

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 properties [3–6] allow for more efficient tomography of states prepared in experiment. However, such approaches still involve explicit reconstruction of local density operators [3, 7], 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 noisy intermediate-scale quantum hardware [8] and simulators that are currently being realized experimentally, a state reconstruction method which can efficiently extract physical quantities of interest directly from raw, noisy experimental datasets is therefore highly desirable.

Neural network-based machine learning has recently emerged as a powerful technique for learning compact representations of high-dimensional data [9–11]. In experimental quantum science, these tools have already been applied profitably to the classification of experimental snapshots [12, 13] and qubit readout [14]. The same data-driven approach can be applied to tomographic tasks. Recent theoretical work [15] 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 data sets generated by numerical algorithms.

In this Letter, we present a proof-of-principle demonstration of neural network quantum state reconstruction from experimental data. Our experimental system con-

sists of a one-dimensional array of strongly interacting Rydberg atoms [16, 17]. Leveraging the high purity and approximate positivity of the experimental state, we train RBMs using a single measurement basis. Specifically, we use a data set 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 data set. 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 [16, 17] involves a programmable Rydberg atom quantum simulator, a flexible neutral-atom system for realizing Ising-type quantum spin models [16, 18–23]. In the experiments utilized in this study (Fig. 1a), a one-dimensional array of N trapped Rubidium atoms is prepared; $N = 8$ atoms are used for the reconstructions presented below, but we have also applied our protocol to arrays of $N = 9$ atoms [24]. 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 in-

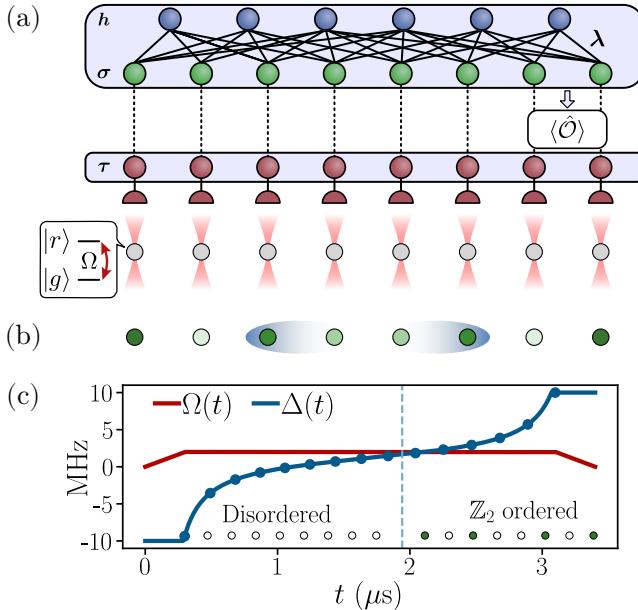


Figure 1. Experiment and reconstruction. (a) Model of the reconstruction process. Individual ^{87}Rb atoms (grey circles) are trapped in an array of optical tweezers and coupled to a Rydberg state with Rabi frequency Ω . Site-resolved fluorescence imaging provides imperfect measurement in the $\hat{\sigma}^z$ 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 $\langle \hat{O} \rangle$ of interest. (b) Representation of the ordered state at the end of the adiabatic sweep – see main text equation (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 $T_{\text{ev}} = 3.4\mu\text{s}$.

teract with a van der Waals-type potential $V(r) \propto r^{-6}$. When subjected to a uniform laser drive, the effective Hamiltonian of the many-body system may be written as [16, 18, 25, 26]

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

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 [16, 25, 27, 28].

We focus on the transition into the \mathbb{Z}_2 phase [16], 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 $|gggggg\dots\rangle$, coinciding with the ground state of Hamiltonian (6) at $t = 0$, and 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 [24]:

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

This state exhibits quantum fluctuations on two pairs of adjacent atoms, as indicated in Fig. 1b. At intermediate times t , in particular when $\Delta(t), \Omega(t)$ approach the quantum critical point (defined in the thermodynamic limit) for the transition into \mathbb{Z}_2 ordering, the ground state cannot be written in such a simple form.

Pure state ansatz - In general, the ground state of the Hamiltonian (6) 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 $\Omega > 0$ [29], which can always be arranged by applying a suitable global unitary [30]. Therefore, if the quantum state of the simulator evolves perfectly adiabatically and with negligible loss of purity, it is uniquely characterized by the probability distribution $p(\sigma)$ it assigns to projective measurements in the $|\sigma\rangle$ basis, and at any time may be written explicitly 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 [31] – 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 local subsystems, to the extent that the corresponding reduced density operators of the true and reconstructed states agree [24], and we adopt the pure, positive state ansatz in all of our reconstruction efforts.

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 severely limited to small systems and datasets with a very large number of samples. In

contrast, *generative models* used in unsupervised machine learning tasks provide the ability to 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 sub-exponentially in the system size, as we demonstrate explicitly in [24], Sec. IV, for data sets of the type used in this manuscript. 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 detailed description below).

We therefore parametrize $p(\sigma)$ with a generative model known as an RBM [32, 33], 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 h captures correlations between visible units. The RBM defines the following probability distribution for the visible layer:

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

$$= \frac{1}{Z_{\lambda}} e^{b \cdot \sigma} \prod_j \left(1 + e^{c_j + \sum_i W_{ji} \sigma_i} \right) \quad (5)$$

where Z_{λ} is a normalization constant, and the real-valued network parameters are $\lambda = \{W, b, c\}$, with W being the weights connecting the two layers and b (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 [15, 34] with positive amplitudes: $\psi_{\lambda}(\sigma) = \langle \sigma | \psi_{\lambda} \rangle = \sqrt{p_{\lambda}(\sigma)}$. We point out that, although pure states with nontrivial phases [15, 35], as well as mixed state models [36, 37], could be applied using similar neural network models, measurements in other bases – not used for the reconstruction process described in this Letter – would be required for this to provide an advantage.

Measurement process and noise layer - We utilize results of measurements taken in the occupation number basis, based on detection of the ground state $|g\rangle$ and Rydberg state $|r\rangle$ [16] at each site in the array. Each individual measurement consists of an N -bit string $\tau = (\tau_1, \dots, \tau_N)$, with $\tau_j = 0, 1$ indicating that atom j was recorded as being in the ground or Rydberg state respectively.

Such measurements are never perfect, and there are small measurement error probabilities $p(1|0) \sim 1\%$, $p(0|1) \sim 4\%$ [31] for an atom in the ground state to be recorded as excited at the time of measurement and vice-versa. These result in experimental data τ that do not correspond to projective measurement. Instead, the measurement process can be described as a positive-operator valued measure (POVM) [38] with measurement operators $\hat{\Pi}_{\tau} = \sum_{\sigma} p(\tau|\sigma) |\sigma\rangle\langle\sigma|$, where $p(\tau|\sigma) =$

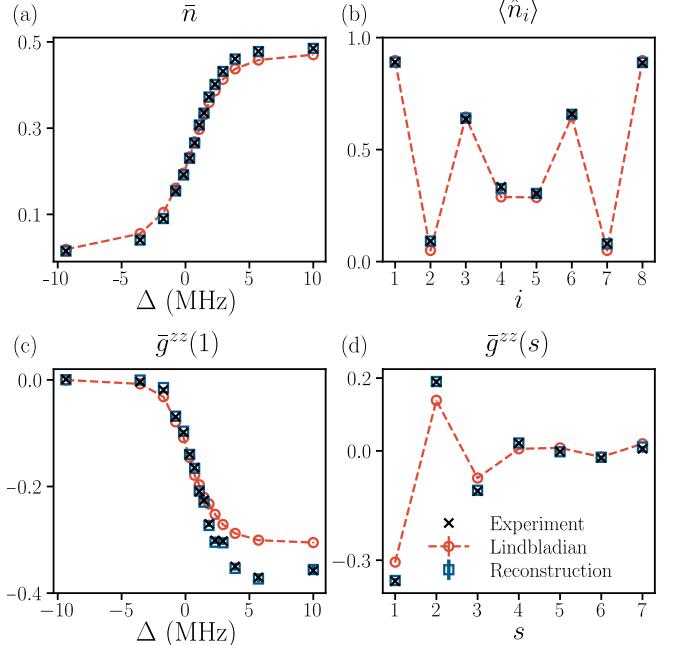


Figure 2. Reconstruction of diagonal observables. Here and throughout the paper, crosses indicate experimental values and blue squares are the RBM values reconstructed from experimental data. Red circles are Lindbladian master equation predictions [24]. Error bars for reconstructions are defined by variation in the final epochs of training across training subsets, and for Lindbladian predictions by variation over different effective Doppler disorder realizations [24]. In order to facilitate comparison with experiment, the values reported here 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) Expected value \bar{n} of the Rydberg population, averaged over all sites, as a function of detuning Δ at which the sweep is halted. (b) Rydberg excitation number at site i for $\Delta = 10$ MHz. (c) Nearest-neighbor correlations $\bar{g}^{zz}(1)$ in the z basis, spatially averaged (see text for definition). (d) Average correlation $\bar{g}^{zz}(s)$ as a function of distance s for $\Delta = 10$ MHz.

$\prod_{j=1}^N p(\tau_j|\sigma_j)$ is the probability of the experimentalist recording bitstring τ if the atoms are prepared in the state $|\sigma\rangle$. The probability distribution sampled in 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 [39], 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 acts as a buffer between the noisy data and the RBM, preventing the parameters from fitting to

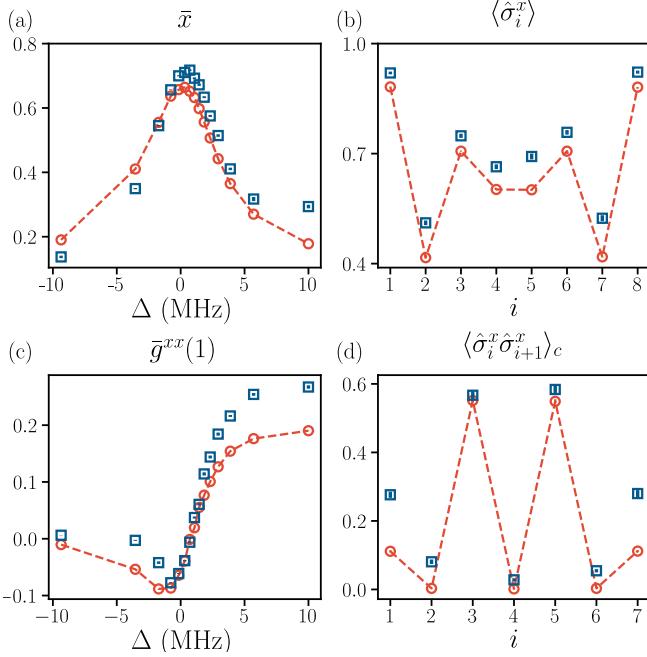


Figure 3. Reconstruction of off-diagonal observables. (a) Spatial average \bar{x} of the transverse field $\langle \hat{\sigma}_i^x \rangle$. (b) Transverse field expectation values at site i for $\Delta \approx 0.7$ MHz. (c) Averaged nearest-neighbor correlations in the x basis. (d) Nearest-neighbor 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.

spurious 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 [24] 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. All reconstructions presented below are obtained in this fashion.

Experimental Reconstruction: accessible observables - 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 single three-layer model with $2N = 16$ hidden units. After training the networks, standard sampling methods can be applied to compute expectation values of observables, with a computational cost scaling polynomially in the network size [24]. 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. 2, 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. Besides the Rydberg excitation profiles $\langle \hat{n}_i \rangle$ (Figs. 2a,b) the networks learn

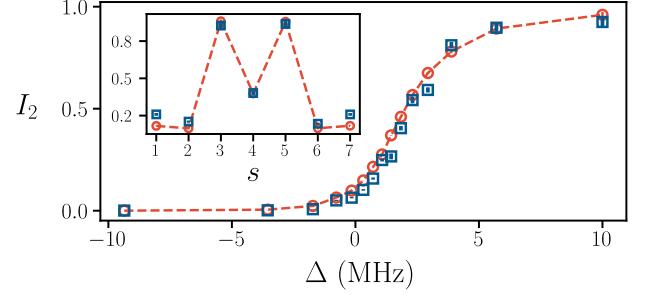


Figure 4. Rényi Mutual Information in the many-body state. The quantum (Rényi) mutual information I_2 , defined as $I_2(s) = S_2(\hat{\rho}_s^A) + S_2(\hat{\rho}_s^B) - S_2(\hat{\rho})$, where $S_2(\hat{\rho}) = -\log \text{Tr} \hat{\rho}^2$ is the second-order Rényi entropy, $\hat{\rho}$ is the (mixed) 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.

the strong two-body correlations $\langle \hat{\sigma}_i^z \hat{\sigma}_j^z \rangle_c$ present in the experimental data (Figs. 2c,d). 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.

Experimental Reconstruction: inaccessible observables - Turning to experimentally inaccessible quantities (Fig. 3), the reconstructions and simulation agree qualitatively in the temporal and spatial profiles of the transverse field $\langle \hat{\sigma}_i^x \rangle$, although the RBMs predict somewhat larger values 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 (Fig. 3d). Training on synthetic data [24] indicates that a large portion of the disagreement between reconstruction and simulation in e.g. the correlations $\bar{g}^{xx}(1)$ (Fig. 3c) is due to the discrepancy between our Lindbladian model and experiment evident in Fig. 2, 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 [40, 41] – may be extracted in a scalable fashion by applying a state-replication and swap procedure virtually [15, 42]. 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 ([43],

[44], see also [24], Sec. VIII), regardless of the sign structure of the true state. We demonstrate a reconstruction of the mutual information defined by the second-order 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 [24].

Conclusions - In this Letter, we have demonstrated neural network reconstruction of experimental quantum states from data produced by a programmable Rydberg-atom quantum simulator. Using measurements in the occupation number basis, we trained a generative model commonly used in unsupervised machine learning tasks, the restricted Boltzmann machine. Once trained, the RBM was sampled 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 takes advantage of the real-positive nature of the ground state wavefunction expected from the effective Hamiltonian, thereby requiring data in only one measurement basis. It provides access to diagonal and off-diagonal observables, and 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 [19, 45–47].

An additional noise layer in the RBM architecture mitigates measurement errors, one of the major challenges in accurate reconstruction from experimental data. We anticipate that such noise layer regularization, readily applicable to other models of mixed states [36] and more general noise processes, will be crucial in future efforts in state reconstruction and snapshot characterization from imperfect measurement data.

In future applications, measurements in rotated bases could be used to certify the reconstructions, by providing direct experimental access to observables which are informationally complete for local subsystems. In addition, access to multiple bases allows for an RBM-based protocol [15] suitable for non-positive states, of immediate relevance to non-equilibrium situations and fermionic quantum simulators [48, 49]. In this context, beamsplitter operations, used to extract Rényi entropies [40, 50], can also provide direct experimental access to quantities which are off-diagonal in the occupation number basis [51], and it may be possible to extract equivalent information from random unitary protocols [41, 52–54]. We envision that raw data from such schemes will be used directly as input for neural network-based quantum state reconstruction.

More generally, 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, and 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. Chekalakar, 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. Riofrío, 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. Holzapfel, 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] M. Ježek, J. Fiurášek, and Z. Hradil, Physical Review A **68**, 012305 (2003).
- [8] J. Preskill, Quantum **2**, 79 (2018).
- [9] G. E. Hinton, Science **313**, 504 (2006).
- [10] A. Graves, A.-r. Mohamed, and G. Hinton, in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing* (IEEE, 2013) pp. 6645–6649.
- [11] Y. LeCun, Y. Bengio, and G. Hinton, Nature **521**, 436 (2015).
- [12] B. S. Rem, N. Kämäring, M. Tarnowski, L. Asteria, N. Fläschner, C. Becker, K. Sengstock, and C. Weitenberg, (2018), arXiv:1809.05519.

- [13] A. Bohrdt, C. S. Chiu, G. Ji, M. Xu, D. Greif, M. Greiner, E. Demler, F. Grusdt, and M. Knap, (2018), arXiv:1811.12425.
- [14] 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.
- [15] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, *Nature Physics* **14**, 447 (2018).
- [16] 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).
- [17] 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).
- [18] P. Schauß, J. Zeiher, T. Fukuhara, S. Hild, M. Cheneau, T. Macrì, T. Pohl, I. Bloch, and C. Gross, *Science* **347**, 1455 (2015).
- [19] H. Labuhn, D. Barredo, S. Ravets, S. de Léséleuc, T. Macrì, T. Lahaye, and A. Browaeys, *Nature* **534**, 667 (2016).
- [20] J. Zeiher, R. van Bijnen, P. Schauß, S. Hild, J.-y. Choi, T. Pohl, I. Bloch, and C. Gross, *Nature Physics* **12**, 1095 (2016).
- [21] J. Zeiher, J.-y. Choi, A. Rubio-Abadal, T. Pohl, R. van Bijnen, I. Bloch, and C. Gross, *Physical Review X* **7**, 041063 (2017).
- [22] E. Guardado-Sánchez, P. T. Brown, D. Mitra, T. Devakul, D. A. Huse, P. Schauß, and W. S. Bakr, *Physical Review X* **8**, 021069 (2018).
- [23] D. Barredo, V. Lienhard, S. de Léséleuc, T. Lahaye, and A. Browaeys, *Nature* **561**, 79 (2018).
- [24] See Supplementary Information.
- [25] T. Pohl, E. Demler, and M. D. Lukin, *Physical Review Letters* **104**, 043002 (2010).
- [26] 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).
- [27] H. Weimer and H. P. Büchler, *Phys. Rev. Lett.* **105**, 230403 (2010).
- [28] E. Sela, M. Punk, and M. Garst, *Physical Review B* **84**, 085434 (2011).
- [29] S. Bravyi, D. P. Divincenzo, R. Oliveira, and B. M. Terhal, *Quantum Info. Comput.* **8**, 361 (2008).
- [30] The exact phases required to render the Rydberg Hamiltonian in the form (6) 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.
- [31] 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).
- [32] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *Cognitive Science* **9**, 147 (1985).
- [33] 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.
- [34] G. Torlai and R. G. Melko, *Physical Review B* **94**, 165134 (2016).
- [35] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
- [36] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, *Nature Machine Intelligence* **1**, 155 (2019).
- [37] G. Torlai and R. G. Melko, *Physical Review Letters* **120**, 240503 (2018).
- [38] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th ed. (Cambridge University Press, Cambridge, 2010).
- [39] Yichuan Tang, R. Salakhutdinov, and G. Hinton, in *2012 IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2012) pp. 2264–2271.
- [40] R. Islam, R. Ma, P. M. Preiss, M. Eric Tai, A. Lukin, M. Rispoli, and M. Greiner, *Nature* **528**, 77 (2015).
- [41] 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.
- [42] M. B. Hastings, I. González, A. B. Kallin, and R. G. Melko, *Physical Review Letters* **104**, 157201 (2010).
- [43] Y. Zhang, T. Grover, and A. Vishwanath, *Physical Review Letters* **107**, 067202 (2011).
- [44] T. Grover and M. P. A. Fisher, *Physical Review A* **92**, 042308 (2015).
- [45] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, *Nature* **462**, 74 (2009).
- [46] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr, *Nature* **471**, 319 (2011).
- [47] A. M. Kaufman, M. E. Tai, A. Lukin, M. Rispoli, R. Schittko, P. M. Preiss, and M. Greiner, *Science* **353**, 794 (2016).
- [48] 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).
- [49] D. Greif, M. F. Parsons, A. Masurenko, C. S. Chiu, S. Blatt, F. Huber, G. Ji, and M. Greiner, *Science* (New York, N.Y.) **351**, 953 (2016).
- [50] A. J. Daley, H. Pichler, J. Schachenmayer, and P. Zoller, *Physical Review Letters* **109**, 020505 (2012).
- [51] A. Bohrdt, C. B. Mendl, M. Endres, and M. Knap, *New Journal of Physics* **19**, 063001 (2017), arXiv:1612.02434.
- [52] M. Ohliger, V. Nesme, and J. Eisert, *New Journal of Physics* **15**, 015024 (2013).
- [53] A. Elben, B. Vermersch, M. Dalmonte, J. I. Cirac, and P. Zoller, *Physical Review Letters* **120**, 050406 (2018).
- [54] B. Vermersch, A. Elben, M. Dalmonte, J. I. Cirac, and P. Zoller, *Phys. Rev. A* **97**, 23604 (2018).
- [55] 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.
- [56] G. E. Hinton, S. Osindero, and Y.-W. Teh, *Neural Computation* **18**, 1527 (2006).
- [57] N. Le Roux and Y. Bengio, *Neural Computation* **20**, 1631 (2008).
- [58] D.-L. Deng, X. Li, and S. Das Sarma, *Physical Review X* **7**, 021021 (2017).
- [59] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, *Physical Review X* **8**, 011006 (2018).
- [60] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, *Physical Review B* **97**, 085104 (2018).
- [61] P. Weinberg and M. Bukov, *SciPost Physics* **2**, 003 (2017).
- [62] D. M. Greig, B. T. Porteous, and A. H. Seheult, *Journal*

- of the Royal Statistical Society: Series B (Methodological) **51**, 271 (1989).
- [63] S. Geman and D. Geman, IEEE Transactions on Pattern Analysis and Machine Intelligence **PAMI-6**, 721 (1984).
- [64] J. Batson and L. Royer, (2019), arXiv:1901.11365.
- [65] J. R. Johansson, P. D. Nation, and F. Nori, Computer Physics Communications **184**, 1234 (2013).

Supplementary Information

In this Supplementary Information, we first provide a derivation of the approximate eight-atom ordered ground state. Next, we discuss how the unsupervised RBM learning process is carried out on experimental datasets, and demonstrate how the networks generalize from the finite datasets used in training. We also detail a regularization method used to mitigate the effect of measurement errors in the training set and provide numerical evidence that this technique significantly improves the fidelity of state reconstruction from noisy data. Finally, we examine how intrinsic decoherence processes impact the quality of the pure-state reconstruction procedure. An appendix provides proofs of two bounds regarding the fidelity and entanglement properties of reconstructions.

I. APPROXIMATE EIGHT-ATOM GROUND STATE

The full Rydberg Hamiltonian is

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

The ground state for $\Delta > 0, V_{nn} \gg \Delta \gg |\Omega|$ lies in the *blockaded subspace* spanned by the maximally-excited basis states

$$|e_1\rangle = |r\ g\ r\ g\ g\ r\ g\ r\rangle \quad (7)$$

$$|e_2\rangle = |r\ g\ g\ r\ g\ r\ g\ r\rangle \quad (8)$$

$$|e_3\rangle = |r\ g\ r\ g\ r\ g\ g\ r\rangle \quad (9)$$

which have no adjacent Rydberg excitations; neglecting interactions beyond second-nearest neighbors (which are small compared to all other relevant frequencies in the sweep), these states are all degenerate under the classical part of the Hamiltonian $-\Delta \sum_{i=1}^N \hat{n}_i + \sum_{i < j} \frac{V_{nn}}{|i-j|^6} \hat{n}_i \hat{n}_j$. This degeneracy is lifted by a nonzero transverse field, which couples the blockaded states at second order. Using perturbation theory, an effective Hamiltonian H_{eff} may be constructed for the blockaded subspace, whose nonzero matrix elements are

$$\langle e_1 | H_{\text{eff}} | e_2 \rangle = \langle e_1 | H_{\text{eff}} | e_3 \rangle = -\frac{\Omega^2}{4\Delta} \quad (10)$$

The corresponding ground state is $|\Psi\rangle = \frac{1}{\sqrt{2}}|e_1\rangle + \frac{1}{2}(|e_2\rangle + |e_3\rangle)$.

II. RECONSTRUCTION METHODS

A. Note on terminology

Below we discuss strategies for training on experimental data which has been corrupted by a fixed, known noise process. σ will denote the variables prior to corruption by measurement errors, while τ will denote those which have been subjected to the noise channel – that is, for a fixed true value σ , the noisy outputs are distributed according to $p(\tau|\sigma)$. In our experiment, τ are the only accessible variables, which yield the bitstrings recorded in each dataset. A model with parameters λ specifies a distribution $p_\lambda(\sigma)$ over the uncorrupted variables σ , and a corresponding corrupted distribution $\tilde{p}_\lambda(\tau) = \sum_\sigma p(\tau|\sigma)p_\lambda(\sigma)$.

B. Standard RBM training method

The standard training method involves fitting the RBM distribution $p_\lambda(\sigma) = \frac{1}{Z_\lambda} \sum_h e^{h^\top W\sigma + b \cdot \sigma + c \cdot h}$ directly to the experimental datasets; in other words, it assumes a noise-free source of data:

$$p(\tau|\sigma) = \delta_{\tau,\sigma} \quad (11)$$

The optimal parameters $\lambda = \{W, b, c\}$ for which the RBM best reproduces the measurement data are found by minimizing the negative log-likelihood

$$\mathcal{L}_\lambda = -\frac{1}{|\mathcal{D}|} \sum_{\tau \in \mathcal{D}} \log p_\lambda(\tau) \quad (12)$$

of the RBM distribution p_λ averaged over the dataset \mathcal{D} ($|\mathcal{D}|$ denotes the size of the dataset). The gradient of the log-likelihood cost function with respect to the trainable parameters λ may be written

$$\nabla_\lambda \mathcal{L}_\lambda = \langle \nabla_\lambda \mathcal{E}_{\text{eff}}(\sigma) \rangle_{p_\lambda(\sigma)} - \frac{1}{|\mathcal{D}|} \sum_{\tau \in \mathcal{D}} \nabla_\lambda \mathcal{E}_{\text{eff}}(\tau) \quad (13)$$

where $\langle \cdot \rangle_{p_\lambda(\sigma)}$ denotes the expectation value with respect to the distribution $p_\lambda(\sigma)$, and the effective energies

$$\mathcal{E}_{\text{eff}}(\sigma) = b \cdot \sigma + \sum_j \log(1 + e^{W_{ji}\sigma_i + c_j}) \quad (14)$$

are defined by $p_\lambda(\sigma) = \frac{1}{Z_\lambda} e^{\mathcal{E}_{\text{eff}}(\sigma)}$. The exact calculation of the first term in this gradient requires knowledge of the partition function $Z_\lambda = \sum_\sigma e^{\mathcal{E}_{\text{eff}}(\sigma)}$, generally not feasible for large system sizes. Instead, it can be approximated by sampling alternately from the conditional

distributions $p_{\lambda}(\sigma|h)$ and $p_{\lambda}(h|\sigma)$. Because of the *restricted* nature of the RBM graph – there are no intra-layer connections – the conditional distributions factorize and each unit in a given layer can be sampled simultaneously (*block Gibbs sampling*). The learning signal is then calculated with the contrastive divergence algorithm [55], where the intractable part of the gradient is approximated as an average over configurations of the visible layer obtained after k steps of block Gibbs sampling (starting from random data samples). In practice, moderate values $k \sim 10$ are sufficient for training with stochastic gradient descent.

C. Noise-regularized training method

In the case where the training set is known to be corrupted by a noise process $p(\tau|\sigma)$, our goal is to learn a model $p_{\lambda}(\sigma)$ whose corresponding *noise-corrupted* distribution $\tilde{p}_{\lambda}(\tau)$ fits the observed data. We therefore define the corresponding log-likelihood cost function

$$\mathcal{L}_{\lambda} = -\frac{1}{|\mathcal{D}|} \sum_{\tau \in \mathcal{D}} \log \tilde{p}_{\lambda}(\tau) \quad (15)$$

and train the network to minimize it on each dataset. The cost gradient takes a form nearly identical to that of the standard training method (13),

$$\nabla_{\lambda} \mathcal{L}_{\lambda} = \langle \nabla_{\lambda} \mathcal{E}_{\text{eff}}(\sigma) \rangle_{p_{\lambda}(\sigma)} - \frac{1}{|\mathcal{D}|} \sum_{\tau \in \mathcal{D}} \langle \nabla_{\lambda} \mathcal{E}_{\text{eff}}(\sigma) \rangle_{\tilde{p}_{\lambda}(\sigma|\tau)} \quad (16)$$

The second term in the gradient update step is now computed not directly from the training set samples $\tau \in \mathcal{D}$, but rather from the Bayesian posterior distribution

$$\tilde{p}_{\lambda}(\sigma|\tau) = \frac{p(\tau|\sigma)p_{\lambda}(\sigma)}{\tilde{p}_{\lambda}(\tau)} \quad (17)$$

which the RBM assigns to visible states σ , given an observation τ in the noisy training set.

This alteration to the cost gradient may be viewed as a regularization of the training based on prior knowledge of the sampling process. Regularization in machine learning generally refers to techniques for improving the generalization performance of a model trained on a particular data set to new datasets drawn from the ‘ground truth’ source. A typical regularization scheme like weight decay does not specify *a priori* how the in-sample and true distributions differ, and therefore typically requires some sort of validation process – testing the model on held-out data – to select good hyperparameters. In contrast, our regularization method is applied in a context where all accessible datasets are corrupted by the same noise process. This makes validation as a means of selecting regularization hyperparameters impossible – but if the

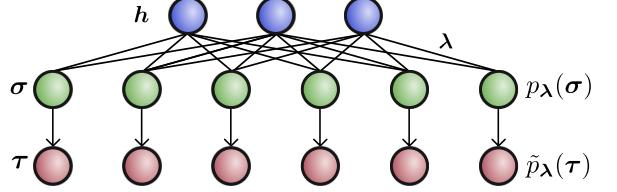


Figure 5. Three layer model. Schematic for how noise-corrupted data is modeled using a three-layer graph. The upper two layers h, σ constitute an RBM with trainable parameters λ , which defines a distribution $p_{\lambda}(\sigma)$ over the uncorrupted variables σ upon tracing out the hidden units h . The corrupted distribution is obtained through the noise process $p(\tau|\sigma)$ as $\tilde{p}_{\lambda}(\tau)$. The noise process is indicated here by arrows which link uncorrupted and corrupted variables at each site.

noise process is known, this is no obstacle as there are no free hyperparameters to select.

In applying equation (16) to the unsupervised training of an RBM, both contributions to the gradient now require computation of expectation values over marginalized distributions $p_{\lambda}(\sigma)$, $\tilde{p}_{\lambda}(\sigma|\tau)$ of the RBM, and are therefore intractable to compute exactly. As in the noise-free training case, this problem may be circumvented using the contrastive divergence method: the first term is estimated by repeated sampling from the conditional distributions $p_{\lambda}(\sigma|h)$, $p_{\lambda}(h|\sigma)$, while the second uses the same alternating sampling from the ‘data-clamped’ distributions $\tilde{p}_{\lambda}(\sigma|h, \tau)$, $\tilde{p}_{\lambda}(h|\sigma, \tau) = p_{\lambda}(h|\sigma)$. As noted above, $p_{\lambda}(\sigma|h)$, $p_{\lambda}(h|\sigma)$ are both efficiently computable due to the restricted structure of the RBM layers σ, h . Similarly, $\tilde{p}_{\lambda}(\sigma|h, \tau)$ is efficiently computable if the error probabilities satisfy a weaker condition, namely factorizing over the uncorrupted variables:

$$p(\tau|\sigma) = \prod_i p(\tau|\sigma_i) \quad (18)$$

In this case, the clamped distribution may be computed explicitly as

$$\tilde{p}_{\lambda}(\sigma|\tau, h) = \prod_i \frac{p(\tau|\sigma_i)p_{\lambda}(\sigma_i|h)}{\sum_{\sigma'_i=0,1} p(\tau|\sigma'_i)p_{\lambda}(\sigma'_i|h)} \quad (19)$$

$$= \prod_i \tilde{p}_{\lambda}(\sigma_i|\tau, h), \quad (20)$$

amenable to efficient block-Gibbs sampling.

Fig. 5 provides an intuitive way to understand the noise regularization – the corrupted variables τ may be included as a third *noise layer* appended to the standard, two-layer RBM graph, with conditional probabilities depending on the σ layer only. These can be interpreted as effective biases for the noise layer, which depend on the uncorrupted variables – for example, the independent bit-flip errors used to model our Rydberg experiment may be

written as

$$\begin{aligned} p(\boldsymbol{\tau}|\boldsymbol{\sigma}) &= \frac{1}{\tilde{Z}} e^{\tilde{b}_{\sigma} \cdot \boldsymbol{\sigma} + \tilde{b}_{\tau} \cdot \boldsymbol{\tau} + \tilde{W} \boldsymbol{\sigma} \cdot \boldsymbol{\tau}} \\ \tilde{W} &= \log \frac{p(1|1)p(0|0)}{p(1|0)p(0|1)} \\ \tilde{b}_{\sigma,i} &= \log \frac{p(0|1)}{p(0|0)} \\ \tilde{b}_{\tau,i} &= \log \frac{p(1|0)}{p(0|0)} \end{aligned}$$

For brevity, we will sometimes refer to RBMs trained with this regularization as ‘three-layer’ machines, as opposed to their ‘two-layer’ counterparts trained in the standard fashion. Similar graphical models known as Deep Belief Nets [56] have previously been used for unsupervised learning tasks, but with a different, layer-wise training algorithm that does not incorporate prior information; a gated RBM architecture similar to the three-layer machine has also been applied to Gaussian noise models in occluded images [39].

D. Sampling from trained RBMs

After an RBM has been trained, new configurations of the uncorrupted variables $\{\boldsymbol{\sigma}\}$ can be drawn from the distribution $p_{\lambda}(\boldsymbol{\sigma})$ using the block-Gibbs sampling techniques discussed above. The expectation value of a generic observable $\hat{\mathcal{O}}$ can then be approximated with a Monte Carlo average $\langle \hat{\mathcal{O}} \rangle_{\psi_{\lambda}} \simeq n_{\text{mc}}^{-1} \sum_{k=1}^{n_{\text{mc}}} \mathcal{O}_L(\boldsymbol{\sigma}_k)$, where the “local estimate” of the observable is $\mathcal{O}_L(\boldsymbol{\sigma}) = \sum_{\boldsymbol{\sigma}'} \frac{\psi_{\lambda}(\boldsymbol{\sigma}')}{\psi_{\lambda}(\boldsymbol{\sigma})} \langle \boldsymbol{\sigma}' | \hat{\mathcal{O}} | \boldsymbol{\sigma} \rangle$, and $\psi_{\lambda}(\boldsymbol{\sigma}) = \sqrt{p_{\lambda}(\boldsymbol{\sigma})}$. In the case of nontrivial noise processes, to sample from the corrupted distribution $\tilde{p}_{\lambda}(\boldsymbol{\tau})$ one may first generate an uncorrupted batch $\{\boldsymbol{\sigma}\}$ of data and then sample once from the conditional distribution $p(\boldsymbol{\tau}|\boldsymbol{\sigma})$ for each uncorrupted configuration.

For an RBM with N visible and N_h hidden units, the times for training and Monte Carlo observable estimation scale as $\mathcal{O}(NN_h)$, or in terms of the model complexity $\alpha = N_h/N$, as $\mathcal{O}(\alpha N^2)$; note that the number of visible units is fixed by the system size. Although arbitrary binary distributions require an exponentially large number of hidden units to model exactly [57], many quantum states relevant to experiment, such as ground states of paradigmatic Hamiltonians and some matrix product states, have been found to admit efficient descriptions [15, 34, 58–60]. In the present work with eight atoms, the Hilbert space is small enough that all amplitudes and expectation values can be computed exactly, providing a valuable check on our procedure. Such a benchmark quickly becomes impossible with current classical hardware when the number of atoms approaches ~ 20 for pure states, and at even smaller chain lengths for the exact evaluation of non-pure states.

III. TRAINING DETAILS

A. Methods

The reconstructions presented in this work were trained using the three-layer scheme detailed above on experimental datasets of $N \approx 3000$ samples each. Training was performed using stochastic gradient descent with a decayed learning rate, the gradients being estimated via contrastive divergence with $k = 30$ sampling steps. Since the visible layers of our machines are relatively small, exact computation of the negative-log-likelihood was possible on each set. Hyperparameters for training were therefore selected by cross-validation on a randomly chosen experimental set; the same hyperparameters were used in training on all datasets. The reconstructions presented in the text were trained on the full datasets; RBMs were also trained on 90/10 splits of each dataset in order to verify that the out-of-sample negative log likelihood did not grow during training. Error bars on reconstructed observables were computed from their variation across these training subsets in the final epochs of training. We found it beneficial to train each machine with the error rates set to zero for the first epoch.

To check that the networks learned a consistent representation of the experimental data, we performed a scaling analysis of the number of hidden units N_h of the RBM when training on experimental data. Increasing the number of hidden units, we found convergence of the observables and log-likelihood for $N_h \sim N$ (see Fig. 6 for examples). The reconstructions presented in this work used RBMs with $N_h = 2N = 16$.

B. Training on larger systems

As a test of the robustness of our reconstruction procedure, we also trained RBMs on a second set of Rydberg atom data sampled from a larger chain of $N = 9$ atoms. The dynamics of this system is governed by a master equation identical in structure to that used for modeling the eight-atom data presented in the main text, but with slightly different detuning and Rabi frequency profiles, and different effective decoherence rates.

Fig. 7 compares the results of this reconstruction to predictions of the relevant Lindbladian model, as well as experimental values where appropriate. Without alteration of the training procedure, the RBMs reconstruct quantum dynamics, as manifested in the transverse field and mutual information, in good agreement with Lindbladian predictions. This is a key benefit conferred to the experimentalist by the RBM reconstruction method. Indeed, given previous knowledge regarding the properties of the quantum state prepared in the experiment, RBM reconstruction of experimentally inaccessible observables

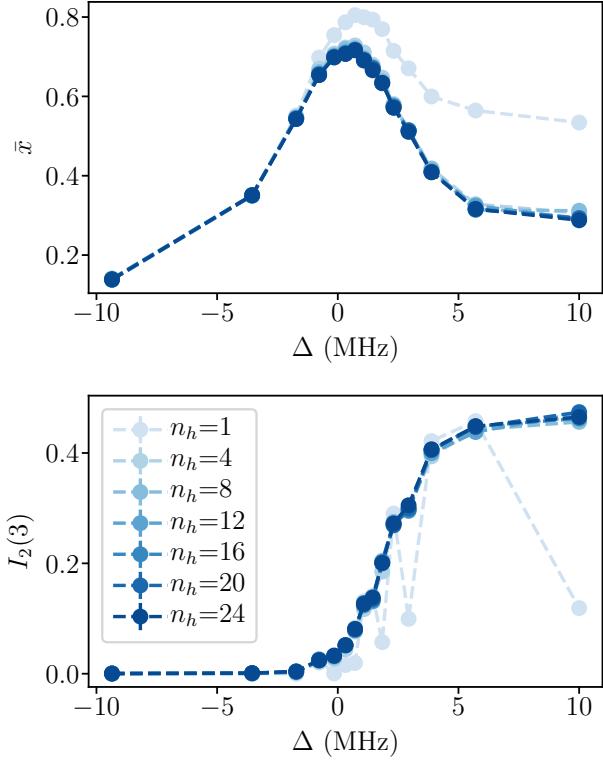


Figure 6. Examples of the scaling of observables with hidden layer size, for RBMs trained on experimental data. Top: spatially averaged transverse field values. Bottom: the Renyi mutual information at bond $s = 3$. Error bars are defined by variation of reconstructed observables in the final epochs of training.

allows for rapid and inexpensive detection of errors in state preparation and manipulation.

IV. GENERALIZATION CAPABILITIES

A generative model is of little use if it merely mimics the statistics of the training set. Successful machine learning applications are built upon the ability to *generalize* from a given dataset, extracting representations of the data that capture relevant features of the ground truth distribution from which it was sampled. This requires some structure in the data for the machine to learn, and the extent to which it succeeds in doing so depends upon the architecture of the machine as well as the size of the dataset.

For relatively small datasets such as those used in this work, it is natural to wonder whether the apparatus of machine learning is necessary at all. In particular, given

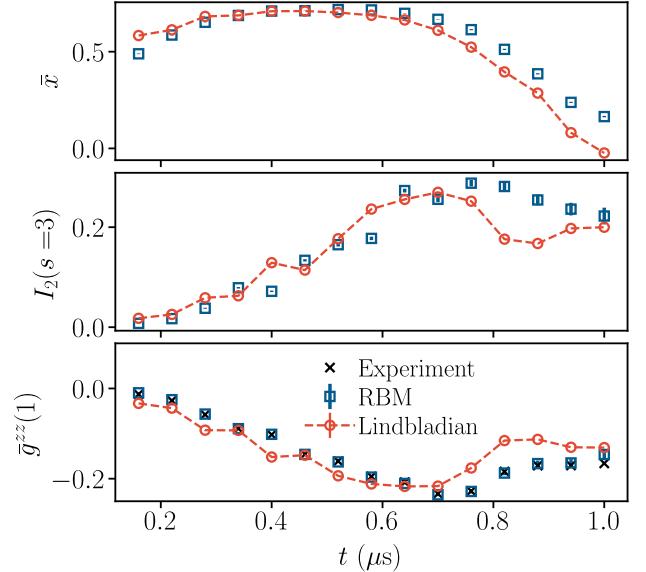


Figure 7. Some examples of observables reconstructed from nine-atom data, plotted as a function of sweep time t . From top to bottom: average transverse field \bar{x} , Renyi mutual information I_2 corresponding to a partition at bond $s = 3$; averaged nearest-neighbor correlations in the measurement basis (including same noise model as in the main text). The machines were trained with the same hyperparameters as in the eight-atom case, using $N_h = 2N = 18$ hidden units.

access to the frequency distribution (FD)

$$P_{\text{FD}}(\tau) = \frac{1}{N_s} \sum_{\tau_i \in \mathcal{D}} \delta_{\tau, \tau_i} \quad (21)$$

defined by a particular dataset \mathcal{D} consisting of N_s samples, one may define a naive *frequency distribution reconstruction* of a pure state corresponding to the data, which simply memorizes the training set:

$$|\Psi\rangle = \sum_{\tau} \sqrt{P_{\text{FD}}(\tau)} |\tau\rangle \quad (22)$$

The FD state model can be computed and stored in a time linear in the size of the dataset, by building a lookup table that associates each observed bitstring τ with its empirical probability in the dataset, and assigning probability zero to all other bitstrings. Such a model may then be used to produce Monte-Carlo estimates of desired observables, in the same fashion as for RBM states.

In general, the FD reconstruction approach cannot scale to high-entropy distributions – if H_2 is the second-order Renyi entropy of the ground-truth distribution $P_{\text{GT}}(\tau)$, the fidelity $F(P_{\text{FD}}, P_{\text{GT}}) = \sum_{\tau} \sqrt{P_{\text{FD}}(\tau)P_{\text{GT}}(\tau)}$ between the frequency distribution and the ground truth obeys the inequality

$$F(P_{\text{FD}}, P_{\text{GT}}) \leq \sqrt{N_s} e^{-H_2/4} \quad (23)$$

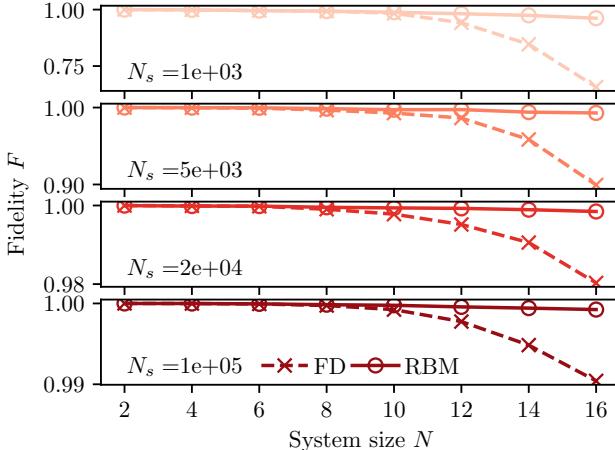


Figure 8. Generalization from ground-state datasets: fidelity improvements conferred by RBMs over frequency-distribution reconstructions, for a selection of dataset sizes N_s . Note the change in scale.

– for a proof, see Section (VII). In particular, if the measurement-basis entropy is proportional to the system size – as is the case in even some very simple states, such as a product state of spins not aligned with the measurement basis – the frequency-distribution fidelity will decay exponentially in system size. The ability to extract a modest number of *physically relevant features* is therefore essential for accurate state reconstruction from generic datasets of realistic size. However, our eight-atom system is small enough compared to the size of the datasets ($\sqrt{N_s} \sim 2^L$) that the FD approach is not *a priori* infeasible.

To quantify the performance of RBM and FD reconstructions in the small-system regime, we sampled synthetic datasets (in the occupation number basis) of size N_s up to 10^5 from ground states of the Rydberg Hamiltonian in equation (6), for a selection of system sizes up to $N = 16$ atoms. Ground state wavefunctions were computed using the QuSpin exact diagonalization package [61]; the Hamiltonian parameters were constant throughout and chosen to place the system near the phase transition into \mathbb{Z}_2 state: $V_{nn} = 30\text{MHz}$, $\Omega = 2\text{MHz}$, $\Delta \approx 1\text{MHz}$. For each dataset, we computed the fidelity F_{FD} of the frequency distribution state onto the ground-truth Rydberg wavefunction; an RBM with $N_h = N$ hidden units was then trained on the same dataset, and its fidelity F_{RBM} onto the true state was also recorded. The RBMs were all trained with the hyperparameters described in section III, but with $k = 10$ contrastive divergence steps. Fig. 8 plots the resulting fidelities achieved by both reconstructions as a function of system size – RBMs of fixed complexity achieve significantly higher fidelities for large systems, with small improvements even at $N = 8$.

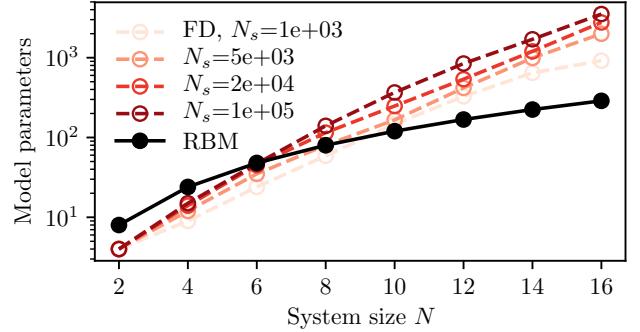


Figure 9. Dependence of model size on physical system size (note the log scale). The solid line indicates the number of parameters required to specify an RBM model with $N_h = N$. The dashed lines indicate the number of parameters required to build a lookup table for the FD model, for various dataset sizes N_s .

Another issue of practical relevance is model size: given a dataset \mathcal{D} of a particular size N_s , how many parameters are required to store each trained model? For the RBM, the number of (real-valued) parameters required to specify the model completely is determined by the size of the bias vectors and weight matrix, $N \cdot N_h + N + N_h$, and therefore quadratic in the system size for a fixed model complexity N_h/N . For the FD model, the number of parameters is determined by the size of the lookup table, i.e. the number of unique samples present in the dataset, and therefore bounded above by the dataset size N_s .

In Fig. 9, the model sizes of the RBM and FD reconstructions from Fig. 8 are compared as a function of system size; for $N \gtrsim 8$ atoms the RBMs are a significantly more efficient (not to mention more accurate) description of the quantum state.

Finally, we note that even for small systems, generative models provide an additional advantage in state reconstruction from noisy data: in the presence of measurement errors, the FD model is not representative of the ground truth for any dataset size, and simply inverting the conditional probabilities will generally result in unphysical prior distributions. Denoising methods for cleaning noisy binary datasets prior to reconstruction [62–64] may be applied, but a model training step is still required.

V. EFFECTS OF DECOHERENCE

For pure state reconstruction to be useful in near-term quantum simulators, realistic decoherence processes must be accounted for. Here, we provide a brief description of the Lindbladian master equation used in our modeling of the experiment, and discuss means of assessing the quality of pure state reconstructions in the presence of

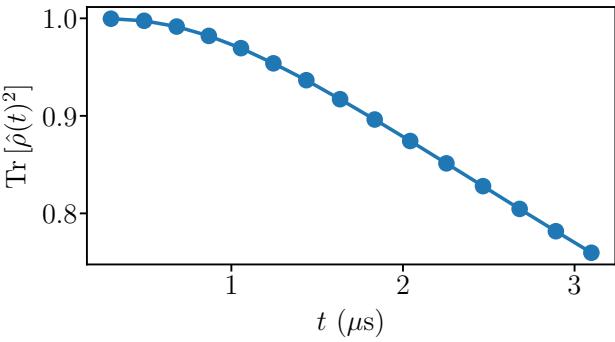


Figure 10. Purity of the master equation solutions as a function of sweep time t .

decoherence.

A. Master equation for the Rydberg machine

To account for decoherence processes quantitatively, we have used a Lindblad model, described in detail in Ref. [31], which includes two jump operators $\tilde{\sigma}_i^{rg} = |g\rangle\langle r|, \tilde{\sigma}_i^{gg} = |g\rangle\langle g|$ to represent decay and dephasing processes acting on atom i . The time evolution of the full state is given by the master equation

$$\frac{d\hat{\rho}}{dt} = -i[\hat{H}(\Omega(t), \Delta(t)) + \hat{H}_{dis}, \hat{\rho}] + \sum_{i=1}^N \sum_{t=rg,gg} \gamma_t \left(\tilde{\sigma}_i^t \hat{\rho} \tilde{\sigma}_i^{t\dagger} - \frac{1}{2} \left\{ \tilde{\sigma}_i^{t\dagger} \tilde{\sigma}_i^t, \hat{\rho} \right\} \right) \quad (24)$$

where $\hat{H}_{dis} = -\sum_{i=1}^N \delta_i \hat{n}_i$ is the static disorder Hamiltonian containing the Doppler shifts δ_i , and $\gamma_t, t = rg, gg$ are decoherence rates estimated from single-atom measurements [31] as $1/\gamma_{rg} = 80\mu\text{s}$, $1/\gamma_{gg} = 40\mu\text{s}$ respectively. The Doppler shifts δ_i were assumed to be Gaussian-distributed with an rms width of $2\pi \cdot 43.5\text{kHz}$. Direct spontaneous decay processes from the Rydberg states, which occur over longer timescales, were neglected. Numerical solutions of the master equation (24) were performed using QuTiP [65], and observables were averaged over 100 disorder realizations $\{\delta_i\}$. Uncertainties in observables were computed from the standard error of the mean of these realizations.

This master equation predicts a substantial loss in purity $\text{Tr} [\hat{\rho}^2]$ for states produced at the end of the sweep (Fig. 10), whose detrimental effects on our pure-state reconstruction process we quantify below.

B. Reconstruction fidelities

To assess the quality of quantum state reconstruction, we consider the fidelity between two states $\hat{\rho}, \hat{\sigma}$,

$$F(\hat{\rho}, \hat{\sigma}) = \text{Tr} \left[\sqrt{\sqrt{\hat{\rho}} \hat{\sigma} \sqrt{\hat{\rho}}} \right] \quad (25)$$

which reduces to the norm of the overlap in the case where $\hat{\rho}, \hat{\sigma}$ are pure states. An ideal state reconstruction $\hat{\sigma}$ of a mixed state $\hat{\rho}$ would yield $F(\hat{\rho}, \hat{\sigma}) = 1$. For pure state reconstructions $\hat{\sigma} = |\psi_\lambda\rangle\langle\psi_\lambda|$, this is not possible if the true state $\hat{\rho}$ is non-pure. However, one may still seek an approximate reconstruction which reproduces the local reduced density operators of $\hat{\rho}$. In particular, specializing to the case of one-dimensional systems, we can consider contiguous subsystems formed from s adjacent sites, $A_i^{(s)} = \{i, i+1, \dots, i+s-1\}$. Given two density operators $\hat{\rho}, \hat{\sigma}$ for the global system \mathcal{S} of size N , the reduced density operators which describe the subsystem in each state are obtained by tracing out the rest of the chain,

$$\begin{aligned} \hat{\rho}_i^{(s)} &= \text{Tr}_{\mathcal{S}/A_i^{(s)}} [\hat{\rho}] \\ \hat{\sigma}_i^{(s)} &= \text{Tr}_{\mathcal{S}/A_i^{(s)}} [\hat{\sigma}] \end{aligned}$$

Then we define a *subsystem averaged fidelity* as the spatial average of the fidelity between these local operators, over all subsystems of a particular size s :

$$F_s(\hat{\rho}, \hat{\sigma}) = \frac{1}{N+s-1} \sum_{i=1}^{N-s+1} F(\hat{\rho}_i^{(s)}, \hat{\sigma}_i^{(s)}) \quad (26)$$

$F_s(\hat{\rho}, \hat{\sigma})$ is a measure of how well, on average, $\hat{\sigma}$ is able to reproduce the s -local physics of $\hat{\rho}$.

To examine the quality of the RBM states $\hat{\rho}_\lambda = |\psi_\lambda\rangle\langle\psi_\lambda|$ in reproducing local density operators, we solved the master equation (24) for set of decay rates $\gamma_{rg} = \alpha\gamma_{rg}^{\text{exp}}, \gamma_{gg} = \alpha\gamma_{gg}^{\text{exp}}$, with $\gamma_{rg}^{\text{exp}}, \gamma_{gg}^{\text{exp}}$ denoting our estimates of the experimental values, and α a dimensionless parameterization of the overall decoherence strength. For each set of decoherence rates, the master equation was solved and synthetic data sampled from the resulting mixed states. Pure state RBMs were trained on each of these datasets, and the resulting averaged fidelities $F_s(\hat{\rho}, \hat{\rho}_\lambda)$ were computed.

As a representative example, Fig. 11 shows how the fidelities computed in the final state of the sweep vary as a function of the average Renyi entropy \bar{s}_2 of subsystems of a given size – one observes a roughly linear decay in the average fidelity with the averaged entropy. These numerical results suggest that pure state reconstruction techniques should focus on few-body operators, where the entropy build-up due to global decoherence process is limited in proportion to the system size.

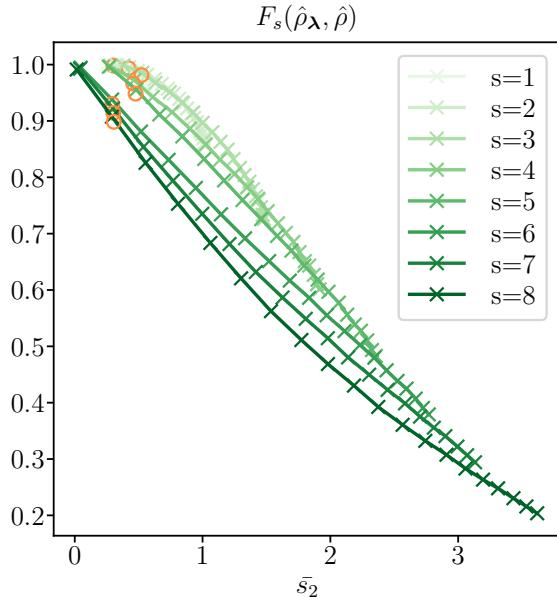


Figure 11. Average subsystem fidelities. For each subsystem size s , the average subsystem fidelity between the reconstructed state $\hat{\rho}_x$ and the state $\hat{\rho}$ from which its training data was sampled is plotted, for varying values of the decoherence rates, as quantified by the average Renyi entropy $\bar{s}_2 = -\frac{1}{N+s-1} \sum_{i=1}^{N-s+1} \log \text{Tr}(\hat{\rho}_i^{(s)2})$ of the local reduced density operators. The data plotted are for states taken at the end of the sweep, at $\Delta = 10\text{MHz}$. Open circles indicate the fidelities obtained using the decoherence rates from the experimental model presented in the main text. The fidelity behavior at other points in the sweep (not shown) is qualitatively similar.

VI. RECONSTRUCTION IMPROVEMENT FROM NOISE LAYER REGULARIZATION

Numerical experiments have demonstrated that noise layer regularization results in higher-fidelity pure state reconstruction when training on uniformly noisy data.

Fig. 12 compares fidelities achieved by regularized and unregularized RBMs, when trained on synthetic datasets subjected to the bitflip error channel described in the main text. The improvement is significant, especially in the ordered phase, where global state purity is lowest. It is important to note that in these experiments the noise process is known ahead of time and built into the three layer networks as in Fig. 5. We have also trained three-layer machines using incorrect values of the error rates on the same noisy synthetic data. Although the fidelity performance varied somewhat, depending most sensitively on $p(0|1)$, the quality of the regularized reconstructions is generally robust, and deep in the ordered phase all three-layer machines exhibited higher fidelities than their two layer counterparts on the corresponding

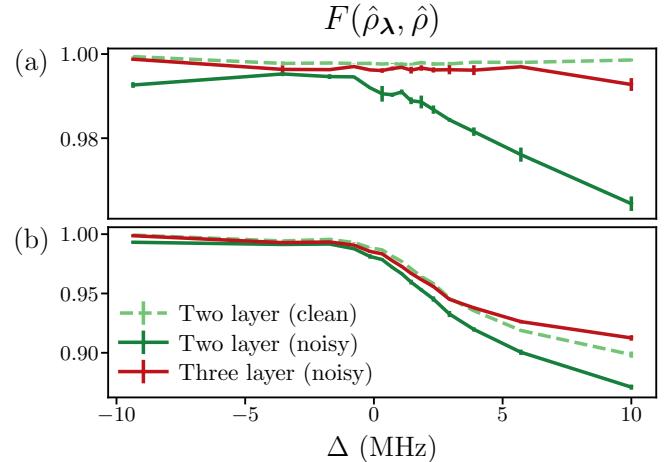


Figure 12. Fidelity improvements from noise layer regularization. As a demonstration of the efficacy of noise layer regularization, we plot the fidelity $F(\hat{\rho}_\lambda, \hat{\rho})$ obtained between the underlying state $\hat{\rho}$ and the reconstruction $\hat{\rho}_\lambda = |\psi_\lambda\rangle\langle\psi_\lambda|$, when training on synthetic data subjected to measurement errors ('noisy' data), as a function of detuning Δ . We compare regularized training (red solid lines, 'Three Layer') with unregularized training (green solid lines, 'Two Layer'), for (a) Data sampled from pure, positive Rydberg ground states, and (b) Data sampled from the mixed states $\hat{\rho}$ predicted by our Lindbladian model. As a benchmark we plot in each case the fidelity obtained by a two-layer RBM training on 'clean' data without measurement errors (green dashed lines). The regularized training leads to higher fidelities for all states sampled.

datasets, for error rates with bounds set by single-atom measurements [31]. Generically, of course, a sufficiently large mismatch between the true and assumed error rates will lead to decreased reconstruction fidelity. Future work will investigate more generally the task of selecting a regularization method for noisy quantum data.

Fig. 13 compares the predictions of these synthetically trained two- and three-layer machines (using the known noise values) for some of the observables discussed in the main text. We find that noise-layer training allows the RBMs to provide much tighter agreement in, for example, values of the transverse field and mutual information. Surprisingly, the three-layer machines actually produce poorer estimates of the transverse field correlator in the ordered phase, despite yielding two-body density operators with higher fidelities for all sampled states. A more detailed analysis of the ordered phase states reveals that regularized training does indeed produce better estimates of the one- and two-body expectation values $\langle \sigma_i^x \rangle$, $\langle \hat{\sigma}_i^x \hat{\sigma}_{i+1}^x \rangle$ for all sites i . However, at bonds (3, 4) and (5, 6), where quantum fluctuations are strongest, the three-layer improvement in the two-body expectation value is relatively small, while the reduction in the one-body expectation value is substantial. Upon computing the connected correlator $\langle \hat{\sigma}_i^x \hat{\sigma}_{i+1}^x \rangle - \langle \sigma_i^x \rangle \langle \sigma_{i+1}^x \rangle$, the overall effect is an overestimate of the true correlation.

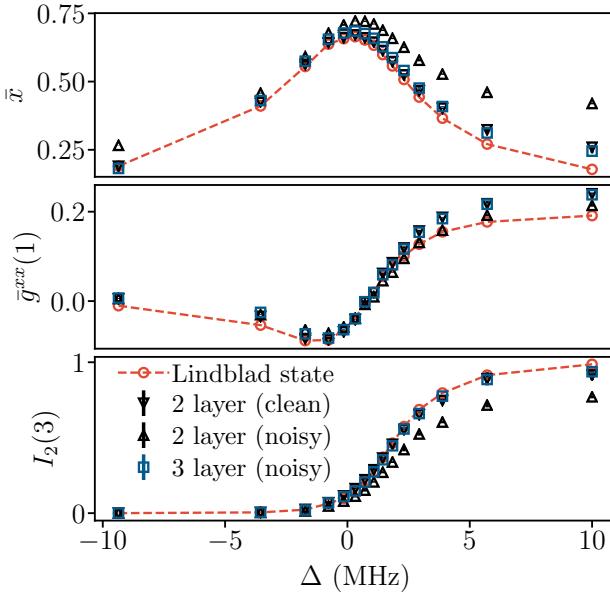


Figure 13. Observable reconstructions from synthetic data. A comparison of two- and three-layer reconstructions of the Lindbladian state when subjected to measurement errors. From top to bottom: average transverse field, average nearest-neighbor XX correlation, and Renyi mutual information at bond 3. ‘noisy’ (‘clean’) indicates training data with (without) measurement errors. Note the close agreement between the three layer machines trained on noisy data (blue squares) and the two-layer machines trained on clean data .

VII. APPENDIX: PROOF OF CLASSICAL FIDELITY BOUND

Inequality (23) is obtained by bounding the probability of the most likely outcome using the Renyi entropy H_2 in the measurement basis. By definition, $H_2 = -\log \sum_{\tau} P_{\text{GT}}(\tau)^2$, and $\sum_{\tau} P_{\text{GT}}(\tau)^2 \geq \max_{\tau} P_{\text{GT}}(\tau)^2$, so $-\log \sum_{\tau} P_{\text{GT}}(\tau)^2 \leq -2 \log \max_{\tau} P_{\text{GT}}(\tau)$. Rearranging, $\max_{\tau} P_{\text{GT}}(\tau) \leq e^{-H_2/2}$. In particular, this bounds the probability of any event in the training set, so

$$\begin{aligned} F(P_{\text{FD}}, P_{\text{GT}}) &= \sum_{\tau} \sqrt{P_{\text{FD}}(\tau) P_{\text{GT}}(\tau)} \\ &\leq \sum_{\tau} \sqrt{P_{\text{FD}}(\tau) e^{-H_2/2}} \\ &= \sqrt{N_s} e^{-H_2/4} \end{aligned}$$

VIII. APPENDIX: RENYI ENTROPY BOUND FROM POSITIVE PURE STATES

The n th order Renyi entropy of a quantum state $\hat{\rho}$ is defined as $S_n[\hat{\rho}] = \frac{1}{1-n} \log \text{Tr} \hat{\rho}^n$.

Consider a system S partitioned into subsets A and B , and a density operator $\hat{\rho}$ defined on S ; its reduced

density operator in the A subsystem is $\hat{\rho}_A = \text{Tr}_B \hat{\rho}$. Let $|i\rangle, |j\rangle$ denote orthonormal bases for A, B respectively, so that the set of product states $|i, j\rangle$ forms an orthonormal basis for the full system S . Let $p_{i,j}$ be the probability assigned by $\hat{\rho}$ to the measurement outcome i, j : $p_{i,j} = \text{Tr}(\hat{\rho}|i, j\rangle\langle i, j|)$. The positive-pure partner to the mixed state is defined as

$$|\Psi^P[\rho]\rangle = \sum_{i,j} \sqrt{p_{i,j}} |i, j\rangle, \quad (27)$$

and the corresponding reduced density operator on A is $\hat{\rho}_A^P = \text{Tr}_B |\Psi^P[\rho]\rangle\langle\Psi^P[\rho]|$.

Theorem: For $n > 1$, the Renyi entropies S_n of the two density operators satisfy the inequality

$$S_n[\hat{\rho}_A^P] \leq S_n[\hat{\rho}_A] \quad (28)$$

As a consequence, in the case of pure states $\hat{\rho}$, where the global Renyi entropy vanishes, the positive-pure partner provides a lower bound on the mutual information:

$$I_n[\hat{\rho}^P] = S_n[\hat{\rho}_A^P] + S_n[\hat{\rho}_B^P] \quad (29)$$

$$\leq S_n[\hat{\rho}_A] + S_n[\hat{\rho}_B] \quad (30)$$

$$= I_n[\hat{\rho}] \quad (31)$$

We note that for the case of the $n = 2$ Renyi entropy and pure states $\hat{\rho}$, this result has been obtained in previous work [43, 44].

Proof: Choose an auxiliary system R to purify $\hat{\rho} = \text{Tr}_R |\Psi\rangle\langle\Psi|$ for some pure state $|\Psi\rangle$ living in $S \otimes R$. If $|\alpha\rangle$ is an orthonormal basis for R , we can expand the larger pure state in the joint basis $|i, j, \alpha\rangle$: $|\Psi\rangle = \sum_{i,j,\alpha} \Psi_{i,j}^\alpha |i, j, \alpha\rangle$ for some complex coefficients $\Psi_{i,j}^\alpha$.

In terms of these amplitudes, the reduced density operator of the mixed state on A is

$$\hat{\rho}_A = \sum_{\alpha, j} \Psi_{i,j}^\alpha \Psi_{i',j}^{\alpha*} |i\rangle\langle i'| \quad (32)$$

and so

$$\text{Tr} \hat{\rho}_A^n = (\Psi_{i_1,j_1}^{\alpha_1} \Psi_{i_2,j_1}^{\alpha_1*}) (\Psi_{i_2,j_2}^{\alpha_2} \Psi_{i_3,j_2}^{\alpha_2*}) \dots (\Psi_{i_n,j_n}^{\alpha_n} \Psi_{i_1,j_n}^{\alpha_n*}) \quad (33)$$

with summation over all indices implied. The reduced density operator for positive-pure partner may be obtained from the definition above:

$$\hat{\rho}_A^P = \sum_{i,i',j} \sqrt{p_{i,j} p_{i',j}} |i\rangle\langle i'| \quad (34)$$

whence

$$\text{Tr}(\hat{\rho}_A^P)^n = (\sqrt{p_{i_1,j_1} p_{i_2,j_1}}) (\sqrt{p_{i_2,j_2} p_{i_3,j_2}}) \dots (\sqrt{p_{i_n,j_n} p_{i_1,j_n}}) \quad (35)$$

(summation implied). Furthermore,

$$p_{i,j} = \sum_{\alpha} \Psi_{i,j}^{\alpha} \Psi_{i,j}^{\alpha*} \quad (36)$$

and so by the Cauchy-Schwartz inequality,

$$\left| \sum_{\alpha} \Psi_{i,j}^{\alpha} \Psi_{i',j'}^{\alpha*} \right| \leq \sqrt{\left(\sum_{\alpha} \Psi_{i,j}^{\alpha} \Psi_{i,j}^{\alpha*} \right) \left(\sum_{\alpha'} \Psi_{i',j'}^{\alpha'} \Psi_{i',j'}^{\alpha'*} \right)} \quad (37)$$

$$= \sqrt{p_{i,j} p_{i',j'}} \quad (38)$$

Therefore, writing $\mathbf{i} = (i_1, \dots, i_n), \mathbf{j} = (j_1, \dots, j_n)$,

$$\text{Tr} \hat{\rho}_A^n = \sum_{\mathbf{i}, \mathbf{j}} \left(\sum_{\alpha_1} \Psi_{i_1, j_1}^{\alpha_1} \Psi_{i_2, j_1}^{\alpha_1*} \right) \left(\sum_{\alpha_2} \Psi_{i_2, j_2}^{\alpha_2} \Psi_{i_3, j_2}^{\alpha_2*} \right)$$

$$\dots \left(\sum_{\alpha_n} \Psi_{i_n, j_n}^{\alpha_n} \Psi_{i_1, j_n}^{\alpha_n*} \right) \quad (39)$$

$$\leq \sum_{\mathbf{i}, \mathbf{j}} \left| \sum_{\alpha_1} \Psi_{i_1, j_1}^{\alpha_1} \Psi_{i_2, j_1}^{\alpha_1*} \right| \left| \sum_{\alpha_2} \Psi_{i_2, j_2}^{\alpha_2} \Psi_{i_3, j_2}^{\alpha_2*} \right| \dots \left| \sum_{\alpha_n} \Psi_{i_n, j_n}^{\alpha_n} \Psi_{i_1, j_n}^{\alpha_n*} \right| \quad (40)$$

$$\leq \sum_{\mathbf{i}, \mathbf{j}} (\sqrt{p_{i_1, j_1} p_{i_2, j_1}}) (\sqrt{p_{i_2, j_2} p_{i_3, j_2}}) \dots (\sqrt{p_{i_n, j_n} p_{i_1, j_n}}) \quad (41)$$

$$= \text{Tr} (\hat{\rho}_A^P)^n \quad (42)$$

Hence $-\log \text{Tr} (\hat{\rho}_A^P)^n \leq -\log \text{Tr} (\hat{\rho}_A)^n$, which means that for $n > 1$, $S_n(\hat{\rho}_A^P) \leq S_n(\hat{\rho}_A)$.