Deep learning Hamiltonians from disordered image data in quantum materials

S. Basak,[1,2] M. Alzate Banguero,[3] L. Burzawa,[4] F. Simmons,[1,2] P. Salev,[5,6] L. Aigouy,[3] M. M. Qazilbash,[7] I. K. Schuller,[6] D. N. Basov,[8] A. Zimmers,[3] and E. W. Carlson[1,2,*]

[1]*Department of Physics and Astronomy, Purdue University, West Lafayette, Indiana 47907, USA*
[2]*Purdue Quantum Science and Engineering Institute, West Lafayette, Indiana 47907, USA*
[3]*Laboratoire de Physique et d'Étude des Matériaux, ESPCI Paris, PSL Université, CNRS, Sorbonne Université, 75005 Paris, France*
[4]*Department of Computer Science, Purdue University, West Lafayette, Indiana 47907, USA*
[5]*Department of Physics and Astronomy, University of Denver, Denver, Colorado 80208, USA*
[6]*Department of Physics and Center for Advanced Nanoscience, University of California–San Diego, La Jolla, California 92093, USA*
[7]*Department of Physics, College of William and Mary, Williamsburg, Virginia 23187, USA*
[8]*Department of Physics, Columbia University, New York, New York 10027, USA*

(Received 27 October 2022; revised 5 April 2023; accepted 6 April 2023; published 10 May 2023)

The capabilities of image probe experiments are rapidly expanding, providing new information about quantum materials on unprecedented length- and timescales. Many such materials feature inhomogeneous electronic properties with intricate pattern formation on the observable surface. This rich spatial structure contains information about interactions, dimensionality, and disorder—a spatial encoding of the Hamiltonian driving the pattern formation. Image recognition techniques from machine learning are an excellent tool for interpreting information encoded in the spatial relationships in such images. Here, we develop a deep learning framework for using the rich information available in these spatial correlations in order to discover the underlying Hamiltonian driving the patterns. We first vet the method on a known case, scanning near-field optical microscopy on a thin film of VO_2. We then apply our trained convolutional neural network architecture to new optical microscope images of a different VO_2 film as it goes through the metal-insulator transition. We find that a two-dimensional Hamiltonian with both interactions and random field disorder is required to explain the intricate, fractal intertwining of metal and insulator domains during the transition. This detailed knowledge about the underlying Hamiltonian paves the way for using the model to control the pattern formation via, e.g., tailored hysteresis protocols. We also introduce a distribution-based confidence measure on the results of a multilabel classifier, which does not rely on adversarial training. In addition, we propose a machine-learning-based criterion for diagnosing a physical system's proximity to criticality.

DOI: 10.1103/PhysRevB.107.205121

I. INTRODUCTION

The types of surface probes, such as atomic force microscopy (AFM), scanning tunneling microscopy (STM), and scattering scanning near-field infrared microscopy (SNIM), among many others [1,2], in addition to the wealth of data they generate, are increasing at a rapid pace. As often happens in science, new experimental frontiers reveal new physics: These scanning and image probe experiments often reveal complex electronic pattern formation spanning multiple lengthscales at the surface of correlated quantum materials, even when they are atomically smooth [3–12]. For example, manganites can have ferromagnetic and antiferromagnetic regions that coexist on multiple lengthscales [9]. In the unidirectional electronic glass in cuprates [4], domains of stripe orientation take fractal form with correlations over four orders of magnitude in lengthscale [13]. Magnetic domains in $NdNiO_3$ were also revealed to have fractal textures [14]. We focus here on VO_2, a material whose metal and insulator domains can show self-similar structure over multiple lengthscales [3,8,15].

Unfortunately, most of our theoretical tools are designed for understanding and describing homogeneous electronic states. Therefore, it is vital that we envision new theoretical frameworks for understanding why the patterns form in strongly correlated materials. The cluster analysis techniques we developed for interpreting these images have already uncovered universal behavior among disparate quantum materials [3,7,8,14,16], but the methods only work on systems near criticality, and for sufficiently large fields of view. Powerful image recognition methods from machine learning (ML) hold the potential to complement and extend these analyses into new regimes.

There has been tremendous growth recently in the application of ML methods to condensed matter. (For reviews, see Refs. [17–20].) ML is being applied as a tool to tackle various problems in condensed-matter physics, including disordered and glassy systems [21–23], quantum many body problems [24], quantum transport [25], renormalization group [26,27], and big data in materials science [28,29]. ML also benefits from physics, an area known as physics-inspired ML theory

*Corresponding author: ewcarlson@purdue.edu

[17]. Applied to experimental data, ML has been used to detect in which phase of matter a physical system is [30], and to aid in the experimental detection of the glass transition temperature [31]. Other common uses of ML for experimental data include the extraction of material parameters from experiment [28,32], or using ML to replace a lengthy and time-consuming fitting procedure [32,33].

Regarding phase transitions, ML has been used to detect in which phase of matter a theoretical configuration is [19,20,34], as well as to identify the transition temperature of a theoretical model [19,34,35], each in cases in which a particular Hamiltonian is already assumed. Relatively little attention has been paid to the critical region [20], where domains display power-law structure across multiple lengthscales. In addition, much of the work done to identify phases or detect phase transitions has been purely computational, with the Hamiltonian assumed [19,20,34,35]. By contrast, our method utilizes the rich spatial correlations available in near-critical configurations to detect which *Hamiltonian* should be used to describe a physical system, and we apply the method to experimentally derived data.

Here we develop a deep learning (DL) classifier to recognize spatial configurations from several different Hamiltonians. We test the DL classifier on experimental image data of VO_2 obtained via SNIM, and then we apply it to new optical microscope images of VO_2. Convolutional neural networks in particular are heavily used in image classification. We have previously shown that with ML, images from simulation can be classified with a very good accuracy of $\sim$97% [36]. Here we show that a DL architecture can classify two-dimensional (2D) surface images into one of seven candidate theoretical models, to even better accuracy (>99%). We introduce a symmetry reduction method that reduces training time over the data augmentation method. In addition, we use the DL model on experimental images derived from SNIM and optical microscope data to discover the underlying Hamiltonian driving pattern formation of metal and insulator puddles in films of VO_2. We also introduce a method for judging the confidence of a multilabel classifier, based on the multivariate distribution of values of the output nodes. We furthermore propose that this confidence measure tracks proximity to criticality.

This article is organized as follows: In Sec. II A, we give an overview of the Hamiltonians of interest from statistical mechanics in this paper. Section III shows the end-to-end deep learning architecture and process. We demonstrate the effectiveness of symmetry reduction to reduce training time as compared with data augmentation, and we develop a confidence criterion to judge the reliability of predictions. In Sec. IV, we make predictions on SNIM and optical microscope data on thin films of VO_2. We show that using only simulated data for training, we have developed a robust deep learning classification model that can learn the Hamiltonian driving pattern formation from experimental surface probe images.

II. DEVELOPING A DEEP LEARNING MODEL TO REVEAL UNDERLYING HAMILTONIANS

We first construct several possible Hamiltonians that could potentially describe the morphology of these metal and insulator domains, including the multiscale behavior. Then, we use numerical simulations to generate thousands of spatial configurations of metal and insulator domains that can arise in these Hamiltonians. Next, we develop and train a deep learning (DL) convolutional neural network (CNN) on a subset of these images in supervised learning mode. After we validate that the DL model can correctly identify the underlying Hamiltonian from a single domain configuration with greater than 99% accuracy, we then apply our trained DL model to experimental data on VO_2 obtained via both SNIM and optical microscopy.

A. Candidate Hamiltonians and the morphologies they produce

We use numerical simulations to generate typical configurations of metal and insulator domains that can arise from various model Hamiltonians that could potentially be controlling the metal-insulator domain structure. Simulation methods are described briefly in this section and in detail in the Supplemental Material (SM), and they are summarized in Table I. As VO_2 undergoes a temperature-driven transition from metal to insulator and vice versa, the macroscopic resistivity changes by four to five orders of magnitude. However, rather than doing so homogeneously, we previously used SNIM to produce spatially resolved images of the metal and insulator domains, which revealed that VO_2 thin films transition *inhomogeneously*, with metal and insulator domains interleaving with each other over a wide range of lengthscales [3,37]. (Note that the voltage-driven transition is also inhomogeneous, as revealed by optical measurements [38].) We introduce a range of possible Hamiltonians that could be responsible for driving the multiscale textures during the metal-insulator transition in VO_2. Domain configurations from these Hamiltonians will be used to train the DL model to identify the underlying physics driving pattern formation in this material.

Because the experimental probes of interest are directly measuring electronic degrees of freedom, we construct Hamiltonians that are about these electronic degrees of freedom. The intricate patterns of metal and insulator domains happen across multiple lengthscales from the resolution of the probes all the way out to the field of view ($\approx$20 nm–4 μm for SNIM and $\approx$370 nm–28 μm for optical microscope data). Therefore, we construct Hamiltonians at the order-parameter level. Because the theories are constructed at the order-parameter level, they are not microscopic, although they can provide constraints on microscopic models.

First, we consider a clean, interacting Hamiltonian. A reasonable ansatz is that the interaction energy between neighboring domains is lower for like domains than for unlike domains. We model this proclivity toward neighboring like domains [8,39] with a nearest-neighbor Ising Hamiltonian:

$$H = -J \sum_{\langle i,j \rangle} \sigma_i \sigma_j - \sum_i h\sigma_i, \tag{1}$$

where $\sigma_i = \pm 1$ is a two-state local order parameter, which in this case tracks metal (e.g., $\sigma = 1$) and insulator (e.g., $\sigma = -1$) domains. In an infinite size system, this model undergoes an equilibrium, second-order phase transition as a function of temperature at a critical temperature of $T_c^{2D} \approx 2.27J$ in two-dimensional systems and $T_c^{3D} \approx 4.51J$ in three-dimensional systems [40,41]. However, the model also undergoes a first-order phase transition as a function of applied field h. This

TABLE I. Parameters of simulations of the statistical mechanics models. In the first part of the table, parameters are in the critical region. In the second part of the table, parameters are not near criticality.

Model	Parameters simulated	Simulation method
2D clean Ising (C-2D)	$T = 2.25 - 2.64$	Monte Carlo
3D clean Ising (C-3D)	$T = 4.45 - 4.65$	Monte Carlo
2D RFIM (RF-2D)	$R = 1.00 - 1.19$	Zero-temperature field sweep
3D RFIM (RF-3D)	$R = 2.25 - 2.29$	Zero-temperature field sweep
2D percolation (P-2D)	$p = 0.57 - 0.61$	Biased coin flip
3D percolation (P-3D)	$p = 0.29 - 0.33$	Biased coin flip
	$p = 0.02 - 0.2$	Biased coin flip
Noncritical percolation (P*)	$p = 0.48 - 0.52$	Biased coin flip
	$p = 0.8 - 0.98$	Biased coin flip

first-order line terminates in the critical end point mentioned above. The phenomenology of a first-order line terminating in a critical end point is why this model is often used in conjunction with the liquid-gas transition. For example, when the liquid-gas transition is approached along the coexistence curve in temperature and pressure, the transition is second order [see the red dotted line in Fig. 2(a)] [42]. The influence of that critical point is felt throughout a critical region [the light green region in Fig. 2(a)], which includes part of the first-order line in the vicinity of the critical end point. Similarly, this model can be used to describe the first-order metal-insulator transition in VO_2, with a critical end point whose influence extends along the first-order line [8]. The physical structure of domains is power law throughout this critical region, when viewed on lengthscales shorter than the correlation length, which diverges as the critical point is approached. In mapping this order-parameter model to the temperature-driven metal-insulator transition in VO_2, we are making the ansatz that a sweep of temperature in the experiment maps to a combination of temperature and field sweep in the model, as in our prior work [8] and Ref. [43].

We simulate configurations near criticality (see Table I), since that is where this Hamiltonian can cause structure over multiple lengthscales. We use Monte Carlo simulations to generate typical examples of multiscale morphologies of insulator and metal domains that can arise from the clean Ising Hamiltonians of Eq. (1). Intricate domain configurations arise near the critical points of this model. Figures 1(a)–1(d) show some configurations near T_c^{2D} on a 100×100 lattice, with periodic boundary conditions. Figures 1(e)–1(h) show some representative configurations near T_c^{3D} on a $100 \times 100 \times 100$ lattice. Further simulation details are in the SM.

The correlation length of a system diverges at criticality, $\xi \propto 1/|T - T_c|^{\nu}$. When viewed on lengthscales $x < \xi$, the system exhibits critical fluctuations, i.e., fluctuations on all lengthscales between the correlation length ξ and the short distance cutoff, which for the lattice models we study is the lattice spacing, and in the real physical system it is the size of a unit cell. Close enough to criticality, this lengthscale will exceed any finite field of view (FOV). Therefore, when observed on a finite FOV (experimentally, or in simulation), there is a finite range of parameters over which the system displays critical pattern formation. For this reason, the entire range of parameters listed in the first part of Table I should be viewed as critical for the FOVs considered in this paper.

In addition to an interaction energy between domains, material disorder also affects the types of shapes that metal and insulator domains take. Because material disorder may make certain regions of the sample more favorable to an insulator, and certain others more favorable to metal, we use a random-field Ising model (RFIM) to simulate the effects of material disorder on the metal and insulator textures [44]:

$$H = -J \sum_{\langle ij \rangle} \sigma_i \sigma_j - \sum_i (h_i + h)\sigma_i. \tag{2}$$

The first term is the clean Ising model of Eq. (1). In the second term, the uniform field h and the local random fields h_i couple directly with the local order parameter. The random fields are chosen from a Gaussian distribution of width R where the probability of h_i is $P(h_i) = \exp(-h_i^2/(2R^2))/\sqrt{2\pi R^2}$. In the physical system, VO_2 changes from insulator to metal as the temperature is changed. Within the model, this physics presents itself as a combination of model temperature *and uniform field h* [8].

The ordered phase corresponds to all metal or all insulator, and the transition is second order when approached as a function of temperature or disorder strength at zero applied field [see the red dotted lines in Fig. 2(b)]. When instead the field is swept across the ordered region, the transition is a first-order change from metal to insulator [see the blue dotted line in Fig. 2(b)]. When temperature and the disorder strength are both nonzero, the behavior of the model in the vicinity of the phase transition is dominated by the random field [45]. That is, the random field is relevant but the temperature is irrelevant in the renormalization-group sense in a broad range around the solid green line in Fig. 2(b). We therefore model the patterns of metal and insulator domains that are possible with this Hamiltonian by generating domain configurations at zero temperature while sweeping the uniform field h. At zero temperature, this model undergoes an equilibrium phase transition at a random field strength of $R_c \approx 2.27J$ in an infinite size three-dimensional system (RF-3D) [46]. In two dimensions (RF-2D), the critical disorder strength is $R_c \to 0$ [47] in the infinite size limit, although in a finite-size system or with finite FOV, $R_c(L) > 0$. For the FOV we consider, $R_c \approx J$ for RF-2D.

When the random field model is near criticality, as the uniform field is swept from low to high or high to low, intricate patterns develop over multiple lengthscales near the

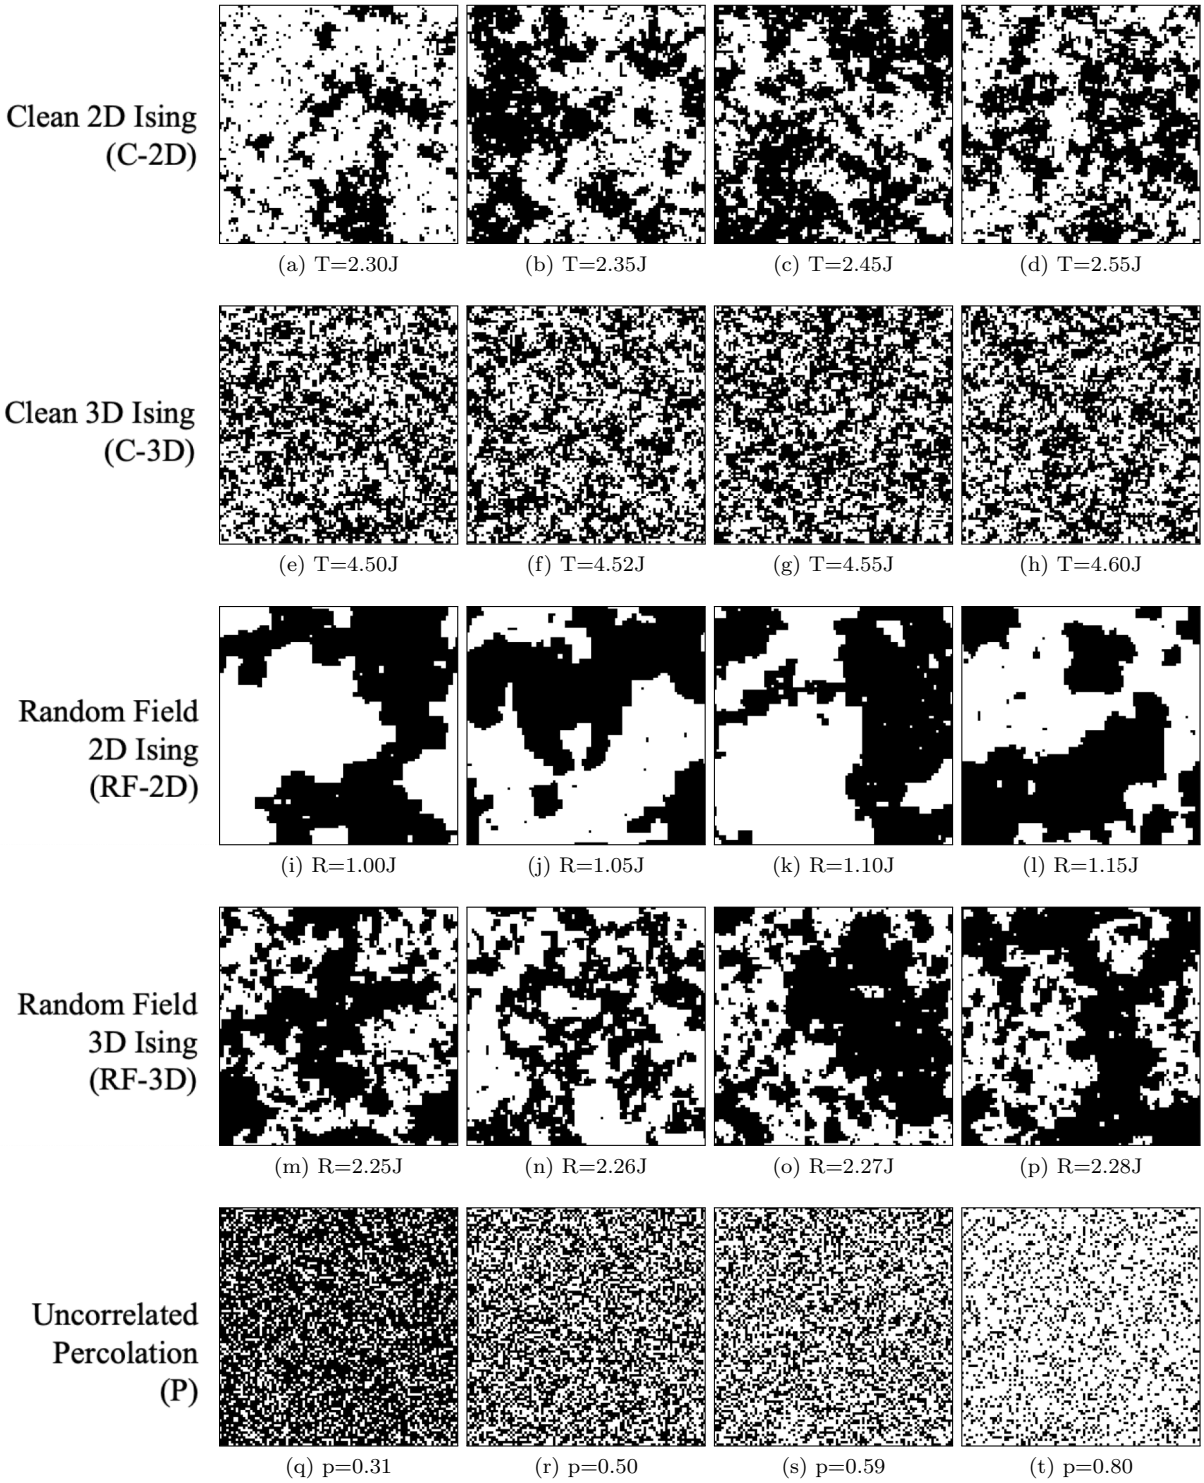

FIG. 1. Typical critical configurations generated from simulations of clean and random field Ising models and percolation models.

coercive field strength, where the metal/insulator domain fraction changes most rapidly with respect to uniform field h. Figures 1(i)–1(l) show representative configurations of RF-2D for a 100×100 lattice. Figures 1(m)–1(p) show representative configurations on the surface of a $100 \times 100 \times 100$ lattice near the 3D critical disorder strength, R_c^{3D}.

There is also the possibility that in fact domains are *not* interacting with each other as in the above Hamiltonians, but rather each domain acts independently. In the corresponding uncorrelated percolation model, a site is labeled "metallic" with a probability p; otherwise it is labeled "insulating." When $p \neq 0.5$, this is like flipping a biased coin, where p is the probability of turning up heads. This model also has a second-order phase transition as a function of p, and it displays structure across multiple lengthscales near its critical point. The critical percolation strength p_c is marked

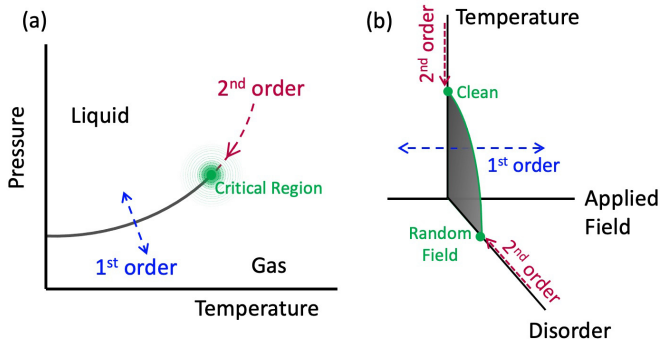

FIG. 2. (a) Generic liquid-gas phase diagram above the triple point. For paths along the dotted blue line, the phase transition is first order. However, the transition is second order when approached along the dotted red line. The critical end point (solid green circle) exerts itself over a critical region (open green circles). (b) Phase diagram of the random-field Ising model. While the clean or random-field Ising model has a second-order phase transition as a function of temperature (the red dotted lines), the transition is first order when approached as a function of applied field, crossing through the ordered region (blue dotted line). Critical behavior is observed in the vicinity of the green critical temperature line, whose critical behavior is controlled by the random-field fixed point for any finite disorder strength. Because the random-field fixed point is a zero-temperature fixed point, the critical region is much broader than in the clean case.

by a percolating cluster spanning the entire system, meaning that it touches one side of the system, and also the opposite side. In a two-dimensional system on an infinite square lattice, this threshold occurs at $p_c^{2D} \approx 0.59$ and in a three-dimensional system on an infinite cubic lattice it occurs at $p_c^{3D} \approx 0.31$ [48]. Figure 1(s) shows a percolation configuration of size 100×100 at $p_c^{2D} = 0.59$. Figure 1(q) shows a percolation configuration $100 \times 100 \times 100$ at $p_c^{3D} = 0.31$.

To further train the DL model to distinguish configurations that are near criticality (such as those described above) from configurations that are not near criticality, we also generate training images on uncorrelated percolation away from any critical point. To avoid the multiscale, fractal textures associated with criticality, in this set of images we use the percolation model in the following ranges: $p = 0.02 - 0.2$; $0.48 - 0.52$; $0.8 - 0.98$. The first range produces images that are mostly black, the second range produces images that are "white noise" [such as Fig. 1(r)], and the third range produces images that are mostly all white [like those in Fig. 1(t)]. Table I summarizes the parameter ranges we use for generating simulated data for training and validation from each of the above Hamiltonians.

III. CUSTOMIZED DEEP LEARNING MODEL

The parameters from Table I are used to generate 8000 images for each model near its transition, with the exception that percolation away from the 2D and 3D critical percolation strengths accounts for 16 000 images, for a total of 64 000 training images of synthetic data. We describe in the following three subsections the three major components of our deep learning model: Sec. III A, data preparation via symmetry reduction; Sec. III B, a CNN with multiple layers; and Sec. III C, our method for judging the confidence of the classifier.

A. Data preparation: Symmetry reduction method

The entire phase space associated with typical configurations generated by the models above satisfies certain symmetries. For example, the clean Ising model [Eq. (1)] satisfies the Z_2 symmetry $\sigma_i \to -\sigma_i$. Similarly, the RFIM [Eq. (1)] is symmetric under the simultaneous operations $\sigma_i \to -\sigma_i$ with $h_i \to -h_i$. Likewise, the percolation model is symmetric under the simultaneous operations $\sigma_i \to -\sigma_i$ with $p \to 1 - p$. In addition, for the square domain configurations we use as training data, the statistical weight of typical configurations in phase space is symmetric under all of the operations of the dihedral group of the square, D_4. Such symmetries are often employed in ML via a technique called data augmentation, in which all of the distinct symmetry operations are applied to specific configurations in order to generate more configurations and thereby augment the training data. When a neural network is trained under this kind of augmented data set, the resulting trained neural network respects all of the symmetries of the underlying models that produced the training data, rather than suffering from accidental asymmetries that mimic the random nature by which the training data are produced. The number of distinct symmetry operations available in our case is that of $Z_2 \otimes D_4$, or $2 \times 8 = 16$. For the square-shaped images of domain patterns that we generate, using this method of data augmentation would increase the training set by a factor of 16.

Rather than employ data augmentation, we introduce an alternative method: *symmetry reduction*. We prepare the data by reducing the symmetry of each configuration as much as possible before feeding them into the neural network. This symmetry reduction is as effective as the data augmentation method, but significantly reduces the time needed to train the neural network. For symmetry reduction to be effective, it is essential that *all* data go through the symmetry reduction before being fed into the CNN (including training, validation, and any subsequent real-world data fed into the classifier).

Let us turn our attention to the $Z_2 \otimes D_4$ symmetry operations in effect. Our models [Eqs. (1) and (2), including the noninteracting percolation limit where $J \to 0$] map metal and insulator domains to Ising spins $\sigma = +1$ for metal, and $\sigma = -1$ for insulator. There are $2 \times 8 = 16$ symmetry operations that can be applied to these spin configurations while preserving the weights of the typical configurations in phase space. We perform the following symmetry operations to each configuration in order to prepare the data:

(i) Ising Z_2 symmetry $\sigma_i \to -\sigma_i$: If a domain configuration has majority spin down, we flip all spins to make it majority spin up.

(ii) Rotations by 0, $\pi/2$, π, $3\pi/2$: The configuration is rotated such that of the four quadrants, quadrant I has the most spins up.

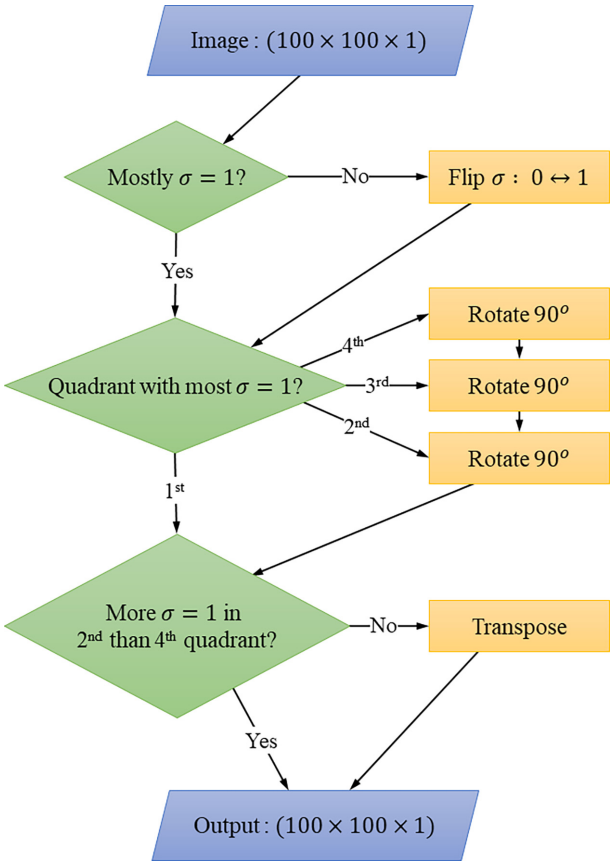

FIG. 3. Symmetry reduction method, as described in the text.

(iii) Transpose (reflection about the *xy* diagonal): If quadrant IV has more spins up than quadrant II, we transpose the configuration to ensure that quadrant II has more spins up than quadrant IV.

All the above operations are performed in the given order, and the logic is summarized in Fig. 3 [49].

B. Convolutional neural net architecture

We pass to the neural network single-channel binary images (i.e., strictly black and white, the same image space as QR codes) of size 100 × 100. The architecture of the CNN is as follows (Fig. 4): We use two sets of convolutional layers interleaved with max pooling layers. The first convolutional layer applies a suite of 32 filters of size 5 × 5 to the image, resulting in an image with 32 channels. (By way of comparison, an RGB image has three channels, so that each pixel is described by three numbers.) These 32 filters have a total of 32 × 5 × 5 parameters to be trained. The subsequent max pooling layer groups successive sets of 2 × 2 pixels, keeping only the largest value in each channel, thus reducing the image size to 50 × 50. The next convolutional layer applies a suite of 64 filters of size 5 × 5 × 32 to the 32-channel image, which was passed from the previous max pooling layer. These 64 filters have a total of 64 × 5 × 5 parameters to be trained. This is followed by another max pooling layer, reducing the image size to 25 × 25 pixels, now with 64 channels.

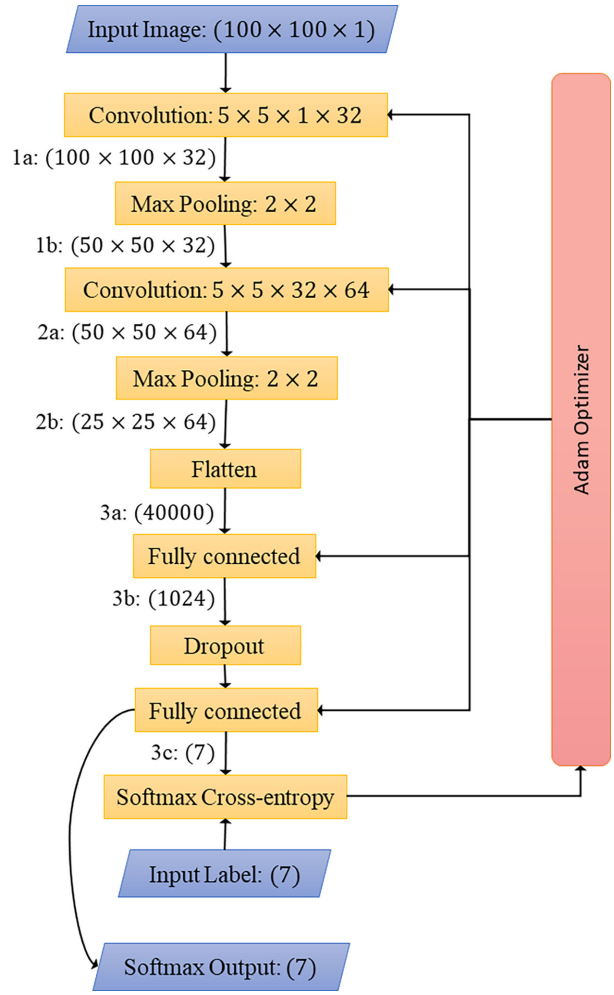

FIG. 4. Convolutional neural network. The input image here is reduced by the symmetry operations given in Fig. 3. The multidimensional output 2b is flattened into a one-dimensional array (3a) before it is fed into the fully connected layer. We use the Adam (adaptive moment estimation) optimization algorithm to train the network [50]. The output label is determined using softmax activation on the output layer.

All of this is followed by a fully connected layer, followed by a dropout of 50% of the connections, followed by a final fully connected layer, resulting in seven-dimensional output for classification (C-2D, C-3D, RF-2D, RF-3D, P-2D, P-3D, and P*). We use a softmax activation in the final output layer, which results in single label classification. If there are n output classes with numbers v_i, the softmax activation function is defined as

$$Y_i^{\text{softmax}} = \exp(v_i) \Big/ \sum_{j=1}^{n} \exp(v_j), \tag{3}$$

where Y_i^{softmax} is the output likelihood estimate.

After the symmetry reduction, 80% of the configurations are used for training; the remaining 20% are used for validation. The training set is used to train the network, whereas the validation set is used to predict the expected error upon generalization beyond the training set. Figure 5 shows how the errors evolve with the training epoch. The epoch at which

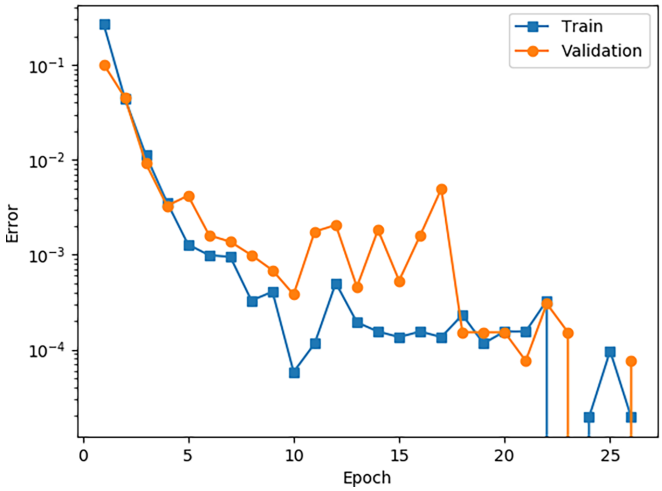

FIG. 5. Error in the training and validation set vs the number of epochs. Epochs correspond to the number of times the training set went through a training process. To prevent overfitting, we chose epoch = 4 for testing with experimental images. Training/validation accuracy = 99.64%/99.67%.

the classification errors in the validation phase deviate from the errors in the training phase roughly marks the onset of overfitting. Figure 5 shows that the classification errors are less than 0.5% at this point.

C. Figure of merit of classifications

Our goal is to use our ML model developed under supervised learning conditions to distinguish among hypotheses about datasets from real experiments. But before applying our trained ML model to images from experiment, it is important to understand that a trained classifier is only as good as its training set. Thus far, we have generated "simulated data" from various theoretical Hamiltonians, and we have trained an ML algorithm as to which sets of simulated data came from which underlying Hamiltonian. A major challenge in going from simulated data to real-world data is how to control for hypotheses that were not originally envisioned. For example, if an ML classifier has been trained to recognize the difference between cats and dogs, what answer will it give when shown a banana? A simple classifier will give a classification from its training set, but ideally the answer should not be "cat" or "dog," but rather "neither." Likewise, if our ML classifier is shown experimental data from a system whose underlying Hamiltonian is sufficiently different that none of our Hamiltonians used in the training process are a good description of the physical system, a simple classifier will still return *some* classification. Therefore, it is necessary to devise a method for flagging potentially dubious classifications. One method is adversarial training, i.e., to train the CNN on images that are not in the set of Hamiltonians comprising the hypothesis. Once again, this is limited by human imagination. For example, how will one know when this process is sufficiently completed, and how can one control for unforeseen image types arising in experiment? It is better to design a neutral method for flagging suspicious classifications, one that is not limited by the adversarial training set.

Therefore, we seek to devise a completely different method for identifying potentially dubious classifications. To do this, we turn our attention to the distribution of values observed right after the last fully connected layer in Fig. 4. Figure 6 shows what the distribution of values looks like at this step, over the entire training set. Since this distribution is well clustered for the seven models of interest, a prediction point lying far from its corresponding cluster should be scrutinized rather than blindly accepted.

For each class, the distribution at the end of the last fully connected layer (see Fig. 6) is generated from the training examples. We form the seven-dimensional standard deviation vector of these clusters about their centers of mass. We subsequently *flag as suspicious* any output in this layer that is a distance in this space of more than one standard deviation vector from all points in the cluster. Setting the cutoff at smaller distances rejects too many correct predictions in the validation set. A generalization of this method would be to use any or all of the intermediate layers for detecting such an anomaly in the input data; see Ref. [53].

IV. APPLICATION TO EXPERIMENTAL IMAGES ON VO_2

A. Testing the CNN on SNIM images of a thin film of VO_2

We next turn our attention to testing the trained CNN on an experimentally derived dataset for which the Hamiltonian underlying the experimentally observed pattern formation is already known, before applying the CNN to a new experimental dataset for which the answer is not previously known. In this section, we consider experimental data taken via SNIM on a thin film of VO_2. VO_2 undergoes a metal-insulator transition just above room temperature, in which the resistivity changes by over five orders of magnitude [54]. Rather than transitioning all at once, we previously showed that there is a finite regime of phase coexistence in which the metal and insulator puddles show significant pattern formation [3]. In fact, the spatial correlations reveal structure on all length-scales measured via SNIM, from the pixel size (20 nm) all the way out to the field of view (4 μm) [8]. The physics driving the pattern formation in this sample is already known via the cluster analysis techniques we recently developed [13,55,56]. By applying these techniques to analyze the metal and insulator puddles in this thin film of VO_2, we showed that the multiscale domains are of a fractal nature, with quantitative geometric characteristics including avalanche statistics matching those of the RF-2D [8].

Figure 7 shows the application of the CNN to experimental data on a thin film of VO_2 as it undergoes the metal-insulator transition. The data were obtained using SNIM, and first reported in Ref. [3]. SNIM measurements return an intensity a as a continuous variable at each pixel, resulting in single channel images. These SNIM images are of a size 256 × 256 px. The SNIM images are converted to black pixels and white pixels by assigning SNIM values of $a < 2.5$, which are insulating, to be white, and SNIM values of and $a > 2.5$, which are metallic, to be black, as discussed in Ref. [8]. These thresholded images are shown in the top row of Fig. 7. Reference [8] showed that the geometric characteristics of the

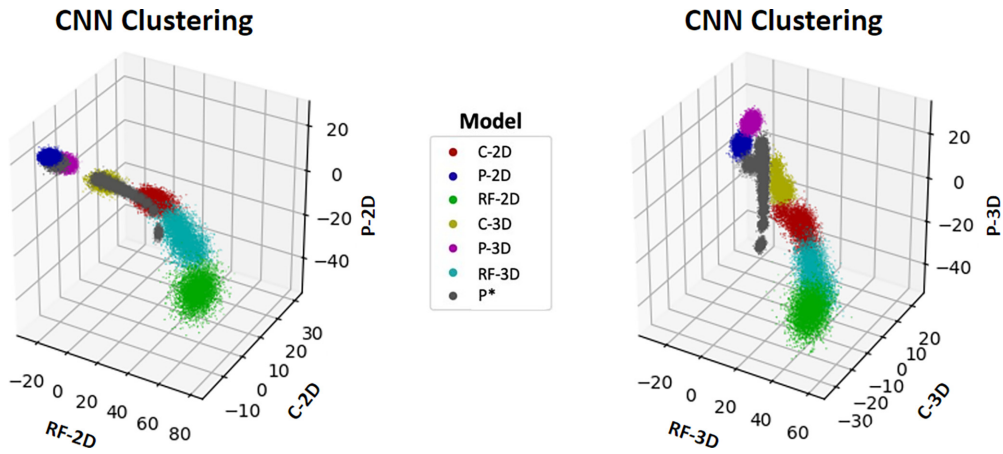

FIG. 6. Distribution of values of the output nodes for each class in the last fully connected layer, for all of the training sets. These clusters inhabit a seven-dimensional "model" space. The two representations in the figure are projections of the same seven-dimensional information onto two different three-dimensional subspaces. An interactive 3D visualization of the full 7D file is available in the Supplemental Material [51] and Ref. [52].

pattern formation are insensitive to changes in the threshold a within about 15% of this threshold value.

As shown in Fig. 8, the CNN takes in the experimentally derived images 100×100 px at a time, in each instance returning a classification indicating which Hamiltonian likely produced the pattern formation. To make full use of the spatial structure in the image, we use a sliding window of size 100×100 px, resulting in $(256 - 100 + 1) \times (256 - 100 + 1) = 157 \times 157$ classifications for each image.

In Fig. 9, we show the distribution of values in the last fully connected layer of the CNN, over the set of sliding windows. The colored dots corresponding to the training set are the same as those shown in Fig. 6. Results of CNN applied to the experimentally derived SNIM images of VO_2 that are within one seven-dimensional standard deviation of a training set are indicated by orange dots. Results of the SNIM data that are farther away are indicated by black dots, as described in Sec. III C.

Figure 7 shows the final results of the classifier applied to the SNIM data. Below each SNIM image (top row), the bar chart indicates the percentage of sliding windows that give a particular classification. Bright bars and numbers in

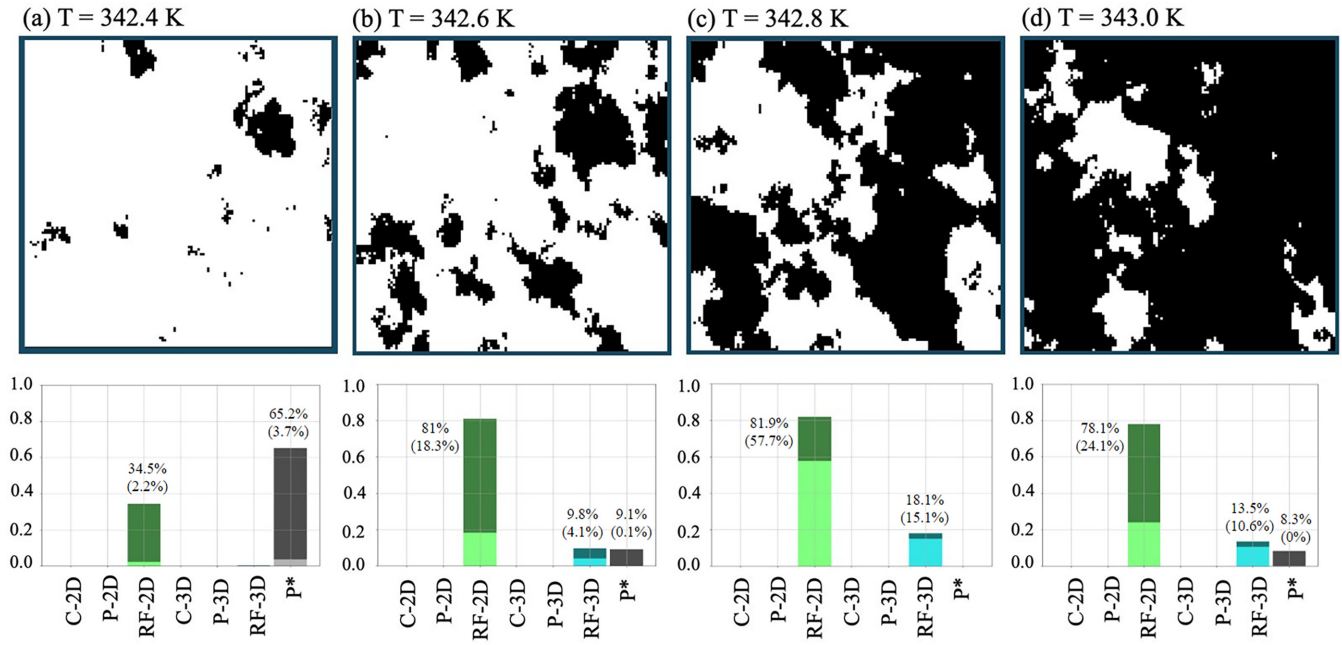

FIG. 7. Classification results of our deep learning model applied to SNIM images on a thin film of VO_2 as described in the text. The top row shows the thresholded data as described in the text. The field of view is 4 μm × 4 μm. White patches are insulating; black patches are metallic. The total percentage of classifications for a particular model are reported in the bar charts of panels (a)–(d). Classification percentages that fall within 1σ of a cluster in the training set are indicated in parentheses. Classifications that fall more than 1σ away from the edge of the corresponding cluster in the training set are colored darker in the bar chart.

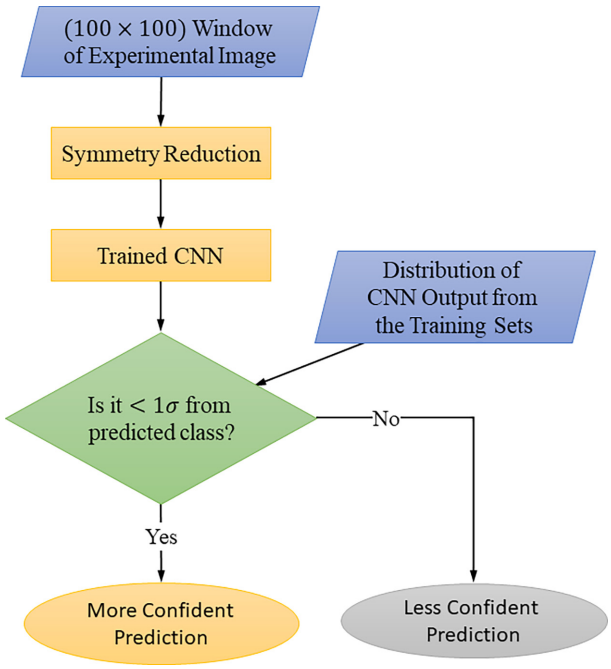

FIG. 8. End-to-end classification flowchart with CNN. We use a sliding window of size 100 × 100 px over the experimentally derived images, and slide the window 1 px at a time in each direction. For each 100 × 100 px image, we first apply symmetry reduction before feeding the image into the trained CNN. We compare the result of the CNN classifier to the entire distribution of CNN output from our training set. If the result is within one seven-dimensional standard deviation of one of the training images, the prediction is considered more confident than if it is farther away.

parentheses correspond to classifications that are within one seven-dimensional standard deviation of the training sets. The darker part of the bar, and the numbers not in parentheses, refer to the total percentage of sliding windows which give the corresponding classification. Thus the overall result of the ML classifier on an experimentally derived dataset is that it agrees with the classification from cluster techniques, with at least 83% confidence.

Notice that in Fig. 9, the distribution of values of the output nodes in the last fully connected layer for the CNN applied to the SNIM data is always close to the RF-2D model. Moreover, the entire set of points moves toward the training set distribution and then away from it, as a function of temperature. The temperature of closest approach is $T = 342.8$ K. This same phenomenon is borne out in the bar charts of Fig. 7, where the height of the bright green bar also peaks at $T = 342.8$ K. This is highly reminiscent of critical behavior, which grows in strength as the system approaches criticality, and diminishes as the system moves away from criticality. We propose that the distance of the center of mass of the SNIM cluster from the training clusters can be used as a measure of proximity to the critical point. Further study is needed to test this idea.

B. Applying the CNN to new optical microscope images of a VO_2 film

The top panels in Fig. 10 show metal and insulator domains in a thin film of VO_2 made at UCSD, taken using a home-built optical microscopy system capable of remaining in focus while temperature is cycled through the full metal-insulator transition. (See the Supplemental Material [51] and Ref. [63] for full details of the sample preparation and experimental setup.) The optical data are taken at a series of temperatures going through the metal-insulator transition. The physical dimensions of the square image sizes in Fig. 10 are all 28 μm × 28 μm, and the pixel size is 50 nm [64]. Both the FOV and the pixel size are larger than those of the SNIM images in Fig. 7.

We apply the same sliding window technique as with the SNIM data to analyze pieces of each image, 100 × 100 px at a time, in each instance returning a classification indicating which Hamiltonian likely produced the pattern formation. Because the optical images in Fig. 10 are 760 × 760 pixels, this results in $(760 - 100 + 1) \times (760 - 100 + 1) = 661^2$ classifications for each image.

The bottom panels in Fig. 10 show the final results of the CNN classifier applied to the optical microscope data. Below each optical microscope image, the bar chart indicates the percentage of sliding windows which give a particular classification. In this case, the images from temperatures $T =$ 339–343 K are each identified as RF-2D with a maximum greater than 89% confidence. In Fig. 11we show the distribution of values in the last fully connected layer of the CNN, over the set of sliding windows. The small circles correspond to the training set, and they are the same as those shown in Fig. 6. We discuss the implications of this identification in Sec. V.

From a theoretical point of view, we do not expect *every* image acquired from the experiments to have significant pattern formation. For example, once the image saturates to metal or insulator, there is no pattern formation left, and consequently there is much less information available in these datasets about the underlying model. Rather, we expect the images to display criticality which reaches peak prominence at a particular temperature. The typical method to discern proximity to criticality is through correlation lengths. The correlation length is expected to blow up as a power law, $\xi \propto 1/|T - T_c|^{\nu}$ in the vicinity of the critical temperature. However, the maximum correlation length our CNN can discern is cut off by the maximum FOV that the CNN is fed from the experimental data. Furthermore, the CNN analysis does not return a lengthscale. Instead, we observe once again the interesting behavior that the proximity of the experimentally derived data's cluster of output values in the last fully connected layer approaches and then retreats from the cluster of output values in the training sets as a function of temperature, as evidenced by the nonmonotonic behavior of the height of the bright green bars with temperature in Fig. 12. This is in line with our previous conjecture that the average distance of the cluster from that of the training set can be used as a measure of proximity to criticality.

V. DISCUSSION

For both experimental datasets, whether from SNIM or from optical microscopy, the deep learning CNN determined that the intricate pattern of metal and insulator patches was being set by the physics of the RF-2D. For the SNIM data,

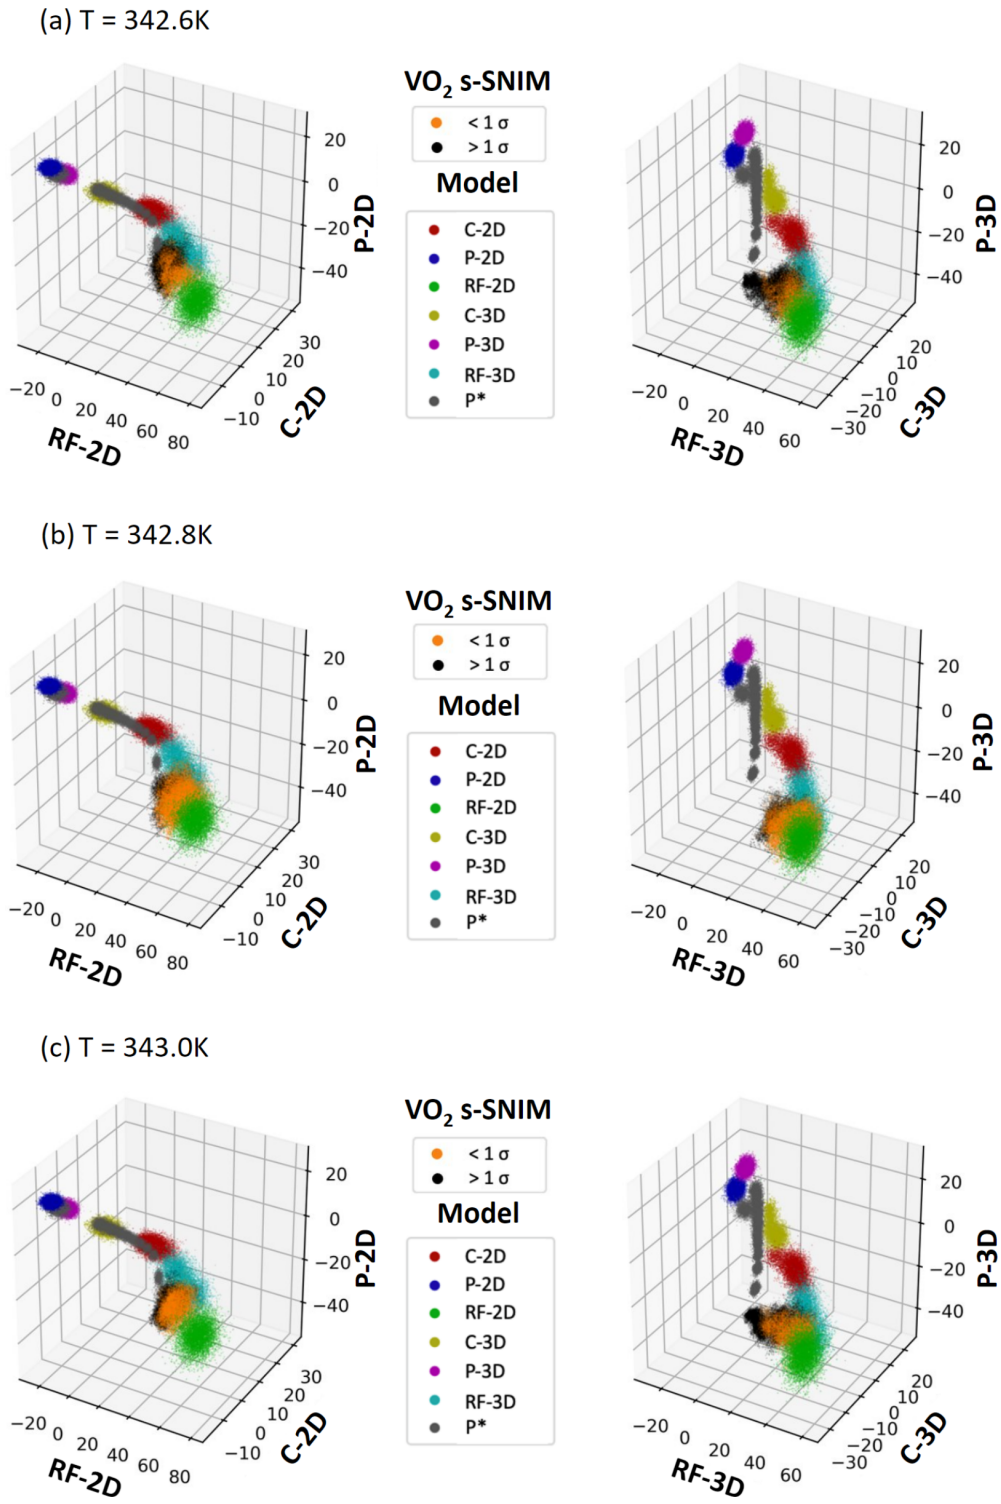

FIG. 9. Distribution of values of the output nodes for each class in the last fully connected layer, for the VO_2 SNIM data, superimposed on the distribution for the training sets shown in Fig. 6. Results for the VO_2 data that are within one seven-dimensional standard deviation of a training set are indicated by orange dots. Results for the VO_2 data that are farther away are indicated by black dots. An interactive 3D visualization of the full 7D file is available in the Supplemental Material [51] and Ref. [52].

this matches our prior identification using cluster methods [8]. For the microscope data, it was already known prior to application of the CNN that the physics driving the pattern formation should be arising from a 2D Hamiltonian. This is because the thickness of the film (≈300 nm) is comparable to the lateral resolution of the instrument. Consequently, the spatial correlations being measured are firmly in the two-dimensional limit. However, the fact that the CNN returned a two-dimensional model and not a three-dimensional model gives us further confidence in the CNN method.

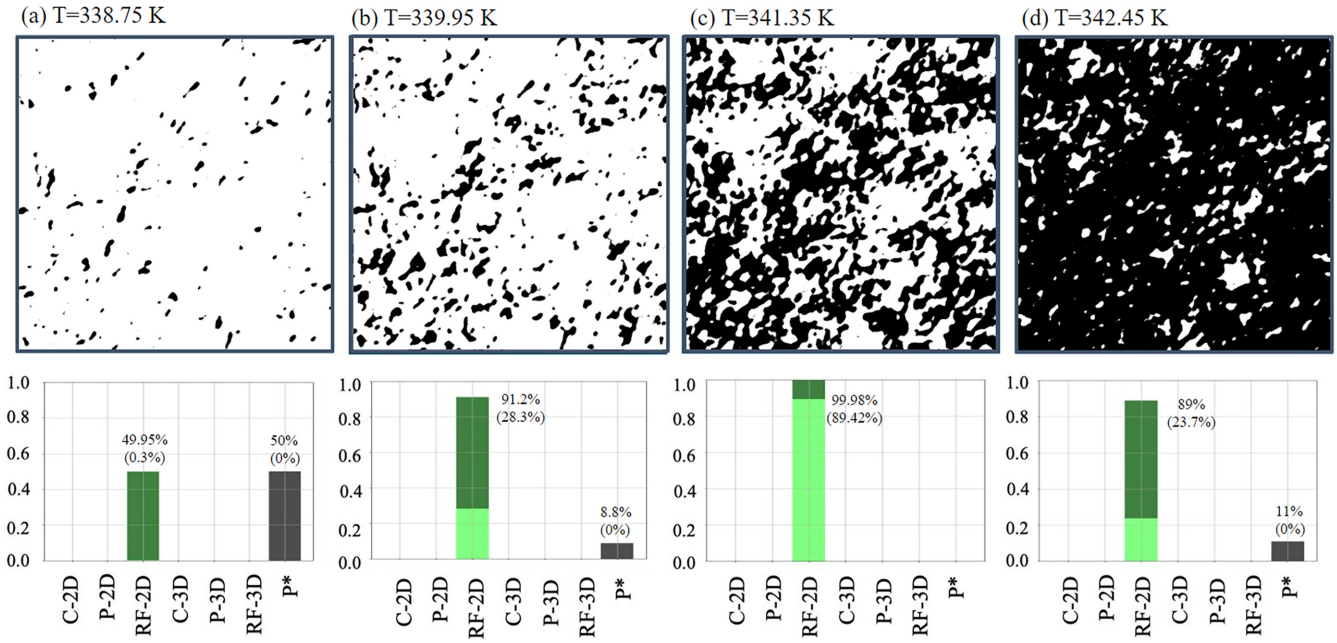

FIG. 10. Classification results of our deep learning model applied to new 28 μm × 28 μm optical microscopy images of a VO_2 thin film as described in the text. White patches are insulating; black patches are metallic. The total percentage of classifications for a particular model is reported in the bar charts of panels (a)–(e). Classification percentages that fall within 1σ of a cluster in the training set are indicated in parentheses. Classifications that fall more than 1σ away from the edge of the corresponding cluster in the training set are colored darker in the bar chart. All CNN predictions from optical data during a full temperature ramp up are presented in Fig. 12 of the Supplemental Material [51] (see also Refs. [57–62] therein).

The identification of the Hamiltonian as RF-2D means that a combination of material disorder and interactions between spatially proximate regions of the sample drives the pattern formation. The fact that interactions must be present rules out a Preisach model [65] of independent hysteretic switchers, as we previously argued based on first-order reversal curve measurements [66] and a cluster analysis of the critical exponents during the transition [8]. The multiscale nature of the pattern formation is driven by proximity to criticality, which can happen even in a first-order phase transition, near a critical end point [8,67]. Consistent with proximity to criticality, we have previously shown that there is significant slowing down of the relaxation time near the phase transition [68,69].

Random-field critical points exhibit extreme critical slowing down: because the barriers to equilibration grow as a power law as the system nears criticality, the characteristic relaxation time grows exponentially as the system approaches criticality [45]. Because of this extreme critical slowing down, the model is notorious for highly nonequilibrium behavior, including hysteresis, glassiness, coarsening, and aging. In addition, the model has an anomalously large region of critical behavior: a system that is 85% away from the critical point can still display 2 decades of scaling [47]. This means that it is fairly easy within this model to get into a regime that displays pattern formation across multiple lengthscales, including fractal textures.

With the model controlling this pattern formation now well established from this study and from our previous work [8], we can make the following statements about VO_2: Increasing disorder is expected to broaden hysteresis curves, and also decrease the slope of the hysteresis curve at its inflection point [70]. Indeed, these expectations are borne out in recent ion irradiation studies of resistivity in VO_2 [71]. In addition, due to the pronounced memory effects with exponentially long equilibration times, exact identification of material properties can be history-dependent, leading to the appearance of nonrepeatability. On the other hand, disorder can ultimately be exploited as another means of control [9,72].

The ML method is complementary to the aforementioned cluster techniques. Whereas the cluster techniques require at least two decades of scaling in the dataset, we have shown here and in Ref. [36] that an ML classifier can make determinations on datasets with smaller FOV. And while the cluster techniques are designed to extract information from datasets in systems that are in the vicinity of a critical point, we expect that the ML methods developed here can be useful farther away from criticality, because they are able to make determinations on smaller FOV, i.e., they do not require that the system have the long correlation length associated with proximity to criticality.

In the same way that the critical exponents are encoded in the shapes and statistics of the fractal electronic textures that arise near a critical point [8,13,55], our ML study reveals that the universal features of the model itself are encoded in the spatial correlations of the textures, without needing the intermediate step of identifying critical exponents. Criticality presents itself even at the moderate (*i.e., not long*) lengthscales our CNN views, which is set by the size of the sliding window we employ on the datasets to be classified.

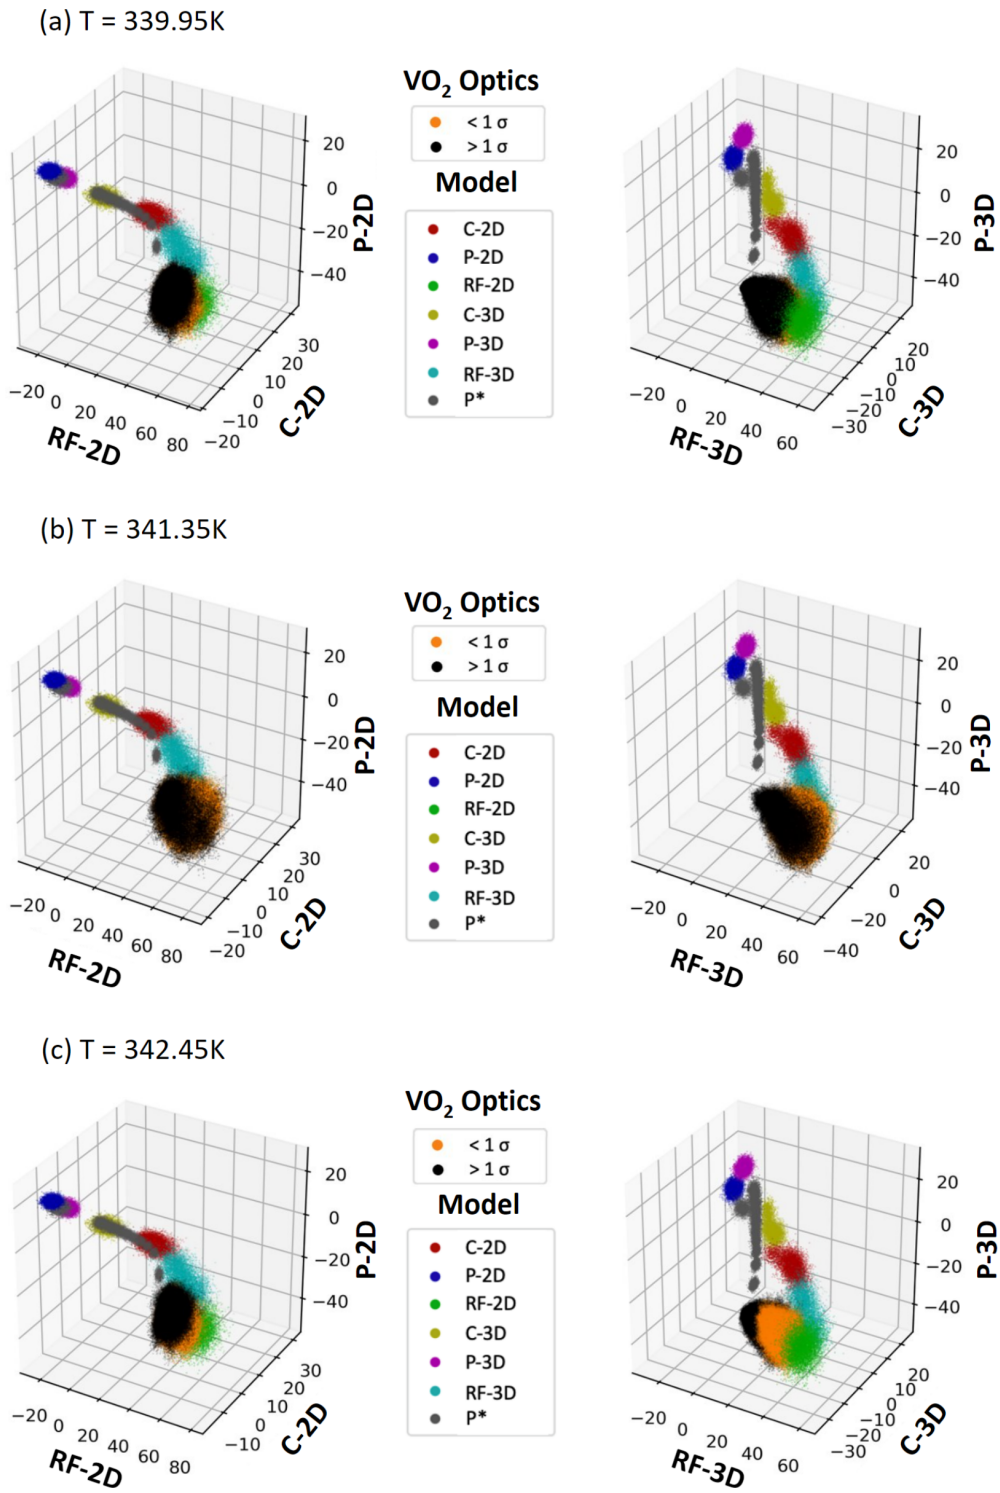

FIG. 11. Distribution of relative weights of each class in the last fully connected layer, for the VO_2 optical data (square sample presented in Fig. 10), superimposed on the distribution for the training sets shown in Fig. 6. Results for the VO_2 data that are within one seven-dimensional standard deviation of a training set are indicated by orange dots. Results for the VO_2 data that are farther away are indicated by black dots. An interactive 3D visualization of the full 7D file is available in the Supplemental Material [51] and Ref. [52].

The method is also potentially extendible to handle nondiscrete order parameters, such as continuum models, which present a challenge for cluster methods. For example, it may be possible to use a similar framework to diagnose pattern formation that reveals an underlying XY model or Heisenberg model. In addition, by using regression, we expect to be able to go beyond criticality to begin to determine the values of parameters in the Hamiltonian.

We have developed this ML method first on critical systems, which have no characteristic lengthscale due to the power-law structure, and therefore display spatial structure on every lengthscale within a correlation length.

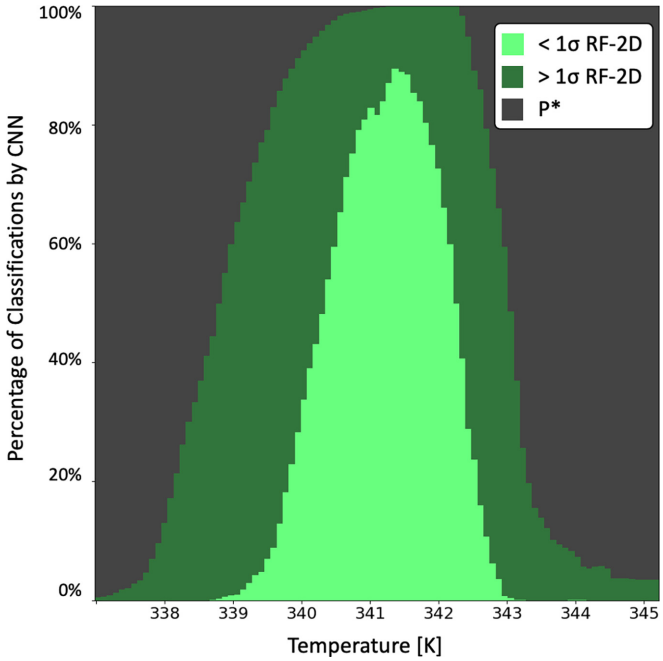

FIG. 12. All CNN predictions from optical data during a temperature ramp up of data presented in Fig. 10. Darker colors denote classifications that are more than 1 standard deviation from the identified training set.

However, we expect this general scheme to also be broadly applicable to systems that have an emergent lengthscale, such as frustrated phase-separation systems in general, such as block copolymers, the mixed phase of type I superconductors, reaction-diffusion systems, and convection rolls [73,74].

Dagotto [9] points out that quenched disorder plays an important role in many strongly correlated materials, and based on this, he argues that for such materials, "it is not sufficient to consider phase diagrams involving only temperature and hole-doping x. A disorder strength axis should be incorporated into the phase diagram of these materials as well." Models incorporating disorder predict that nonequilibrium behavior including glassiness (multiple nearby local energy minima) and hysteresis are prominent features when electronic phase separation occurs in the presence of quenched disorder [8,13,14,16,55,70,72,75–77]. The methods we have employed here, which identify the terms in the Hamiltonian, when extended to include a regression analysis to identify the values of the parameters in those terms, have the potential to identify the disorder strength. Mapping out this disorder strength axis in strongly correlated phase diagrams has the potential to help disentangle some of the ambiguities and apparent inconsistencies heretofore reported in the literature of these systems [78–81].

Future work on this type of classifier will also benefit from (i) generalizing the CNN to handle input images of any size, (ii) developing a learning-based optimization for the rejection classifier, and (iii) handling grayscale images without the need to threshold them.

VI. CONCLUSION

In conclusion, we have extended machine learning methods to be able to identify the Hamiltonian driving pattern formation in complex electronic matter. We have shown that the accuracy that can be achieved by using a CNN to classify synthetic data is better than 99%, and about 83–89 % accurate on experimental data. We introduce a symmetry reduction method, which significantly lowers the training time over data reduction without reducing accuracy. In addition, we introduce a distribution-based method for quantifying confidence of multilabel classifier predictions, without the problems associated with introducing adversarial training sets. We also propose a machine learning based criterion for diagnosing proximity to criticality.

We have also demonstrated that this framework can be successfully applied to real experimental images by using it to classify the Hamiltonian of SNIM data on a thin film of VO_2, for which the answer was already known from a complementary theoretical method. Having thus vetted our ML model, we applied it to optical microscope data on a different sample of VO_2. In each case, we find that the pattern formation of metal-insulator domains in thin films of VO_2 is driven by proximity to a critical point of the two-dimensional random field Ising model. Further tests of this model include hysteresis protocols in the presence of a series of engineered disorder strengths.

ACKNOWLEDGMENTS

We thank M. J. Carlson for technical assistance with image stabilization, and we acknowledge helpful conversations with A. El Gamal and K. A. Dahmen. The work at ESPCI (M.A.B., L.A., and A.Z.) was supported by Cofund AI4theSciences hosted by PSL Université, through the European Union's Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie Grant No. 945304. The work at UCSD (P.S. and I.K.S.) was supported by the AFOSR Award No. FA9550-20-1-0242. M.M.Q. acknowledges support from the National Science Foundation (NSF) via Grant No. IIP-1827536. D.N.B acknowledges support by the Center on Precision-Assembled Quantum Materials, funded through the US National Science Foundation (NSF) Materials Research Science and Engineering Centers (Award No. DMR-2011738). S.B., F.S., and E.W.C. acknowledge support from NSF Grant No. DMR-2006192, NSF XSEDE Grants No. TG-DMR-180098 and No. DMR-190014, and the Research Corporation for Science Advancement Cottrell SEED Award. F.S. acknowledges support from the COVID-19 Research Disruption Fund at Purdue through the U.S. Department of Education HEERF III (ARP) Award No. P425F204928. S.B. acknowledges support from a Bilsland Dissertation Fellowship. E.W.C. acknowledges support from a Fulbright Fellowship and from DOE BES Award No. DE-SC0022277. L.B. acknowledges support from a Summer Undergraduate Research Fellowship at Purdue. This research was supported in part through computational resources provided by Research Computing at Purdue, West Lafayette, Indiana [82].

[1] D. N. Basov, R. D. Averitt, D. van der Marel, M. Dressel, and K. Haule, Rev. Mod. Phys. **83**, 471 (2011).

[2] G. Binnig and H. Rohrer, Rev. Mod. Phys. **71**, S324 (1999).

[3] M. M. Qazilbash, M. Brehm, B.-G. Chae, P.-C. Ho, G. O. Andreev, B.-J. Kim, S. J. Yun, A. V. Balatsky, M. B. Maple, F. Keilmann, H.-T. Kim, and D. N. Basov, Science **318**, 1750 (2007).

[4] Y. Kohsaka, C. Taylor, K. Fujita, A. Schmidt, C. Lupien, T. Hanaguri, M. Azuma, M. Takano, H. Eisaki, H. Takagi, S. Uchida, and J. C. Davis, Science **315**, 1380 (2007).

[5] S. Guénon, S. Scharinger, S. Wang, J. G. Ramírez, D. Koelle, R. Kleiner, and I. K. Schuller, Europhys. Lett. **101**, 57003 (2013).

[6] M. Lange, S. Guénon, Y. Kalcheim, T. Luibrand, N. M. Vargas, D. Schwebius, R. Kleiner, I. K. Schuller, and D. Koelle, Phys. Rev. Appl. **16**, 054027 (2021).

[7] B. Phillabaum, *Using Anisotropy as a Probe for Nematic Order in the Cuprates* (Purdue University Press, West Lafayette, IN, 2012).

[8] S. Liu, B. Phillabaum, E. W. Carlson, K. A. Dahmen, N. S. Vidhyadhiraja, M. M. Qazilbash, and D. N. Basov, Phys. Rev. Lett. **116**, 036401 (2016).

[9] E. Dagotto, Science **309**, 257 (2005).

[10] A. Moreo, M. Mayr, A. Feiguin, S. Yunoki, and E. Dagotto, Phys. Rev. Lett. **84**, 5568 (2000).

[11] D. Cho, K. M. Bastiaans, D. Chatzopoulos, G. D. Gu, and M. P. Allan, Nature (London) **571**, 541 (2019).

[12] I. Battisti, K. M. Bastiaans, V. Fedoseev, A. de la Torre, N. Iliopoulos, A. Tamai, E. C. Hunter, R. S. Perry, J. Zaanen, F. Baumberger, and M. P. Allan, Nat. Phys. **13**, 21 (2016).

[13] B. Phillabaum, E. W. Carlson, and K. A. Dahmen, Nat. Commun. **3**, 915 (2012).

[14] J. Li, J. Pelliciari, C. Mazzoli, S. Catalano, F. Simmons, J. T. Sadowski, A. Levitan, M. Gibert, E. Carlson, J.-M. Triscone, S. Wilkins, and R. Comin, Nat. Commun. **10**, 4568 (2019).

[15] A. Sharoni, J. G. Ramírez, and I. K. Schuller, Phys. Rev. Lett. **101**, 026404 (2008).

[16] K. W. Post, A. S. McLeod, M. Hepting, M. Bluschke, Y. Wang, G. Cristiani, G. Logvenov, A. Charnukha, G. X. Ni, P. Radhakrishnan, M. Minola, A. Pasupathy, A. V. Boris, E. Benckiser, K. A. Dahmen, E. W. Carlson, B. Keimer, and D. N. Basov, Nat. Phys. **14**, 1056 (2018).

[17] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Rev. Mod. Phys. **91**, 045002 (2019).

[18] J. Carrasquilla, Adv. Phys.: X **5**, 1797528 (2020).

[19] E. Bedolla, L. C. Padierna, and R. Castañeda-Priego, J. Phys.: Condens. Matter **33**, 053001 (2021).

[20] P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, Phys. Rep. **810**, 1 (2019).

[21] P. Ronhovde, S. Chakrabarty, D. Hu, M. Sahu, K. K. Sahu, K. F. Kelton, N. A. Mauro, and Z. Nussinov, Eur. Phys. J. E **34**, 105 (2011).

[22] S. S. Schoenholz, E. D. Cubuk, D. M. Sussman, E. Kaxiras, and A. J. Liu, Nat. Phys. **12**, 469 (2016).

[23] L.-F. Arsenault, A. Lopez-Bezanilla, O. A. von Lilienfeld, and A. J. Millis, Phys. Rev. B **90**, 155136 (2014).

[24] G. Carleo and M. Troyer, Science **355**, 602 (2016).

[25] A. Lopez-Bezanilla and O. A. von Lilienfeld, Phys. Rev. B **89**, 235411 (2014).

[26] P. Mehta and D. J. Schwab, arXiv:1410.3831.

[27] C. Beny, arXiv:1301.3124.

[28] J. Schmidt, M. R. G. Marques, S. Botti, and M. A. L. Marques, npj Comput. Mater. **5**, 83 (2019).

[29] L. M. Ghiringhelli, J. Vybiral, S. V. Levchenko, C. Draxl, and M. Scheffler, Phys. Rev. Lett. **114**, 105503 (2015).

[30] Y. Zhang, A. Mesaros, K. Fujita, S. D. Edkins, M. H. Hamidian, K. Chng, H. Eisaki, S. Uchida, J. C. S. Davis, E. Khatami, and E.-A. Kim, Nature (London) **570**, 484 (2019).

[31] Y. Zhang and X. Xu, Heliyon **6**, e05055 (2020).

[32] S. Xu, A. S. McLeod, X. Chen, D. J. Rizzo, B. S. Jessen, Z. Yao, Z. Wang, Z. Sun, S. Shabani, A. N. Pasupathy, A. J. Millis, C. R. Dean, J. C. Hone, M. Liu, and D. N. Basov, ACS Nano **15**, 18182 (2021).

[33] X. Chen, Z. Yao, S. Xu, A. S. McLeod, S. N. G. Corder, Y. Zhao, M. Tsuneto, H. A. Bechtel, M. C. Martin, G. L. Carr, M. M. Fogler, S. G. Stanciu, D. N. Basov, and M. Liu, ACS Photonics **8**, 2987 (2021).

[34] J. Carrasquilla and R. G. Melko, Nat. Phys. **13**, 431 (2017).

[35] L. Wang, Phys. Rev. B **94**, 195105 (2016).

[36] L. Burzawa, S. Liu, and E. W. Carlson, Phys. Rev. Mater. **3**, 033805 (2019).

[37] A. Sohn, T. Kanki, K. Sakai, H. Tanaka, and D.-W. Kim, Sci. Rep. **5**, 10417 (2015).

[38] N. J. McLaughlin, Y. Kalcheim, A. Suceava, H. Wang, I. K. Schuller, and C. R. Du, Adv. Quantum Technol. **4**, 2000142 (2021).

[39] S. Papanikolaou, R. M. Fernandes, E. Fradkin, P. W. Phillips, J. Schmalian, and R. Sknepnek, Phys. Rev. Lett. **100**, 026408 (2008).

[40] L. Onsager, Phys. Rev. **65**, 117 (1944).

[41] M. E. Fisher, Rev. Mod. Phys. **46**, 597 (1974).

[42] C. Kittel and H. Kroemer, *Thermal Physics* (Freeman, San Francisco, 1980).

[43] P. Limelette, A. Georges, D. Jerome, P. Wzietek, P. Metcalf, and J. M. Honig, Science **302**, 89 (2003).

[44] Disorder can also cause spatial variations in the coupling J. This random bond disorder is irrelevant in the renormalization-group sense when random field disorder is present.

[45] D. S. Fisher, Phys. Rev. Lett. **56**, 416 (1986).

[46] A. A. Middleton and D. S. Fisher, Phys. Rev. B **65**, 134411 (2002).

[47] O. Perković, K. Dahmen, and J. P. Sethna, Phys. Rev. Lett. **75**, 4528 (1995).

[48] Y. Deng and H. W. J. Blöte, Phys. Rev. E **72**, 016126 (2005).

[49] There are some configurations that cannot be mapped to a unique configuration using the symmetry reduction operations described in the text. For example, if exactly half the domains are metal, the first step cannot reduce it down to a unique configuration. But the average frequency with which that occurs is $\frac{^{N}C_{N/2}}{2^{N}} \approx 1.592 \times 10^{3008}/1.995 \times 10^{3010} \approx 0.8\%$, where $N = 100 \times 100$.

[50] D. P. Kingma and J. Ba, arXiv:1412.6980.

[51] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevB.107.205121 for interactive, 3D visualizations of the results of the ML classifier predictions in the 7D classifier space, for codes, for theoretical methods, and for experimental methods including sample preparation as well as data on a second sample measured by optical methods.

[52] S. Basak, M. A. Banguero, L. Burzawa, F. Simmons, P. Salev, L. Aigouy, M. M. Qazilbash, I. K. Schuller, D. N. Basov, A. Zimmers, and E. W. Carlson, Purdue University Research Repository (2023), doi: 10.4231/966Q-6F95.
[53] K. S. Sivamani, R. Sahay, and A. E. Gamal, IEEE Lett. Comput. Soc. **3**, 25 (2020).
[54] F. J. Morin, Phys. Rev. Lett. **3**, 34 (1959).
[55] E. W. Carlson, S. Liu, B. Phillabaum, and K. A. Dahmen, J. Supercond. Novel Magn. **28**, 1237 (2015).
[56] S. Liu, E. W. Carlson, and K. A. Dahmen, Condens. Matter **6**, 39 (2021).
[57] S. Liu, M. Liu, and Z. Yang, EURASIP J. Adv. Sign. Proc. **2016**, 70 (2016).
[58] S. Basak, Y. Sun, M. Alzate Banguero, P. Salev, I. K. Schuller, L. Aigouy, E. W. Carlson, and A. Zimmers, Adv. Electron. Mater. (to be published).
[59] J. Ziv and A. Lempel, IEEE Trans. Inf. Theor. **24**, 530 (1978).
[60] T. Welch, Computer **17**, 8 (1984).
[61] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).
[62] U. Wolff, Phys. Rev. Lett. **62**, 361 (1989).
[63] M. Alzate Banguero, S. Basak, N. Raymond, F. Simmons, P. Salev, I. K. Schuller, L. Aigouy, E. W. Carlson, and A. Zimmers, arXiv:2301.04220.
[64] The far-field optical resolution is 300–500 nm.
[65] L. A. L. de Almeida, G. S. Deep, A. M. N. Lima, H. F. Neff, and R. C. S. Freire, IEEE Trans. Instrum. Meas. **50**, 1030 (2001).
[66] J.-G. Ramírez, A. Sharoni, Y. Dubi, M. E. Gómez, and I. K. Schuller, Phys. Rev. B **79**, 235110 (2009).
[67] Interestingly, our identification of RF-2D means that in fact the first-order line has gone away, since the critical disorder strength in 2D is 0. However, hysteresis loops remain open because of the extreme critical slowing down of random-field Ising models, which generate exponentially large timescales for equilibration.
[68] S. Wang, J. G. Ramírez, J. Jeffet, S. Bar-Ad, D. Huppert, and I. K. Schuller, Europhys. Lett. **118**, 27005 (2017).
[69] J. del Valle, P. Salev, F. Tesler, N. M. Vargas, Y. Kalcheim, P. Wang, J. Trastoy, M.-H. Lee, G. Kassabian, J. G. Ramírez, M. J. Rozenberg, and I. K. Schuller, Nature (London) **569**, 388 (2019).
[70] J. P. Sethna, K. Dahmen, S. Kartha, J. A. Krumhansl, B. W. Roberts, and J. D. Shore, Phys. Rev. Lett. **70**, 3347 (1993).
[71] J. G. Ramirez, T. Saerbeck, S. Wang, J. Trastoy, M. Malnou, J. Lesueur, J.-P. Crocombette, J. E. Villegas, I. K. Schuller, J. Villegas, and I. Schuller, Phys. Rev. B **91**, 205123 (2015).
[72] E. W. Carlson and K. A. Dahmen, Nat. Commun. **2**, 379 (2011).
[73] M. Seul and D. Andelman, Science **267**, 476 (1995).
[74] V. J. Emery and S. A. Kivelson, Physica C **209**, 597 (1993).
[75] E. W. Carlson, K. A. Dahmen, E. Fradkin, and S. A. Kivelson, Phys. Rev. Lett. **96**, 097003 (2006).
[76] Y. L. Loh, E. W. Carlson, and K. A. Dahmen, Phys. Rev. B **81**, 224207 (2010).
[77] K. Dahmen and J. P. Sethna, Phys. Rev. B **53**, 14872 (1996).
[78] T. Timusk and B. Statt, Rep. Prog. Phys. **62**, 61 (1999).
[79] N. P. Armitage, P. Fournier, and R. L. Greene, Rev. Mod. Phys. **82**, 2421 (2010).
[80] B. Keimer, S. A. Kivelson, M. R. Norman, S. Uchida, and J. Zaanen, Nature (London) **518**, 179 (2015).
[81] I. Bozovic and J. Levy, Nat. Phys. **16**, 712 (2020).
[82] T. Hacker, B. Yang, and G. McCartney, Educause Review (2014), https://er.educause.edu/articles/2014/7/empowering-faculty-a-campus-cyberinfrastructure-strategy-for-research-communities.