

Classifying Snapshots of the Doped Hubbard Model with Machine Learning

Annabelle Bohrdt,^{1,2} Christie S. Chiu,² Geoffrey Ji,² Muqing Xu,²
Daniel Greif,² Markus Greiner,² Eugene Demler,² Fabian Grusdt,^{1,2} and Michael Knap¹

¹*Department of Physics and Institute for Advanced Study,
Technical University of Munich, 85748 Garching, Germany*

²*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*

(Dated: December 3, 2018)

Quantum gas microscopes for ultracold atoms can provide high-resolution real-space snapshots of complex many-body systems. We implement machine learning to analyze and classify such snapshots of ultracold atoms. Specifically, we compare the data from an experimental realization of the two-dimensional Fermi-Hubbard model to two theoretical approaches: a doped quantum spin liquid state of resonating valence bond type, and the geometric string theory, describing a state with hidden spin order. This approach considers all available information without a potential bias towards one particular theory by the choice of an observable and can therefore select the theory which is more predictive in general. Up to intermediate doping values, our algorithm tends to classify experimental snapshots as geometric-string-like, as compared to the doped spin liquid. Our results demonstrate the potential for machine learning in processing the wealth of data obtained through quantum gas microscopy for new physical insights.

Introduction.— The phase diagram of the Fermi Hubbard model and its connection to high-temperature superconductivity have been the subject of a vast amount of theoretical and experimental studies in the past decades [1, 2]. While a large number of theories exist, each with its own merits, a unifying analytic understanding is nonetheless still lacking. In the regime of low temperatures and finite doping even numerical simulations become increasingly difficult. In recent years, tremendous progress has been made in using ultracold atoms to study quantum magnetism in the Fermi-Hubbard model [3–11]. These ultracold atom experiments are now exploring finite doping regimes of the phase diagram where no consensus on a theoretical description and the most appropriate way to experimentally characterize the system exists.

All information about the quantum state of the system is contained in the many-body density matrix, where the number of degrees of freedom scales exponentially with the system size. A measurement collapses the quantum state, such that only a projection of it can be accessed. Repeated projective measurements provide a plethora of data, which in the past has mostly been analyzed to obtain conventional observables such as one- and two-point correlation functions, that are also traditionally measured in solids. However, measurements performed in quantum gas microscopes contain considerably more information. Therefore, the need arises for new methods to analyze the data that take all available information into consideration and hence use the capabilities of quantum gas microscopes to their full extent.

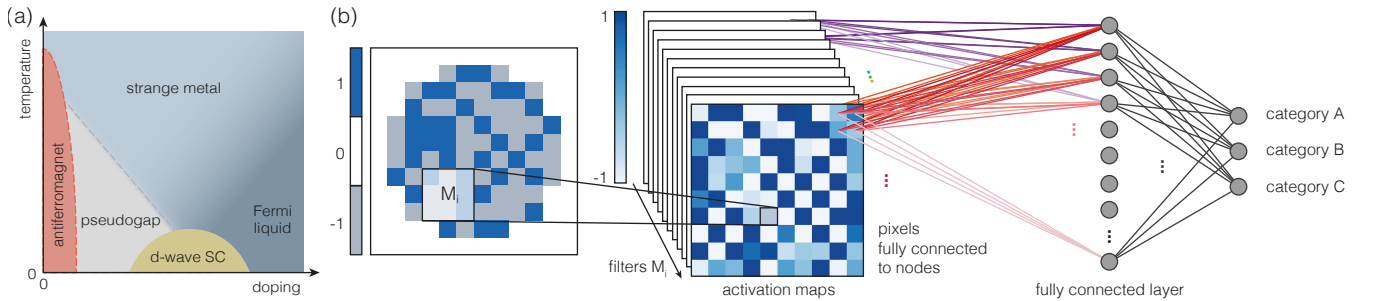


FIG. 1. **Classifying quantum gas microscope snapshots of the doped Fermi-Hubbard model with convolutional neural networks (CNNs).** (a) Schematic phase diagram of the 2D Fermi-Hubbard model. We use snapshots of the many-body quantum state at fixed doping and temperature as input data for the CNN. (b) The main building block of CNNs, which are conventionally used to analyze visual imagery, is the convolutional layer with a set of learnable filters M_i as parameters [12]. At each possible position of a given filter in the input image, the inner product between the filter and the input data is computed. This yields a two dimensional activation map of the filter. During training, the network learns to set the entries of the filters such that the corresponding value in the activation map is high when specific types of patterns are detected. The convolutional layer is followed by a fully connected layer, which then sorts the data into the different categories.

In this work, we employ machine learning techniques to characterize the large amount of data produced by quantum gas microscopy of the doped Fermi-Hubbard model. Recently, machine learning has emerged as a tool in condensed matter physics. Its main applications so far include representing the wave functions of correlated many-body states [13–15], the determination of a phase transition [16–24] and recently the analysis of scanning tunneling microscopy images [25]. Here, we take a complementary route and use a CNN to classify experimental data at finite filling into different theoretical categories in order to determine which theory describes the system best on the microscopic level, see Fig. 1. This approach provides insight into the underlying microscopic structures of the system which may be inaccessible to conventional observation but can be essential for gaining a deeper understanding of the emergent physics.

Physical system.— The experimental data that we analyze with our machine-learning algorithm has been measured with a quantum gas microscope for ultracold Lithium atoms in an optical lattice and is available in Ref. 11. This system is modeled by the Fermi-Hubbard Hamiltonian

$$\hat{H} = -t \sum_{\sigma=\uparrow,\downarrow} \sum_{\langle i,j \rangle} (\hat{c}_{i,\sigma}^\dagger \hat{c}_{j,\sigma} + \text{h.c.}) + U \sum_j \hat{c}_{j,\uparrow}^\dagger \hat{c}_{j,\uparrow} \hat{c}_{j,\downarrow}^\dagger \hat{c}_{j,\downarrow}$$

where the first term describes tunneling with amplitude t of spin-1/2 fermions between nearest-neighbor sites on a two-dimensional square lattice. The second term corresponds to on-site interactions of strength U between fermions with opposite spin; $U \approx 8t$ in the experiment [11]. The half-filling limit of the two-dimensional Hubbard model is comparably well understood and can be approximately described for large interactions by the Heisenberg Hamiltonian with superexchange coupling $J = 4t^2/U$ [26]. Starting from high temperature and antiferromagnetic (AFM) correlations with increasing correlation length, AFM correlations emerge. We now investigate the decrease of AFM correlations with doping by comparing the snapshots obtained from the quantum gas microscope to two different theories, a doped resonating valence bond (RVB) liquid [27, 28] and the geometric string theory [11, 29, 30] over a wide range of dopings. Before presenting our results, we provide a brief account of the two theories from which we numerically sample snapshots of the many-body density matrix.

π -flux theory.— In the RVB picture, the ground state of the doped Hubbard model is described as a superposition of different spin-singlet coverings of the lattice, through which deconfined chargons can move freely. Our simulations for this theory are based on a mean-field parton Hamiltonian \hat{H}_{MF} with free spin-1/2 fermions hopping on a square lattice with a magnetic flux of π per plaquette

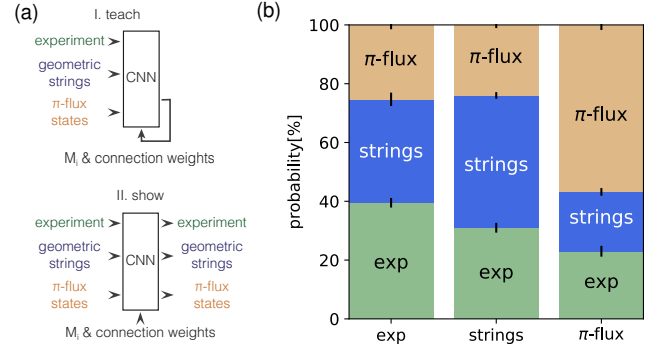


FIG. 2. **Classifying single snapshots of the many-body density matrix.** (a) The convolutional neural network is trained to identify to which dataset category any given snapshot belongs. Here, we consider (i) experimental data, (ii) geometric strings and (iii) snapshots from π -flux theory, all at 9% doping. The probabilities in (b) show how a sample of 400 snapshots which have not been used during training is classified. While the π -flux theory is recognized comparably well, a clear separation between experimental data and geometric strings is not possible. We repeat the process ten times to average out fluctuations in the results, see Supplementary Material for details [31].

[27, 28, 32]. A Gutzwiller projection of the corresponding thermal density matrix $\hat{\rho} = e^{-\beta \hat{H}_{\text{MF}}}$ removes double occupancies in accordance with a large on-site interaction $U \gg t$.

We use Monte Carlo sampling techniques to generate snapshots in the Fock basis of the projected mean-field density matrix. To take into account virtual charge fluctuations present in the larger physical Hilbert space, we introduce doublon-hole pairs into the snapshots on neighboring sites with probability $4t^2/U^2$ determined by second order perturbation theory. The overall energy scale in the mean-field Hamiltonian is fixed such that the nearest-neighbor spin correlator at half-filling matches the experimental value. This approach has been shown in Ref. 11 to lead to good agreement of spin correlations for all relevant doping values. Our results are robust under small variations in the overall energy scale.

Geometric string theory.— In the underdoped regime, this theory describes the fermionic charge carriers as bound states of two partons [33–35]: a heavy spinon and a light chargon; see also Refs. [36–38]. Their internal structure is described by a fluctuating geometric string of displaced spins connecting the spinon to the chargon [29, 30]. In order to derive the properties of this string, the frozen spin approximation is assumed, in which the spin background does not change with doping but the anti-ferromagnetic order is hidden by the hole motion. Each hole displaces the spins along the string by one site, which leads to an increase in spin interaction energy proportional to the spin correlations in the undoped system and a decrease of the overall staggered magnetization.

The distribution of the geometric string length is obtained from a microscopic calculation of the motion of a single hole at a given temperature and Hubbard parameter U/t [11].

To generate snapshots for the geometric string theory, we start from the experimental data at half-filling and for each doping value place the corresponding number of holes independently into the snapshots. The holes are then moved independently from one another in random directions through the anti-ferromagnet for a number of sites which is sampled from the theoretical string length distribution.

The experimental images only contain information about one spin species, while the other spin species as well as doublons and holes are detected as empty sites. Hence, before comparing our theoretical images to experimental results, the second spin species and doubly occupied sites are converted to empty sites in the theoretical data.

Classifying snapshots.— We now train a convolutional neural network to distinguish snapshots from the following classes: (i) experimental data, (ii) geometric string theory and (iii) π -flux theory, all at 9% doping. In all snapshots, the underlying $SU(2)$ symmetry of the Fermi-Hubbard Hamiltonian leads to fluctuations of the Néel ordering vector. To simplify the pattern recognition, we use $\sim 30\%$ of the snapshots with the highest absolute values of the staggered magnetization. Importantly, this still corresponds to a large number of snapshots, which can be considered as a representative sample of the quantum state, see Supplementary Material [31].

The performance of our neural network is visualized in Fig. 2. In this plot, the x -axis displays the actual class of a snapshot and the y -axis shows the probability for the neural network to sort it into the different classes. The accuracy for the classification of images, which corresponds to the weighted average of the diagonal entries, is 47.1%. The main source of confusion for the CNN is the similarity between the experimental and the geometric string theory data, while a differentiation of the π -flux theory snapshots is more successful. Taking the first two categories together, the accuracy of the classification increases to 69.2%. This is a first indication that the geometric string theory resembles the experimental data at 9% doping more closely than π -flux theory.

Sorting experimental data into theory categories.— One of the most powerful features of neural networks is their ability to generalize to new situations not encountered during training. We make use of this property by first training a CNN to distinguish between snapshots from π -flux and geometric string theory at a fixed doping value; a task for which the CNN achieves a precision above 70%. Subsequently we show experimental data to the CNN to sort it into one of the two theory categories. The classification of experimental data then reveals how similar these snapshots are to the theoretically simulated data.

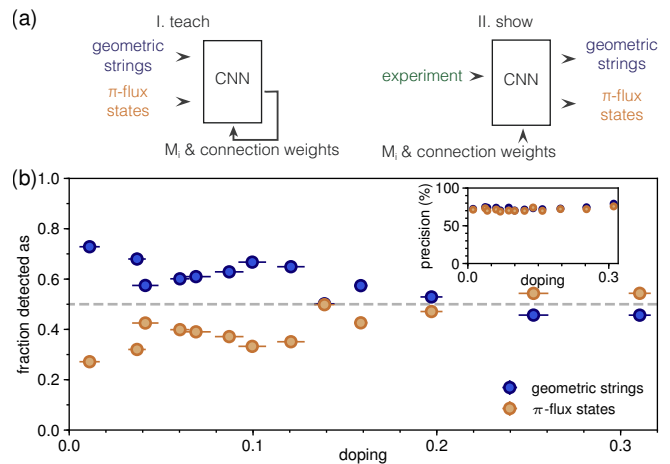


FIG. 3. **Classifying experimental data.** (a) The network is trained to distinguish snapshots from geometric string theory (blue) and π -flux theory (orange) at each doping value. After training, experimental images at the same doping are shown to the network. (b) The average of the resulting classification of the experimental data at the corresponding doping value. The inset shows the precision for the trained classes on a subset of data not used for training. We repeat the process ten times to average out fluctuations in the results.

As shown in Fig. 3, the neural network classifies a majority of the experimental snapshots as geometric string theory over a broad range of doping values up to about 15%, even though conventional spin and charge correlation functions coincide equally well with experimental results in that regime for both theories [11]. For larger dopings, the experimental data is more often sorted into the π -flux theory class, see also Supplementary Material [31].

The ability of the neural network to distinguish π -flux from geometric string theory on the level of individual images indicates that the physical structure of these states is different. We can further improve the accuracy of our classification by taking into account the information that an entire set of measurements belongs to the same physical state. When the CNN sorts each snapshot into one of the two categories with probabilities p and $1 - p$, the entire sample is classified by the category in which the majority is sorted. Because the number of shots in each category follows a binomial distribution, the probability to make a wrong classification of the entire sample decays exponentially with the number of snapshots. Therefore, the entire experimental dataset at any doping value below $\sim 15\%$ would be classified as geometric string theory data with almost 100% probability [31]. When the input to the network consists of four snapshots from the same category, the precision is above 80% already and the fraction of experimental images classified as string theory increases significantly [31].

Moreover, our algorithm also classifies the low-temperature experimental snapshots to geometric string

theory rather than experimental high temperature data, again up to doping values of about 15%; see Supplementary Material for details.

Summary and Outlook.— We have demonstrated that convolutional neural networks provide a powerful tool to analyze the large amount of experimental data obtained from quantum gas microscopes. Individual snapshots can be classified to theoretical predictions and we can thereby determine which theory fits best. We apply this method to the Fermi-Hubbard model on a square lattice and find that on a microscopic level the experimental data more closely resembles the geometric string theory with short-range hidden order than the π -flux RVB theory in the regime of low doping. Our analysis suggests a qualitative change of this behavior between 15% to 20% doping, beyond which the π -flux RVB data resembles the experiment more closely.

Conventional observables, such as the staggered magnetization or two-point spin correlation functions, hardly allow for a distinction between the theories under consideration and it depends on the chosen observable which theory will be favored [11]. By contrast, the CNN searches for patterns in the collection of snapshots in an unbiased way without specifying certain physical observables and with that searches for structure in the many-body density matrix. Turning this argument around, it remains an interesting open challenge to understand how the CNN classifies the snapshots, which we plan to address in a future work.

In this work, we compared two theories out of many potential candidates to the experimental data. In future work, the investigation of a larger class of theories will provide us with further information about the structure of the quantum state of the two-dimensional Fermi Hubbard model. Straightforward extensions include the comparison of snapshots from the Fermi-Hubbard model to different resonating valence bond states or predictions by quantum dimer models [35]. Examining completely different parameter regimes or even models could reveal additional insights. Current experiments have been performed at comparably high temperatures, where no d-wave pairing or charge order is expected. Once colder temperatures are achievable, it will be interesting to compare geometric string theory to theoretical models with different types of order parameters built in.

The analysis of snapshots from quantum gas microscopy with machine learning techniques has the capability to reveal microscopic mechanisms and hidden order in the considerable amount of available data. Machine learning of quantum many-body states, perhaps possible through experimental snapshots, offers prospects to find the most predictive theory among a multitude of competing theories.

Acknowledgments.— We thank Ehud Altman, Juan Carrasquilla, Ehsan Khatami, Frank Pollmann, Achim Rosch, Subir Sachdev, Richard Schmidt and Dries Sels

for insightful discussions and Márton Kanász-Nagy additionally for his Heisenberg QMC code. We acknowledge support from Harvard-MIT CUA, NSF Grant No. DMR-1308435, AFOSR-MURI Quantum Phases of Matter (grant FA9550-14-1-0035), AFOSR grant No. FA9550-16-10323; DoD NDSEG; the Gordon and Betty Moore Foundation EPIQS program and grant No. 6791; NSF GRFP and grant Nos. PHY-1506203 and PHY-1734011; ONR grant No. N00014-18-1-2863; SNSF; Studienstiftung des deutschen Volkes; and the Technical University of Munich - Institute for Advanced Study, funded by the German Excellence Initiative and the European Union FP7 under grant agreement 291763, from the DFG grant No. KN 1254/1-1, and DFG TRR80 (Project F8) and the German Excellence Strategy MCQST.

Supplementary Material: Classifying Snapshots of the Doped Hubbard Model with Machine Learning

Data

The experimental data is obtained from Refs. 7 and 11. All snapshots used in this work consist of a circular region of interest (ROI) with 80 sites. In Table S1 we list the amount of data available at $T = 0.6J \pm 0.1J$ at various dopings, as well as at half-filling for different temperatures. Conventional observables such as the staggered magnetization and spin correlators are very similar to the experiment for both the geometric string and the π -flux theories over a wide range of dopings [11]. A distinction of the theories by the neural network solely based on such 'trivial' quantities can therefore be excluded. Furthermore, for the nearest-neighbor spin correlator, which could probably be used to differentiate single snapshots, the experiment seems to be better described by π -flux theory than by geometric string theory [11]. However, geometric strings do describe the experimentally measured staggered magnetization more closely; as such we separately calculate the staggered magnetizations of the experimental data classified into geometric-string theory and into pi-flux theory in Fig. S1. While the part of the dataset classified as π -flux theory has a slightly lower average staggered magnetization than the set of snapshots sorted into the geometric string theory class, see Fig. S1, the values for both categories are comparably close to each other. The similarity of staggered magnetization values for both theories indicates that the CNN is classifying on the basis of more involved patterns.

As discussed in the main text, we post-select the snapshots in all cases to a high value of the staggered magnetization. For experimental as well as theory data, we

TABLE S1. Number of snapshots from the experiment available at different doping values at $T = 0.6J \pm 0.1J$ (left) and at different temperature values at half-filling (right).

doping (%)	number of snapshots	T/J	number of snapshots
0	5326		
2	480	0.5	65
3	248	0.6	1776
4.5	253	0.7	4657
6	1588	0.8	356
7	687	0.9	346
9	1205	1.0	317
10	483	1.1	313
12.5	185	1.2	475
14	383	1.3	151
17	372	1.4	91
20	335	1.5	62
25	148	1.6	811
32	332	1.8	183

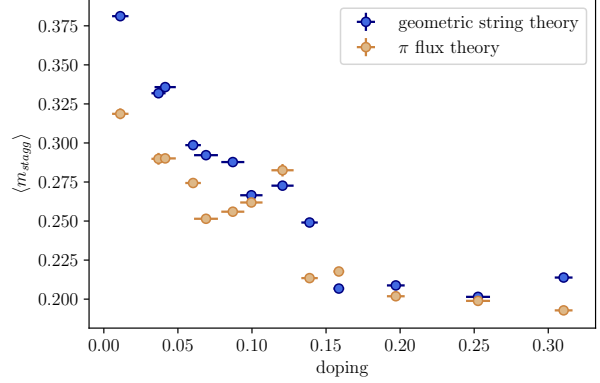


FIG. S1. **Staggered magnetization of the two classes after sorting.** After training the CNN to distinguish geometric string theory and π -flux theory, we show it experimental data, which is then classified by the network as one of the two theories. As a result, we obtain a subset of experimental data sorted into the string theory class and accordingly a second complementary subset classified as π -flux theory. For these two new datasets, we now determine the average value of the staggered magnetization. The similarity of the resulting values for the different theories shows that a classification by the CNN solely based on the staggered magnetization is not convincingly possible.

calculate the absolute value of the staggered magnetization $|m_z|$ for every snapshot and use only the $\sim 30\%$ of all snapshots with the highest value of $|m_z|$. The post-selection is necessary to avoid confusion due to images with Néel ordering vector not aligned with the measurement axis, in which case a single snapshot alone does not contain any useful information. We performed the analysis shown in the main text for different percentages of the total amount of data and found that our results are independent of the details of the post-selection.

Errors

Before training the neural network, we randomly select 400 snapshots from the full data for testing. These snapshots are not used during training, but only to test the performance of the network and determine the accuracy. In order to average out fluctuations in the results, we repeat the training with a different choice of the test data set and sorting of experimental data 10 times. The error bars for the different accuracies as well as the percentage of experimental snapshots detect as one of the two theories shown in the figures represent the standard deviation over the different runs,

$$\Delta x = \frac{\sqrt{\sum_{i=1}^{\mathcal{N}} (x_i - \bar{x})^2}}{\mathcal{N} - 1}, \quad (\text{S1})$$

TABLE S2. Network parameters for Fig. 2, Fig. 3 and S5.

	Fig. 2	Fig. 3	Fig. S5
available snapshots	2059	5326	2848
post-selection (%)	30	30	40
batch size	200	250	250
conv. layer 1	20	24	16
conv. layer 2	-	-	4
dropout	0.8	0.55	0.7
fully conn. layer 1	40	100	40
fully conn. layer 2	30	-	-
dropout	0.7	0.5	0.6
output layer	3	2	2
learning rate	$4 \cdot 10^{-4}$	$1 \cdot 10^{-3}$	$3 \cdot 10^{-4}$
learning rate decay	0.6	0.5	0.5
iterations	1000	1000	1000

where \mathcal{N} is the number of runs, x_i is the quantity under consideration in the i -th run and \bar{x} is the average value of the said quantity over all \mathcal{N} runs.

Network

The convolutional neural network used in this work was implemented using TensorFlow [39]. We tested the performance of the network with different architectures and while there are some quantitative fluctuations, the qualitative result is robust under changes of the architecture. Since the learning tasks as well as the amount of training data available are different for Fig. 2, Fig. 3 and Fig. S5, the optimal performance in terms of overall test accuracy is achieved for slightly different network architectures. The concrete parameters used for the different figures are listed in Table S2.

Distinguishing theories

Geometric string theory based on Heisenberg QMC

In the main text, we train the network to distinguish between geometric string theory and π -flux theory and subsequently sort experimental data into one of those categories. One potential bias of this approach is that the geometric string theory snapshots are based on experimental data at half-filling, whereas the π -flux theory data is purely theoretical apart from a fitted overall energy scale. This bias is to some extent intrinsic, because the geometric string theory does not make a statement about the state at half-filling itself, but only about how the introducing holes changes it. However, we can also generate purely theoretical images for the geometric string theory category by using Heisenberg Quantum Monte Carlo (QMC) snapshots as starting point. After adding the strings by hand as described above we also

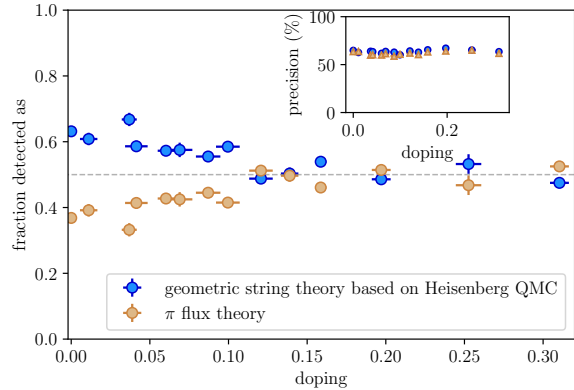


FIG. S2. **Distinguishing geometric string theory from π -flux theory.** The network is trained to distinguish snapshots from geometric string theory (blue) and π -flux theory (orange) at each doping value. As opposed to Fig. 3 in the main text, the geometric string theory data is based on Heisenberg quantum Monte Carlo simulations to avoid a potential bias by using experimental half-filling data. After training, experimental images at the same doping are shown to the network. The plot shows the average of the resulting classification over all available experimental data at the corresponding doping value. The inset shows the precision that is achieved for the trained classes on a subset of data not used for training. We repeat the process ten times in order to average out fluctuations in the results.

add doublon hole pairs in the same way as for π -flux theory. We then train the CNN to distinguish the geometric string theory data based on Heisenberg QMC from π -flux theory data and subsequently sort the experimental data into the two categories. As can be seen from Fig. S2, the qualitative result is very similar to what we obtained in Fig. 3 in the main text where geometric string theory data is based on experimental half-filling snapshots.

Distinguishing geometric string theory from random images

The precision achieved by the CNN during training for the different tasks so far has been around 75%. In Fig. S3 we show the precision as well as the resulting classification of experimental data after training the network to distinguish geometric string theory from random data, where the ratio of occupied to empty sites in every image is chosen according to the doping value under consideration. For doping below 10%, the precision is close to 100%, demonstrating that the network is capable of learning to classify different theories almost perfectly.

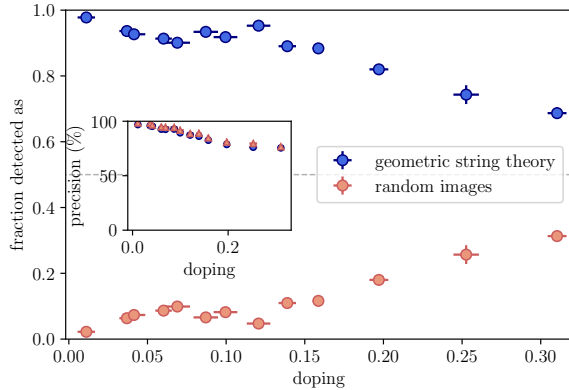


FIG. S3. **Distinguishing geometric string theory from random snapshots.** The network is trained to distinguish snapshots from geometric string theory (blue) and random images (salmon) at each doping value. After training, experimental images at the same doping are shown to the network. The plot shows the average of the resulting classification over all available experimental data at the corresponding doping value. The inset shows the precision that is achieved for the trained classes on a subset of data not used for training. We repeat the process ten times to average out fluctuations in the results.

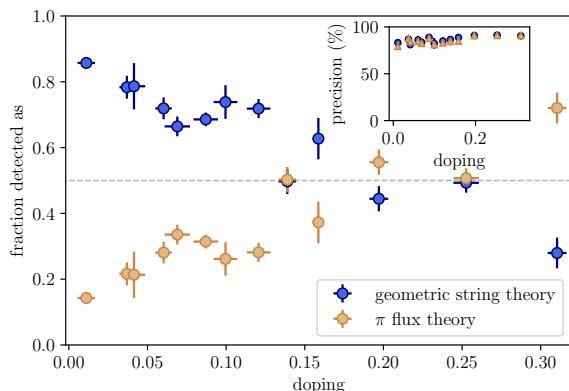


FIG. S4. **Distinguishing geometric string theory from π -flux theory, four snapshots at a time.** The network is trained to distinguish snapshots from geometric string theory (blue) and π -flux theory (orange) at each doping value. Each input corresponds to four randomly chosen images from one theory. After training, experimental images at the same doping are shown to the network. Again, we show the network four images at a time and the network has to classify them into the same category. The plot shows the average of the resulting classification over all available experimental data at the corresponding doping value. The inset shows the precision that is achieved for the trained classes on a subset of data not used for training. We repeat the process five times to average out fluctuations in the results.

In Fig. 3, the neural network sorts about 60% of the experimental snapshots into the geometric string theory category in the low doping regime. As stated in the main text, the accuracy of our classification can be increased further by taking into account the information that an entire set of measurements belongs to the same physical state. Here, we take one step in that direction and group four snapshots of the same category at a time. Additionally, for each of the four snapshots we choose a random number $r = 0, 1, 2, \text{ or } 3$ and rotate the snapshot $r \cdot 90$ degrees. Using the same network architecture, the precision increases to about 85% throughout all doping values, see Fig. S4. After training, we show the network four experimental snapshots at a time and ask it to assign a single label to them. Again, we apply a random number of rotations to each individual snapshot. For dopings below 14%, the fraction of experimental images classified as string theory is between 65 and 85%.

Comparing finite doping to high temperature

Starting at half-filling and cold temperatures around $0.6J$ the AFM correlations vanish both with increasing doping or increasing temperature [11]. We approach the question of how similar these regimes are on the microscopic level. At each doping value we train the CNN to distinguish geometric string theory data and experimental data at high temperature and half-filling. We randomly add holes to the hot data according to the respective doping level to prevent the network from distinguishing the two theories trivially by the density. In order to obtain a sufficiently large training set, we use half-filling data for temperatures between $T = 0.9J$ and $1.8J$ for the high temperature class.

For doping below $\sim 15\%$, the network classifies the experimental data as geometric string theory. However, the precision decreases with increased doping, see inset of Fig. S5. This shows that it is difficult for the CNN to distinguish the theories, indicating that the classification of experimental pictures around 15% doping is not very reliable.

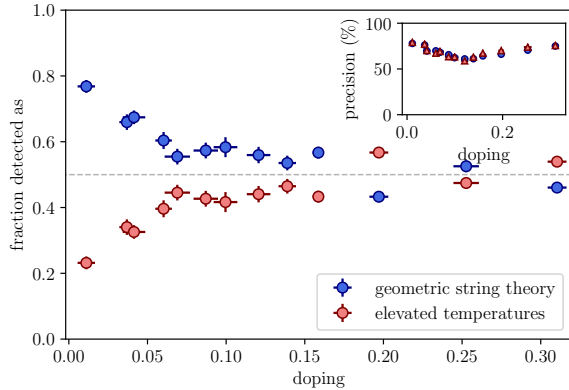


FIG. S5. **Distinguishing finite doping from high temperatures.** The CNN is trained to distinguish geometric string theory at a given doping value from experimental images at half-filling and temperatures $T \geq 0.9J$ with the corresponding number of holes randomly sprinkled into the pictures to prevent a distinction based on the filling. After training, experimental images at low temperatures, $T \approx 0.6J$, and the corresponding doping value are shown to the network. The plot shows the average of the resulting classification. The inset shows the precision that is achieved for the trained classes on a subset of data not used for training. The process is repeated ten times to average out fluctuations in the results.

- [1] P. A. Lee, N. Nagaosa, and X.-G. Wen, *Rev. Mod. Phys.* **78**, 17 (2006).
- [2] B. Keimer, S. A. Kivelson, M. R. Norman, S. Uchida, and J. Zaanen, *Nature* **518**, 179 (2015).
- [3] D. Greif, T. Uehlinger, G. Jotzu, L. Tarruell, and T. Esslinger, *Science* **340**, 1307 (2013).
- [4] R. A. Hart, P. M. Duarte, T.-L. Yang, X. Liu, T. Paiva, E. Khatami, R. T. Scalettar, N. Trivedi, D. A. Huse, and R. G. Hulet, *Nature* **519**, 211 EP (2015).
- [5] L. W. Cheuk, M. A. Nichols, K. R. Lawrence, M. Okan, H. Zhang, E. Khatami, N. Trivedi, T. Paiva, M. Rigol, and M. W. Zwierlein, *Science* **353**, 1260 (2016).
- [6] T. A. Hilker, G. Salomon, F. Grusdt, A. Omran, M. Boll, E. Demler, I. Bloch, and C. Gross, *Science* **357**, 484 (2017).
- [7] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kánász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner, *Nature* **545**, 462 (2017).
- [8] P. T. Brown, D. Mitra, E. Guardado-Sanchez, R. Nourafkan, A. Reymbaut, S. Bergeron, A.-M. S. Tremblay, J. Kokalj, D. A. Huse, P. Schauss, and W. S. Bakr, ArXiv e-prints (2018), [arXiv:1802.09456 \[cond-mat.quant-gas\]](#).
- [9] M. A. Nichols, L. W. Cheuk, M. Okan, T. R. Hartke, E. Mendez, T. Senthil, E. Khatami, H. Zhang, and M. W. Zwierlein, ArXiv e-prints (2018), [arXiv:1802.10018 \[cond-mat.quant-gas\]](#).
- [10] G. Salomon, J. Koepsell, J. Vijayan, T. A. Hilker, J. Nespolo, L. Pollet, I. Bloch, and C. Gross, ArXiv e-prints (2018), [arXiv:1803.08892 \[cond-mat.quant-gas\]](#).
- [11] C. S. Chiu, G. Ji, A. Bohrdt, M. Xu, M. Knap, E. Demler, F. Grusdt, M. Greiner, and D. Greif, ArXiv e-prints (2018), [arXiv:1810.03584 \[cond-mat.quant-gas\]](#).
- [12] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016).
- [13] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
- [14] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, *Physical Review X* **8**, 011006 (2018), [arXiv:1710.04045 \[quant-ph\]](#).
- [15] S. Lu, X. Gao, and L.-M. Duan, ArXiv e-prints (2018), [arXiv:1810.02352 \[quant-ph\]](#).
- [16] J. Carrasquilla and R. G. Melko, *Nature Physics* **13**, 431 EP (2017).
- [17] B. S. Rem, N. Käming, M. Tarnowski, L. Asteria, N. Fläschner, C. Becker, K. Sengstock, and C. Weitenberg, ArXiv e-prints (2018), [arXiv:1809.05519 \[cond-mat.quant-gas\]](#).
- [18] P. Broecker, J. Carrasquilla, R. G. Melko, and S. Trebst, *Scientific Reports* **7**, 8823 (2017).
- [19] K. Ch'ng, J. Carrasquilla, R. G. Melko, and E. Khatami, *Phys. Rev. X* **7**, 031038 (2017).
- [20] M. J. S. Beach, A. Golubeva, and R. G. Melko, *Phys. Rev. B* **97**, 045207 (2018), [arXiv:1710.09842 \[cond-mat.stat-mech\]](#).
- [21] X.-Y. Dong, F. Pollmann, and X.-F. Zhang, ArXiv e-prints (2018), [arXiv:1806.00829 \[cond-mat.dis-nn\]](#).
- [22] J. Greitemann, K. Liu, and L. Pollet, ArXiv e-prints (2018), [arXiv:1804.08557 \[cond-mat.stat-mech\]](#).
- [23] K. Liu, J. Greitemann, and L. Pollet, ArXiv e-prints (2018), [arXiv:1810.05538 \[cond-mat.stat-mech\]](#).
- [24] M. Koch-Janusz and Z. Ringel, *Nature Physics* **14**, 578 (2018), [arXiv:1704.06279 \[cond-mat.dis-nn\]](#).
- [25] Y. Zhang, A. Mesaros, K. Fujita, S. D. Edkins, M. H. Hamidian, K. Ch'ng, H. Eisaki, S. Uchida, J. C. Séamus Davis, E. Khatami, and E.-A. Kim, ArXiv e-prints (2018), [arXiv:1808.00479 \[cond-mat.dis-nn\]](#).
- [26] A. Auerbach, *Interacting Electrons and Quantum Magnetism* (Springer, Berlin, 1998).
- [27] P. W. Anderson, *Science* **235**, 1196 (1987).
- [28] G. Baskaran, Z. Zou, and P. Anderson, *Solid State Communications* **63**, 973 (1987).
- [29] F. Grusdt, M. Kánász-Nagy, A. Bohrdt, C. S. Chiu, G. Ji, M. Greiner, D. Greif, and E. Demler, *Phys. Rev. X* **8**, 011046 (2018).
- [30] F. Grusdt, Z. Zhu, T. Shi, and E. Demler, ArXiv e-prints (2018), [arXiv:1806.04426 \[cond-mat.str-el\]](#).
- [31] See supplementary online material.
- [32] J. B. Marston and I. Affleck, *Phys. Rev. B* **39**, 11538 (1989).
- [33] P. Beran, D. Poilblanc, and R. Laughlin, *Nuclear Physics B* **473**, 707 (1996).
- [34] G. Baskaran, [arXiv:0709.0902](#) (2007).
- [35] M. Punk, A. Allais, and S. Sachdev, *Proceedings of the National Academy of Sciences* **112**, 9552 (2015).
- [36] L. Bulaevskii, E. Nagaev, and D. Khomskii, *JETP*, Vol. 27, No. 5, p. 836 (November 1968).
- [37] S. A. Trugman, *Phys. Rev. B* **37**, 1597 (1988).
- [38] E. Manousakis, *Phys. Rev. B* **75**, 035106 (2007).
- [39] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray,

C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, “[TensorFlow: Large-](#)

[scale machine learning on heterogeneous systems,](#)” (2015), software available from [tensorflow.org](https://www.tensorflow.org).