nature physics

LETTERS

https://doi.org/10.1038/s41567-019-0565-x

Classifying snapshots of the doped Hubbard model with machine learning

Annabelle Bohrdt[1,2,3], Christie S. Chiu[2], Geoffrey Ji[2], Muqing Xu[2], Daniel Greif[2], Markus Greiner[2], Eugene Demler[2], Fabian Grusdt[1,2,3] and Michael Knap[1,3*]

Quantum gas microscopes for ultracold atoms can provide high-resolution real-space snapshots of complex many-body systems. We implement machine learning to analyse 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[1,2], and the geometric string theory[3,4], describing a state with hidden spin order. This technique considers all available information without a potential bias towards one particular theory by the choice of an observable and can therefore select the theory that 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.

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[5,6]. 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[7–15]. 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 have mostly been analysed to obtain conventional observables such as one- and two-point correlation functions, which 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 analyse the data that take all available information into consideration and hence use the capabilities of quantum gas microscopes to their full extent.

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 new tool in condensed matter physics. Its main applications so far include representing the wavefunctions of correlated many-body states[16–19], the determination and characterization of a phase transition[20–29], quantum state tomography[30] and, most recently, a mode of hypothesis testing for experimental scanning tunnelling microscopy data[31]. Here, we take an alternative route and use a convolutional neural network (CNN) to classify experimental data at finite doping into different theoretical categories to determine which theory describes the system best on the microscopic level (see Fig. 1). This approach provides insights into the underlying microscopic structures of the state, which may be inaccessible to conventional observables but can be essential for gaining a deeper understanding of the emergent physics.

The experimental data that we analyse with our machine-learning algorithm have been measured with a quantum gas microscope for ultracold lithium atoms in an optical lattice and are available in ref. [32]. This system is modelled by the Fermi–Hubbard Hamiltonian

$$\hat{\mathcal{H}} = -t \sum_{\sigma=\uparrow,\downarrow} \sum_{\langle i,j \rangle} (\hat{c}^{\dagger}_{i,\sigma}\hat{c}_{j,\sigma} + \text{h.c.}) + U \sum_{j} \hat{c}^{\dagger}_{j,\uparrow}\hat{c}_{j,\uparrow}\hat{c}^{\dagger}_{j,\downarrow}\hat{c}_{j,\downarrow} \quad (1)$$

where the first term describes tunnelling with amplitude t of spin-1/2 fermions between nearest-neighbour sites of a two-dimensional (2D) square lattice. The second term corresponds to on-site interactions of strength U between fermions with opposite spin; $U \approx 8t$ in the experiment[15]. The half-filling limit of the 2D Hubbard model is comparably well understood and can be approximately described for large interactions by the Heisenberg Hamiltonian with superexchange coupling $J = 4t^2/U$ (ref. [33]). Starting from high temperatures $T > J$, on decreasing the temperature, antiferromagnetic correlations with increasing correlation length emerge. We now investigate the decrease of antiferromagnetic correlations with doping by comparing the snapshots obtained from the quantum gas microscope to two different theories, a doped resonating valence bond (RVB) liquid[1,2] and the geometric string theory[3,15,34] 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.

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 π-flux theory are based on a mean-field parton Hamiltonian $\hat{\mathcal{H}}_{\text{MF}}$ with free spin-12 fermions hopping on a square lattice with a magnetic flux of π per plaquette[1,2,35]. A Gutzwiller projection of the corresponding thermal density matrix

[1]Department of Physics and Institute for Advanced Study, Technical University of Munich, Garching, Germany. [2]Department of Physics, Harvard University, Cambridge, MA, USA. [3]Munich Center for Quantum Science and Technology, Munich, Germany. *e-mail: michael.knap@ph.tum.de

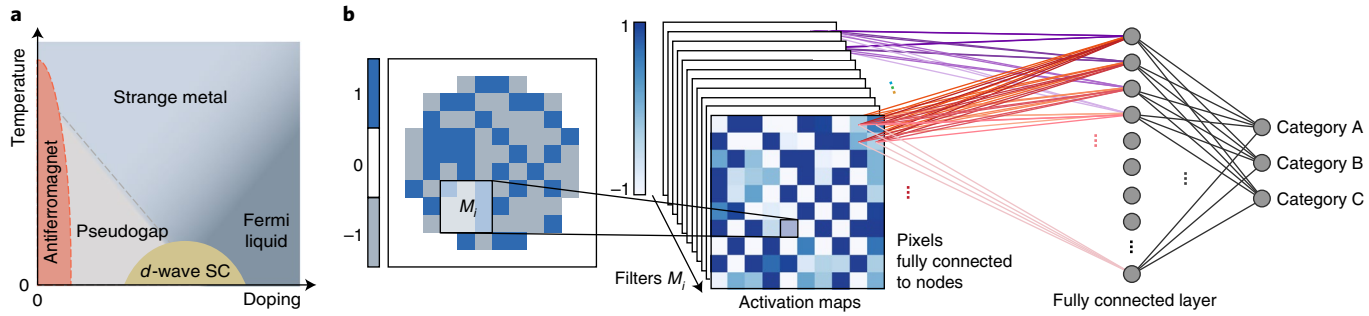

Fig. 1 | Classifying quantum gas microscope snapshots of the doped Fermi–Hubbard model with CNNs. a, A schematic of the conjectured phase diagram of the finite-size 2D Fermi–Hubbard model. We use snapshots of the many-body quantum state at fixed doping and temperature as input data for the CNN. **b**, The main building block of CNNs, which are conventionally used to analyse visual imagery, is the convolutional layer with a set of learnable filters M_i as parameters[42]. 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 2D 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 pattern are detected. The convolutional layer is followed by a fully connected layer, which then sorts the data into the different categories.

$\hat{\rho} \cong e^{-\beta \hat{\mathcal{H}}_{\mathrm{MF}}}$ 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 neighbouring 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-neighbour spin correlator at half-filling matches the experimental value. This approach has been shown in ref. [15] 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.

In the underdoped regime, the geometric string theory describes the fermionic charge carriers as bound states of two partons[36–38]: a heavy spinon and a light chargon (see also refs. [4,39–41]). Their internal structure is described by a fluctuating geometric string of displaced spins connecting the spinon to the chargon[3,34]. 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 (ref. [15]).

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 antiferromagnet for a number of sites that is sampled from the theoretical string length distribution.

The experimental images contain information about only 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. All data used in this analysis are obtained for a temperature of $T = 0.6J \pm 0.1J$, which corresponds to the currently lowest temperatures available in the experiment.

We now train a CNN to distinguish snapshots from the following classes: experimental data, geometric string theory and π-flux theory, all at 9% doping.

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%. This indicates that a classification of the experimental snapshots as one of the theories is in principle possible, since otherwise the CNN would be able to distinguish experimental data from either theory with a high accuracy. 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%. This is a first indication that the geometric string theory resembles the experimental data at 9% doping more closely than π-flux theory.

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%. Here, the precision is defined as the percentage of snapshots classified as a theory that actually belong to the said theory. The precision of the CNN can be further improved by increasing the system size, detecting holes and increasing the size of the training set (see Supplementary Information). Subsequently, we show experimental data to the CNN to sort them into one of the two theory categories. The classification of experimental data then reveals how similar these snapshots are to the theoretically simulated data.

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[15]. For larger dopings, the experimental data cannot be unambiguously classified (see also Supplementary Information).

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. As the number of shots in each category follows a binomial distribution,

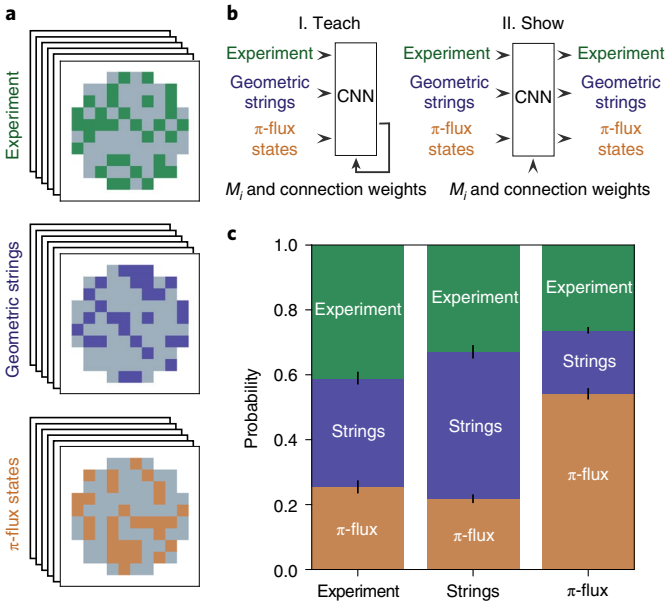

Fig. 2 | Classifying single snapshots of the many-body density matrix. **a**, Randomly chosen snapshots from the experiment and the two theories. Coloured sites are occupied by one spin species; grey sites represent the other spin species, holes and doublons. **b**, The CNN is trained to identify to which dataset category any given snapshot belongs. Here, we consider experimental data, geometric strings and snapshots from π-flux theory, all at 9% doping. **c**, The probabilities show how a sample of 400 snapshots that 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. The standard errors of the mean over ten repetitions of the process are shown by the error bars (see Supplementary Information for details).

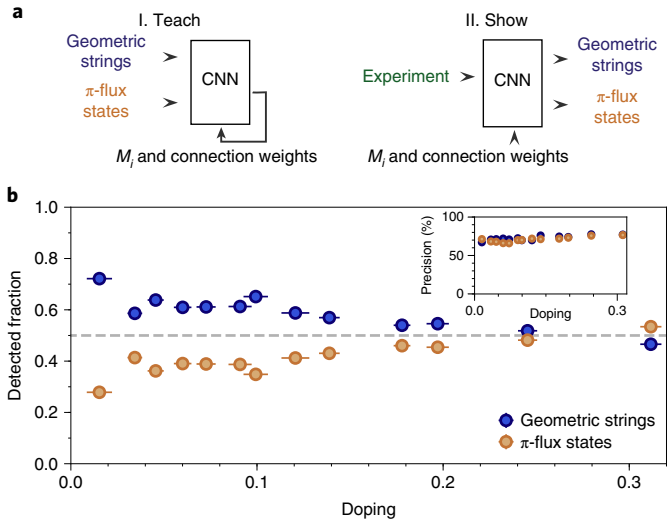

Fig. 3 | Classifying experimental data. **a**, The network is trained to distinguish snapshots from geometric string theory (blue) and π-flux theory (orange) at each doping value. After the training sequence, experimental images at the same doping are shown to the network. **b**, The average of the resulting 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. The vertical error bars show the standard errors of the mean over ten repetitions of the process and are smaller than the plot markers. The horizontal error bars for the doping are obtained as described in the Supplementary Information.

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 ~15% would be classified as geometric string theory data with almost 100% probability (see Supplementary Information). When the input to the network consists of four snapshots from the same category, the precision is already above 80% and the fraction of experimental images classified as string theory at low dopings increases significantly (see Supplementary Information).

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 Information for details).

In summary, we have demonstrated that CNNs provide a powerful tool to analyse 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 resemble 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 behaviour between 15 to 20% doping.

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 favoured[15]. 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 2D Fermi–Hubbard model. Straightforward extensions include the comparison of snapshots from the Fermi–Hubbard model to different RVB states or predictions by quantum dimer models[38]. 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 parameter built.

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.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, statements of code and data availability and associated accession codes are available at https://doi.org/10.1038/s41567-019-0565-x.

Data availability

The data that support the plots within this paper and other findings of this study are available from the corresponding author upon reasonable request. The raw data are available in ref. [32].

Code availability

The computer codes used to generate the results of this paper are available from the corresponding author upon reasonable request.

Received: 1 December 2018; Accepted: 17 May 2019;
Published online: 01 July 2019

References

1. Anderson, P. W. The resonating valence bond state in La_2CuO_4 and superconductivity. *Science* **235**, 1196–1198 (1987).
2. Baskaran, G., Zou, Z. & Anderson, P. The resonating valence bond state and high-T_c superconductivity—a mean field theory. *Solid State Commun.* **63**, 973–976 (1987).
3. Grusdt, F. et al. Parton theory of magnetic polarons: mesonic resonances and signatures in dynamics. *Phys. Rev. X* **8**, 011046 (2018).
4. Grusdt, F., Bohrdt, A. & Demler, E. Microscopic spinon-chargon theory of magnetic polarons in the t-J model. Preprint at https://arxiv.org/abs/1901.01113 (2019).
5. Lee, P. A., Nagaosa, N. & Wen, X.-G. Doping a Mott insulator: physics of high-temperature superconductivity. *Rev. Mod. Phys.* **78**, 17–85 (2006).
6. Keimer, B., Kivelson, S. A., Norman, M. R., Uchida, S. & Zaanen, J. From quantum matter to high-temperature superconductivity in copper oxides. *Nature* **518**, 179–186 (2015).
7. Greif, D., Uehlinger, T., Jotzu, G., Tarruell, L. & Esslinger, T. Short-range quantum magnetism of ultracold fermions in an optical lattice. *Science* **340**, 1307–1310 (2013).
8. Hart, R. A. et al. Observation of antiferromagnetic correlations in the Hubbard model with ultracold atoms. *Nature* **519**, 211–214 (2015).
9. Cheuk, L. W. et al. Observation of spatial charge and spin correlations in the 2D Fermi–Hubbard model. *Science* **353**, 1260–1264 (2016).
10. Hilker, T. A. et al. Revealing hidden antiferromagnetic correlations in doped Hubbard chains via string correlators. *Science* **357**, 484–487 (2017).
11. Mazurenko, A. et al. A cold-atom Fermi–Hubbard antiferromagnet. *Nature* **545**, 462–466 (2017).
12. Brown, P. T. et al. Bad metallic transport in a cold atom Fermi–Hubbard system. *Science* **363**, 379–382 (2019).
13. Nichols, M. A. et al. Spin transport in a Mott insulator of ultracold fermions. *Science* **363**, 383–387 (2019).
14. Salomon, G. et al. Direct observation of incommensurate magnetism in Hubbard chains. *Nature* **565**, 56–60 (2019).
15. Chiu, C. S. et al. String patterns in the doped Hubbard model. Preprint at https://arxiv.org/abs/1810.03584 (2018).
16. Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).
17. Glasser, I., Pancotti, N., August, M., Rodriguez, I. D. & Cirac, J. I. Neural-network quantum states, string-bond states, and chiral topological states. *Phys. Rev. X* **8**, 011006 (2018).
18. Choo, K., Carleo, G., Regnault, N. & Neupert, T. Symmetries and many-body excitations with neural-network quantum states. *Phys. Rev. Lett.* **121**, 167204 (2018).
19. Lu, S., Gao, X. & Duan, L.-M. Efficient representation of topologically ordered states with restricted Boltzmann machines. *Phys. Rev. B* **99**, 155136 (2019).
20. Carrasquilla, J. & Melko, R. G. Machine learning phases of matter. *Nat. Phys.* **13**, 431–434 (2017).
21. van Nieuwenburg, E. P. L., Liu, Y.-H. & Huber, S. D. Learning phase transitions by confusion. *Nat. Phys.* **13**, 435–439 (2017).
22. Rem, B. S. et al. Identifying quantum phase transitions using artificial neural networks on experimental data. *Nat. Phys.* https://doi.org/10.1038/s41567-019-0554-0 (2019).
23. Broecker, P., Carrasquilla, J., Melko, R. G. & Trebst, S. Machine learning quantum phases of matter beyond the fermion sign problem. *Sci. Rep.* **7**, 8823 (2017).
24. Ch'ng, K., Carrasquilla, J., Melko, R. G. & Khatami, E. Machine learning phases of strongly correlated fermions. *Phys. Rev. X* **7**, 031038 (2017).
25. Beach, M. J. S., Golubeva, A. & Melko, R. G. Machine learning vortices at the Kosterlitz–Thouless transition. *Phys. Rev.* **97**, 045207 (2018).
26. Dong, X.-Y., Pollmann, F. & Zhang, X.-F. Machine learning of quantum phase transitions. *Phys. Rev. B* **99**, 121104 (2019).
27. Greitemann, J., Liu, K. & Pollet, L. Probing hidden spin order with interpretable machine learning. *Phys. Rev. B* **99**, 060404 (2019).
28. Liu, K., Greitemann, J. & Pollet, L. Learning multiple order parameters with interpretable machines. *Phys. Rev. B* **99**, 104410 (2019).
29. Koch-Janusz, M. & Ringel, Z. Mutual information, neural networks and the renormalization group. *Nat. Phys.* **14**, 578–582 (2018).
30. Torlai, G. et al. Neural-network quantum state tomography. *Nat. Phys.* **14**, 447–450 (2018).
31. Zhang, Y. et al. Machine learning in electronic-quantum-matter imaging experiments. *Nature* https://doi.org/10.1038/s41586-019-1319-8 (2019).
32. Chiu, C. S. et al. Data for 'String patterns in the doped Hubbard model' (Harvard Dataverse, 2019); https://doi.org/10.7910/DVN/1CSVBV
33. Auerbach, A. *Interacting Electrons and Quantum Magnetism* (Springer, 1998).
34. Grusdt, F., Zhu, Z., Shi, T. & Demler, E. Meson formation in mixed-dimensional t–J models. *SciPost Phys.* **5**, 057 (2018).
35. Marston, J. B. & Affleck, I. Large-*n* limit of the Hubbard-Heisenberg model. *Phys. Rev. B* **39**, 11538–11558 (1989).
36. Beran, P., Poilblanc, D. & Laughlin, R. Evidence for composite nature of quasiparticles in the 2d *t-J* model. *Nucl. Phys. B* **473**, 707–720 (1996).
37. Baskaran, G. 3/2-Fermi liquid: the secret of high-Tc cuprates. Preprint at https://arxiv.org/abs/0709.0902 (2007).
38. Punk, M., Allais, A. & Sachdev, S. Quantum dimer model for the pseudogap metal. *Proc. Natl Acad. Sci. USA* **112**, 9552–9557 (2015).
39. Bulaevskii, L., Nagaev, E. & Khomskii, D. A new type of auto-localized state of a conduction electron in an antiferromagnetic semiconductor. *J. Exp. Theor. Phys.* **27**, 836 (1968).
40. Trugman, S. A. Interaction of holes in a Hubbard antiferromagnet and high-temperature superconductivity. *Phys. Rev. B* **37**, 1597–1603 (1988).
41. Manousakis, E. String excitations of a hole in a quantum antiferromagnet and photoelectron spectroscopy. *Phys. Rev. B* **75**, 035106 (2007).
42. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* (MIT Press, 2016).

Acknowledgements

We thank E. Altman, I. Bloch, J. Carrasquilla, M. Kanász-Nagy, E. Khatami, F. Pollmann, A. Rosch, S. Sachdev, R. Schmidt and D. Sels for insightful discussions and M. Kanász-Nagy in addition 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 programme 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, the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy–EXC-2111–390814868, the DFG grant no. KN1254/1-1, and DFG TRR80 (Project F8).

Author contributions

A.B., F.G. and M.K. devised the method. A.B. carried out the numerical simulations and analysis. F.G. and E.D. developed the geometric string theory. C.S.C., G.J., M.X. and D.G. performed the experiments. M.G., E.D., F.G. and M.K. supervised the work. All authors contributed to the writing of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at https://doi.org/10.1038/s41567-019-0565-x.

Reprints and permissions information is available at www.nature.com/reprints.

Correspondence and requests for materials should be addressed to M.K.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.