communications physics

A Nature Portfolio journal

Article

https://doi.org/10.1038/s42005-025-02226-7

Recurrent neural network wave functions for Rydberg atom arrays on kagome lattice

Mohamed Hibat-Allah[1,2,3,4] ✉, Ejaaz Merali[2,3], Giacomo Torlai[5], Roger G. Melko[2,3] & Juan Carrasquilla[3,4,6]

Rydberg atom array experiments have demonstrated the ability to act as powerful quantum simulators, preparing strongly-correlated phases of matter which are challenging to study for conventional computer simulations. A key direction has been the implementation of interactions on frustrated geometries, in an effort to prepare exotic many-body states such as spin liquids and glasses. In this paper, we apply two-dimensional recurrent neural network (RNN) wave functions to study the ground states of Rydberg atom arrays on the kagome lattice. We implement an annealing scheme to find the RNN variational parameters in regions of the phase diagram where exotic phases may occur, corresponding to rough optimization landscapes. For Rydberg atom array Hamiltonians studied previously on the kagome lattice, our RNN ground states show no evidence of exotic spin liquid or emergent glassy behavior. In the latter case, we argue that the presence of a non-zero Edwards-Anderson order parameter is an artifact of the long autocorrelations times experienced with quantum Monte Carlo (QMC) simulations, and we show that autocorrelations can be systematically reduced by increasing numerical effort. This result emphasizes the utility of autoregressive models, such as RNNs, in conjunction with QMC, to explore Rydberg atom array physics on frustrated lattices and beyond.

Rydberg atom arrays have emerged as a rich playground for quantum simulation of many-body problems[1]. A key property of these arrays is their high degree of programmability, which enables the realization of multiple Hamiltonians on different lattice geometries and parameter ranges. This programmability facilitates the simulation of a wide array of phases of matter[2,3] and enables the solution to challenging combinatorial optimization problems[3–5]. The preparation of spin liquid phases—disordered phases of matter characterized by the presence of anyonic excitations, topological invariants, and long-range entanglement—has been demonstrated in programmable Rydberg arrays, potentially serving as building blocks of future generation of fault-tolerant qubits[6–8].

Recent numerical studies have investigated the physics of the ground state of Rydberg atom arrays in different lattice geometries, in particular in one[9] and two spatial dimensions in various geometries[10–16]. In lattices such as ruby and honeycomb lattices, strong numerical evidence favors the existence of a spin liquid phase in agreement with experiments[14,16]. Another recent example is the kagome lattice, where Density Matrix Renormalization Group (DMRG)[17,18] studies provided evidence that Rydberg atom arrays host a liquid-like regime[13], while Quantum Monte Carlo (QMC) simulations predicted the existence of a spin glass phase[19]. These systems display frustration arising from lattice geometry and Hamiltonian interactions, leading to the existence of a large number of quantum states with nearly degenerate energies but markedly different properties. This makes it computationally difficult to accurately approximate the ground state of these systems.

Here we focus on applying recurrent neural network (RNNs) wave functions[20,21] to a Rydberg array of atoms on the kagome lattice. The effectiveness of RNNs and Transformer language models has already been demonstrated in Rydberg atom arrays on the square lattice[22–24]. RNNs possess two key properties that make them particularly well-suited for studying frustrated systems. Firstly, their ability to perform exact sampling helps mitigate frustration-induced ergodicity issues in quantum Monte Carlo. Secondly, the ability to define them in any spatial dimension without incurring additional computational intractability helps address challenges faced by techniques like DMRG, such as the increased computational cost stemming from increased entanglement in higher dimensions[20,25].

Our findings reveal that in the highly frustrated and highly entangled regimes of the system, the RNN predicts a paramagnetic phase without

[1]Department of Applied Mathematics, University of Waterloo, Waterloo, ON, Canada. [2]Perimeter Institute for Theoretical Physics, Waterloo, ON, Canada. [3]Department of Physics and Astronomy, University of Waterloo, Waterloo, ON, Canada. [4]Vector Institute, Toronto, ON, Canada. [5]AWS Center for Quantum Computing, Pasadena, CA, USA. [6]Institute for Theoretical Physics, ETH Zürich, Zurich, Switzerland. ✉e-mail: mhibatallah@uwaterloo.ca

topological order, consistent with earlier QMC simulations[19]. However, in contrast to the QMC results in ref. 19, the RNN suggests the absence of a spin-glass phase. Nevertheless, in agreement with QMC, our numerical simulations indicate the emergence of a rugged optimization landscape, necessitating more optimization steps and thermal-like fluctuations to mitigate local minima in the RNN's parameter landscape.

Overall, our results showcase the remarkable applicability and advantages of machine learning-based wave functions, particularly RNNs, in tackling challenging problems at the forefront of Rydberg atom array physics. These findings pave the way for further exploration of exotic phases and phenomena in highly frustrated quantum systems, harnessing the power of modern machine learning techniques to advance our understanding in this field.

Methods

We focus our attention on an array of neutral atoms on the kagome lattice, interacting via laser excitation to atomic Rydberg states. We consider a lattice with periodic boundary conditions (PBC). The Hamiltonian of this system is given by refs. 1,13:

$$\hat{H} = -\sum_{i=1}^{N}\frac{\Omega}{2}\left(|g\rangle_i\langle r|_i + |r\rangle_i\langle g|_i\right) - \delta\sum_{i=1}^{N}|r\rangle_i\langle r|_i + \frac{1}{2}\sum_{i,j}V(||\boldsymbol{x}_i - \boldsymbol{x}_j||)|r\rangle_i\langle r|_i \otimes |r\rangle_j\langle r|_j.$$

Here $|g\rangle_i, |r\rangle_i$ are respectively the ground and excited states of the Rydberg atom i. Ω is the Rabi frequency and δ is the laser detuning. $V(R) = C/R^6$ is the repulsive potential due to the dipole-dipole interaction between Rydberg atoms, which is responsible for the blockade mechanism[1]. In practice, we define a blockade radius R_b such that $V(R_b/a) = \Omega$, where a is the distance between two neighboring Rydberg atoms. Finally, we note that the sum over all possible pairs is truncated to a sum over neighbors separated by a distance cutoff $R_c = 2$ or $R_c = 4$. The choice $R_c = 2$ is taken to compare with the DMRG results reported in ref. 13 as well as with the QMC findings in ref. 19

Two dimensional RNNs

The Rydberg Hamiltonian is stoquastic in nature[26], which implies that the ground-state wave function contains only positive amplitudes. This offers the opportunity to model the ground state with an RNN wave function with only positive amplitudes[20] which we adopt below. Complex extensions of RNN wave functions for non-stoquastic Hamiltonians have been explored in refs. 20,21,25. To model a positive RNN wave function, we can express our ansatz in the computational basis as:

$$\Psi_{\boldsymbol{\theta}}(\boldsymbol{\sigma}) = \sqrt{p_{\boldsymbol{\theta}}(\boldsymbol{\sigma})},$$

such that $\boldsymbol{\theta}$ corresponds to the variational parameters of the ansatz $|\Psi_{\boldsymbol{\theta}}\rangle$, and $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ is a configuration of the Rydberg atoms. The main advantage of using RNN wave functions is the possibility of estimating observables through autoregressive sampling, which allows obtaining uncorrelated samples by construction[20]. To do so, we model the joint probability $p_{\boldsymbol{\theta}}(\boldsymbol{\sigma})$ by constructing the conditionals $p_{\boldsymbol{\theta}}(\sigma_i|\sigma_{<i})$ by taking advantage of the probability chain rule

$$p_{\boldsymbol{\theta}}(\boldsymbol{\sigma}) = p_{\boldsymbol{\theta}}(\sigma_1)p_{\boldsymbol{\theta}}(\sigma_2|\sigma_1)\cdots p_{\boldsymbol{\theta}}(\sigma_N|\sigma_{N-1},\ldots,\sigma_2,\sigma_1).$$

These conditional probabilities are obtained through a Softmax layer as follows:

$$p_{\boldsymbol{\theta}}(\sigma_i|\sigma_{<i}) = \boldsymbol{y}_i \cdot \boldsymbol{\sigma}_i.$$

Here $\boldsymbol{y}_i = \mathrm{Softmax}(U\boldsymbol{h}_i + \boldsymbol{c})$ where U and $\boldsymbol{c}$ are, respectively, trainable weights and biases, and 'Softmax' corresponds to the normalizing Softmax activation function. Additionally, the memory (hidden) state $\boldsymbol{h}_i$ is obtained recursively as ref. 27:

$$\boldsymbol{h}_i = f(W[\boldsymbol{h}_{i-1};\boldsymbol{\sigma}_{i-1}] + \boldsymbol{b}), \quad (1)$$

such that $[.\,;.\,]$ is a concatenation of two vectors, while $\boldsymbol{\sigma}_{i-1}$ is a one-hot encoding of σ_{i-1}. These computations are illustrated in Fig. 1a. W and $\boldsymbol{b}$ are also trainable weights and biases, and f is a user-defined activation function.

By virtue of the 'Softmax' activation function, the conditionals $p_{\boldsymbol{\theta}}(\sigma_i|\sigma_{<i})$ are normalized to one. This property implies that the RNN joint probability $p_{\boldsymbol{\theta}}$ is also normalized[20]. Furthermore, by sampling the conditionals $p_{\boldsymbol{\theta}}(\sigma_i|\sigma_{<i})$ sequentially, as illustrated in Fig. 1b, we can extract exact samples from the joint RNN probability $p_{\boldsymbol{\theta}}$. An attractive property of this scheme is the possibility to efficiently generate uncorrelated samples from different modes present in $p_{\boldsymbol{\theta}}$, whereas traditional Metropolis sampling scheme may get stuck in only one mode.

The atom configurations of a Rydberg atom array on a kagome lattice can be seen as an $L \times L \times 3$ array of binary degrees of freedom where L is the size of each side of the lattice. As illustrated in Fig. 1c, we can map our kagome lattice with a local Hilbert space of 2 to a square lattice with an enlarged Hilbert space of size $2^3 = 8$ which we can study using our two-dimensional (2D) RNN wave function[23,28].

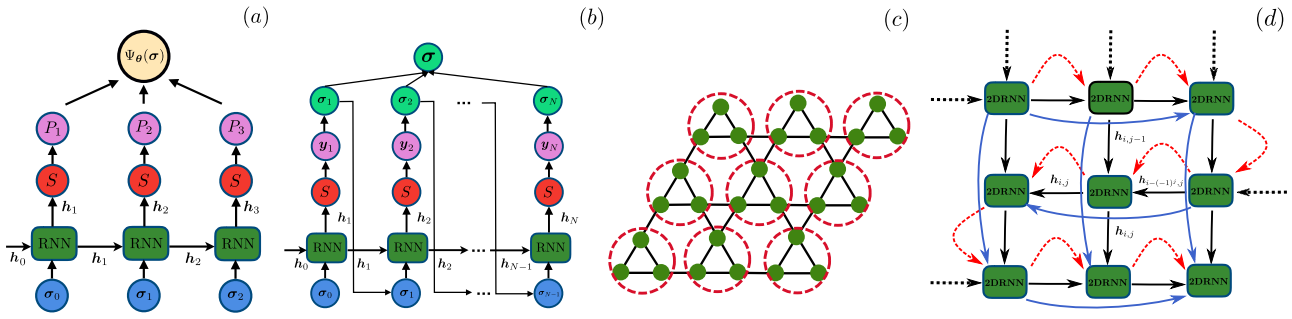

Fig. 1 | RNN wave functions architecture details. a An illustration of a positive RNN wave function. Each RNN cell receives an input $\boldsymbol{\sigma}_{n-1}$ and a hidden state $\boldsymbol{h}_{n-1}$ and outputs a new hidden state $\boldsymbol{h}_n$. This vector is taken as an input to the Softmax layer (denoted S) that computes the conditional probability P_i. **b** RNN autoregressive sampling scheme: after obtaining the probability vector $\boldsymbol{y}_i$ from the Softmax layer (S) in step i, we sample it to produce $\boldsymbol{\sigma}_i$. The latter is taken again as an input to the RNN along with the hidden state $\boldsymbol{h}_i$ to sample the following degree of freedom $\boldsymbol{\sigma}_{i+1}$. **c** Mapping of a Kagome lattice to a square lattice by embedding three atoms in a larger local Hilbert space. **d** A two-dimensional (2D) RNN with periodic boundary conditions for a 3×3 lattice for illustration purposes. A bulk RNN cell receives two hidden states $\boldsymbol{h}_{i,j-1}$ and $\boldsymbol{h}_{i-(-1)^j,j}$, as well as two input vectors $\boldsymbol{\sigma}_{i,j-1}$ and $\boldsymbol{\sigma}_{i-(-1)^j,j}$ (not shown) as illustrated by the black solid arrows. RNN cells at the boundary receive additional hidden states $\boldsymbol{h}_{i,j+1}$ and $\boldsymbol{h}_{i+(-1)^j,j}$, as well as two input vectors $\boldsymbol{\sigma}_{i,j+1}$ and $\boldsymbol{\sigma}_{i+(-1)^j,j}$ (not shown), as demonstrated by the blue curved and solid arrows. The sampling path is taken as a zigzag path, as demonstrated by the dashed red arrows. The initial memory states of the 2D RNN and the initial inputs are null vectors, as indicated by the dashed black arrows.

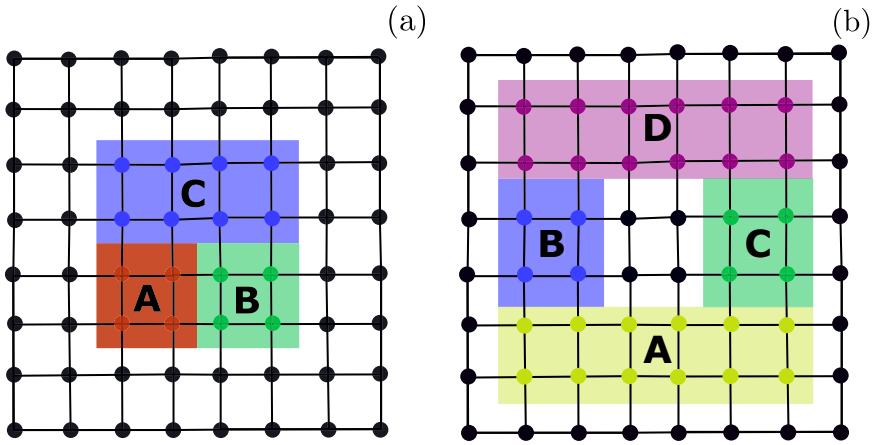

Fig. 2 | Topological entanglement entropy constructions. a A sketch of the parts A, B, and C that we use for Kitaev-Preskill construction to compute the TEE. **b** Levin-Wen construction using the regions A, B, C and D. For the Rydberg atoms Hamiltonian on a kagome lattice, each dot on the square lattice corresponds to a block of three binary degrees of freedom, as shown in Fig. 1c.

To construct a 2D RNN ansatz that can handle PBC, we modify our RNN recursion in Eq. (1) to a two-dimensional recursion relation as:

$$\boldsymbol{h}_{i,j} = f\left(W[\text{Neighbours}\,(\boldsymbol{h}_{i,j});\ \text{Neighbours}\,(\boldsymbol{\sigma}_{i,j})] + \boldsymbol{b}\right). \tag{2}$$

$\boldsymbol{h}_{i,j}$ is a memory state with two indices for each atom in the two-dimensional lattice. Here 'Neighbours($\boldsymbol{\sigma}_{i,j}$)' returns a concatenation of the neighbors of $\boldsymbol{\sigma}_{i,j}$. The same observation goes for 'Neighbours($\boldsymbol{h}_{i,j}$)'. These neighbors correspond to incoming vectors indicated by the black and blue arrows as illustrated in Fig. 1d. More specifically, we define

$$\text{Neighbours}\,(\boldsymbol{h}_{i,j}) \equiv [\boldsymbol{h}_{i-(-1)^j,j}; \boldsymbol{h}_{i,j-1}; \boldsymbol{0}; \boldsymbol{0}]$$

on the bulk. On the boundaries, we take

$$\text{Neighbours}\,(\boldsymbol{h}_{i,j}) \equiv \left[\boldsymbol{h}_{i-(-1)^j,j}; \boldsymbol{h}_{i,j-1}; \boldsymbol{h}_{i+(-1)^j,j}; \boldsymbol{h}_{i,j+1}\right].$$

Note that PBC on the indices is assumed. The additional inputs $\boldsymbol{\sigma}_{i+(-1)^j,j}$, $\boldsymbol{\sigma}_{i,j+1}$ and hidden states $\boldsymbol{h}_{i+(-1)^j,j}$, $\boldsymbol{h}_{ij+1}$ allow to take PBC into account and to introduce correlations between degrees of freedom at the boundaries. During the autoregressive sampling procedure, the input and hidden vectors are initialized to a null vector, if not previously defined, to preserve the autoregressive nature of our scheme, as illustrated in Fig. 1b. Also, note that the particular choice of the indices is motivated by the zigzag sampling path. In this study, we use an advanced version of 2D RNNs incorporating the gating mechanism as previously done in refs. 25,29,30. More details can be found in Supplementary Note 1. Finally, since $\boldsymbol{h}_{i,j}$ is a summary of the history of the generated $\sigma_{<i,j}$, it is used to compute the conditional probabilities as follows:

$$p_{\boldsymbol{\theta}}(\sigma_{i,j}|\sigma_{<i,j}) = \text{Softmax}\,(U\boldsymbol{h}_{i,j} + \boldsymbol{c}) \cdot \boldsymbol{\sigma}_{i,j}. \tag{3}$$

Supplementing RNN optimization with annealing

To reach the ground state of the Rydberg atoms array Hamiltonian on the kagome lattice, we minimize the energy expectation value $E_{\boldsymbol{\theta}} = \langle\Psi_{\boldsymbol{\theta}}|\hat{H}|\Psi_{\boldsymbol{\theta}}\rangle$ using the Variational Monte Carlo (VMC) scheme[31] (see Supplementary Note 2). Due to the frustrated nature of the kagome lattice which can induce local minima in the VMC scheme, we leverage annealing with thermal-like fluctuations to mitigate local minima. This technique has been suggested and implemented in refs. 21,25,28,32–34. In this case, we obtain a free-energy like cost function, defined as

$$F_{\boldsymbol{\theta}}(n) = E_{\boldsymbol{\theta}} - T(n)S_{\text{classical}}(p_{\boldsymbol{\theta}}), \tag{4}$$

where $F_{\boldsymbol{\theta}}$ is a variational pseudo Free energy and $S_{\text{classical}}$ is the classical Shannon entropy:

$$S_{\text{classical}}(p_{\boldsymbol{\theta}}) = -\sum_{\boldsymbol{\sigma}} p_{\boldsymbol{\theta}}(\boldsymbol{\sigma})\log\left(p_{\boldsymbol{\theta}}(\boldsymbol{\sigma})\right). \tag{5}$$

The previous sum goes over all classical Rydberg configurations $\{\boldsymbol{\sigma}\}$ in the computational z-basis. Note that $S_{\text{classical}}$ is a pseudo-entropy that can be efficiently estimated using our RNN wave function as opposed to the quantum von Neumann entropy. Additionally, $T(n)$ is a pseudo-temperature that is annealed from some initial value T_0 to zero as follows: $T(n) = T_0(1 - n/N_{\text{a}})$ where $n \in [0,\ N_{\text{a}}]$ and N_{a} is the total number of annealing steps. We note that for each annealing step, we train our RNN for 5 training steps. We present more details about the hyperparameters of our training scheme in Supplementary Note 3.

Topological entanglement entropy

To investigate the existence of a topological property in the Rydberg atom arrays on the kagome lattice, we compute the topological entanglement entropy (TEE)[35–41]. For a gapped phase of matter, where the area law is satisfied, the Renyi-2 entanglement entropy follows the scaling law $S_2(A) = aL - \gamma + \mathcal{O}(L^{-1})$, assuming A and B is partition of the system, L is the size of the boundary between A and B and $S_2(A) \equiv -\log(\text{Tr}\,(\rho_A^2))$. In this case, γ is the so-called TEE. In this paper, we use the swap trick with our RNN wave function ansatz[20,42,43] to calculate the second Renyi entropy S_2 to extract the TEE γ.

We extract γ using two different strategies, namely the Kitaev-Preskill construction[38] and the Levin-Wen construction[37], illustrated in Fig. 2.

The Kitaev-Preskill construction consists of choosing three subregions A, B, C with geometries as shown in Fig. 2a. The TEE can be then obtained by computing

$$\begin{aligned}\gamma \quad &= -S_2(A) - S_2(B) - S_2(C) + S_2(AB)\\ &\quad + S_2(AC) + S_2(BC) - S_2(ABC),\end{aligned}$$

where $S_2(A)$ is the second Renyi entropy of the subsystem A, and AB is the union of A and B and similarly for the other terms. It is worth mentioning that finite size effects on γ can be reduced by extrapolating the size of the subregions[38,44]. Finally, note that this approach combined with RNN wave functions was successful in extracting a non-zero TEE on the toric code and the hard-core Bose-Hubbard model on the kagome lattice[28].

The Levin-Wen construction allows to extract the TEE γ by constructing four different subsystems $A_1 = A \cup B \cup C \cup D$, $A_2 = A \cup C \cup D$, $A_3 = A \cup B \cup D$ and $A_4 = A \cup D$ as illustrated in Fig. 2b such that[40]:

$$\gamma = \frac{-S_2(A_1) + S_2(A_2) + S_2(A_3) - S_2(A_4)}{2}.$$

Note that finite size effects on γ can be reduced by extrapolating the width and thickness of A_1, A_2, A_3 and A_4[40,44].

Finally, we would like to highlight that our ability to study quantum systems with fully periodic boundary conditions is key to mitigating boundary effects, as opposed to cylinders used in DMRG[45,46], which can introduce boundary effects in the TEE value[13].

Results

According to the RNN numerics, our results show that the ground state at $R_b = 1.95$ and $\delta = 3.3$, which is suggested to be in the spin-liquid phase according to ref. 13, is rather a disordered state with no topological order. We first plot the correlations $\langle n_0 n_r \rangle$ in Fig. 3a. The results indicate that the extracted state has short-range correlations. To confirm the correctness of our variational implementation, we perform a sanity check and compare our ground state energies with QMC and DMRG as shown in Supplementary Note 4. We found a good agreement between our RNN energies and QMC as well as DMRG energies. Most importantly, we observe that our RNN results using only $d_h = 100$ are more accurate compared to DMRG with a bond dimension $\chi = 1000$ in the highly entangled regime at $R_b = 1.95$ and $\delta = 3.3$.

To investigate the existence of a spin liquid in this regime, we calculate the TEE γ using the Kitaev-Preskill construction[38] for a system size $L = 8$ (see Fig. 2a), and for different values of $\delta \in [2.0, 3.7]$ and $R_c = 2$, 4 at $R_b = 1.95$. We also do the same using the Levin-Wen construction[40] in Fig. 2b. Our results, illustrated in Fig. 3b suggest that the TEE extracted by the RNN is consistent with zero and different from ln(2) within error bars. These results suggest the non-existence of a spin liquid within our settings and also suggest that the state we find in this regime is a disordered state. Our findings are further corroborated by a recent QMC study[19] and also by previous results in the literature suggesting that the paramagnetic 'liquid' phase in Ising systems on the kagome lattice is not exotic[47–49].

To address the finite-size scaling of the TEE, we compute the TEE at $R_b = 1.95$, $\delta = 3.3$ for $N = 10 \times 10 \times 3$ by pre-training from the RNN optimized at $L = 8$ and also for $L = 12$ by starting from the optimized parameters at $L = 10$. We follow a similar KP construction to Fig. 2 of the regions A, B and C where the size of each subregion is given as $3 \times 3 \times 3$, $3 \times 3 \times 3$, and $6 \times 3 \times 3$ for $L = 6$ respectively, and $4 \times 4 \times 3$, $4 \times 4 \times 3$, and $8 \times 4 \times 3$ for $L = 8$. Our estimates $\gamma_{\mathrm{RNN}} = 0.10 \pm 0.26$, -0.1 ± 0.22 for $L = 6$ and $L = 8$ respectively. These values, which are consistent with zero TEE, corroborate the absence of a Z_2 spin-liquid according to the RNN variational calculations. Note that the error bars can be systematically reduced by increasing the number of samples in the swap trick calculations[20,28].

In this QMC study[19], it was suggested that the region, around $R_b = 1.95$ and the values of δ used in our study contains an emergent spin-glass phase instead of a paramagnetic state. To verify this claim, we compute the Edwards-Anderson (EA) order parameters[50,51], defined as:

$$q_{\mathrm{EA}} = \frac{\sum_{i=1}^{N} \langle n_i - \rho \rangle^2}{N\rho(1-\rho)}, \tag{6}$$

where N is the system size, n_i is the occupation number of site i and $\rho = (\sum_{i=1}^{N} n_i)/N$. Deviations of this order parameter from zero values are signals of the existence of a spin-glass phase. In Fig. 3c, we plot this order parameter as a function of δ with $R_c = 2$, 4 and $R_b = 1.95$. We find that the values of the order parameter are consistent with zero, as opposed to the results of QMC in ref. 19. Furthermore, we report in Fig. 3d the density-density overlap $\frac{1}{N}\sum_{i=1}^{N} n_i^{(1)} n_i^{(2)}$ and the spin-spin overlap $\frac{1}{N}\sum_{i=1}^{N} S_i^{(1)} S_i^{(2)}$ between different RNN samples at $R_b = 1.95$, $\delta = 3.3$, and $R_c = 2$. Here labels (1) and (2) correspond to two independent sets of samples, which are obtained from optimized RNNs with 10 different training seeds. The Gaussian nature of the overlap distribution in both representations is another indicator that there is no static signature of a spin-glass order[52].

The discrepancy in our results and previous QMC findings[19] could be related to emergent glassy dynamics in the QMC simulations, which results in very long auto-correlations times and thus in a non-ergodic behavior. To corroborate our findings, we run QMC simulations[53], based on Stochastic Series Expansion (SSE)[54,55], for larger inverse temperatures compared to ref. 19, namely for $\beta \geq 200$ and using 2.2×10^6 Monte Carlo samples. We find that the QMC prediction for the EA order parameter is given as $q_{\mathrm{EA}}^{\mathrm{QMC}} = 0.0000018(5)$ for $R_b = 1.95$, $\delta = 3.3$, a system size $8 \times 8 \times 3$, and for a radius cutoff $R_c = 2$. The previous result agrees very well with our RNN findings in Fig. 3c. This result is also confirmed by the good agreement

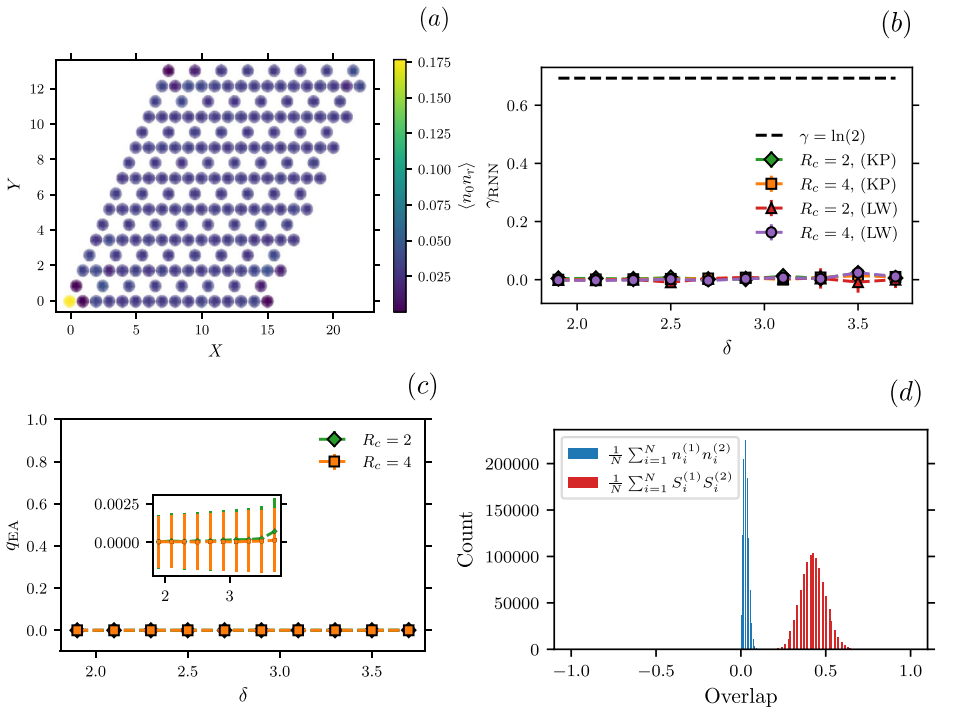

Fig. 3 | Numerical results of 2DRNNs applied to Rydberg atom arrays on kagome lattice. In these panels, we focus on the Blockade radius $R_b = 1.95$. **a** Plot of two point correlations $\langle n_0 n_r \rangle$ with $\delta = 3.3$ and for a system size $N = 8 \times 8 \times 3$ and $R_c = 2$. **b** Plots of the topological entanglement entropy versus δ for two different values of the cutoff radius R_c, using the Levin-Wen (LW) construction and the Kitaev-Preskill (KP) construction, for $N = 8 \times 8 \times 3$. **c** A histogram of the Edwards-Anderson order parameter q_{EA} defined in Eq. (6) as a function of δ for $N = 8 \times 8 \times 3$. The inset provides a zoomed-in view close to zero. **d** A plot of the density overlaps $\frac{1}{N}\sum_{i=1}^{N} n_i^{(1)} n_i^{(2)}$ and the spin overlaps $\frac{1}{N}\sum_{i=1}^{N} S_i^{(1)} S_i^{(2)}$ at $\delta = 3.3$. Here $S_i = 2n_i - 1$, and (1) and (2) are labels for two sets of samples obtained from our optimized RNN, that are aggregated from 10 different training seeds, for $N = 6 \times 6 \times 3$ and $R_c = 2$. For each seed, we generate 2×10^5 independent samples and divide them into two sets. The error bars indicate the statistical uncertainty of one standard deviation, calculated across different samples. In this and other plots throughout the Article, error bars may be hidden if they are smaller than the symbol size.

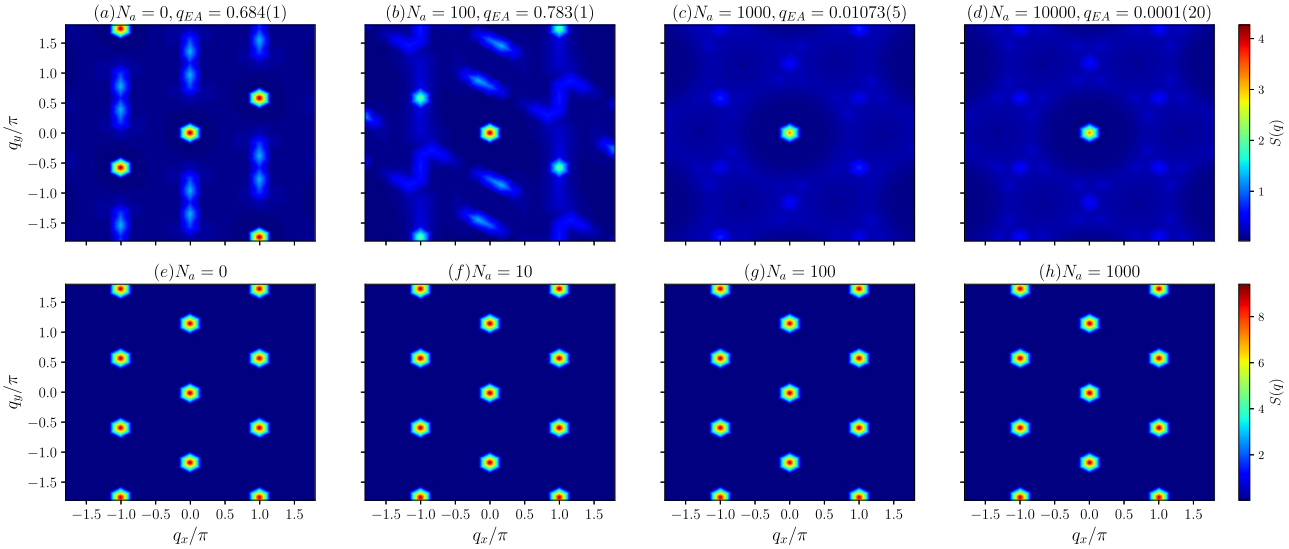

Fig. 4 | Structure factor evolution with the number of annealing steps. We illustrate the structure factor for different values of the number of annealing steps N_a at $R_b = 1.95$, $\delta = 3.3$ in (**a**–**d**) and at $R_b = 1.7$, $\delta = 3.3$ in (**e**–**h**). The color bars represent the magnitude of the structure factor $S(\boldsymbol{q})$. Additionally, these results correspond to a system size $N = 6 \times 6 \times 3$. Note that we observe convergence to a paramagnetic state without ordering peaks beyond $N_a = 1000$ for $R_b = 1.95$, $\delta = 3.3$. For $R_b = 1.7$, $\delta = 3.3$, we find that the nematic state is not affected by the choice of N_a and can be obtained without a need for annealing.

between the RNN energies and the QMC energies as shown in Supplementary Note 4. Our findings are further supported by the results of ref. 56, which suggests the possibility of transition in a quantum dimer model between nematic to paramagnetic to staggered states. In conclusion, our numerical investigation suggests that the long auto-correlation time could be a limiting factor in the QMC results reported in ref. 19.

We note that the emergence of a long autocorrelation time in QMC coincides with the emergence of a rugged optimization landscape, which in our simulations implies a longer number of annealing steps in our RNN simulations to achieve convergence. To demonstrate this point, we compute the structure factor

$$S(\boldsymbol{q}) = \frac{1}{N}\sum_{i,j}\langle n_i n_j\rangle\, e^{i\boldsymbol{q}\cdot(\boldsymbol{x}_i-\boldsymbol{x}_j)} \tag{7}$$

to extract the nature of the states obtained by our RNN ansatz and investigate their dependence on the number of annealing steps N_a. We expect the optimization landscape to be rougher as N_a required to converge increases. Figure 4a–d at $R_b = 1.95$ and $\delta = 3.3$ show that the RNN finds different states for different numbers of annealing steps N_a, until it converges to a state without ordering peaks, i.e., the paramagnetic state. In contrast, the nematic state at $R_b = 1.7$, $\delta = 3.3$ can be reached without the need for annealing, as illustrated by the structure factors at different N_a in Fig. 4e–h. These observations suggest an emergent rugged optimization landscape when optimizing our ansatz in the highly entangled regime. Finally, to find the optimal number of annealing steps N_a in the highly entangled regime, we note that we conduct a scaling study as shown in Supplementary Note 5.

To indicate the quality of our variational calculations, we use the v-score as a metric[57], which we report in Supplementary Note 6. Furthermore, to investigate the effect of parameter sharing in the RNN, illustrated in Eqs. (2), (3), we optimized our RNN using site-dependent parameters at $R_b = 1.95$, $\delta = 3.3$ and for $L = 8$ and we find that both TEE and EA order parameters are consistent with zero within errors bars. This result confirms that parameter-sharing in our RNN ansatz does not bias our findings. More details are shared in Supplementary Note 7.

Finally, we report a comparison between RNN and QMC runtimes in Supplementary Note 8. In particular, we find a significant speed-up when training RNNs on an A100 GPU and QMC on a single CPU. This speedup is enabled by GPU hardware in addition to the ability of the 2DRNN to do transfer learning from smaller lattices to larger ones[21,25,28,30,58,59].

Conclusion

In this paper, we demonstrate a successful application of recurrent neural network (RNN) wave functions to the task of investigating topological order on Rydberg atom arrays on kagome lattice. We use these architectures to estimate the second Renyi entropies using the swap trick[20]. The latter allows us to compute the TEEs using the Kitaev-Preskill[38] and the Levin-Wen[37] constructions. Furthermore, with the possibility of handling periodic boundary conditions in RNNs, the boundary effects on the TEE are reduced compared to DMRG, which has challenges with boundary effects on cylinders[46].

Our main finding, suggested by the two-dimensional RNN wave functions results, points out that Rydberg atom arrays on the kagome lattice do not establish a Z_2 spin liquid in the highly entangled regime. This observation is also consistent with previous QMC studies[19]. Our RNN numerics also suggest that the highly entangled region corresponds to a trivial paramagnetic state and that there is no signature for spin glass order as opposed to the observations outlined in ref. 19. We believe that the ability of RNNs to generate uncorrelated samples from a multimodal distribution is a crucial factor for our numerics to indicate the non-existence of the spin-glass phase. Additionally, we find that autocorrelation could be the main factor behind the spin glass phase observed in previous QMC simulations[19]. In particular, our QMC numerics with more numerical effort compared to ref. 19. suggest the absence of a spin glass phase. Furthermore, supplementing RNNs with annealing turns out to be a valuable tool for mitigating local minima induced by the frustrated nature of the kagome lattice in the highly entangled regime. We highlight that advanced optimization schemes, such as minimal Stochastic Reconfiguration (minSR)[60,61], is a potential avenue for enhancing the optimization of our 2DRNN ansätzes. Refs. 62,63 reported that Stochastic Reconfiguration (SR) is not as effective as Adam optimizer when applied to RNN wave functions. As a result, we believe that reconciling SR and RNN wave function optimization is an interesting research direction that deserves a thorough study in the future. Furthermore, studying dynamic properties of Rydberg atoms arrays is another promising research direction that can corroborate our findings by performing time-evolution on our RNN ansatz[64,65].

Finally, we note that our method can be generalized to study other systems with potential topological order, such as the Rydberg atom arrays on the Ruby lattice[14,66,67]. One could also use quantum state tomography

with RNNs[68] in a wide variety of quantum simulators and also combine data from QMC or quantum simulators with VMC to improve the variational results[22,24,69,70]. We also believe in the potential of RNN wave functions ansätzes in the discovery of new phases of matter with topological order. Overall, these results highlight the promising future of RNN wave functions[20,21], language-model based wave functions, and neural quantum states[71], in general, for investigating open questions and discovering new physics within the condensed matter community and beyond.

Data availability

The data generated in this study is available from the corresponding author upon reasonable request.

Code availability

Our code is made publicly available at "http://github.com/mhibatallah/RNNWavefunctions". The hyperparameters we use are given in Supplementary Note 3.

Received: 21 October 2024; Accepted: 10 July 2025;
Published online: 25 July 2025

References

1. Browaeys, A. & Lahaye, T. Many-body physics with individually controlled Rydberg atoms. *Nat. Phys.* **16**, 132–142 (2020).
2. Ebadi, S. et al. Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature* **595**, 227–232 (2021).
3. Wurtz, J. et al. Aquila: Quera's 256-qubit neutral-atom quantum computer. https://doi.org/10.48550/arXiv.2306.11727 (2023).
4. Ebadi, S. et al. Quantum optimization of maximum independent set using Rydberg atom arrays. *Science* **376**, 1209–1215 (2022).
5. Nguyen, M.-T. et al. Quantum optimization with arbitrary connectivity using Rydberg atom arrays. *PRX Quantum* **4**, 010316 (2023).
6. Dennis, E., Kitaev, A., Landahl, A. & Preskill, J. Topological quantum memory. *J. Math. Phys.* **43**, 4452–4505 (2002).
7. Kitaev, A. Anyons in an exactly solved model and beyond. *Ann. Phys.* **321**, 2–111 (2006).
8. Kitaev, A. & Laumann, C. Topological phases and quantum computation. https://doi.org/10.48550/arXiv.0904.2771 (2009).
9. Samajdar, R., Choi, S., Pichler, H., Lukin, M. D. & Sachdev, S. Numerical study of the chiral $\mathbb{Z}_3$ quantum phase transition in one spatial dimension. *Phys. Rev. A* **98**, 023614 (2018).
10. Samajdar, R., Ho, W. W., Pichler, H., Lukin, M. D. & Sachdev, S. Complex density wave orders and quantum phase transitions in a model of square-lattice Rydberg atom arrays. *Phys. Rev. Lett.* **124**, 103601 (2020).
11. Kalinowski, M. et al. Bulk and boundary quantum phase transitions in a square Rydberg atom array. *Phys. Rev. B* **105**, 174417 (2022).
12. Li, C.-X., Yang, S. & Xu, J.-B. Quantum phases of Rydberg atoms on a frustrated triangular-lattice array. *Opt. Lett.* **47**, 1093–1096 (2022).
13. Samajdar, R., Ho, W. W., Pichler, H., Lukin, M. D. & Sachdev, S. Quantum phases of Rydberg atoms on a kagome lattice. *Proc. Natl Acad. Sci. USA* **118**, e2015785118 (2021).
14. Verresen, R., Lukin, M. D. & Vishwanath, A. Prediction of toric code topological order from Rydberg blockade. *Phys. Rev. X* **11**, 031005 (2021).
15. Yang, S. & Xu, J.-B. Density-wave-ordered phases of rydberg atoms on a honeycomb lattice. *Phys. Rev. E* **106**, 034121 (2022).
16. Kornjača, M. et al. Trimer quantum spin liquid in a honeycomb array of Rydberg atoms. *Commun. Phys.* **6**, 358 (2023).
17. White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **69**, 2863–2866 (1992).
18. Schollwöck, U. The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.* **326**, 96–192 (2011).
19. Yan, Z., Wang, Y.-C., Samajdar, R., Sachdev, S. & Meng, Z. Y. Emergent glassy behavior in a kagome Rydberg atom array. *Phys. Rev. Lett*. **130**, 206501 (2023).
20. Hibat-Allah, M., Ganahl, M., Hayward, L. E., Melko, R. G. & Carrasquilla, J. Recurrent neural network wave functions. *Phys. Rev. Res.* **2**, 023358 (2020).
21. Roth, C. Iterative retraining of quantum spin models using recurrent neural networks. https://doi.org/10.48550/arXiv.2003.06228 (2020).
22. Moss, M. S. et al. Enhancing variational Monte Carlo using a programmable quantum simulator. *Phys. Rev. A* **109**, 032410 (2024).
23. Sprague, K. & Czischek, S. Variational Monte Carlo with large patched transformers. *Commun. Phys.* **7**, 90 (2024).
24. Czischek, S., Moss, M. S., Radzihovsky, M., Merali, E. & Melko, R. G. Data-enhanced variational Monte Carlo simulations for Rydberg atom arrays. *Phys. Rev. B* **105**, 205108 (2022).
25. Hibat-Allah, M., Melko, R. G. & Carrasquilla, J. Supplementing recurrent neural network wave functions with symmetry and annealing to improve accuracy. https://doi.org/10.48550/arXiv.2207.14314 (2022).
26. Bravyi, S. Monte Carlo simulation of stoquastic hamiltonians. *Quant. Inf. Comp.* **15**, 1122–1140 (2015).
27. Lipton, Z. C., Berkowitz, J. & Elkan, C. A critical review of recurrent neural networks for sequence learning. https://doi.org/10.48550/arXiv.1506.00019 (2015).
28. Hibat-Allah, M., Melko, R. G. & Carrasquilla, J. Investigating topological order using recurrent neural networks. *Phys. Rev. B* **108**, 075152 (2023).
29. Casert, C., Vieijra, T., Whitelam, S. & Tamblyn, I. Dynamical large deviations of two-dimensional kinetically constrained models using a neural-network state ansatz. *Phys. Rev. Lett.* **127**, 120602 (2021).
30. Luo, D. et al. Gauge-invariant and anyonic-symmetric autoregressive neural network for quantum lattice models. *Phys. Rev. Res.* **5**, 013216 (2023).
31. Becca, F. & Sorella, S. *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, 2017).
32. Hibat-Allah, M., Inack, E. M., Wiersema, R., Melko, R. G. & Carrasquilla, J. Variational neural annealing. *Nat. Mach. Intell.* **3**, 952–961 (2021).
33. Roth, C., Szabó, A. & MacDonald, A. High-accuracy variational Monte Carlo for frustrated magnets with deep neural networks. *Phys. Rev. B* **108**, 054410 (2023).
34. Khandoker, S. A., Abedin, J. M. & Hibat-Allah, M. Supplementing recurrent neural networks with annealing to solve combinatorial optimization problems. *Mach. Learn. Sci. Technol.* **4**, 015026 (2023).
35. Hamma, A., Ionicioiu, R. & Zanardi, P. Bipartite entanglement and entropic boundary law in lattice spin systems. *Phys. Rev. A* **71**, 022315 (2005).
36. Hamma, A., Ionicioiu, R. & Zanardi, P. Ground state entanglement and geometric entropy in the Kitaev model. *Phys. Lett. A* **337**, 22–28 (2005).
37. Levin, M. & Wen, X.-G. Detecting topological order in a ground state wave function. *Phys. Rev. Lett.* **96**, 110405 (2006).
38. Kitaev, A. & Preskill, J. Topological entanglement entropy. *Phys. Rev. Lett.* **96**, 110404 (2006).
39. Flammia, S. T., Hamma, A., Hughes, T. L. & Wen, X.-G. Topological entanglement rényi entropy and reduced density matrix structure. *Phys. Rev. Lett.* **103**, 261601 (2009).
40. Isakov, S. V., Hastings, M. B. & Melko, R. G. Topological entanglement entropy of a Bose-Hubbard spin liquid. *Nat. Phys.* **7**, 772–775 (2011).
41. Wildeboer, J., Seidel, A. & Melko, R. G. Entanglement entropy and topological order in resonating valence-bond quantum spin liquids. *Phys. Rev. B* **95**, 100402 (2017).
42. Hastings, M. B., González, I., Kallin, A. B. & Melko, R. G. Measuring renyi entanglement entropy in quantum Monte Carlo simulations. *Phys. Rev. Lett.* **104**, 157201 (2010).

43. Wang, Z. & Davis, E. J. Calculating rényi entropies with neural autoregressive quantum states. *Phys. Rev. A* **102**, 062413 (2020).
44. Furukawa, S. & Misguich, G. Topological entanglement entropy in the quantum dimer model on the triangular lattice. *Phys. Rev. B* **75**, 214407 (2007).
45. Stoudenmire, E. & White, S. R. Studying two-dimensional systems with the density matrix renormalization group. *Annu. Rev. Condens. Matter Phys.* **3**, 111–128 (2012).
46. Gong, S.-S., Zhu, W., Sheng, D. N., Motrunich, O. I. & Fisher, M. P. A. Plaquette ordered phase and quantum phase diagram in the spin-$\frac{1}{2}$ $J_1 - J_2$ square Heisenberg model. *Phys. Rev. Lett.* **113**, 027201 (2014).
47. Nikolić, P. & Senthil, T. Theory of the kagome lattice Ising antiferromagnet in weak transverse fields. *Phys. Rev. B* **71**, 024401 (2005).
48. Moessner, R. & Sondhi, S. L. Ising models of quantum frustration. *Phys. Rev. B* **63**, 224401 (2001).
49. Moessner, R., Sondhi, S. L. & Chandra, P. Two-dimensional periodic frustrated Ising models in a transverse field. *Phys. Rev. Lett.* **84**, 4457–4460 (2000).
50. Edwards, S. F. & Anderson, P. W. Theory of spin glasses. *J. Phys. F: Met. Phys.* **5**, 965 (1975).
51. Richards, P. M. Spin-glass order parameter of the random-field Ising model. *Phys. Rev. B* **30**, 2955–2957 (1984).
52. Castellani, T. & Cavagna, A. Spin-glass theory for pedestrians. *J. Stat. Mech.: Theory Exp.* **2005**, P05012 (2005).
53. Merali, E., Vlugt, I. J. S. D. & Melko, R. G. Stochastic series expansion quantum Monte Carlo for Rydberg arrays. *SciPost Phys. Core* **7**, 016 (2024).
54. Sandvik, A. W. & Kurkijärvi, J. Quantum Monte Carlo simulation method for spin systems. *Phys. Rev. B* **43**, 5950–5961 (1991).
55. Sandvik, A. W. Stochastic series expansion method with operator-loop update. *Phys. Rev. B* **59**, R14157–R14160 (1999).
56. Yan, Z., Samajdar, R., Wang, Y.-C., Sachdev, S. & Meng, Z. Y. Triangular lattice quantum dimer model with variable dimer density. *Nat. Commun.* **13**, 5799 (2022).
57. Wu, D. et al. Variational benchmarks for quantum many-body problems. *Science* **386**, 296–301 (2024).
58. Moss, M. S., Wiersema, R., Hibat-Allah, M., Carrasquilla, J. & Melko, R. G. Leveraging recurrence in neural network wavefunctions for large-scale simulations of Heisenberg antiferromagnets: the square lattice. https://arxiv.org/abs/2502.17144 (2025).
59. Moss, M. S., Wiersema, R., Hibat-Allah, M., Carrasquilla, J. & Melko, R. G. Leveraging recurrence in neural network wavefunctions for large-scale simulations of Heisenberg antiferromagnets: the triangular lattice. https://arxiv.org/abs/2505.20406 (2025).
60. Chen, A. & Heyl, M. Empowering deep neural quantum states through efficient optimization. *Nat. Phys.* **20**, 1476–1481 (2024).
61. Rende, R., Viteritti, L. L., Bardone, L., Becca, F. & Goldt, S. A simple linear algebra identity to optimize large-scale neural network quantum states. *Commun. Phys.* **7**, 260 (2024).
62. Donatella, K., Denis, Z., Le Boité, A. & Ciuti, C. Dynamics with autoregressive neural quantum states: Application to critical quench dynamics. *Phys. Rev. A* **108**, 022210 (2023).
63. Lange, H., Döschl, F., Carrasquilla, J. & Bohrdt, A. Neural network approach to quasiparticle dispersions in doped antiferromagnets. *Commun. Phys.* **7**, 187 (2024).
64. Sinibaldi, A., Hendry, D., Vicentini, F. & Carleo, G. Time-dependent neural galerkin method for quantum dynamics https://arxiv.org/abs/2412.11778.(2024).
65. de Walle, A. V., Schmitt, M. & Bohrdt, A. Many-body dynamics with explicitly time-dependent neural quantum states https://arxiv.org/abs/2412.11830.(2024).
66. Semeghini, G. et al. Probing topological spin liquids on a programmable quantum simulator. *Science* **374**, 1242–1247 (2021).
67. Giudici, G., Lukin, M. D. & Pichler, H. Dynamical preparation of quantum spin liquids in Rydberg atom arrays. *Phys. Rev. Lett.* **129**, 090401 (2022).
68. Carrasquilla, J., Torlai, G., Melko, R. G. & Aolita, L. Reconstructing quantum states with generative models. *Nat. Mach. Intell.* **1**, 155–161 (2019).
69. Bennewitz, E. R., Hopfmueller, F., Kulchytskyy, B., Carrasquilla, J. F. & Ronagh, P. Neural error mitigation of near-term quantum simulations. *Nat. Mach. Intell.* **4**, 618–624 (2022).
70. Lange, H. et al. Transformer neural networks and quantum simulators: a hybrid approach for simulating strongly correlated systems. *Quantum* **9**, 1675 (2025).
71. Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).
72. Abadi, M. et al. TensorFlow: Large-scale machine learning on heterogeneous systems. https://www.tensorflow.org/. Software available from tensorflow.org (2015).
73. Harris, C. R. et al. Array programming with numpy. *Nature* **585**, 357–362 (2020).

Acknowledgements

We would like to thank Subir Sachdev, Anders Sandvik, and Arun Paramekanti for their helpful and inspiring discussions. Our RNN implementation is based on Tensorflow[72] and NumPy[73]. Computer simulations were made possible thanks to the Vector Institute computing cluster and the Digital Research Alliance of Canada cluster. We acknowledge support from Natural Sciences and Engineering Research Council of Canada (NSERC), the Shared Hierarchical Academic Research Computing Network (SHARCNET), Compute Canada, and the Canadian Institute for Advanced Research (CIFAR) AI chair program. This work is not related to the research being performed at AWS. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Colleges and Universities. This research was supported in part by grant NSF PHY-2309135 to the Kavli Institute for Theoretical Physics (KITP).

Author contributions

All the authors contributed significantly to the design of this study and the writing of this manuscript. M.H. launched and analyzed the RNN simulations. E.M launched and analyzed the QMC simulations. G.T. contributed to the Density Matrix Renormalization Group (DMRG) simulations. J.C. and R.G.M. were advising this project.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s42005-025-02226-7.

Correspondence and requests for materials should be addressed to Mohamed Hibat-Allah.

Peer review information *Communications Physics* thanks Jonas Rigo and the other, anonymous, reviewers for their contribution to the peer review of this work.