

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

Integrating Neural Networks with a Quantum Simulator for State Reconstruction

Giacomo Torlai, Brian Timar, Evert P. L. van Nieuwenburg, Harry Levine, Ahmed Omran, Alexander Keesling, Hannes Bernien, Markus Greiner, Vladan Vuletić, Mikhail D. Lukin, Roger G. Melko, and Manuel Endres

Phys. Rev. Lett. **123**, 230504 — Published 6 December 2019

DOI: [10.1103/PhysRevLett.123.230504](https://doi.org/10.1103/PhysRevLett.123.230504)

Integrating Neural Networks with a Quantum Simulator for State Reconstruction

Giacomo Torlai,^{1,2,3,*} Brian Timar,^{4,*} Evert P.L. van Nieuwenburg,⁴ Harry Levine,⁵ Ahmed Omran,⁵ Alexander Keesling,⁵ Hannes Bernien,⁶ Markus Greiner,⁵ Vladan Vuletić,⁷ Mikhail D. Lukin,⁵ Roger G. Melko,^{2,3} and Manuel Endres⁴

¹*Center for Computational Quantum Physics, Flatiron Institute, New York, New York 10010, USA*

²*Department of Physics and Astronomy, University of Waterloo, Ontario N2L 3G1, Canada*

³*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*

⁴*Division of Physics, Mathematics and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA*

⁵*Department of Physics, Harvard University, Cambridge, MA 02138, USA*

⁶*Institute for Molecular Engineering, University of Chicago, Chicago, IL 60637, USA*

⁷*Department of Physics and Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

We demonstrate quantum many-body state reconstruction from experimental data generated by a programmable quantum simulator, by means of a neural network model incorporating known experimental errors. Specifically, we extract restricted Boltzmann machine (RBM) wavefunctions from data produced by a Rydberg quantum simulator with eight and nine atoms in a single measurement basis, and apply a novel regularization technique to mitigate the effects of measurement errors in the training data. Reconstructions of modest complexity are able to capture one- and two-body observables not accessible to experimentalists, as well as more sophisticated observables such as the Rényi mutual information. Our results open the door to integration of machine learning architectures with intermediate-scale quantum hardware.

Quantum state tomography [1] is an important tool for reconstructing generic quantum states, but traditional techniques require a number of measurements scaling exponentially in the system size [2]. In certain cases, methods that exploit particular entanglement or symmetry properties [3–7] allow for more efficient tomography of states prepared in experiments. However, such approaches still involve explicit reconstruction of local density operators [3, 8], incurring a significant computational overhead in the estimation of nontrivial observables from experimental data – especially in the presence of measurement errors introduced by realistic experimental hardware. In order to facilitate the characterization of near-term quantum hardware [9], a state reconstruction method which can efficiently extract physical quantities of interest directly from noisy experimental datasets is highly desirable.

Neural network-based machine learning has recently emerged as a powerful technique for learning compact representations of high-dimensional data [10–12]. In experimental quantum science, these tools have already been applied profitably to the classification of experimental snapshots [13, 14] and qubit readout [15]. The same data-driven approach can be applied to tomographic tasks. Recent theoretical work has demonstrated that a generative model called a restricted Boltzmann machine (RBM) is capable of accurate reconstruction of quantum states and observables directly from synthetic datasets generated by numerical algorithms [16].

In this Letter, we present a proof-of-principle demonstration of neural network quantum state reconstruction from experimental data. Our experimental system consists of a one-dimensional array of strongly interacting

Rydberg atoms [17, 18]. Leveraging the high purity and approximate positivity of the experimental state, we train RBMs using single measurement basis data consisting of bit-strings obtained via repeated, simultaneous single-shot readout of the ground and Rydberg populations of all atoms. The RBMs learn a higher-fidelity and more efficient representation of the underlying bit-string probability distributions than standard inference from the limited size training dataset. This approach also enables us to implement an efficient procedure for denoising the full probability distribution from bit-flip-type measurement errors, by incorporating a dedicated “noise layer” in the network architecture. We test the validity of our approach by comparing predictions of the trained RBMs with numerical results for observables that are off-diagonal in the measurement basis, including the quantum mutual information. These results demonstrate the utility of RBMs in reconstructing approximately pure, positive states from experimental data, and pave the way to further integration of neural network models with quantum hardware.

Experimental system. Our experimental approach [17, 18] involves a programmable Rydberg atom quantum simulator, a flexible neutral-atom system for realizing Ising-type quantum spin models [17, 19–24]. In the present experiments (Fig. 1a), a one-dimensional array of N trapped Rubidium atoms is prepared; $N = 8$ atoms are used below, but we have also applied our protocol to arrays of $N = 9$ atoms [25]. Each atom can occupy a ground state $|g\rangle$ or an excited (Rydberg) state $|r\rangle$, and two atoms excited to the Rydberg state at a distance r interact with a van der Waals-type potential $V(r) \propto r^{-6}$. When subjected to a uniform laser drive,




Figure 1. Experiment and reconstruction. (a) Model of the reconstruction process. Individual ^{87}Rb atoms (grey circles) are trapped in an array of optical tweezers and coupled to a Rydberg state with Rabi frequency Ω . Site-resolved fluorescence imaging provides imperfect measurement in the $\hat{\sigma}^z$ basis. Our neural network model describes the true quantum state as an RBM (blue and green neurons), while the binary data τ accessible to the experimentalist are included as an auxiliary ‘noise’ layer (red neurons). By training on this data, the network learns parameters λ describing the experimental quantum state, which are subsequently used to compute observables $\langle \hat{O} \rangle$. (b) Representation of the ordered state at the end of the adiabatic sweep – see Eq (2). Darker circles represent a higher probability of Rydberg excitation, and the shading indicates quantum fluctuations localized at bonds (3,4) and (5,6). (c) The effective laser detuning Δ and Rabi frequency Ω as a function of sweep time t . Circular markers indicate the times at which the sweep was halted to collect data. Vertical line: approximate transition to ordering in the finite system. The nearest-neighbor interaction is $V_{nn} = 30$ MHz, the final detuning is 10 MHz, and the peak Rabi frequency is 2 MHz; the total sweep time is $T_{\text{ev}} = 3.4\mu\text{s}$.

the effective Hamiltonian of the many-body system can be written as [17, 19, 26, 27]

$$\hat{H}(\Omega, \Delta) = -\Delta \sum_{i=1}^N \hat{n}_i - \frac{\Omega}{2} \sum_{i=1}^N \hat{\sigma}_i^x + \sum_{i < j} \frac{V_{nn}}{|i-j|^6} \hat{n}_i \hat{n}_j, \quad (1)$$

where V_{nn} is the interaction strength between Rydberg atoms at adjacent sites, $\hat{\sigma}_i^\alpha$, with $\alpha = x, y, z$, are the Pauli pseudo-spin operators at site i (defined as $\hat{\sigma}_i^z = |r_i\rangle\langle r_i| - |g_i\rangle\langle g_i|$, $\hat{\sigma}_i^x = |r_i\rangle\langle g_i| + h.c.$, etc), and $\hat{n}_i = \frac{1}{2}(1 + \hat{\sigma}_i^z)$ projects onto the Rydberg state at site i . The parameters Ω, Δ denote the effective Rabi frequency and detuning, respectively, which characterize the laser drive, and can be varied in time as $\Omega(t), \Delta(t)$ to drive the system into nontrivial ordered phases [17, 26, 28, 29].

We focus on the transition into the $\mathbb{Z}_2$ phase [17], where a high density of Rydberg excitations is energetically favorable, subject to the constraint that no two adjacent atoms are excited. The atoms are initially pumped into the fiducial state $|ggggg\dots\rangle$, coinciding with the ground state of Hamiltonian (1) at $t = 0$. They then evolve adiabatically under a ‘sweep’ of the laser parameters $\Omega(t), \Delta(t)$ for a time T_{ev} , with $\hat{H}(\Omega(T_{\text{ev}}), \Delta(T_{\text{ev}}))$ lying deep in the $\mathbb{Z}_2$ phase (Fig. 1c). For our eight-atom system, the final $\mathbb{Z}_2$ -ordered state at $t = T_{\text{ev}}$ is well approximated by the ground state of the Rydberg Hamiltonian with a small transverse field and short-range interactions only [25]:

$$|\psi\rangle = \frac{1}{\sqrt{2}}|rgrggrgr\rangle + \frac{1}{2}|rgrggrgr\rangle + \frac{1}{2}|rggrgrgr\rangle. \quad (2)$$

This state exhibits quantum fluctuations on two pairs of adjacent atoms, as indicated in Fig. 1b.

Pure state ansatz. The ground state of the Hamiltonian (1) has real, positive amplitudes in the occupation number basis $|\sigma\rangle = |\sigma_1, \dots, \sigma_N\rangle$ – defined as the simultaneous eigenstates of $\hat{n}_1, \dots, \hat{n}_N$ – as long as $\Omega > 0$ [30], which can always be arranged by applying a suitable global unitary [31]. Therefore, if the quantum state of the simulator evolves perfectly adiabatically and with negligible loss of purity, it is uniquely characterized by its probability distribution $p(\sigma)$ over projective measurements in the $|\sigma\rangle$ basis, and at any time may be written as the pure state

$$|\psi\rangle = \sum_{\sigma} \sqrt{p(\sigma)} |\sigma\rangle. \quad (3)$$

Of course, some loss of purity is inevitable – in our experiments, due primarily to single-atom decay and dephasing processes [32] – and the true state is described by a mixed density operator $\hat{\rho}$. Although this pure state approximation cannot capture all of the physics of the experimental state, it can in principle accurately describe reduced density operators of the true and reconstructed states agree [25]. We adopt the pure, positive state ansatz in all of our reconstruction efforts below.

Neural network model. While the quantum state (3) can in principle be inferred directly from a set of raw measurements (i.e. by inverting the measurement counts of each configuration to estimate $p(\sigma)$), such an approach is limited to small systems and very large datasets. In contrast, *generative models* used in unsupervised machine learning tasks can capture the structure of the distribution $p(\sigma)$, generalizing beyond a limited set of training samples. This results in a higher-fidelity reconstruction and a model size scaling polynomially in the system size (Fig. 2). Moreover, using a generative model rather than direct inference from the data enables automatic correction of this distribution for known measurement errors using a ‘noise layer’ (see Fig. 1a and description below).

We parametrize $p(\sigma)$ with a generative model known as an RBM [33, 34], a stochastic neural network with two layers of binary units. The “visible” layer σ describes the atomic states of the Rydberg chain in the occupation number basis, while a hidden layer $\mathbf{h}$ captures correlations between visible units. The RBM defines the following probability distribution for the visible layer:

$$p_{\lambda}(\sigma) = \frac{1}{Z_{\lambda}} \sum_{\mathbf{h}} e^{\mathbf{h}^T \mathbf{W} \sigma + \mathbf{b} \cdot \sigma + \mathbf{c} \cdot \mathbf{h}}, \quad (4)$$

where Z_{λ} is a normalization constant, and the real-valued network parameters are $\lambda = \{\mathbf{W}, \mathbf{b}, \mathbf{c}\}$, with $\mathbf{W}$ being the weights connecting the two layers and $\mathbf{b}$ ($\mathbf{c}$) the visible (hidden) bias vectors. We use the visible layer of the RBM to define the projective measurement distribution $p(\sigma)$ of the pure state (3), resulting in an RBM wavefunction with positive amplitudes [35]: $\psi_{\lambda}(\sigma) = \langle \sigma | \psi_{\lambda} \rangle = \sqrt{p_{\lambda}(\sigma)}$. We have numerically verified that this RBM wavefunction can accurately describe states relevant to our experiment, with a number of parameters scaling only quadratically in system size (Fig. 2 and [25], Sec. IV), in accordance with recent scaling studies for quantum Ising ground states [36]. We point out that, although pure states with nontrivial phases [37, 38], as well as mixed state models [39, 40], could be applied using similar neural network models, measurements in other bases would be required.

Measurement process and noise layer. Measurement data consists of a collection of N -bit strings $\tau = (\tau_1, \dots, \tau_N)$, with $\tau_j = 0, 1$ indicating that atom j was recorded as being in the ground $|g\rangle$ or Rydberg state $|r\rangle$ respectively [17]. Such measurements are never perfect, and there are small measurement error probabilities $p(1|0) \sim 1\%$, $p(0|1) \sim 4\%$ [32] for an atom in the ground state to be recorded as excited and vice-versa. These result in experimental data τ that do not correspond to projective measurements. Instead, the measurement process can be described as a positive-operator valued measure (POVM) [41] with measurement operators $\hat{\Pi}_{\tau} = \sum_{\sigma} p(\tau|\sigma) |\sigma\rangle\langle\sigma|$, where $p(\tau|\sigma) = \prod_{j=1}^N p(\tau_j|\sigma_j)$ is the probability of the experimentalist recording τ if the atoms are prepared in the state $|\sigma\rangle$. The probability distribution sampled in the experiment is then $P_{\text{exp}}(\tau) = \text{Tr}[\hat{\rho} \hat{\Pi}_{\tau}]$.

The experimental measurement process is incorporated into our model via a third binary layer, the so-called *noise layer* (Fig. 1a), which represents the observed POVM outcomes τ . The measurement error rates $p(\tau|\sigma)$ are included as connections between the visible and noise layers [42], by assigning a probability $\tilde{p}_{\lambda}(\tau) = \sum_{\sigma} p(\tau|\sigma) p_{\lambda}(\sigma)$ to the measurement result τ .

The full three-layer network is trained to learn parameters λ which maximize the log-likelihood of the recorded POVM outcomes under $\tilde{p}_{\lambda}(\tau)$. During training, the noise layer prevents the parameters λ from fitting to spurious




Figure 2. Benchmarking RBM reconstruction. (a) Fidelity of reconstruction. We sample synthetic datasets from states obtained by exact time-evolution under the Hamiltonian (1) without decoherence. The exact quantum state fidelity F between the true state $\hat{\rho}$ and the reconstruction $\hat{\rho}_{\lambda} = |\psi_{\lambda}\rangle\langle\psi_{\lambda}|$ is plotted as a function of detuning Δ . Training standard RBMs on datasets without measurement noise (green dashed line), we achieve uniformly high fidelities, demonstrating that the RBM wavefunction ansatz is capable of representing states relevant to our experiment. Training on datasets with measurement noise with (red solid line) and without (green solid line) noise-layer regularization shows how the modified training improves reconstruction. Inset: same data, for time-evolution including a realistic decoherence model. (b) Model size. Here we compare the number of parameters N_p required to specify an RBM wavefunction with N hidden units with the size of the frequency-distribution (FD) model required to perform direct inference (i.e. number of different configurations in the dataset), for a typical Rydberg ground state, as a function of system size N and for several dataset sizes N_s . Note that the FD model size depends on N_s , while the RBM size does not. For further discussion, see [25].

features in the data produced by measurement errors. This *noise layer regularization* significantly improves the fidelity between $|\psi_{\lambda}\rangle$ and the state $\hat{\rho}$ underlying the data; numerical tests (Fig. 2) based on Lindbladian simulation of our experiment result in fidelities greater than 90% for the full many-body state at the end of the sweep, even when decoherence processes are included. All reconstructions presented below are obtained in this fashion.

Experimental reconstruction. In the experiment, at fifteen subsequent time-steps t (Fig. 1c), the sweep is halted and measurements τ are sampled from the state $\hat{\rho}(t)$. At each time-step, a dataset of around 3,000 samples is collected and used to train a three-layer model with $2N = 16$ hidden units. After training the networks,




Figure 3. Few-body observables. Comparison of the RBM reconstruction (squares) with the experiment results (crosses) and the predictions from the Lindbladian master equation (circles) [25]. In order to facilitate comparison with experiment, the values reported in (a) and (b) for the RBM and Lindbladian observables are computed including the known measurement error rates $p(0|1) = 0.04$, $p(1|0) = 0.01$. (a) Nearest-neighbor correlations $\bar{g}^{zz}(1)$ in the z basis, spatially averaged (see text for definition). (b) Average correlation $\bar{g}^{zz}(s)$ as a function of distance s for $\Delta = 10$ MHz. (c) Spatial average $\bar{x}$ of the transverse field $\langle \hat{\sigma}_i^x \rangle$. (d) Nearest-neighbor correlation $\langle \hat{\sigma}_i^x \hat{\sigma}_{i+1}^x \rangle_c$ as a function of position i for $\Delta = 10$ MHz. The two peaks correspond to the bonds highlighted in Fig. 1b.

standard sampling methods can be applied to compute expectation values of observables, with a computational cost scaling polynomially in the network size [25]. We consider in particular the connected correlation functions $\langle \hat{\sigma}_i^\alpha \hat{\sigma}_j^\alpha \rangle_c = \langle \hat{\sigma}_i^\alpha \hat{\sigma}_j^\alpha \rangle - \langle \hat{\sigma}_i^\alpha \rangle \langle \hat{\sigma}_j^\alpha \rangle$ for $\alpha = x, y, z$, and their spatial averages, $\bar{g}^{\alpha\alpha}(s) = \frac{1}{N-s} \sum_{i=1}^{N-s} \langle \hat{\sigma}_i^\alpha \hat{\sigma}_{i+s}^\alpha \rangle_c$.

In Fig. 3a-b, we verify that our reconstructions learn to represent their training sets, by examining their ability to accurately reproduce observables which are diagonal in the occupation number basis. The networks learn the strong two-body correlations $\langle \hat{\sigma}_i^z \hat{\sigma}_j^z \rangle_c$ present in the experimental data. We compare the results of the reconstruction process to the exact solutions of a Lindblad master equation for the full many-body evolution. Our Lindbladian simulation predicts Rydberg excitation profiles in excellent agreement with experiment, but its significantly weaker correlations suggest our model for the sweep dynamics is partially incomplete.

Turning to experimentally inaccessible quantities (Fig. 3c-d), the reconstructions and simulation agree qualitatively in the temporal and spatial profiles of the transverse field $\langle \hat{\sigma}_i^x \rangle$ and its two-point correlation function, although the RBMs predict somewhat larger val-




Figure 4. Rényi Mutual Information. The quantum (Rényi) mutual information I_2 , defined as $I_2(s) = S_2(\hat{\rho}_s^A) + S_2(\hat{\rho}_s^B) - S_2(\hat{\rho})$, where $S_2(\hat{\rho}) = -\log \text{Tr} \hat{\rho}^2$ is the second-order Rényi entropy, $\hat{\rho}$ is the (mixed) state of the whole system, and $\hat{\rho}_s^A, \hat{\rho}_s^B$ are the reduced density matrices for the subsystems $A_s = \{1, \dots, s\}$, $B_s = \{s+1, \dots, N\}$ respectively, defined by a partitioning of the system at bond $(s, s+1)$. The mutual information is plotted for a partition at bond $(3, 4)$, as a function of detuning. Inset: The mutual information $I_2(s)$ as a function of the cut bond s for $\Delta = 10$ MHz.

ues in the ordered phase. Note that the distinct spatial variation of the transverse field correlations, a signature of quantum fluctuations captured in the approximate state (2), is reconstructed directly from our experimental data. Training on synthetic data [25] indicates that a large portion of the disagreement between reconstruction and simulation is due to the discrepancy between our Lindbladian model and experiment evident in Fig. 3a-b, not the RBM model itself.

Beyond few-body observables, an important question is whether entanglement properties are reproduced accurately in reconstruction. From our RBMs, the Rényi entropy – which requires specialized or hardware-specific protocols to access directly in experiment [43, 44] – may be extracted in a scalable fashion by applying a state-replication and swap procedure virtually [37, 45]. In fact, for pure experimental states, positive-pure ansatzes such as the RBM wavefunction provide a lower bound on the mutual information defined by the Rényi entropy ([46], [47], see also [25], Sec. VIII), regardless of the sign structure of the true state. We demonstrate a reconstruction of the mutual information defined by the Rényi entropy in Fig. 4, finding that the RBM values are in remarkable agreement with the results of numerical simulation. Reconstructions on experimental states of $N = 9$ capture a similar buildup in the mutual information during the sweep predicted by Lindbladian simulation [25].

Conclusions. In this Letter, we have demonstrated neural-network reconstruction of experimental quantum states from data produced by a programmable Rydberg-atom quantum simulator. By leveraging the real-positive nature of the ground state wavefunction expected from the effective Hamiltonian, we trained restricted Boltzmann machines on measurements in the occupation basis

only. An additional noise layer was added to the standard RBM architecture to mitigate measurement errors. Once trained, the RBM was queried to produce a variety of observables not accessible in the original experimental setup, including the Rényi entropy - a basis independent measure of the quantum entanglement of the wavefunction.

Our approach can be integrated without alteration into existing platforms where a positive wavefunction ansatz is a valid approximation, such as Bose-Hubbard experiments and some non-frustrated quantum spin simulators [20, 48–50]. Access to multiple measurement bases would allow enhanced certification of the reconstruction, by providing direct experimental access to observables which are informationally complete for local subsystems. Also, with access to different bases the RBM protocol can be easily adapted to reconstruct non-positive and complex wavefunctions [37]. In this case, the reconstruction cannot immediately accommodate the noise-layer regularization. In turn, the denoising can be implemented in any scheme where the quantum state is uniquely specified by a classical probability distribution, such as in generative modeling of POVM measurements [40]. For non-positive wavefunctions, identifying the minimal set of measurement bases and the optimal protocol to collect the statistics represents a crucial step towards reconstruction of quantum states prepared by fermionic quantum simulators and non-equilibrium dynamics [51, 52].

In conclusion, machine learning techniques offer a means of increasing the amount of useful information that can be extracted from experiments, especially when hardware constrains the quantity or quality of accessible measurements. They can be used to offload the burden of technically expensive – or fundamentally impossible – measurements from experimental platforms in a noise-resilient fashion. We expect experimentalists will profit from deeper integration of machine learning architectures with quantum devices.

We thank Dmitry Abanin for helpful discussions, and Soonwon Choi and Hannes Pichler for pointing out the bound on Rényi entropies. M.E. and B.T. acknowledge funding provided by the Institute for Quantum Information and Matter, an NSF Physics Frontiers Center (NSF Grant PHY-1733907), as well as the NSF CAREER award (1753386), and the AFOSR YIP (FA9550-19-1-0044). The Flatiron Institute is supported by the Simons Foundation. R.G.M. was supported by NSERC of Canada, a Canada Research Chair, and the Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported through Industry Canada and by the Province of Ontario through the Ministry of Research & Innovation. This research was supported in part by NSF Grant No. PHY-1748958, NIH Grant No. R25GM067110, and the Gordon and Betty Moore Foundation Grant No. 2919.01.

* These authors contributed equally.

- [1] K. Banaszek, M. Cramer, and D. Gross, *New Journal of Physics* **15**, 125020 (2013).
- [2] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chekalkar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt, *Nature* **438**, 643 (2005).
- [3] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, *Nature Communications* **1**, 149 (2010).
- [4] J. Y. Lee and O. Landon-Cardinal, *Phys. Rev. A* **91**, 062128 (2015).
- [5] C. A. Riofrio, D. Gross, S. T. Flammia, T. Monz, D. Nigg, R. Blatt, and J. Eisert, *Nature Communications* **8**, 15305 (2017).
- [6] B. P. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. S. Buyskikh, A. J. Daley, M. Cramer, M. B. Plenio, R. Blatt, and C. F. Roos, *Nature Physics* **13**, 1158 (2017).
- [7] G. Tóth, W. Wieczorek, D. Gross, R. Krischek, C. Schwemmer, and H. Weinfurter, *Phys. Rev. Lett.* **105**, 250403 (2010).
- [8] M. Ježek, J. Fiurášek, and Z. Hradil, *Physical Review A* **68**, 012305 (2003).
- [9] J. Preskill, *Quantum* **2**, 79 (2018).
- [10] G. E. Hinton, *Science* **313**, 504 (2006).
- [11] A. Graves, A.-r. Mohamed, and G. Hinton, in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing* (IEEE, 2013) pp. 6645–6649.
- [12] Y. LeCun, Y. Bengio, and G. Hinton, *Nature* **521**, 436 (2015).
- [13] B. S. Rem, N. Käming, M. Tarnowski, L. Asteria, N. Fläschner, C. Becker, K. Sengstock, and C. Weitenberg, *Nature Physics* **15**, 917 (2019).
- [14] A. Bohrdt, C. S. Chiu, G. Ji, M. Xu, D. Greif, M. Greiner, E. Demler, F. Grusdt, and M. Knap, *Nature Physics* **15**, 921 (2019).
- [15] A. Seif, K. A. Landsman, N. M. Linke, C. Figgatt, C. Monroe, and M. Hafezi, *Journal of Physics B: Atomic, Molecular and Optical Physics* **51**, 174006 (2018), arXiv:1804.07718.
- [16] G. Torlai and R. G. Melko, arXiv e-prints, arXiv:1905.04312 (2019), arXiv:1905.04312 [quant-ph].
- [17] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, *Nature* **551**, 579 (2017).
- [18] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, *Science* **354**, 1024 (2016).
- [19] P. Schauß, J. Zeiher, T. Fukuhara, S. Hild, M. Cheneau, T. Macrì, T. Pohl, I. Bloch, and C. Gross, *Science* **347**, 1455 (2015).
- [20] H. Labuhn, D. Barredo, S. Ravets, S. de Léséleuc, T. Macrì, T. Lahaye, and A. Browaeys, *Nature* **534**, 667 (2016).
- [21] J. Zeiher, R. van Bijnen, P. Schauß, S. Hild, J.-y. Choi, T. Pohl, I. Bloch, and C. Gross, *Nature Physics* **12**, 1095 (2016).
- [22] J. Zeiher, J.-y. Choi, A. Rubio-Abadal, T. Pohl, R. van Bijnen, I. Bloch, and C. Gross, *Physical Review X* **7**,

- 041063 (2017).
- [23] E. Guardado-Sanchez, P. T. Brown, D. Mitra, T. Devakul, D. A. Huse, P. Schauß, and W. S. Bakr, *Physical Review X* **8**, 021069 (2018).
 - [24] D. Barredo, V. Lienhard, S. de Léséleuc, T. Lahaye, and A. Browaeys, *Nature* **561**, 79 (2018).
 - [25] See Supplemental Material, which includes Refs. [53-66].
 - [26] T. Pohl, E. Demler, and M. D. Lukin, *Physical Review Letters* **104**, 043002 (2010).
 - [27] R. M. W. van Bijnen, S. Smit, K. A. H. van Leeuwen, E. J. D. Vredenbregt, and S. J. J. M. F. Kokkelmans, *Journal of Physics B: Atomic, Molecular and Optical Physics* **44**, 184008 (2011).
 - [28] H. Weimer and H. P. Büchler, *Phys. Rev. Lett.* **105**, 230403 (2010).
 - [29] E. Sela, M. Punk, and M. Garst, *Physical Review B* **84**, 085434 (2011).
 - [30] S. Bravyi, D. P. Divincenzo, R. Oliveira, and B. M. Terhal, *Quantum Info. Comput.* **8**, 361 (2008).
 - [31] The exact phases required to render the Rydberg Hamiltonian in the form (1) will vary between different experimental realizations, but as the final measurements are always taken in the occupation number basis this has no effect on observables, provided no variation in the laser phase occurs during evolution.
 - [32] H. Levine, A. Keesling, A. Omran, H. Bernien, S. Schwartz, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, *Physical Review Letters* **121**, 123603 (2018).
 - [33] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *Cognitive Science* **9**, 147 (1985).
 - [34] P. Smolensky, in *Parallel Distributed Processing*, edited by D. E. Rumelhart, J. L. McClelland, and C. PDP Research Group (MIT Press, Cambridge, MA, USA, 1986) Chap. Information Processing in Dynamical Systems: Foundations of Harmony Theory, pp. 194–281.
 - [35] G. Torlai and R. G. Melko, *Physical Review B* **94**, 165134 (2016).
 - [36] D. Sehayek, A. Golubeva, M. S. Albergo, B. Kulchytskyy, G. Torlai, and R. G. Melko, arXiv:1908.07532 [quant-ph] (2019), arXiv: 1908.07532.
 - [37] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, *Nature Physics* **14**, 447 (2018).
 - [38] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
 - [39] G. Torlai and R. G. Melko, *Physical Review Letters* **120**, 240503 (2018).
 - [40] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, *Nature Machine Intelligence* **1**, 155 (2019).
 - [41] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th ed. (Cambridge University Press, Cambridge, 2010).
 - [42] Yichuan Tang, R. Salakhutdinov, and G. Hinton, in *2012 IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2012) pp. 2264–2271.
 - [43] R. Islam, R. Ma, P. M. Preiss, M. Eric Tai, A. Lukin, M. Rispoli, and M. Greiner, *Nature* **528**, 77 (2015).
 - [44] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, (2018), arXiv:1806.05747.
 - [45] M. B. Hastings, I. González, A. B. Kallin, and R. G. Melko, *Physical Review Letters* **104**, 157201 (2010).
 - [46] Y. Zhang, T. Grover, and A. Vishwanath, *Physical Review Letters* **107**, 067202 (2011).
 - [47] T. Grover and M. P. A. Fisher, *Physical Review A* **92**, 042308 (2015).
 - [48] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, *Nature* **462**, 74 (2009).
 - [49] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr, *Nature* **471**, 319 (2011).
 - [50] A. M. Kaufman, M. E. Tai, A. Lukin, M. Rispoli, R. Schittko, P. M. Preiss, and M. Greiner, *Science* **353**, 794 (2016).
 - [51] L. W. Cheuk, M. A. Nichols, M. Okan, T. Gersdorf, V. V. Ramasesh, W. S. Bakr, T. Lompe, and M. W. Zwierlein, *Physical Review Letters* **114**, 193001 (2015).
 - [52] D. Greif, M. F. Parsons, A. Mazurenko, C. S. Chiu, S. Blatt, F. Huber, G. Ji, and M. Greiner, *Science (New York, N.Y.)* **351**, 953 (2016).
 - [53] S. Sachdev, *Quantum Phase Transitions*, 2nd ed. (Cambridge University Press, 2011).
 - [54] G. E. Hinton, in *Neural Networks: Tricks of the Trade: Second Edition*, edited by G. Montavon, G. B. Orr, and K.-R. Müller (Springer Berlin Heidelberg, Berlin, Heidelberg, 2012) pp. 599–619.
 - [55] G. E. Hinton, *Neural Comput.* **14**, 1771 (2002).
 - [56] M. J. S. Beach, I. D. Vlugt, A. Golubeva, P. Huembeli, B. Kulchytskyy, X. Luo, R. G. Melko, E. Merali, and G. Torlai, *SciPost Phys.* **7**, 9 (2019).
 - [57] G. E. Hinton, S. Osindero, and Y.-W. Teh, *Neural Computation* **18**, 1527 (2006).
 - [58] N. Le Roux and Y. Bengio, *Neural Computation* **20**, 1631 (2008).
 - [59] D.-L. Deng, X. Li, and S. Das Sarma, *Physical Review X* **7**, 021021 (2017).
 - [60] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, *Physical Review X* **8**, 011006 (2018).
 - [61] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, *Physical Review B* **97**, 085104 (2018).
 - [62] P. Weinberg and M. Bukov, *SciPost Physics* **2**, 003 (2017).
 - [63] D. M. Greig, B. T. Porteous, and A. H. Seheult, *Journal of the Royal Statistical Society: Series B (Methodological)* **51**, 271 (1989).
 - [64] S. Geman and D. Geman, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-6**, 721 (1984).
 - [65] J. Batson and L. Royer, (2019), arXiv:1901.11365.
 - [66] J. R. Johansson, P. D. Nation, and F. Nori, *Computer Physics Communications* **184**, 1234 (2013).