme allows us to explore 2D parameter spaces, where we utilized the snake model from computer vision. Our method is, in spirit, similar to the actor-critic scheme for reinforcement learning [57] and the adversarial training scheme for generative models [46].

The major limitation for the snake is the need for an initial state that has overlap with the desired features to be detected, so that it is able to probe a gradient. This was also true for their use in computer vision. In applications to phase diagrams, we have the clear advantage of some known extreme limits at which we can fix the snake. We

can also overcome this problem by scaling or moving the snake.

Ch'ng *et al.* have proposed to train neural networks deep inside the known phases with supervision, and then use them to extrapolate the whole phase diagram [3]. Such a method is, compared to our method, simpler and faster. However, the data for supervised training have to be carefully chosen, otherwise interpolation of the phase boundary could be qualitatively incorrect [3]. On the contrary, the snake can actively explore a much larger area in the parameter space. For general phase transition problems, one could use both methods complementarily.

Machine-learning applications usually assume the existence of big data. However, in science it might be expensive to obtain these data. With the DCN scheme, it is possible for a machine-learning agent to suggest parameters for the physicist to carry out experiments or simulations, and rapidly locate interesting phenomena. In this Letter we put forward a proposal to realize this scheme.

The authors thank L. Wang, S.D. Huber, S. Trebst, K. Hepp, M. Sigrist, and T.M. Rice for reading the manuscript and helpful suggestions. Y.-H.L. thanks G. Sordi and A.-M. Tremblay for stimulating discussions. Y.-H.L. is supported by ERC Advanced Grant SIMCOFE and the Canada First Research Excellence Fund. E. P. L. v. N. gratefully acknowledges financial support from the Swiss National Science Foundation (SNSF) through Grant No. P2EZP2-172185. The authors used TensorFlow [58] for machine learning.

*yehua.liu@usherbrooke.ca

†evert@caltech.edu

- [1] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, *Nat. Phys.* **13**, 431 (2017).
- [2] L. Wang, Discovering phase transitions with unsupervised learning, *Phys. Rev. B* **94**, 195105 (2016).
- [3] K. Ch'ng, J. Carrasquilla, R. G. Melko, and E. Khatami, Machine Learning Phases of Strongly Correlated Fermions, *Phys. Rev. X* **7**, 031038 (2017).
- [4] 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).
- [5] E. P. L. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Learning phase transitions by confusion, *Nat. Phys.* **13**, 435 (2017).
- [6] F. Schindler, N. Regnault, and T. Neupert, Probing many-body localization with neural networks, *Phys. Rev. B* **95**, 245134 (2017).
- [7] 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).
- [8] S. J. Wetzel, Unsupervised learning of phase transitions: From principal component analysis to variational autoencoders, *Phys. Rev. E* **96**, 022140 (2017).
- [9] T. Ohtsuki and T. Ohtsuki, Deep learning the quantum phase transitions in random two-dimensional electron systems, *J. Phys. Soc. Jpn.* **85**, 123706 (2016).
- [10] P. Broecker, F. Assaad, and S. Trebst, Quantum phase recognition via unsupervised machine learning, [arXiv:1707.00663](https://arxiv.org/abs/1707.00663).
- [11] N. C. Costa, W. Hu, Z. J. Bai, R. T. Scalettar, and R. R. P. Singh, Principal component analysis for fermionic critical points, *Phys. Rev. B* **96**, 195138 (2017).
- [12] 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).
- [13] W.-J. Rao, Z. Li, Q. Zhu, M. Luo, and X. Wan, Identifying product order with restricted Boltzmann machines, [arXiv:1709.02597v1](https://arxiv.org/abs/1709.02597v1) [*Phys. Rev. B* (to be published)].
- [14] Z. Li, M. Luo, and X. Wan, Extracting critical exponent by finite-size scaling with convolutional neural networks, [arXiv:1711.04252v1](https://arxiv.org/abs/1711.04252v1).
- [15] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [16] D.-L. Deng, X. Li, and S. D. Sarma, Machine learning topological states, *Phys. Rev. B* **96**, 195145 (2017).
- [17] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, On the equivalence of restricted Boltzmann machines and tensor network states, *Phys. Rev. B* **97**, 085104 (2018).
- [18] X. Gao and L.-M. Duan, Efficient representation of quantum many-body states with deep neural networks, *Nat. Commun.* **8**, 662 (2017).
- [19] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Many-body quantum state tomography with neural networks, [arXiv:1703.05334](https://arxiv.org/abs/1703.05334) [*Nat. Phys.* (to be published)].
- [20] Z. Cai and J. Liu, Approximating quantum many-body wave functions using artificial neural networks, *Phys. Rev. B* **97**, 035116 (2018).
- [21] 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).
- [22] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, Neural-Network Quantum States, String-Bond States, and Chiral Topological States, *Phys. Rev. X* **8**, 011006 (2018).
- [23] L. Wang, Exploring cluster Monte Carlo updates with Boltzmann machines, *Phys. Rev. E* **96**, 051301 (2017).
- [24] L. Huang, Y.-F. Yang, and L. Wang, Recommender engine for continuous-time quantum Monte Carlo methods, *Phys. Rev. E* **95**, 031301 (2017).
- [25] L. Huang and L. Wang, Accelerated Monte Carlo simulations with restricted Boltzmann machines, *Phys. Rev. B* **95**, 035105 (2017).
- [26] H. Fujita, Y. O. Nakagawa, S. Sugiura, and M. Oshikawa, Construction of Hamiltonians by machine learning of energy and entanglement spectra, *Phys. Rev. B* **97**, 075114 (2018).
- [27] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, Self-learning Monte Carlo method, *Phys. Rev. B* **95**, 041101 (2017).

- [28] J. Liu, H. Shen, Y. Qi, Z. Y. Meng, and L. Fu, Self-learning Monte Carlo method and cumulative update in fermion systems, *Phys. Rev. B* **95**, 241104 (2017).
- [29] X. Y. Xu, Y. Qi, J. Liu, L. Fu, and Z. Y. Meng, Self-learning quantum Monte Carlo method in interacting fermion systems, *Phys. Rev. B* **96**, 041119 (2017).
- [30] Y. Nagai, H. Shen, Y. Qi, J. Liu, and L. Fu, Self-learning Monte Carlo method: Continuous-time algorithm, *Phys. Rev. B* **96**, 161102 (2017).
- [31] M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, Fast and Accurate Modeling of Molecular Atomization Energies with Machine Learning, *Phys. Rev. Lett.* **108**, 058301 (2012).
- [32] F. A. Faber, A. Lindmaa, O. A. von Lilienfeld, and R. Armiento, Machine Learning Energies of 2 Million Elpasolite (ABC_2D_6) Crystals, *Phys. Rev. Lett.* **117**, 135502 (2016).
- [33] G. Pilania, C. Wang, X. Jiang, S. Rajasekaran, and R. Ramprasad, Accelerating materials property predictions using machine learning, *Sci. Rep.* **3**, 2810 (2013).
- [34] L.-F. Arsenault, A. Lopez-Bezanilla, O. A. von Lilienfeld, and A. J. Millis, Machine learning for many-body physics: The case of the Anderson impurity model, *Phys. Rev. B* **90**, 155136 (2014).
- [35] L.-F. Arsenault, O. A. von Lilienfeld, and A. J. Millis, Machine learning for many-body physics: efficient solution of dynamical mean-field theory, [arXiv:1506.08858](https://arxiv.org/abs/1506.08858).
- [36] A. P. Bartók, S. De, C. Poelking, N. Bernstein, J. R. Kermode, G. Csányi, and M. Ceriotti, Machine learning unifies the modeling of materials and molecules, *Sci. Adv.* **3**, e1701816 (2017).
- [37] F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K.-R. Müller, Bypassing the Kohn-Sham equations with machine learning, *Nat. Commun.* **8**, 872 (2017).
- [38] O. Landon-Cardinal and D. Poulin, Practical learning method for multi-scale entangled states, *New J. Phys.* **14**, 085004 (2012).
- [39] C. Bény, Deep learning and the renormalization group, [arXiv:1301.3124](https://arxiv.org/abs/1301.3124).
- [40] P. Mehta and D. J. Schwab, An exact mapping between the Variational Renormalization Group and Deep Learning, [arXiv:1410.3831](https://arxiv.org/abs/1410.3831).
- [41] H. W. Lin, M. Tegmark, and D. Rolnick, Why does deep and cheap learning work so well? *J. Stat. Phys.* **168**, 1223 (2017).
- [42] D. Rolnick and M. Tegmark, The power of deeper networks for expressing natural functions, [arXiv:1705.05502](https://arxiv.org/abs/1705.05502).
- [43] M. Koch-Janusz and Z. Ringel, Mutual information, neural networks and the renormalization group, [arXiv:1704.06279](https://arxiv.org/abs/1704.06279).
- [44] M. Kass, A. Witkin, and D. Terzopoulos, Snakes: Active contour models, *Int. J. Comput. Vis.* **1**, 321 (1988).
- [45] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Learning representations by back-propagating errors, *Nature (London)* **323**, 533 (1986).
- [46] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative adversarial networks, [arXiv:1406.2661](https://arxiv.org/abs/1406.2661).
- [47] The source code can be found in <https://github.com/rhinech/snake>.
- [48] We have preprocessed each configuration by flipping all its spins when the net magnetization $\sum_i s_i$ is negative.
- [49] S. Ruder, An overview of gradient descent optimization algorithms, [arXiv:1609.04747](https://arxiv.org/abs/1609.04747).
- [50] S. Ioffe and C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, [arXiv:1502.03167](https://arxiv.org/abs/1502.03167).
- [51] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher, Boson localization and the superfluid-insulator transition, *Phys. Rev. B* **40**, 546 (1989).
- [52] W. Krauth, M. Caffarel, and J.-P. Bouchaud, Gutzwiller wave function for a model of strongly interacting bosons, *Phys. Rev. B* **45**, 3137 (1992).
- [53] Additionally, we have tested that the snake is capable of finding the lobe from an initially circular (periodic) configuration.
- [54] D. P. Kingma and J. L. Ba, Adam: A method for stochastic optimization, in *International Conference on Learning Representations, San Diego (2015)*, <https://iclr.cc/archive/www/doku.php%3Fid=iclr2015:accepted-main.html>.
- [55] F. Pollmann, A. M. Turner, E. Berg, and M. Oshikawa, Entanglement spectrum of a topological phase in one dimension, *Phys. Rev. B* **81**, 064439 (2010).
- [56] G. Vidal, Classical Simulation of Infinite-Size Quantum Lattice Systems in One Spatial Dimension, *Phys. Rev. Lett.* **98**, 070201 (2007).
- [57] V. R. Konda and J. Tsitsiklis, Actor-critic algorithms, *SIAM J. Control Optim.* **42**, 1143 (2003).
- [58] M. Abadi *et al.*, TensorFlow: Large-scale machine learning on heterogeneous distributed systems, [arXiv:1603.04467](https://arxiv.org/abs/1603.04467).