

Entanglement and optimization within autoregressive neural quantum states

Andrew Jreissaty^{1,2,*}, Hang Zhang,¹ Jairo C. Quijano³, Juan Carrasquilla^{1,2} and Roeland Wiersema⁴

¹*Institute for Theoretical Physics, ETH Zurich, CH-8093 Zurich, Switzerland*

²*Vector Institute, Toronto, Ontario M5G 0C6, Canada*

³*Politécnico Colombiano JIC, Medellín, Colombia*

⁴*Center for Computational Quantum Physics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA*



(Received 22 September 2025; accepted 12 January 2026; published 10 February 2026)

Neural quantum states (NQSs) are powerful variational ansätze capable of representing highly entangled quantum many-body wavefunctions. While the average entanglement properties of ensembles of restricted Boltzmann machines are well understood, the entanglement structure of autoregressive NQSs such as recurrent neural networks and transformers remains largely unexplored. We perform large-scale simulations of ensembles of random autoregressive wavefunctions for chains of up to 256 spins and uncover signatures of transitions in their average entanglement scaling, entanglement spectra, and correlation functions. We show that the standard softmax normalization of the wavefunction suppresses entanglement and fluctuations, and introduce a square modulus normalization function that restores them. Finally, we connect the insights gained from our entanglement and activation function analysis to initialization strategies for finding the ground states of strongly correlated Hamiltonians via variational Monte Carlo.

DOI: [10.1103/t2cg-kr7y](https://doi.org/10.1103/t2cg-kr7y)

I. INTRODUCTION

Efficient simulation of strongly correlated quantum systems in and out of equilibrium is a fundamental problem of significant interest to the condensed matter community. While quantum Monte Carlo approaches are considered state of the art in sign-problem-free systems for the simulation of equilibrium states, they continue to struggle to overcome the sign problem in fermionic and frustrated spin models [1–10]. Likewise, matrix product state (MPS) and tensor network techniques combined with algorithms such as the density matrix renormalization group (DMRG) outperform many approaches in one dimension for systems with area-law entanglement, but start to suffer from efficiency issues owing to exponential bond dimension growth in more strongly entangled and higher-dimensional systems [11–15].

While there is likely more to come from such approaches as recent successful tensor network simulations of dynamics of quantum spin glasses in higher dimensions show [16], the present gaps in these methods have opened the door for machine-learning-based neural network quantum states (NQSs) as variational ansätze for the targeting of ground states (including volume-law states) of strongly correlated Hamiltonians [17–24], phase classification across critical points [25,26], quantum state tomography [27], and the simulation of quantum dynamics [28–33].

Autoregressive NQSs have emerged in recent years as ansätze from which independent samples can be efficiently obtained. By outputting normalized conditional probabilities that depend only on previously sampled degrees of freedom, successive independent samples can be generated without having to deal with the autocorrelation issues that can inhibit Markov chain Monte Carlo sampling required by unnormalized NQS approaches [22,23,34–36].

We focus on two primary autoregressive architectures that have been adopted to tackle the quantum many-body problem, namely, the recurrent neural network (RNN) [22,37–42] and the autoregressive transformer (ATF) [42–48]. And while there is evidence that these architectures can suffer from expressivity issues in the context of fermionic and strongly correlated spin Hamiltonians [42,49], clever design choices have allowed these models to overcome these obstacles and thrive, and in many cases provide state-of-the-art results, such as in the recent large-scale ground-state simulations of the antiferromagnetic Heisenberg model on the triangular lattice (single-layer RNN) and of the Rydberg Hamiltonian on the square lattice (ATF) [41,43]. Important developments have also been made in the use of autoregressive models for the simulation of quench dynamics [30], though issues with the regularization of the quantum geometric tensor in time-dependent variational Monte Carlo (VMC) persist [30,50].

Despite this exciting progress, many open questions remain. While NQSs have proven to be highly compelling variational wavefunctions for the representation of a variety of strongly correlated ground states, strong performance hinges on the specific choice of architecture and extensive hyperparameter tuning [20]. It is therefore important to develop a clearer understanding of the physical properties and capabilities of the various NQS models, in order to better inform the choice of architecture and hyperparameters for the purposes of solving a given quantum many-body problem.

*Contact author: ajreissaty@phys.ethz.ch

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

To that end, the entanglement properties of ensembles of restricted Boltzmann machines (RBMs) have already been studied extensively [51,52]. Deng *et al.* [51] showed that while RBMs with only short-range connections obey area-law entanglement, fully connected RBMs follow a volume-law scaling. Building on Ref. [51], Sun *et al.* [52] mapped fully connected random RBMs to statistical mechanics models and computed their average entanglement phase diagram, which features generic volume-law behavior and entanglement spectra that in certain regimes exhibit statistics following the Gaussian Orthogonal Ensemble (GOE), similar to thermalizing quantum many-body systems [53].

Motivated by these observations, in this article we consider the RNN of Ref. [22] and ATF of Ref. [43], and study their average entanglement properties. We compute entanglement entropy phase diagrams, uncovering a contrast between peak regions of entropy and lower entanglement states similar to the observations in the random RBM case. We study how tuning nonlinearities and parameter choices impact entanglement scaling, correlations, and entanglement spectrum level statistics, reaching system sizes of up to $L = 256$ spins in the process. While our primary goal is to understand the extent to which transitions in entanglement scaling and the other relevant quantities can be induced by the above-listed hyperparameters, including the autoregressive property itself, we also develop optimized initialization strategies for ground-state search of various strongly correlated Hamiltonians.

The paper is organized as follows. In Sec. II, we introduce the RNN and ATF architectures and the methods used for our calculations. We examine the average entanglement properties of random RNN wavefunctions in Sec. III, and present evidence of transitions in the relevant quantities as induced by the hyperparameters of interest, including the given nonlinear activation functions. We also consider VMC initialization strategies for the transverse-field Ising model (TFIM) and quantum Heisenberg model. In Sec. IV, we carry out the entire analysis for the autoregressive transformer and conclude with a final discussion and outlook in Sec. V.

II. ARCHITECTURES AND METHODS

A. Autoregressive neural quantum states

The variational state for a general autoregressive spin model can be expressed as

$$|\Psi_\lambda\rangle = \sum_{\sigma} \Psi_\lambda(\sigma) |\sigma\rangle, \quad (1)$$

$$\Psi_\lambda(\sigma) = \sqrt{P_\lambda(\sigma)} e^{i\phi_\lambda(\sigma)}, \quad (2)$$

$$P_\lambda(\sigma) = \prod_n P_\lambda(\sigma_n | \sigma_{<n}), \quad (3)$$

where $\Psi_\lambda(\sigma)$ is the probability amplitude corresponding to a given spin-1/2 configuration $\sigma \in \{0, 1\}^L$ and parametrized by a set of weights $\{\lambda\}$, $P_\lambda(\sigma) = |\Psi_\lambda(\sigma)|^2$ is the probability associated with that configuration, expressed in Eq. (3) as a product of conditional probabilities that depend only on previously sampled spins, and $\phi_\lambda(\sigma)$ is the phase, calculated by combining the local phases $\phi_\lambda(\sigma_n | \sigma_{<n})$ at each site and writing $\phi_\lambda(\sigma) = \sum_n \phi_\lambda(\sigma_n | \sigma_{<n})$. In certain cases, we may be interested in building a fully real and positive variational




FIG. 1. (a) RNN and (b) ATF architectures. (a) The RNN cell (green) applies a generally nonlinear activation function f to a linearly transformed hidden state h_{n-1} and input vector σ_{n-1} [Eq. (4)]. The new hidden state is transformed in the output layer to produce the conditional probability (blue) and phase (gray) vectors, as per Eqs. (5) and (6). (b) The ATF directly accepts as input all previously sampled spins $\sigma_1, \sigma_2, \dots, \sigma_{n-1}$. Each spin is embedded separately into a d_{emb} -dimensional space, and from there undergoes the transformations detailed in Sec. II C to produce the final conditional probability (blue) and phase (gray) vectors. We note that r_1 denotes the argument of f in Eq. (4), while r_2 and r_3 denote the arguments of the output layer activation functions of the RNN and ATF.

wavefunction, e.g., for the simulation of the ground states of stoquastic Hamiltonians [54]. For this, we would set $\Psi_\lambda(\sigma) = \sqrt{P_\lambda(\sigma)}$ and avoid calculating a phase.

B. Recurrent neural network

First, we focus on the one-dimensional (1D) RNN wavefunction of Ref. [22]. The vanilla RNN, depicted in Fig. 1(a), operates as follows. At each site n , a fully connected layer (the RNN cell) accepts the $(n-1)$ th spin σ_{n-1} and the corresponding hidden state h_{n-1} , and outputs the next hidden state

$$h_n = f(W[h_{n-1}; \sigma_{n-1}] + b), \quad (4)$$

with the semicolon representing a concatenation, and f is applied elementwise. Next, h_n is passed to the output layer, producing the conditionals and phase components as

$$P_\lambda(\sigma_n | \sigma_{n-1}) = g(Uh_n + c), \quad (5)$$

$$\Phi_\lambda(\sigma_n | \sigma_{n-1}) = \mathbb{1}(Vh_n + d) = Vh_n + d. \quad (6)$$

Here, g is applied elementwise, $\mathbb{1}$ is the identity function officially defined in Eq. (11), and $P_\lambda(\sigma_n | \sigma_{n-1})$ and $\Phi_\lambda(\sigma_n | \sigma_{n-1})$ are, respectively, the conditional probability and the phase component vectors whose elements correspond to the two levels of a spin-1/2 degree of freedom, easily generalizable to a d -dimensional quantum spin. The hidden state $h_n \in \mathbb{R}^{d_h}$ is a vector that stores information about previously sampled spins and encodes correlations between the sites. As such, d_h

is referred to as the number of memory units, and it is one of the key hyperparameters we tune in Sec. III. We note that the initial vectors $\mathbf{h}_0$ and $\mathbf{\sigma}_0$ are chosen to be zero vectors, which is standard for RNN wavefunctions [22]. The network parameters $\{\lambda\} = \{W, U, V, \mathbf{b}, \mathbf{c}\}$ are such that $W \in \mathbb{R}^{d_h \times (d_h+2)}$ and $U, V \in \mathbb{R}^{2 \times d_h}$, as well as the bias vectors $\mathbf{b} \in \mathbb{R}^{d_h}$ and $\mathbf{c}, \mathbf{d} \in \mathbb{R}^2$. The hidden dimension d_h controls the total number of parameters N_p in the model. Weights are shared among all sites, which allows one to use the same parameters for different system sizes [22,40,41]. Importantly for our analysis, we explicitly define the activation functions to be tuned as follows:

$$f \rightarrow \text{vanilla RNN cell activation function}, \quad (7)$$

$$g \rightarrow \text{output layer activation function}. \quad (8)$$

Throughout this work, f and g are functions that we allow to vary. For the model to be autoregressive, g must act on the linearly transformed hidden state to produce normalized conditionals. It is usually chosen to be the nonlinear softmax function $\text{SM}(\mathbf{r})$ [22,23,40–42], which acts on the elements of a vector $\mathbf{r}$ to produce

$$\text{SM}(r_j) = \frac{e^{r_j}}{\sum_i e^{r_i}} \forall j, \quad (9)$$

the elements of which sum to one. In Sec. III, we will propose and explore two alternatives to the softmax: a square modulus function $g(\mathbf{r}) = \text{MOD}(\mathbf{r})$ that preserves the normalization,

$$\text{MOD}(r_j) = \frac{r_j^2}{\sum_i r_i^2} \forall j, \quad (10)$$

and the identity function $g(\mathbf{r}) = \mathbb{1}(\mathbf{r})$, which removes a non-linearity, renders the RNN unnormalized, and removes the autoregressive property:

$$\mathbb{1}(r_j) = r_j \forall j. \quad (11)$$

As for f , in a vanilla RNN cell it is generally taken to be the hyperbolic tangent function $f(\mathbf{r}) = \tanh(\mathbf{r})$ [22]. With the universal approximation theorem for multilayer feedforward neural networks [55] depending explicitly on the presence of nonlinearities in the fully connected layers of those models, a commonly held view in the community is that increasing the number of nonlinear functions in an NQS model is likely to improve its expressive power [56] in terms of its ability to simulate a wider range of highly entangled wavefunctions, since linear networks can only approximate a limited subset of functions.

C. Autoregressive transformer

Unlike the RNN that processes one spin at a time, the ATF [43] processes the entire spin configuration in parallel using masked self-attention layers. This design allows every spin's representation to immediately incorporate information from all preceding spins, enabling the ATF to capture long-range correlations [57–59].

As illustrated in Fig. 1(b), the ATF takes as input the configuration of all previously sampled spins $\sigma_1, \sigma_2, \dots, \sigma_{n-1}$, as opposed to the RNN cell, which only takes in σ_{n-1} directly

and encodes the dependence on spins $\sigma_1, \sigma_2, \dots, \sigma_{n-2}$ into the $(n-1)$ th hidden state.

We consider the multihead ATF, where each spin is embedded into a d_{emb} -dimensional space via a linear transformation, with fixed sinusoidal encodings, which are added to the result to retain positional information [60]. The resulting input matrix $X \in \mathbb{R}^{L \times d_{\text{emb}}}$ is then used to calculate query, key, and value matrices for each of the h heads:

$$Q_i = XW_i^Q, \quad K_i = XW_i^K, \quad V_i = XW_i^V, \quad (12)$$

where $W_i^Q, W_i^K, W_i^V \in \mathbb{R}^{d_{\text{emb}} \times d_K}$ with $i = 1, \dots, h$ are trainable weights and $d_K = d_{\text{emb}}/h$ is the per-head dimension. Recent studies have introduced diverse strategies for parametrizing the attention in transformers [60–63]; here, we choose to investigate two distinct attention mechanisms: the standard softmax attention [60] and a masked circulant attention kernel that we propose, motivated by the idea of relative positional encoding [58,63]. We will denote this attention scheme by $\mathcal{A}$.

In the softmax attention scheme $\mathcal{A}_{\text{SM}}$, we calculate an attention score matrix $R_i^{\text{ASM}} \in \mathbb{R}^{L \times L}$ as

$$R_i^{\text{ASM}} = \frac{Q_i K_i^\top}{\sqrt{d_K}} + M, \quad (13)$$

where M is a lower-triangular mask matrix with $M_{ij} = -\infty$ for $i < j$ and $M_{ij} = 0$ otherwise that encodes the autoregressive structure. We then apply the softmax attention to obtain

$$A_i = \text{SM}(R_i^{\text{ASM}}). \quad (14)$$

Here, $\text{SM}(R_i^{\text{ASM}})$ normalizes the rows of R_i^{ASM} as per Eq. (9).

In the circulant attention scheme $\mathcal{A}_{\text{Circ}}$, we generate a matrix $R_{\text{Circ}} \in \mathbb{R}^{L \times L}$ as per

$$R_i^{\text{Circ}} \equiv \text{Circ}(\mathbf{r}_i), \quad (15)$$

where $\text{Circ}(\mathbf{r}_i)$ is the circulant matrix produced from a trainable kernel vector $\mathbf{r}_i \in \mathbb{R}^L$. The masked circulant attention is calculated as

$$A_i = R_i^{\text{Circ}} \odot M', \quad (16)$$

where $\odot$ is the elementwise Hadamard product, and M' is a lower-triangular binary matrix defined elementwise by $M'_{ij} = 1$ for $j \leq i$ and $M'_{ij} = 0$ otherwise, which once again restricts each position i to only attend to positions $j \leq i$. Circulant self-attention reduces the number of trainable parameters from $\mathcal{O}(L^2)$ to $\mathcal{O}(L)$, while naturally inducing the translational invariance particularly well suited for systems with periodic boundary conditions [63]. Although the mask breaks the exact circulant symmetry, it preserves some shift invariance locally while enabling autoregressive conditioning.

Finally, for both types of attention, we calculate the output for each head:

$$\text{head}_i = A_i V_i \in \mathbb{R}^{L \times d_K}. \quad (17)$$

Within multihead attention (MHA), the attention mechanism is performed in parallel across h heads. Each head transforms the input using separate learned projections, which are then finally concatenated as

$$\text{MHA}(X) = [\text{head}_1; \dots; \text{head}_h], \quad (18)$$

yielding the final output of the attention function $\text{MHA}(X) \in \mathbb{R}^{L \times d_{\text{emb}}}$.

Each transformer block is equipped with residual connections and layer normalization [64] around the MHA function and subsequent feedforward layers (FFLs), following standard architectural conventions, such that

$$X' = \text{LayerNorm}(X + \text{MHA}(X)) \in \mathbb{R}^{L \times d_{\text{emb}}}, \quad (19)$$

$$X'' = \text{LayerNorm}(X' + \text{FFLs}(X')) \in \mathbb{R}^{L \times d_{\text{emb}}}, \quad (20)$$

with the pointwise FFLs defined as

$$\text{FFLs}(X') = f_{\text{FL}}(X'W_1 + \mathbf{b}_1)W_2 + \mathbf{b}_2, \quad (21)$$

and $f_{\text{FL}}(\mathbf{r})$ commonly taken to be the rectified linear unit activation function $\text{ReLU}(\mathbf{r})$. Here, we have $W_1 \in \mathbb{R}^{d_{\text{emb}} \times d_{\text{FL}}}$, $W_2 \in \mathbb{R}^{d_{\text{FL}} \times d_{\text{emb}}}$, $\mathbf{b}_1 \in \mathbb{R}^{d_{\text{FL}}}$, and $\mathbf{b}_2 \in \mathbb{R}^{d_{\text{emb}}}$ as trainable weight tensors, with d_{FL} the width of the feedforward layer. After the MHA and FFLs, the ATF outputs the conditional probability and phase vectors, as per

$$P_\lambda(\sigma_n | \sigma_{n-1}) = g(W_3 X' + \mathbf{c}')_n, \quad (22)$$

$$\Phi_\lambda(\sigma_n | \sigma_{n-1}) = \mathbb{1}(W_4 X' + \mathbf{d}')_n, \quad (23)$$

where $W_3, W_4 \in \mathbb{R}^{d_{\text{emb}} \times 2}$ and $\mathbf{c}', \mathbf{d}' \in \mathbb{R}^2$ are additional trainable weights and biases, and where $P_\lambda(\sigma_n | \sigma_{n-1})$ and $\Phi_\lambda(\sigma_n | \sigma_{n-1})$, respectively, correspond to the n th row of the $L \times 2$ matrices $g(W_3 X' + \mathbf{c}')$ and $\mathbb{1}(W_4 X' + \mathbf{d}')$. In our simulations, we fix $h = 2$ and the number of feedforward layers $N_{\text{FFLs}} = 1$, and we allow $\mathcal{A}$, f_{FL} , and g to vary as we investigate the impact of nonlinearity on the entanglement properties of the model.

D. Methodology

Our ensemble of wavefunctions is constructed by drawing the weights and biases of the neural networks from a Gaussian distribution of zero mean and fixed width σ , i.e., $\mathcal{N}(0, \sigma^2)$. Unless otherwise specified, every calculation of an average quantity for our random wavefunctions goes as follows:

- (1) Select desired NQS architecture: RNN versus ATF.
- (2) Select hyperparameters:
 - (a) Complex versus Positive wavefunction
 - (b) Activation functions: RNN $\rightarrow f, g$; ATF $\rightarrow \mathcal{A}, f_{\text{FL}}$, g ;
 - (c) RNN: (d_h, σ) ; ATF: (d_{emb}, σ) .
- (3) Draw network weights from $\mathcal{N}(0, \sigma^2)$.
- (4) Calculate desired properties of resulting random state using direct Monte Carlo sampling of the autoregressive model [22,35,65]. When the model is not autoregressive, e.g., when $g(\mathbf{r}) = \mathbb{1}(\mathbf{r})$ for the RNN, Markov chain Monte Carlo is used instead for sample generation.
- (5) Repeat steps (3) and (4) for N_{init} random wavefunctions, and calculate averages for the relevant quantities (e.g., entanglement entropy) over the wavefunction ensemble.
- (6) Repeat steps (1)–(5) for all desired hyperparameter combinations.

We focus on random Gaussian ensembles because Gaussian initialization is common in machine learning [66,67], and as a result, has also been widely used by the NQS community [52].

III. RNN: ENTANGLEMENT AND VARIATIONAL MONTE CARLO

A. Entanglement phase diagrams

We begin by examining the bipartite Rényi entanglement entropy $S_2^A = -\ln[\text{Tr}(\hat{\rho}_A^2)]$, calculated using the swap operator technique of Ref. [68] with a Monte Carlo estimator [22,69], for an equal bipartition of the system. In Fig. 2, we consider an $L = 20$ 1D spin chain and compute the ratio of S_2^A and the maximal entropy $S_2^A/(L_A \ln 2)$, with $L_A \equiv L/2$. We, respectively, average over $N_{\text{init}} = 100$ and $N_{\text{init}} = 20$ random RNN wavefunctions in the autoregressive [Figs. 2(a1)–2(c1) and Figs. 2(a2)–2(c2)] and unnormalized [Figs. 2(d1) and 2(d2)] cases, and plot this ratio as a function of the number of memory units d_h and the width of the parameter distribution function σ , for a variety of combinations of the f and g activation functions, and for both complex and fully positive architectures.

We first consider the case of complex wavefunctions and specifically Fig. 2(a1), which shows the results for the autoregressive vanilla RNN of Ref. [22] ($f = \tanh$, $g = \text{SM}$), perhaps the most prototypical iteration of the RNN architecture for the purposes of solving problems in quantum many-body physics. For fixed $d_h > 0$, the S_2^A versus σ curves exhibit a peak that both grows and shifts to the left with increasing d_h [see Fig. 2(a1) in the main text and Fig. 12 in Appendix C]. Together, these curves combine to form a narrow region of peak entanglement as seen in Fig. 2(a1), with the peak generating approximately half the maximal entropy near $d_h = 50$, but also widening as it declines in value and shifts to the right in the limit $d_h \rightarrow 0$. Calculations in the $d_h \in [51, 200]$ region (see Fig. 13 in Appendix C) show that as d_h increases, the position of the peak appears to approach $\sigma = 0$ at a polynomially decreasing rate (at least), while its value also appears to saturate to some maximal value slightly larger than $S_2^A/(L_A \ln 2) \sim 0.5$. We expect the peak to only ever reach $\sigma = 0$ in the limit $d_h \rightarrow \infty$; for finite d_h , it will always occur at finite $\sigma > 0$. In comparison, while the corresponding RBM entanglement phase diagram of Ref. [52] (S_2^A as a function of λ_{RBM} , the ratio of hidden-to-visible neurons, and σ), which was estimated analytically in the limit $N \rightarrow \infty$, also contains a region of higher entropy that is narrow in width at small σ and large λ_{RBM} , a region of near maximal entanglement emerges at lower λ_{RBM} , occurring between $\lambda_{\text{RBM}} \sim 0.3$ and $\lambda_{\text{RBM}} \sim 1.2$ and stretching from small σ all the way to $\sigma \rightarrow \infty$. This is a first indication that the autoregressive nature of the RNN may be suppressing entanglement at large σ .

A Gaussian width of $\sigma = 0$ forces all parameters to vanish and produces real probability amplitudes satisfying $\Psi_\lambda(\sigma) = 1/\sqrt{2^N}$ for all spin configurations σ . The resulting state is a product state, $|\Psi_\lambda\rangle_{\sigma=0} = |+\rangle^{\otimes N} = [\frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle)]^{\otimes N}$, an eigenstate of $\hat{S}_j^x \forall j$. Thus, the entanglement effectively vanishes when $\sigma \sim 0$, irrespective of the choice of f and g , as seen in Fig. 2. The more interesting limit is $\sigma \rightarrow \infty$. Here, for $(f, g) = (\tanh, \text{SM})$, the entanglement entropy is clearly suppressed. Upon drawing multiple wavefunctions for large σ , Monte Carlo sampling produces a single dominant configuration in the $\hat{S}^z$ basis. An example of this is shown in Table I(b) in Appendix C, where samples are drawn from a random




FIG. 2. Average bipartite entanglement entropy divided by the maximal entropy $\overline{S}_2^A / (L_A \ln 2)$ across N_{init} random 20-spin RNN wavefunctions as a function of the hidden unit dimension d_h and the Gaussian distribution width σ . The different panels explore various combinations of the activation functions f and g [as per Eqs. (7) and (8)], namely, (a) $f = \tanh$, $g = \text{SM}$, (b) $f = \mathbb{1}$, $g = \text{SM}$, (c) $f = \tanh$, $g = \text{MOD}$, and (d) $f = \tanh$, $g = \mathbb{1}$. The wavefunctions are complex in the first row [panels (a1)–(d1)] and positive in the second row [panels (a2)–(d2)], while the averages are taken over $N_{\text{init}} = 100$ random states for the autoregressive architectures [panels (a)–(c)] (10^5 samples used for the MC estimation), and $N_{\text{init}} = 20$ otherwise [panel (d)] (96 000 samples used here). In the presence of the softmax, adding a hyperbolic tangent in the fully connected linear on top of the linear transformation induces a richer structure of entanglement entropy across the random wavefunction ensemble, while a comparison of panel (a1) with panels (c1) and (d1) highlights the suppression of entanglement induced by the softmax in the large σ regime for most values of $d_h > 0$.

$(f, g) = (\tanh, \text{SM})$ wavefunction initialized at $\sigma = 50$, and only one configuration is ever produced. We note that this configuration is entirely random, with a new dominant configuration emerging as new parameters are drawn from the same Gaussian and the wavefunction is re-initialized. As long as there are no correlations between successive sets of RNN weights, successive dominant configurations are never correlated.

The random RBM wavefunctions explored by Ref. [52] do not experience this large- σ entanglement suppression, which suggests that the autoregressive nature of the RNN causes it. We investigate this in Fig. 2(d1), replacing the softmax in the output layer with the identity function $g = \mathbb{1}$, thereby producing unnormalized conditionals and making sampling possible only via Markov chain Monte Carlo (negative signs are incorporated into the complex phase). While the peak entanglement region is relatively unchanged along its curved left outline, the large- σ S_2^A suppression is removed, and the peak entanglement region extends indefinitely in the limit $\sigma \rightarrow \infty$. These observations motivate us to replace the softmax with an alternative that retains the autoregressive property without relying on exponentials, namely, the square modulus function $\text{MOD}(\mathbf{r})$ introduced in Eq. (10).

The $(f, g) = (\tanh, \text{MOD})$ results are shown in Fig. 2(c1). While the peak average entropy values are lower than

those produced by the $(f, g) = (\tanh, \mathbb{1})$ wavefunctions of Fig. 2(d1) ($S_2^A / L_A \ln 2 \sim 0.4$ compared to 0.5), there is no entanglement suppression in the $\sigma \rightarrow \infty$ limit. Thus, any vanishing of the entanglement in that limit cannot be a result of the autoregressive property per se—it must be due to the softmax function itself, as the numerical results confirm. As σ increases, the resulting network weights are larger in magnitude and more widely scattered about the zero mean. This ensures that the unnormalized two-component vector $\mathbf{r}$ produced by a forward pass through the layers and which enters $g(\mathbf{r})$ as an argument in Eq. (5) is more likely to have two elements whose difference is significant enough that the softmax will exponentially suppress one at the expense of the other. The larger σ is, the more configurations σ and their corresponding $|\Psi_\lambda(\sigma)|^2$ are suppressed.

We note that when $\sigma \sim 0$, this phenomenon initially amplifies the correlations, because the starting point is an equal superposition of all computational basis states, a fully uncorrelated wavefunction, and initial cancellations of configurations are needed to generate entanglement. This is why an entanglement peak emerges in the initial increase of the Gaussian width in Fig. 2(a1). Beyond the peak, continued increase in σ causes entanglement to decrease, as not enough configurations make strong enough contributions to maintain the correlations between the relevant subsystems. In the large




FIG. 3. Ratio of the average entanglement entropy and the maximal entropy $\overline{S}_2^A / (L_A \ln 2)$ as a function of σ at $d_h = 40$ for the four complex RNN wavefunction architectures of Fig. 2. The error bars correspond to the variance of the entropies in the random ensemble, and do not include the minimal-by-comparison RNN sampling errors induced by the Monte Carlo estimation.

σ limit, the state eventually collapses to a random single dominant product state configuration.

Instead, when $g = \text{MOD}$, there is no exponential suppression, only polynomial, and our numerics show that even in the limit of very large σ , this does not suppress the correlations developed at small σ . This point is illustrated in Appendix A (see Table I), which compares random samples drawn from two wavefunctions initialized at $\sigma = 0.4$: one with $g = \text{SM}$ and the other with $g = \text{MOD}$. The same comparison is repeated at $\sigma = 50$. At small σ , neither wavefunction is dominated by a single configuration, and this behavior persists for the modulus-based RNN even at $\sigma = 50$, in contrast to its softmax counterpart. Additional numerical results indicate that the modulus function continues to behave similarly in the limit $\sigma \rightarrow \infty$, highlighting the extent of its scale invariance. We construct a simple argument in Appendix B that justifies the behavior of the entanglement in the small and large σ regimes.

Beyond the output layer, we also examine the role of the fully connected nonlinearity f in Eq. (4). As seen in Figs. 2(a1) and 2(b1), replacing the identity with $\tanh$ under the softmax yields a much richer entanglement structure in the (d_h, σ) diagram, lending credence to the idea that nonlinearities enhance neural network expressivity [55]. However, this does not imply that simply stacking nonlinearities will extend correlations indefinitely—the softmax suppression at large σ shows otherwise. Instead, it is the interplay of linear and nonlinear transformations across layers that drives greater expressivity (see Fig. 3 for cross-section comparisons). We expect that stronger nonlinearities in the fully connected layers of an already normalized random wavefunction should further aid expressivity.

In order to bring this idea of expressivity into sharper focus, we calculate the standard deviation of the ensemble entanglement entropy values used to compute the averages of Fig. 2, and plot the results in Appendix C (Fig. 14). These lead

us to conclusions largely analogous to those gleaned from the analysis of the averages, for example that the softmax also suppresses fluctuations, one of many possible indicators of wavefunction expressivity.

The impact of sign structure on the entanglement of random wavefunctions has already been partially investigated [70,71]. In Ref. [71], Grover and Fisher proved that general wavefunctions $|\Psi\rangle = \sum_{\sigma} \Psi(\sigma)|\sigma\rangle$ where $\Psi(\sigma) = \pm 1/\sqrt{2^N}$ with equal probability exhibit on average volume-law scaling for the 2-Rényi entropy, while random positive wavefunctions encode area-law scaling for all α -Rényi entropies with $\alpha > 1$ [70]. We would thus expect the enforcing of a random complex phase on our random positive RNNs to at the very least boost the average entanglement entropies that they can encode, and this is exactly what our results in Fig. 2 show, with the impact most pronounced when $(f, g) = (\tanh, \text{SM})$ and $(f, g) = (\tanh, \mathbb{1})$.

The ensemble of wavefunctions generated in the large σ with the SM activation can be paralleled with eigenstates of the prototypical Hamiltonian for the study of the many-body localization (MBL) phenomenon, namely, the disordered Heisenberg model, $\hat{H}_{\text{MBL}} = -J \sum_i \hat{S}_i^x \cdot \hat{S}_{i+1}^x - \sum_i h_i \hat{S}_i^z$, where the set of $\{h_i\}$ are drawn independently from the uniform distribution spanning $[-W, W]$ [53,72,73]. In the strong-disorder regime ($W \gg J$), eigenstates essentially reduce to a single random configuration with perturbative corrections that vanish as $W \rightarrow \infty$. Analogously, large σ drives our softmax-based RNN wavefunctions toward a single dominant configuration, where σ acts as a disorderlike localization parameter. In contrast, when $\sigma \rightarrow 0$, the wavefunctions collapse to the trivial all-up $\hat{S}^x$ eigenstate, which is localized but not disorder induced.

This motivates two questions. First, given the similarity between the large σ limit and an MBL regime, to what extent can the peak-entanglement states in Fig. 2 be regarded as “thermal” or ergodic [72]? Second, do the (d_h, σ) diagrams capture transitions in physically relevant quantities—such as entanglement scaling, level statistics, or correlations—and thus serve as genuine phase diagrams? We aim at shedding light on these questions in the following subsections.

B. Entanglement scaling

At least two exact mappings between specific matrix product state and RNN constructions are known. In Ref. [74], Wu *et al.* constructed an exact mapping from a generic matrix product state to a specific single-layer RNN with a linear memory update, i.e., a hidden state update that sets $f = \mathbb{1}$ in Eq. (4). Such a linear RNN, referred to as an MPS-RNN, must exhibit an area law in entanglement scaling just like its equivalent MPS with bond dimension χ , which encodes a maximum bipartite entanglement entropy of $S_{\text{max}} \sim \ln \chi$ [11,12]. However, because the mapping leads to a specific form for the linear RNN, the possibility of a reverse mapping from a generic linear RNN to an MPS remains unclear.

Another known mapping was established in a series of influential works by Levine and co-workers on the relation between network depth and quantum entanglement [75,76]. There, the authors connected a shallow recurrent network

known as a Recurrent Arithmetic Circuit (RAC) to an equivalent matrix product state. The RAC architecture is an RNN of the form described in Sec. II B with $f \equiv f(\mathbf{h}_{n-1}, \boldsymbol{\sigma}_{n-1}) = W_h \mathbf{h}_{n-1} \odot W_\sigma \boldsymbol{\sigma}_{n-1}$ (with W_h and W_σ defining linear transformations) and $g = \mathbb{1}$. The MPS mapping proves that the shallow, one-layer RAC encodes an area law of entanglement. The authors later showed that including an identical second fully connected layer in the RAC produces a network that can admit logarithmic corrections to the area law, i.e. $S_{\max} \sim \ln L$. In yet another recent work by Yang *et al.* [77], sign-free RNNs with a specific property known as conditional independence were shown to also satisfy an area law with possible logarithmic corrections.

We now seek to numerically study the scaling properties of our normalized autoregressive single-layer RNNs. Intuitively, we expect any RNN architecture of the form shown in Fig. 1(a) to find it difficult to encode longer range correlations and thus to encode volume-law or sub-volume-law scaling, due to the absence of nonlocal connections between cells. We focus on the softmax-based model, and study the dependence of the scaling on σ at a fixed hidden state size $d_h = 40$. We ask if increasing nonlinearity in f can induce entanglement transitions in the thermodynamic limit, on top of the entanglement boost incurred by the move $f = \mathbb{1} \rightarrow f = \tanh$.

In Fig. 4(a), we fix $(f, g) = (\mathbb{1}, \text{SM})$ and plot the bipartite entanglement entropy as a function of system size L , for a number of values of $\sigma \in [0.1, 10.1]$, with the twin goals of cutting across the entanglement peaks in Figs. 2(a1) and 2(b1) and testing the large- σ regime. Each data point represents an average over an $N_{\text{init}} = 400$ random complex wavefunction ensemble. We achieve maximal system sizes of $L = 256$, and fit the points to curves described by $S_{\text{fit}}(L) = aL^\nu + b \ln L + c$. As per the parameter values in Table II in Appendix. H, we find that the curves at $\sigma = 0.1$ and $\sigma \gtrsim 2.1$ exhibit little to no entanglement and saturate to constants, which is strongly indicative of an area law in those regimes. However, near the entanglement peak of Fig. 2(b1), the $\sigma = 0.3$ and $\sigma = 0.4$ curves appear to exhibit logarithmic growth, with the $\mathcal{O}(\ln L)$ term dominant in S_{fit} at least up to $L = 256$, while at the actual peak ($\sigma \sim 0.2$), $\overline{S}_2^A$ appears to have a weak power-law dependence on top of the stronger logarithmic response. These results are further highlighted in Fig. 15(a) in Appendix D, where the semilog equivalent of Fig. 4(a) is shown. It is important to note that for all σ , any growth observed here is relatively weak, and could yet come to a halt in the limit $L \rightarrow \infty$, thus potentially producing a true area law.

Next, we introduce nonlinearity setting $f = \tanh$ while keeping everything else unchanged, and display the results in Fig. 4(b). Up to $L = 128$, and in most cases $L = 256$, the presence of nonlinearity induces much stronger entanglement scaling across the board, except at $\sigma = 0.1$, where the state remains close to $|\Psi\rangle = [\frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle)]^{\otimes L}$ for all L . As shown in Table II and emphasized in the semilog plot of Fig. 15(b), for $\sigma \in [0.2, 0.4]$, we see strong power-law dependence with $\nu \in [0.61, 0.78]$, while for $\sigma = 2.1, 5.1$, and 10.1 ($\sigma \gtrsim 0.6$), the scaling appears remarkably linear, with ν approaching 1. For some of the higher entropy curves, we are forced to discard data points at $L > 128$ because the averages start diverging exponentially (not shown), due to the




FIG. 4. Average bipartite entanglement entropy $\overline{S}_2^A$ as a function of system size L for (a) $(f, g) = (\mathbb{1}, \text{SM})$ and (b) $(f, g) = (\tanh, \text{SM})$ complex random RNN wavefunction ensembles, at various values of $\sigma \in [0.1, 10.1]$ and fixed $d_h = 40$. In panel (a), 2×10^4 Monte Carlo samples are used for the estimation, while 10^5 samples are used in panel (b). Each data point represents an average over $N_{\text{init}} = 400$ states. The dashed curves represent curves of best fit described by $S_{\text{fit}}(L) = aL^\nu + b \ln L + c$, with the fitted parameters displayed in Table II in Appendix H. The error bars shown are the errors associated with the random ensemble, with the RNN sampling errors negligible by comparison for the data points we include. We discard points in panel (b) for which the sampling error starts to dominate, which occurs at larger entropies and large L , causing the average entropy values to blow up and produce unphysical results. The inset in panel (b) zooms into the results at $L \leq 32$ (and fits S_{fit} to only those data points) and shows that if we had restricted our analysis to smaller systems, we may have ended up making confident but insufficiently backed claims about volume laws for $\sigma \lesssim 0.6$.

fact that the sampling error associated with the Monte Carlo estimation becomes large. Controlling this error would require orders of magnitude more samples, which is computationally prohibitive given the ensemble size.

We note that if we had relied on data from small system sizes alone ($L \lesssim 32$), we might have been led to believe that the RNN encodes a volume law at all σ , as the inset in Fig. 4(b) shows. Any claim based only on this would overlook the clear $\nu < 1$ power law or saturating behavior of the $\sigma \lesssim 0.6$ curves at larger L . The approximately linear, volume-law-like scaling observed for larger σ up to $L = 256$




FIG. 5. Entanglement spectrum level statistics presented in the form of the distribution $P(r)$ of the ratio of adjacent energy gaps for 20-spin random complex RNN wavefunction ensembles with (a) $(f, g) = (\tanh, \text{SM})$, (b) $(f, g) = (1, \text{SM})$, and (c) $(f, g) = (\tanh, \text{MOD})$. Each ensemble consists of 5000 random states, and the reduced density matrices are diagonalized exactly. Only eigenvalues of $\hat{\rho}_A$ satisfying $\lambda_{\rho_A}^n \geq 10^{-10}$ were included in the calculation, to avoid possible numerical precision issues brought about by even smaller values. The GOE, GUE, GSE, Poisson, and semi-Poisson distributions of random matrix theory [81] are plotted for comparison. The RNN produces GUE statistics in the peak regions of entanglement entropy [$\sigma \in \sim [0.2, 0.7]$ in panel (a), $\sigma \gtrsim 0.2$ in panel (c)], while the statistics tend to Poisson as we move away from those regions. In the inset of panel (a), the average eigenvalues of the reduced density matrices corresponding to the six wavefunction ensembles that reproduce GUE statistics in panel (a) ($\sigma = 0.2, 0.3, \dots, 0.7$) are displayed alongside the Marchenko-Pastur distribution of random matrix theory [82].

does not necessarily guarantee a volume law either in the thermodynamic limit, since the curves may eventually saturate. Nevertheless, our simulations show that the $\tanh$ nonlinearity substantially enhances entanglement scaling—power law at small σ and nearly linear at larger σ —across the system sizes typically used in finite-size scaling studies of one-dimensional quantum critical points [78–80].

C. Level statistics of entanglement spectra

We have touched on the similarities between random RNN wavefunctions and eigenstates relevant to the ergodic-MBL transition. In Sec. III A, we identified the $\sigma \rightarrow \infty$ limit of softmax-based RNNs as the MBL regime of the (d_h, σ) entanglement diagrams, while Sec. III B showed that nonlinearities in the fully connected layer may drive entanglement from area-law scaling to power-law or even potentially the volume-law scaling characteristic of thermal states. In Appendix E, we further demonstrate that ensembles of random RNN wavefunctions at $L = 40$ realize both decaying and nondecaying correlations, reminiscent of localized and delocalized eigenstates of the random-field Heisenberg model [72], with possible transitions between both regimes induced by the Gaussian width σ and the nonlinearity of the fully connected layer. Taken together, these results indicate that random RNN wavefunctions interpolate between the “MBL-like” and “thermal-like” behavior depending on σ and network nonlinearity.

To probe this connection further, we analyze the ratio of adjacent gaps in the entanglement spectrum of the ensemble of RNN wavefunctions, a standard metric for quantifying quantum chaos. Given a reduced density matrix $\hat{\rho}_A$, the

entanglement Hamiltonian is

$$\hat{H}_{\text{ent}} = -\ln \hat{\rho}_A, \quad (24)$$

with entanglement spectrum $\{E_n = -\ln \lambda_n^{(\rho_A)}\}$ from the eigenvalues $\{\lambda_n^{(\rho_A)}\}$ of $\hat{\rho}_A$. A useful diagnostic is the adjacent-gap ratio,

$$r_n \equiv \frac{\Delta_{n+1}}{\Delta_n} \equiv \frac{E_{n+1} - E_n}{E_n - E_{n-1}}, \quad (25)$$

or its related form

$$r_n^{\min} = \frac{\min(\Delta_{n+1}, \Delta_n)}{\max(\Delta_{n+1}, \Delta_n)}, \quad (26)$$

widely used to probe the ergodic-MBL transition [72,73].

The distribution $P(r)$, defined on the set $\{r_n\}$ in Eq. (25) and originally applied to the spectra of physical Hamiltonians, distinguishes thermalizing from nonthermalizing dynamics. Nonintegrable systems encode the random matrix statistics of Gaussian orthogonal ensembles (GOE; real symmetric), Gaussian unitary ensembles (GUE; complex Hermitian), or Gaussian symplectic ensembles (GSE; quaternionic Hermitian) [51,81,83–88], while integrable ones yield Poisson statistics [83,89,90]. Importantly, $P(r)$ is also relevant for entanglement spectra: Thermal eigenstates encode GOE, GUE, or GSE statistics, whereas localized ones exhibit semi-Poisson behavior [53].

In Fig. 5, we construct large ensembles of 20-spin random complex RNN states with fixed $d_h = 40$ and three (f, g) combinations that preserve the autoregressive property of the models. With these, we compute $P(r)$ for various values of σ in the vicinity of the entanglement peaks of Fig. 2. For $(f, g) = (\tanh, \text{SM})$, the statistics in the range $\sigma \in [0.2, 0.7]$ —around the peak at $\sigma_{\text{peak}} \sim 0.4$ —closely follow the GUE distribution, a striking outcome given the local

structure of the RNN. As we vary σ and move away from the peak region, most of the weight of $P(r)$ shifts to smaller values r and closely reaches a Poissonian distribution at $\sigma \gtrsim 2.0$. Similarly, close-to-Poisson statistics are produced at $\sigma = 0.1$, with a better fit expected for $\sigma \lesssim 0.1$. This again suggests that upon varying σ , RNN wavefunctions interpolate between thermal-like and localizedlike states. In the limits $\sigma \rightarrow 0$ or $\sigma \rightarrow \infty$, the RNN wavefunctions reduce to product states, making $P(r)$ ill defined, though we expect the Poissonian regime to persist for $\sigma \gtrsim 2.0$ for a considerable range.

When the hyperbolic tangent is turned off ($f = \mathbb{1}$), the entanglement peak, now found at $\sigma_{\text{peak}} \sim 0.2$, is significantly narrower and weaker [recall Fig. 2(B1)]. The corresponding statistics in the vicinity of this peak are shown in Fig. 5(b). The $\sigma = 0.2$ curve betrays a hint of Gaussian statistics in terms of its shape but is nowhere near the statistics of any of the Gaussian ensembles under consideration. Away from the peak, at $\sigma = 0.1$ and $\sigma = 0.3$, the statistics clearly tend to a Poissonian. Thus, the Gaussian statistics of quantum chaos emerge near entanglement entropy peaks of softmax-based random RNN states endowed with a strong nonlinearity, where the switch $f = \mathbb{1}$ to $f = \tanh$ induces a transition from nonthermal to thermal-like GUE level statistics.

If the softmax is replaced by the modulus function, entanglement is no longer suppressed, with the peak at small σ appearing to extend indefinitely in the limit $\sigma \rightarrow \infty$ [recall Fig. 2(c1)]. The level statistics of the corresponding entanglement Hamiltonian reflect this—the $P(r)$ curves in Fig. 5(c) follow the GUE distribution almost exactly, but this time for a much wider range of σ values, here shown up to $\sigma \sim 10$. Below $\sigma = 0.2$, the statistics drift from GUE toward the left (see, e.g., $\sigma = 0.1$), though without clearly reaching Poissonian behavior, highlighting the extent to which the modulus function rids the RNN of the entanglement suppression brought about by the softmax.

We note that the results (not shown) for the nonautoregressive RNN of Fig. 2(d1) where $(f, g) = (\tanh, \mathbb{1})$ are almost identical to the $(f, g) = (\tanh, \text{MOD})$ results, cementing the connection between the entanglement entropy maxima of our (d_h, σ) diagrams and the GUE statistics of random matrix theory.

The final result of this subsection is shown in the inset of Fig. 5(a). We plot the average eigenvalues of the reduced density matrices associated with the ensembles in Fig. 5(a) that encode the observed GUE-like statistics near the entanglement peak ($\sigma \in [0.2, 0.7]$). While the curve at the $\sigma = 0.4$ peak comes closest to the Marchenko-Pastur distribution of random matrix theory [82], none of the curves are close to it [similar behavior, not shown, for Figs. 5(b) and 5(c) architectures]. The Marchenko-Pastur distribution describes the average behavior of the eigenvalues of random rectangular matrices in the asymptotic limit, in other words, the eigenvalues of Wishart matrices $Y = XX^\dagger$ with X an $m \times n$ matrix and its elements *i.i.d.* Given that excited states of disordered Heisenberg Hamiltonians in the thermal phase have been shown to encode Marchenko-Pastur behavior in full while the corresponding eigenstates in the MBL phase follow the tails of the distribution [93], this further indicates that our random states are unique, and are different to known localized or

delocalized states, despite some shared properties as we have already seen.

D. Variational Monte Carlo

We now investigate the average behavior of the RNN within the context of variational Monte Carlo. We consider the TFIM in 1D,

$$\hat{H}_{\text{TFIM}} = -J \sum_{i=1}^{L-1} \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z - h \sum_{i=1}^L \hat{\sigma}_i^x, \quad (27)$$

as well as the Heisenberg model

$$\hat{H}_{\text{Heis}} = -J \sum_{i=1}^{L-1} \hat{\vec{S}}_i \cdot \hat{\vec{S}}_{i+1}, \quad (28)$$

both with open boundary conditions. We consider an ensemble of wavefunctions randomly initialized at various points in the (d_h, σ) diagram, perform a ground-state optimization of the selected Hamiltonian up to a predefined accuracy, and compute the number of iterations $\overline{\tau}_{\text{conv}}$ required to reach convergence. We wish to determine whether there is any correlation between the initial state average entanglement entropies and $\overline{\tau}_{\text{conv}}$ for the above Hamiltonians.

We begin with a small test system, the $L = 10$ 1D Heisenberg model. For a given (d_h, σ) pair, we fix $(f, g) = (\tanh, \text{SM})$, construct a random ensemble consisting of $N_{\text{init}} = 100$ complex autoregressive RNN wavefunctions, and perform the optimization as described above, using the Adam optimizer [94] for the parameter updates combined with the Monte Carlo approach outlined in Ref. [22] for the gradient and variational energy estimation. We terminate each simulation when we reach a relative error of $\epsilon_{\text{rel}} \equiv |E_{\text{RNN}} - E_0|/|E_0| = 10^{-3}$ and an energy variance per spin of $\sigma_H^2 \equiv (\langle \hat{H}^2 \rangle - \langle \hat{H} \rangle^2)/L = 10^{-3}$, with $E_{\text{RNN}} \equiv \langle \Psi_\lambda | \hat{H} | \Psi_\lambda \rangle$ and E_0 the reference ground-state energy (computed by exact diagonalization when $L = 10$, and DMRG later when $L = 40$). In Fig. 6(a), we plot $\overline{\tau}_{\text{conv}}$ as a function of σ for various values of $d_h \in [20, 100]$. As d_h increases, the position of the minimum convergence time $\sigma_{\text{min}}^{\overline{\tau}_{\text{conv}}}(d_h)$ shifts left and satisfies an approximate power law

$$\sigma_{\text{min}}^{\overline{\tau}_{\text{conv}}}(d_h) \sim d_h^{-\gamma}, \quad (29)$$

appearing to decrease at an increasing rate ($\gamma > 0$), an effect we showcase in Fig. 6(g). This is precisely how the entanglement entropy peaks of the RNN wavefunctions of Fig. 2(a) behave: Their σ positions $\sigma_{\text{max}}^{\overline{S}_2^A}$ decay polynomially as $d_h \rightarrow \infty$ and appear to satisfy

$$\sigma_{\text{max}}^{\overline{S}_2^A}(d_h) \sim d_h^{-\nu}, \quad (30)$$

with $\nu > 0$, as seen for comparison in Figs. 6(a) and 6(g). This does not imply that the left shift of $\sigma_{\text{max}}^{\overline{S}_2^A}$ is the actual cause of the corresponding $\sigma_{\text{min}}^{\overline{\tau}_{\text{conv}}}$ left shift. As the number of parameters in any neural network layer increases, a lower variance of parameter values is needed to prevent vanishing or exploding gradients, allowing for more efficient optimization and lower convergence times [92]. At the same time, we have already seen that with a rise in the total number of




FIG. 6. Average RNN convergence time $\overline{\tau}_{\text{conv}}$ as a function of σ and d_h for panel (a) the $L = 10$ quantum Heisenberg model and panel (d) the $L = 40$ critical TFIM. Convergence is satisfied when the relative error ϵ_{rel} and variance per spin σ_H^2 reach values of 10^{-3} . The $\overline{\tau}_{\text{conv}}$ averages for both models are taken over an ensemble of $(f, g) = (\tanh, \text{SM})$ RNN wavefunctions, with the states complex in the case of the Heisenberg model and positive in the case of the TFIM. The RNNs in panel (d) are constructed using the gated recurrent unit of Refs. [22,91], while those in panel (a) retain the vanilla RNN structure explored in this manuscript. The corresponding entanglement entropy curves are plotted in panels (a) and (d) for comparison, with the $L = 10$ model entropies computed exactly for $N_{\text{init}} = 10^4$ states, while the $L = 40$ curves are MC estimated with $N_{\text{init}} = 100$. In every optimization, the learning rate was fixed to $\eta = 5 \times 10^{-3}$, while 500 samples were used for the gradient estimation at each step. In panels (b) and (e), enlarged pictures of the $\overline{\tau}_{\text{conv}}$ results of panels (a) and (d) are shown, plotted against the convergence time results associated with the uniform Xavier-Glorot (XG) initialization scheme of Ref. [92]. In panels (c) and (f), the best-in-class Gaussian initialization $\overline{\tau}_{\text{conv}}$ results are plotted against the corresponding XG results, with the error bars calculated as the standard error of the ensemble convergence time results. Finally, panels (g) and (h) depict the approximate σ position of the $\overline{\tau}_{\text{conv}}$ minima, $\sigma_{\text{min}}^{\overline{\tau}_{\text{conv}}}$, as well as the position of the $\overline{S}_2^A$ maxima, $\sigma_{\text{max}}^{\overline{S}_2^A}$, corresponding to the results of panels (a) and (d), respectively, as a function of d_h , alongside dashed curves representing fits described by $ad_h^{-\delta} + b$.

parameters, smaller weight values are enough to guarantee the necessary number of nonzero probability amplitudes needed to maximize correlations and entanglement. In other words, if the forward pass that produces the amplitudes outputs fewer zeros, the backward pass that builds the gradients will as well, since backpropagation relies on the activations produced by a forward pass through the network [95]. Thus, there is likely no causation—only correlation.

The equivalent results for the $L = 40$ critical TFIM ($h/J = 1$) are shown in Figs. 6(d) and 6(h). Because the system is larger, we require a more realistic setting [22] for the optimization, hence we replace the vanilla cell with the gated recurrent unit of Refs. [22,91], and construct the ensemble using positive wavefunctions. Enforcing positivity simplifies the optimization given that the ground state of stoquastic Hamiltonians is real and positive [54]. We observe in Fig. 6(h) that the entanglement peak and the convergence time minimum both shift to the left once again with increasing number of parameters, although the entanglement results at $N_{\text{init}} = 100$

are noisier given the larger system under consideration. We also note that the $\sigma_{\text{min}}^{\overline{\tau}_{\text{conv}}}$ left shift only seems to satisfy polynomial decay with increasing rate at larger d_h . At smaller d_h , the decay appears polynomial with decreasing rate, but this is likely due to the noise inherent in the identification of the exact minimum.

In Figs. 6(b), 6(c), 6(e), and 6(f), we compare our $\overline{\tau}_{\text{conv}}$ results to the average convergence times produced by the XG initialization scheme of Ref. [92] in which the weights of each linear transformation $\{W_{ij}\}$ are drawn from the uniform distribution $U \sim [-\frac{1}{\sqrt{n}}, \frac{1}{\sqrt{n}}]$, where n is the size of the input to the transformation. The XG scheme is one of the most commonly used techniques for neural network initialization, including in the context of the quantum many-body problem [96,97]. Our bar graph comparisons of Figs. 6(c) and 6(f) suggest at the very least that a best-in-class Gaussian initialization scheme can outperform the XG scheme in terms of $\overline{\tau}_{\text{conv}}$ for the specific NQS, physical models and system sizes under consideration. Such a Gaussian scheme may not be as




FIG. 7. Average value of the ratio of the average bipartite entanglement entropy and the maximal entropy $\overline{S_2^A}/(L_A \ln 2)$ across N_{init} random 20-spin ATF wavefunctions as a function of embedding dimension d_{emb} , the Gaussian distribution width σ , and various combinations of the activation functions $\mathcal{A}$ and g [as per Eqs. (17) and (8)], namely, (a) $\mathcal{A} = \text{SM}$, $g = \text{SM}$, (b) $\mathcal{A} = \text{Circ}$, $g = \text{SM}$, (c) $\mathcal{A} = \text{SM}$, $g = \text{MOD}$, and (d) $\mathcal{A} = \text{SM}$, $g = 1$. The wavefunctions are complex, and the averages are taken over $N_{\text{init}} = 100$ random states for the autoregressive architectures [panels (a)–(c)] (2×10^5 Monte Carlo samples), and $N_{\text{init}} = 20$ otherwise [panel (d)] (204 800 samples). For all panels, we set $f_{\text{FL}} = \text{ReLU}$. The softmax attention in panel (a) yields consistently lower entanglement than the circulant attention in panel (b) across all values of σ and d_{emb} , indicating a general suppressive effect on quantum correlations induced by the attention mechanism. A comparison of panel (a) with panels (c) and (d) further reveals that the softmax in the output layer also suppresses entanglement, with the effect becoming particularly pronounced at large σ .

difficult to design as it may seem. Since $\sigma_{\text{min}}^{\overline{\tau_{\text{conv}}}}$ always seems to lie to the left of the $\overline{S_2^A}$ peaks, even cursory knowledge of the nature of the maximal entropy regions is likely to be a good indicator for the approximate $\sigma_{\text{min}}^{\overline{\tau_{\text{conv}}}}$ position. We suspect this effect might generalize to other NQS architectures such as the RBM, whose equivalent entanglement entropy diagrams are well established [52].

It is important to note that many of the above conclusions are model and phase dependent. For example, in the case of the TFIM, in the limit of $h/J \rightarrow \infty$ (deep in the disordered phase), the ground state tends to the lowest-lying eigenstate of the local $\hat{\sigma}^x$ operators, which happens to be the exact $\sigma = 0$ state of our (d_h, σ) diagrams as we have already seen. Here, any initialization close to $\sigma \sim 0$, regardless of d_h , is likely to minimize $\overline{\tau_{\text{conv}}}$, implying the absence of any notable left shift. Instead, we believe the effect is more likely to emerge in the context of critical Hamiltonians, the ground states of which encode greater complexity and are more difficult to simulate.

IV. AUTOREGRESSIVE TRANSFORMER: ENTANGLEMENT AND VARIATIONAL MONTE CARLO

A. Entanglement phase diagrams

We evaluate the bipartite entanglement entropy of ATF wavefunctions with exactly the same protocol used for our RNN ansatz. In Fig. 7, we show results for an $L = 20$ spin chain, averaging over $N_{\text{init}} = 100$ random initializations to compute the average entanglement entropy of complex-valued wavefunctions across the parameter space defined by d_{emb} and σ . Our primary objective is to investigate how nonlinear activation functions in the ATF architecture influence quantum entanglement. To this end, we compare two attention kernels, $\mathcal{A} = \text{SM}$ and $\mathcal{A} = \text{Circ}$, as defined in Sec. II C. For the output layer, we vary g as we did for the RNN. The role

of nonlinearities f_{FL} in the feedforward layers is discussed separately in Appendix F.

Our analysis reveals that the two attention mechanisms yield distinct entanglement characteristics, as shown in Figs. 7(a) and 7(b) where the comparison is made in the presence of the final layer softmax ($g = \text{SM}$). While both plots showcase the entanglement peaks characteristic of the softmax-based RNN, the ATF with circulant attention generates ~ 5 times more entanglement at the peak than an $\mathcal{A} = \text{SM}$ transformer, showing that a softmax away from the output layer can also contribute to entanglement suppression [98].

We also observe a general feature of softmax-based autoregressive architectures in Figs. 7(a) and 7(b). As σ increases, the weight distribution becomes broader, and the final layer softmax increasingly produces highly localized outputs. This behavior is analogous to that seen in the RNN, where $\sigma \rightarrow 0$ and $\sigma \rightarrow \infty$ lead to product states: At low variance, the system approaches the $+$ product state with vanishing entanglement; at high variance, a single dominant configuration emerges in the MBL limit.

In Figs. 7(c) and 7(d), we vary g while fixing $\mathcal{A} = \text{SM}$, and in doing so, further similarities with the RNN emerge. With $g = 1$ and $g = \text{MOD}$, the peak entanglement region extends indefinitely as $\sigma \rightarrow \infty$. This observation demonstrates that the large- σ entanglement suppression is primarily driven by the final softmax layer. We note, however, that a softmax in the attention layer clearly attenuates the peak $\overline{S_2^A}$ values produced by the transformer in the large σ limit, as compared to the maximal values generated by the $(f, g) = (\tanh, \text{MOD})$ and $(f, g) = (\tanh, 1)$ RNNs of Figs. 2(c1) and 2(d1).

Finally, although the nonlinear effects introduced by the layer normalizations in the ATF architecture are not analyzed in this study, their potential to suppress extreme values has been noted in a recent work [99].




FIG. 8. Average bipartite entanglement entropy $\overline{S}_2^A$ as a function of system size L for three combinations of nonlinearities: (a) $(\mathcal{A}, g) = (\text{Circ}, \text{SM})$, (b) $(\mathcal{A}, g) = (\text{SM}, \text{SM})$, and (c) $(\mathcal{A}, g) = (\text{SM}, \mathbb{1})$, evaluated over ensembles of complex random ATF wavefunctions. In all models, $f_{\text{FL}} = \text{ReLU}$. Each curve corresponds to a fixed embedding dimension $d_{\text{emb}} = 20$ and a different value of the Gaussian width $\sigma \in \{0.1, 0.2, 0.4, 0.5, 0.8\}$. In panels (a) and (b), we use 2×10^5 Monte Carlo samples for entropy estimation, while in panel (c) 204 800 samples are used. Each data point represents an average over $N_{\text{init}} = 500$ independently initialized wavefunctions. The dashed lines denote best-fit curves of the form $S_{\text{fit}}(L) = aL^\nu + b \ln L + c$, with the fitted parameters displayed in Table III in Appendix H. Error bars represent standard error of the mean over the initialization ensemble; ATF sampling errors are negligible for the data shown.

B. Entanglement scaling

Next, we investigate how the entanglement entropies of random ATF wavefunctions scale with system size L across different transformer architectures, as shown in Fig. 8. We adopt the same fitting form as for the RNN, $S_{\text{fit}}(L) = aL^\nu + b \ln L + c$, with the parameters values displayed in Table III in Appendix H. At $\sigma = 0.1$, all models exhibit small and almost constant values of $\overline{S}_2^A(L)$, indicating near-product state behavior and area-law scaling, similar to the RNNs. For all $(\mathcal{A}, g)$ combinations and all values of $\sigma \neq 0.1$ shown, the scaling appears either power law ($\nu < 1$), logarithmic, or linear ($\nu \sim 1$) up to $L \sim 100$.

We first investigate the impact of varying $\mathcal{A}$ in the presence of the final layer softmax. Comparing Figs. 8(a) and 8(b), we find that the circulant attention induces significantly stronger entanglement entropies for all $L \lesssim 100$ when compared to $\mathcal{A} = \text{SM}$. Both architectures produce power-law or logarithmic scaling at least up to $L \sim 100$, but the circulant-based ATF exhibits clearer volume-law-like growth in the entropy at both $\sigma = 0.2$ and $\sigma = 0.8$, respectively, to the left and right of the entanglement peak ($\sigma \sim 0.4$). For $(\mathcal{A}, g) = (\text{SM}, \text{SM})$, at $\sigma = 0.4$, we find $a \sim 0.06$, $\nu \approx 0.6$, and $b \approx 0$. At $\sigma = 0.5$, the fitted exponents ($a \sim 0.002$, $\nu \approx 1.3$, and $b \approx 0.08$) suggest a possible crossover between power-law and critical-like logarithmic [51, 100–102] scaling. Such a progression is not as evident in the $(\mathcal{A}, g) = (\text{Circ}, \text{SM})$ case, where the stronger power-law dependence remains clear across the peak. This tallies with our understanding of the softmax-based attention, which disproportionately amplifies large attention scores, while the circulant attention maintains a more uniform distribution over positions, preserving longer-range correlations [103]. An attention-induced entanglement scaling transition in random ATF wavefunctions in the thermodynamic limit is not impossible to conceive, especially in the vicinity of the

$\overline{S}_2^A$ peak and to its right, but more data at larger system sizes $L \gg 100$ are needed to strengthen this argument.

In Fig. 8(c), we study unnormalized wavefunctions by setting $(\mathcal{A}, g) = (\text{SM}, \mathbb{1})$, where we find that in the large σ regime, entanglement is no longer suppressed, as per Fig. 7(d). The model generates substantial entropy at large $\sigma \sim 0.8$, significantly more than both the $(\mathcal{A}, g) = (\text{SM}, \text{SM})$ and $(\mathcal{A}, g) = (\text{Circ}, \text{SM})$ models, and while there is no evidence of a logarithmic crossover as σ increases, the scaling remains at best power law with $\nu < 1$ (no apparent $\nu \sim 1$ volume-law behavior). For an analysis of the feedforward nonlinearity f_{FL} and its impact on scaling, we refer the interested reader to Appendix F. Of course, all curves shown in Fig. 8 may eventually saturate to a constant at larger L . We believe that the strong scaling up to at least $L \sim 100$ generated with a relatively low number of parameters enforces the belief that, like the RNN, the ATF can be used to simulate highly entangled ground states of models of interest, at the large enough sizes needed to adequately characterize their critical behavior [43, 58].

C. Level statistics of entanglement spectra

Next, we investigate the entanglement spectrum level statistics of random ATF wavefunctions. Here, we restrict ourselves to two autoregressive architectures with softmax-based attention at fixed $d_{\text{emb}} = 20$, $(\mathcal{A}, g) = (\text{SM}, \text{SM})$, and $(\mathcal{A}, g) = (\text{SM}, \text{MOD})$, showcasing the results in Fig. 9. We find that in the regions of maximal entanglement, ATF wavefunctions remarkably produce GUE level statistics, just like their autoregressive RNN counterparts in the maximally entangled regions of their own (d_h, σ) phase diagrams. As Fig. 9(a) shows, when $g = \text{SM}$, the GUE statistics are clear only at the $\overline{S}_2^A$ maxima of Fig. 7(a), near $\sigma \sim 0.5$; away from those peaks, the statistics become more Poisson-like, though the Poisson behavior is not as clear at small gap ratio r as it is at larger r .

When $g = \text{MOD}$ in Fig. 9(b), there is no final layer softmax suppression, and the statistics are GUE at almost all σ values shown, including at the larger σ values at which the $g = \text{SM}$ wavefunctions of Fig. 9(a) have already crossed into Poisson-like statistics. In other words, moving from the softmax to a square modulus norm induces GUE statistics at large σ . At the same time, transitions from Poisson to GUE and vice versa are also induced by the Gaussian width itself when $g = \text{SM}$.

While the GUE statistics indicate thermal-like behavior in maximally entangled random ATF wavefunctions, the eigenvalues of the reduced density matrix indicate nonthermal behavior, deviating from the Marchenko-Pastur distribution [82]. Thus, random ATF wavefunctions, like random RNN states, are neither fully thermal nor fully localized; instead they are “thermal-like” in some regimes, “MBL-like” in others.

D. Variational Monte Carlo

In Fig. 10(a), we present average convergence time results as a function of σ for random ATF wavefunction ensembles with fixed $(\mathcal{A}, f_{\text{FL}}, g) = (\text{SM}, \text{ReLU}, \text{SM})$, following the equivalent RNN analysis of Sec. III D. The optimizations are performed on a chain of $L = 100$ spins for the critical 1D




FIG. 9. Entanglement spectrum level statistics presented in terms of the distribution $P(r)$ of the ratio of adjacent energy gaps for 20-spin random complex ATF wavefunction ensembles with (a) $(\mathcal{A}, g) = (\text{SM}, \text{SM})$ and (b) $(\mathcal{A}, g) = (\text{SM}, \text{MOD})$. Each ensemble contains 10000 states, and their respective reduced density matrices are diagonalized exactly. Only eigenvalues $\lambda_{\rho_A}^{(n)} \geq 10^{-9}$ are retained to ensure numerical stability. The reference distributions from random matrix theory (GOE, GUE, GSE, Poisson, and semi-Poisson) [81] are shown for comparison. GUE-like statistics emerge in the maximal entanglement regions of the corresponding (d_{emb}, σ) diagrams ($\sigma \in [0.2, 0.5]$ in panel (a), $\sigma \gtrsim 0.2$ in panel (b)), while Poisson-like behavior appears to occur elsewhere. Panel (c) displays the average eigenvalues of the reduced density matrices associated with the six ensembles of panel (a) corresponding to $\sigma < 1.0$, and compares them with the Marchenko-Pastur distribution [82].

TFIM, and averaged over $N_{\text{init}} = 40$ simulations. Just like for the RNN, a minimum $\bar{\tau}_{\text{conv}}$ emerges at all d_{emb} values that we test, indicating the existence of an optimal σ for VMC. The position of that optimum $\sigma_{\text{min}}^{\bar{\tau}_{\text{conv}}}(d_{\text{emb}})$ appears to decrease at an

increasing rate and satisfy $\sigma_{\text{min}}^{\bar{\tau}_{\text{conv}}}(d_h) \sim d_{\text{emb}}^{-\gamma}$ for some $\gamma > 0$, notwithstanding the effects of noise. Of course the $\bar{\tau}_{\text{conv}}$ values improve with increasing d_{emb} , reflecting the greater expressivity induced by having more parameters in the model.

Perhaps more importantly, we find that, like the RNN, the best-in-class Gaussian initialization performs remarkably better than a uniform Xavier-Glorot scheme, reducing $\bar{\tau}_{\text{conv}}$ by factor of $\sim 2-3$ across the board. While the above effects may well be model and phase dependent, we expect to observe a left shift in $\sigma_{\text{min}}^{\bar{\tau}_{\text{conv}}}(d_{\text{emb}})$ for critical Hamiltonians, mirroring the expected left shift at $L = 100$ of the entanglement entropy maxima for the reasons discussed in Sec. III D.

V. CONCLUSION AND OUTLOOK

In this article, we have studied the entanglement features of Gaussian-initialized random autoregressive neural network wavefunctions, focusing on the recurrent neural network of Ref. [22] and the autoregressive transformer of Ref. [43]. For the RNN (ATF), we found that narrow entanglement peaks emerge in the $\bar{S}_2^A$ phase diagram at optimal values of the hidden state dimension d_h (embedding dimension d_{emb}) and Gaussian width σ , with the peak shifting leftward with increasing d_h (d_{emb}). We uncovered a so-called MBL limit, $\sigma \rightarrow \infty$, in which σ plays the role of the disorder strength in a many-body localized Hamiltonian with random field, and in which the final-layer softmax suppresses entanglement by generating wavefunctions that comprise a single random dominant configuration and corrections, akin to the eigenstates of MBL Hamiltonians in the localized phase [72]. We find that the suppression is caused not by the autoregressive property per se, but by the softmax itself due to its exponential dampening properties, and we prove this by proposing a square modulus normalization function (MOD) that preserves autoregression and for which the peak entanglement region extends indefinitely in the infinite σ regime. Remarkably, we find that within these maximal entanglement regions, both




FIG. 10. (a) Average ATF convergence time $\bar{\tau}_{\text{conv}}$ as a function of Gaussian width σ for various embedding dimensions d_{emb} in ground-state VMC optimizations of the critical 1D TFIM with $L = 100$ spins. Convergence is satisfied when the relative error ϵ_{rel} and variance per spin σ_H^2 reach values of 10^{-3} . The averages are taken over $N_{\text{init}} = 40$ initializations for an $(\mathcal{A}, f_{\text{EL}}, g) = (\text{SM}, \text{ReLU}, \text{SM})$ ensemble of ATF wavefunctions constrained to be positive. In all simulations, a fixed learning rate $\eta = 10^{-3}$ and 1024 samples for the gradient estimation were used. Gray “x” markers indicate the optimal σ minimizing $\bar{\tau}_{\text{conv}}$ for each d_{emb} , which shifts to smaller values as d_{emb} increases. In panel (b), the best-in-class Gaussian initialization results are compared with the results corresponding to a uniform Xavier-Glorot initialization scheme [92]. The results make it clear once again that for the $L = 100$ TFIM simulations using the ATF, the Gaussian initialization scheme performs remarkably better than the XG approach. In both panels, the error bars are calculated as the standard error of the ensemble convergence time results.

architectures encode the level statistics of the Gaussian Unitary Ensemble of random matrix theory [81], and that moving away from those peaks induces a transition to Poisson-like statistics, an effect that is especially clear for the RNN.

Next, we considered the impact of the various nonlinearities on display on the entanglement scaling behavior of the RNN and ATF. We found evidence of scaling transitions induced by specific nonlinearities such as the hyperbolic tangent in the fully connected layer of the vanilla RNN, as well as by the Gaussian width σ in both the RNN and ATF, at least up to the finite system sizes studied ($L = 256$ for the RNN, $L = 128$ for the ATF). Although this evidence does not necessarily imply scaling transitions in the thermodynamic limit, the large finite-size system results further reinforce the idea that these models can be used to target highly entangled ground states at the finite system sizes required for the characterization of quantum critical points, as recent works emphasize [41,43].

There are several avenues of future investigation. Naturally it would be interesting to extend the work done here to two-dimensional (2D) architectures, such as the 2D RNN of Refs. [22,23,38,40,41]. A “Goldilocks” zone of neural network initialization—linked to Hessian curvature—has been proposed [104] but debated [105]. The emergence of best-in-class Gaussian widths $\sigma_{\min}^{\tau_{\text{conv}}}$ suggests future work should probe how initialization curvature relates to VMC performance. Similarly, the anticipated increased expressivity brought about by the square modulus function compared to the softmax may offer improvements to the performance of VMC. Preliminary ground-state simulations of the Heisenberg and modified Haldane-Shastry models suggest that the RNN architecture coupled with the modulus normalization might perform better than softmax-based RNN wavefunctions for optimizations that use stochastic reconfiguration [106,107], in specific regimes of Gaussian initialization. Further testing is required to build and characterize any VMC approach that replaces the final layer SM with MOD, the success of which is likely to depend on the model at hand.

Away from VMC, we anticipate that further analytical work might result in an effective model for these autoregressive architectures, one that might elucidate the underlying reason behind their emergent GUE level statistics at the entanglement peaks, especially when the RNN itself is a locally connected model. We also expect that the RNN (ATF) may encode universal properties such as the exponent ν of Eq. (30), governing the left shift of the entanglement entropy peaks with increasing d_h (d_{emb}).

ACKNOWLEDGMENTS

We thank Matteo D’Anna, Jannes Nys, Matija Medvidovic, Zahra Farahmand, Alev Orfi, M. Schuyler Moss, Matthew Duschenes, Anna Dawid, Hannah Lange, and Yuxuan Zhang for valuable insight and discussion. We also would like to thank Mark Coatsworth, Sebastian Burlacu, Daniel Oltianu, and the entire engineering and computing teams at the Vector Institute for their generous help, guidance, and expertise on the simulation front. The sampling and variational Monte Carlo simulations were conducted using NetKet [108,109], JAX [110], and Flax [111]. We acknowledge the support of the Natural Sciences and Engineering Research

Council of Canada (NSERC). R.W. acknowledges support from the Flatiron Institute. The Flatiron Institute is a division of the Simons Foundation. J.C. acknowledges support from the Shared Hierarchical Academic Research Computing Network (SHARCNET), Compute Canada, and the Canadian Institute for Advanced Research (CIFAR) AI chair program. Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute [112]. This work was also supported as part of the “Swiss AI initiative” by a grant from the Swiss National Supercomputing Centre (CSCS) under project ID a05 on Alps.

DATA AVAILABILITY

Our code for the calculation of properties (e.g., entanglement entropy) of randomly initialized RNN wavefunctions as well as for their VMC optimization is made publicly available in Ref. [113], while the equivalent codes for the ATF are available in Ref. [114]. The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX A: ENTANGLEMENT SUPPRESSION DETAILS

In Table I, random samples drawn from four random RNN wavefunctions are shown, allowing us to compare the softmax normalization of Eq. (9) with the square modulus function of Eq. (10) in the limit of small ($\sigma = 0.4$) and large ($\sigma = 50$) Gaussian distribution widths. From the results, the softmax exponential suppression of all configurations but one in the $\sigma \rightarrow \infty$ regime is clear, with no such suppression present when $g = \text{MOD}$.

APPENDIX B: RNN OUTPUT FUNCTIONS

1. RNN probability distribution

We can understand the behavior of the entanglement in the small and large σ limits as follows. Consider a joint distribution $P_\lambda(\mathbf{x}) \equiv P(\mathbf{x}, \lambda)$ over binary random variables $\mathbf{x} \in \{0, 1\}^L$ and a set of weights $\{\lambda\}$ that form the vector λ . For notational clarity, we use $\mathbf{x}$ instead of σ here in order to clearly distinguish between the binary variables and the Gaussian width σ . According to the chain rule of probability,

$$P(\mathbf{x}, \lambda) = P(\mathbf{x}|\lambda)p(\lambda). \quad (\text{B1})$$

We are interested in the case where $p(\lambda) = \mathcal{N}(\mu, \sigma I)$ (I being the identity matrix) and where $P(\mathbf{x}|\lambda)$ corresponds to an autoregressive RNN:

$$P(\mathbf{x}|\lambda) = \prod_n P_\lambda(x_n|\mathbf{x}_{<n}).$$

Here, $P_\lambda(x_n|\mathbf{x}_{<n})$ is given by a Bernoulli distribution [115]

$$P_\lambda(x_n|\mathbf{x}_{<n}) = \text{Ber}(x_n|s(a_n)), \quad (\text{B2})$$

$$= x_n s(a_n) + (1 - x_n)(1 - s(a_n)), \quad (\text{B3})$$

TABLE I. Random samples produced by four complex random RNN wavefunctions with $L = 20$ spins in both the small and large network parameter regimes, constructed and initialized with the following hyperparameters: (a) $g = \text{SM}$, $\sigma = 0.4$, (b) $g = \text{SM}$, $\sigma = 50$, (c) $g = \text{MOD}$, $\sigma = 0.4$, and (d) $g = \text{MOD}$, $\sigma = 50$. In all four cases, a hyperbolic tangent is chosen as the activation function of choice in the fully connected layer. Clearly, the softmax suppresses all configurations but one in the limit of large Gaussian width σ , but the square modulus function does not.

sample	(a) $\sigma = 0.4$ (SM)	(b) $\sigma = 50$ (SM)	(c) $\sigma = 0.4$ (MOD)	(d) $\sigma = 50$ (MOD)
1	↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑
2	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑
3	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑
4	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑
5	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑	↓↑↑↑↑↑↑↑↑↑↑↑↑↑↑↑
...	...	...	...	...
1000	↑↓↑↓↑↑↑↑↑↑↓↑↓↑↑↑↑↓	↓↑↑↑↓↑↑↑↓↑↑↑↓↑↑↑↑↑	↓↑↑↓↑↑↑↑↑↑↓↑↑↓↑↑↑↓	↓↑↑↓↑↑↑↑↑↑↓↑↑↓↑↑↑↓
...	...	...	...	...

where we used the sigmoid activation function (effectively a softmax function for binary classification) regime, i.e.,

$$s(z) = \frac{1}{1 + e^{-z}}. \quad (\text{B4})$$

The input $a_n \in \mathbb{R}$ is given by

$$a_n = U\mathbf{h}_n + c, \quad (\text{B5})$$

where f is some elementwise smooth activation function and $\mathbf{h}_n$ is as defined in Eq. (4). The recursive nature of this distribution, combined with the nonlinear activation functions, makes this model hard to study.

2. Limiting behavior of the RNN ensemble

In the limit $\sigma \rightarrow 0$, all the network weights vanish and $a_n = 0$, $s(a_n = 0) = 1/2$ and thus all probabilities are given by $P_\lambda(x_n|\mathbf{x}_{<n}) = 1/2$, resulting in the state $|+\rangle^{\otimes L}$. We now consider the limit $\sigma \rightarrow \infty$. Based on our large σ results in Sec. III A and elaborated on in Appendix A, we assume that the random binary variables $\{x_n\}$ are uncorrelated in this

$$P(\mathbf{x}|\boldsymbol{\lambda}) = \prod_n P_\lambda(x_n), \quad (\text{B6})$$

where now

$$P_\lambda(x_n) = \text{Ber}(x_n|s(\theta_n)).$$

Here, we have made the assumption that in the limit $\sigma \rightarrow \infty$, the binary variables $\{x_n\}$ and any information about them encoded in the hidden states $\{\mathbf{h}_n\}$ get completely washed out by the large and normally distributed parameters $\{\lambda\}$ governing the linear transformations in Eqs. (B5) and (4). Thus, we assume a_n is effectively a sum of $d_h + 1$ Gaussian random variables that we define as θ_n . By the central limit theorem [116], the distribution of a sum of Gaussian random variables of infinite width σ is itself a Gaussian of infinite width, so we define the argument of the sigmoid, θ_n , as a single Gaussian random variable. We expect each RNN cell to produce a different sum, hence θ_n is specific to the n th cell. This gives




FIG. 11. Probability density function of the logit-normal distribution of Eq. (B11) with mean $\mu = 0$ and standard deviation σ .




FIG. 12. Ratio of the average entanglement entropy and the maximal entropy $\bar{S}_2^A / (L_A \ln 2)$ as a function of σ for varying values of d_h from 10 to 50, for a complex RNN wavefunction ensemble with $(f, g) = (\tanh, \text{SM})$ and $L = 20$ spins. The curves are cross sections of the results in Fig. 2(a1).




FIG. 13. Ratio of the average entanglement entropy and the maximal entropy $S_2^A / (L_A \ln 2)$ as a function of σ for varying values of $d_h \in [0, 200]$, for complex RNN wavefunction ensembles with $(f, g) = (\tanh, \text{SM})$ and $L = 20$ spins. The results are calculated in increments of $\Delta_{d_h} = 1$ and $\Delta\sigma = 0.1$, and each point represents an average over $N_{\text{init}} = 100$ random wavefunctions.

the joint probability distribution

$$P(\mathbf{x}, \boldsymbol{\lambda}) = P(\mathbf{x}, \boldsymbol{\theta}_\lambda) = \prod_n [\text{Ber}(x_n | \theta_n) p(\theta_n)], \quad (\text{B7})$$

where $\boldsymbol{\theta}_\lambda \equiv (\theta_1, \theta_2, \dots, \theta_L)$.

a. The single-spin case

Consider the case where $L = 1$ and $\lambda = \theta_1 \equiv \theta \in \mathbb{R}$ (we drop the subscript for convenience). Because we are studying ensembles of random RNN wavefunctions, we want to understand their average behavior, so we marginalize the joint distribution $P(\mathbf{x}, \boldsymbol{\theta}) = \text{Ber}(x | \theta) p(\theta)$ to obtain

$$P(x) = \frac{1}{\sigma} \int_{-\infty}^{\infty} \text{Ber}(x | s(\theta)) \phi\left(\frac{\theta - \mu}{\sigma}\right) d\theta, \quad (\text{B8})$$

where $\phi(t) = \mathcal{N}(0, 1) = \frac{1}{\sqrt{2\pi}} e^{-t^2/2}$ is the probability density function of the standard normal distribution. This gives

$$p(x = 1) = \frac{1}{\sigma} \int_{-\infty}^{\infty} s(\theta) \phi\left(\frac{\theta - \mu}{\sigma}\right) d\theta. \quad (\text{B9})$$

We substitute $s(\theta) = y$, with $\theta = \ln(y/(1-y))$ and

$$\begin{aligned} \frac{d\theta}{dy} &= \frac{d}{dy} \ln(y/(1-y)) \\ &= \frac{1-y}{y} \frac{1}{(1-y)^2} \\ &= \frac{1}{y(1-y)}, \end{aligned}$$

giving us

$$p(x = 1) = \frac{1}{\sigma} \int_0^1 y \phi\left(\frac{\ln(y/(1-y)) - \mu}{\sigma}\right) \frac{1}{y(1-y)} dy. \quad (\text{B10})$$

Hence, we can view the discrete probability of sampling $x = 1$ as a random variable itself, whose mean value is equal to $\mathbb{E}[y]$ under the density

$$p(y) = \frac{1}{\sigma} \phi\left(\frac{\ln(y/(1-y)) - \mu}{\sigma}\right) \frac{1}{y(1-y)}. \quad (\text{B11})$$

This function is known as the logit-normal distribution. Because of the marginalization procedure of Eq. (B8) and the definition $y \equiv s(\theta)$, we can interpret y as representing the probability of sampling $x = 1$ for a single random RNN wavefunction. One can verify that the cumulative density function (CDF) is given by the antiderivative of $p(y)$:

$$P(Y \leq y) = \Phi\left(\frac{\ln(y/(1-y)) - \mu}{\sigma}\right), \quad (\text{B12})$$

where Φ is the CDF of the standard normal distribution.

As can be seen in Fig. 11, the logit-normal distribution has sharp peaks at the points 0 and 1 if $\mu = 0$ and $\sigma \gg 1$. To formalize this, we can consider the probability that Y is not ϵ close to either 0 or 1 with $0 < \epsilon < 1$, given this distribution. With $\epsilon \ll 1$, we find

$$\begin{aligned} P(\epsilon \leq Y \leq 1 - \epsilon) &= P(Y \leq 1 - \epsilon) - P(Y \leq \epsilon) \\ &= \Phi\left(\frac{\ln((1-\epsilon)/\epsilon)}{\sigma}\right) \\ &\quad - \Phi\left(\frac{\ln(\epsilon/(1-\epsilon))}{\sigma}\right) \\ &= 1 - 2\Phi\left(\frac{\ln(\epsilon/(1-\epsilon))}{\sigma}\right), \end{aligned}$$

where we used $\Phi(-x) = 1 - \Phi(x)$. In the limit of large σ , we thus get

$$P(\epsilon \leq Y \leq 1 - \epsilon) \approx 1 - 2\Phi(0) = 1 - 2 \times \frac{1}{2} = 0,$$

hence all the probability mass concentrates around 0 and 1. In other words, we can rewrite $p(y)$ as a sum of two delta functions:

$$\lim_{\sigma \rightarrow \infty} p(y) = \frac{1}{2} \delta(y) + \frac{1}{2} \delta(y - 1). \quad (\text{B13})$$

Thus, for a random single-spin RNN wavefunction, $p(x = 1)$ is essentially always 0 or 1 with equal probability, so the state is $|0\rangle$ or $|1\rangle$, with the entirely random outcome dependent on the values of the Gaussian weights at initialization.

b. The multispin case

In the case where $L > 1$, a similar analysis holds. We have

$$P(\mathbf{x}) = \prod_{n=1}^L \left[\frac{1}{\sigma} \int_{-\infty}^{\infty} \text{Ber}(x_n | s(\theta_n)) \phi\left(\frac{\theta_n - \mu}{\sigma}\right) d\theta_n \right].$$

This gives

$$P(\mathbf{x} = (1, \dots, 1)) = \prod_{n=1}^L \left[\frac{1}{\sigma} \int_{-\infty}^{\infty} s(\theta_n) \phi\left(\frac{\theta_n - \mu}{\sigma}\right) d\theta_n \right].$$

Substituting $y_n = s(\theta_n)$ again gives

$$P(\mathbf{x} = (1, \dots, 1)) = \prod_{n=1}^L \left[\int_0^1 y_n p(y_n) dy_n \right],$$




FIG. 14. Standard deviation of the bipartite entanglement entropy across the random RNN wavefunction ensembles of Fig. 2 as a function of d_h , σ , and various combinations of the activation functions f and g . The averages of Fig. 2 and the standard deviations here are all calculated from the same set of entanglement entropy values, for each panel.

where $p(y_n)$ is the logit-normal distribution of Eq. (B11). If we consider the limiting distribution of $p(y_n)$ in the $\sigma \rightarrow \infty$ limit [Eq. (B13)], we find

$$\lim_{\sigma \rightarrow \infty} P(\mathbf{x} = (1, \dots, 1)) = \prod_{n=1}^L \left[\int_0^1 y_n \left(\frac{1}{2} \delta(y_n) + \frac{1}{2} \delta(y_n - 1) \right) dy_n \right]. \quad (\text{B14})$$

For every multispin random RNN wavefunction, each binary variable x_n is thus 1 or 0 with equal probability in the limit of large σ , resulting in a product state comprising a single configuration in the computational basis with random, uncorrelated spin orientations.

3. Summary and note on the modulus

The analysis above tells us that, in both regimes $\sigma \rightarrow 0$ and $\sigma \rightarrow \infty$, a random RNN wavefunction with a sigmoid (softmax) activation function in the output layer is a product state. This explains the lack of entanglement in Fig. 2 in the relevant limits.

To understand the behavior of the case $g = \text{MOD}$, we can see that it is scale invariant. Consider a Gaussian random variable $\theta \sim \mathcal{N}(0, \sigma I)$ that goes through the modulus function as

$$\text{MOD}(\theta) = \frac{\theta_i^2}{\sum_i \theta_i^2} = \frac{(z_i \sigma)^2}{\sum_i (z_i \sigma)^2} = \frac{z_i^2}{\sum_i z_i^2}, \quad (\text{B15})$$

where $\mathbf{z} \sim \mathcal{N}(0, I)$. In the limit $\sigma \rightarrow 0$, the input into the normalization function g in Eq. (5) is a two-component vector $U\mathbf{h}_n + \mathbf{c}$ with both elements almost identically near 0, so the

$|+\rangle^{\otimes L}$ state results just like in the case $g = \text{SM}$. From there, as σ starts to increase, the modulus does amplify some configurations at the expense of others in such a way to increase correlations and entanglement, but the scale invariance of the modulus described in Eq. (B15) appears to emerge clearly in the limit $\sigma \rightarrow \infty$, where the entanglement entropy is not suppressed as seen in Fig. 2(c1).

APPENDIX C: RNN ENTROPY AND FLUCTUATIONS

In Fig. 12, we display a plot of the entanglement entropy as a function of σ taken from various d_h cross sections of Fig. 2(a1). Clearly, the entropy peaks grow and shift to the left with increasing d_h , while the softmax suppresses entanglement in the $\sigma \rightarrow \infty$ regime. Beyond $d_h = 50$, the peak appears to approach $\sigma = 0$ at a polynomially decreasing rate in d_h , but its value seems to saturate at some maximal value near $\overline{S_2^A}/(L_A \ln 2) \approx 0.55$, as can be seen in Fig. 13.

In the main text, we touched on the idea of ansatz expressivity. We explore this idea further in Fig. 14, where we calculate and plot the standard deviations of the ensemble entropies of Fig. 2. We interpret these fluctuations as one of many possible measures of expressivity, and in that sense, we see the softmax as suppressing the RNN's expressive potential, the modulus as retaining it, and the unnormalized RNN as the most expressive of the lot. When the wavefunction is normalized, we expect increasing nonlinearity in the fully connected layer to increase the expressive power of the architecture, a conclusion that the results of Figs. 14(a1) and 14(a2) support. Finally, the fluctuations are boosted by the inclusion of the complex phase, as expected [70].




FIG. 15. Semilog version of Fig. 4, where the average bipartite entanglement entropy $\overline{S}_2^A$ as a function of system size L is shown for two RNN architectures at various values of σ (logarithmic scale on the x axis). The dashed curves represent $S_{\text{fit}}(L) = aL^\nu + b \ln L + c$, with the fitted parameters displayed in Table II in Appendix H. The error bars shown are the errors associated with the random ensemble.

APPENDIX D: RNN SCALING: SEMILOG PLOT

In Fig. 15, we showcase the semilog version of the RNN entanglement scaling results of Fig. 4, alongside the curves of best fit. The corresponding fit parameter values are displayed in Table II in Appendix H.

APPENDIX E: RNN CORRELATIONS

We now turn our attention to correlations. In Fig. 16, we construct random softmax-based RNN wavefunction ensembles of $L = 40$ spins, and calculate the connected correlation function $\ln |\langle \hat{\sigma}_i^z \hat{\sigma}_j^z \rangle_c|$ as a function of $|i - j|$, with $\langle \hat{\sigma}_i^z \hat{\sigma}_j^z \rangle_c \equiv \langle \hat{\sigma}_i^z \hat{\sigma}_j^z \rangle - \langle \hat{\sigma}_i^z \rangle \langle \hat{\sigma}_j^z \rangle$. The average is taken over all spin pairs corresponding to a fixed separation $|i - j| \equiv d$ as well as all $N_{\text{init}} = 1000$ wavefunctions in the random ensemble. When $f = 1$, the wavefunctions in the vicinity and to the right of the entanglement peak ($\sigma > \sim 0.2$) have correlations that exhibit minimal-to-no decay with increasing distance, at least up to the error bars generated by the random ensemble, akin to the eigenstates of the quantum Heisenberg Hamiltonian with random field in the thermal, delocalized




FIG. 16. Average connected correlation function $\ln |\langle \hat{\sigma}_i^z \hat{\sigma}_j^z \rangle_c|$ as a function of separation $|i - j|$ for complex random $L = 40$ RNN wavefunction ensembles corresponding to (a) $f = 1$, $d_h = 20$, (b) $f = 1$, $d_h = 50$, (c) $f = \tanh$, $d_h = 20$, and (d) $f = \tanh$, $d_h = 50$. The averages are taken over all spin pairs satisfying a fixed separation $|i - j| \equiv d$, as well as over all $N_{\text{init}} = 1000$ wavefunctions in a given ensemble. The error bars depicted are the errors produced by the randomness of the ensemble, with the RNN sampling errors minimal by comparison (2×10^4 samples used for the MC estimation). In all panels, the correlation functions are plotted for various Gaussian widths satisfying $\sigma \in [0.1, 10.1]$.

phase [72]. Turning on the hyperbolic tangent ($f = \tanh$) then induces a much wider range of Gaussian widths σ over which the wavefunctions encode objectively stronger




FIG. 17. Average bipartite entanglement entropy $\overline{S}_2^A$ as a function of system size L for two combinations of nonlinearities: (a) $(\mathcal{A}, f_{\text{FL}}) = (\text{Circ}, \text{ReLU})$ and (b) $(\mathcal{A}, f_{\text{FL}}) = (\text{Circ}, 1)$, evaluated over ensembles of complex random ATF wavefunctions. In all ensembles, $g = \text{SM}$. Each curve corresponds to a fixed embedding dimension $d_{\text{emb}} = 20$ and a different value of the Gaussian width $\sigma \in \{0.1, 0.2, 0.4, 0.5, 0.8\}$. For all data points, 2×10^5 MC samples are used for the entropy estimations. The averages are taken over $N_{\text{init}} = 500$ independently initialized wavefunctions, and the error bars specify the corresponding standard error of the mean. Dashed lines represent fits to the form $S_{\text{fit}}(L) = aL^\nu + b \ln L + c$.




FIG. 18. Average connected correlation function $\ln |\langle \hat{\sigma}_i^z \hat{\sigma}_j^z \rangle_c|$ as a function of spin separation $|i - j|$ for complex random $L = 20$ ATF wavefunction ensembles with circulant attention and a final layer softmax, corresponding to (a) $f_{\text{FL}} = \text{ReLU}$, $d_{\text{emb}} = 10$, (b) $f_{\text{FL}} = \text{ReLU}$, $d_{\text{emb}} = 30$, (c) $f_{\text{FL}} = \mathbb{1}$, $d_{\text{emb}} = 10$, and (d) $f_{\text{FL}} = \mathbb{1}$, $d_{\text{emb}} = 30$. Each data point represents an average over all spin pairs with fixed separation $|i - j| \equiv d$, as well as over an ensemble of $N_{\text{init}} = 100$ independently initialized transformer wavefunctions. Correlation curves are shown for various Gaussian widths $\sigma \in [0.1, 0.7]$. Error bars indicate the standard error of the mean with respect to the ensemble, with the ATF sampling errors minimal by comparison (2×10^6 samples used for the MC estimation).

correlations ($\ln |\langle \hat{\sigma}_i^z \hat{\sigma}_j^z \rangle_c| \in [-3, -2]$), but strangely enough, the correlations here all exhibit clearer nonzero decay, possibly exponential (characteristic of the behavior of more localized eigenstates), especially in and around the entanglement peak in the corresponding (d_h, σ) diagram. Given the possible stronger entanglement scaling induced by the move from $f = \mathbb{1}$ to $f = \tanh$ discussed in Sec. III B, this transition in decay behavior is unexpected, highlighting the nontrivial impact of nonlinearities on the average properties of the RNN.

We note that the left shift of the S_2^A peak with increasing d_h discussed in Sec. III A is visible in the form of stronger correlations at lower values of σ in Figs. 16(b) and 16(d) as compared to Figs. 16(a) and 16(c). We also note that there is a clear difference in average correlation behavior between wavefunctions at $\sigma < 0.2$ and those at $\sigma \gtrsim 0.2$, for both $f = \mathbb{1}$ and $f = \tanh$. This effect is particularly clear in Fig. 16(d), where at $\sigma = 0.1$, there is strong exponential decay up to $|i - j| \sim 5$ followed by small but constant correlations ($\ln |\langle \hat{\sigma}_i^z \hat{\sigma}_j^z \rangle_c| \approx -5.5$) for $|i - j| > 5$, as compared to the curves at $\sigma > 0.2$, where the correlations are much larger but decay over the entire $|i - j| \in [0, 20]$ range, with the decay significantly weaker than that observed initially at $\sigma = 0.1$.

In short, possible transitions in correlation decay behavior are induced by the crossing of the entanglement peak via a changing Gaussian width σ , as well as by a linear-to-nonlinear move in the activation function of the fully connected layer.

TABLE II. Values of a , ν , and b corresponding to the $S_{\text{fit}}(L) = aL^\nu + b \ln L + c$ curves of best fit for the RNN entanglement scaling data of Fig. 4. Values are reported to an accuracy of $\mathcal{O}(10^{-4})$.

(f, g)	σ	a	ν	b
($\mathbb{1}$, SM)	0.1	0.0019	0.0001	0.0003
($\mathbb{1}$, SM)	0.2	0.0119	0.6960	0.4327
($\mathbb{1}$, SM)	0.3	0.0000	0.0000	0.2492
($\mathbb{1}$, SM)	0.4	0.0000	0.0000	0.1553
($\mathbb{1}$, SM)	2.1	0.0000	0.0000	0.0055
($\mathbb{1}$, SM)	5.1	0.0000	0.0086	0.0007
($\mathbb{1}$, SM)	10.1	0.0000	1.7821	0.0000
(tanh, SM)	0.1	0.0018	0.0003	0.0003
(tanh, SM)	0.2	0.0416	0.6059	0.0825
(tanh, SM)	0.3	0.2297	0.6254	0.3036
(tanh, SM)	0.4	0.2780	0.7840	0.0000
(tanh, SM)	2.1	0.0415	0.9544	0.0000
(tanh, SM)	5.1	0.0162	0.9735	0.0000
(tanh, SM)	10.1	0.0087	0.9576	0.0000

APPENDIX F: ATF FEEDFORWARD LAYER NONLINEARITIES

When considering other sources of nonlinearity in the transformer architecture and their effects on the average entanglement entropy of random wavefunctions, we find that a move from $f_{\text{FL}} = \text{ReLU}$ to $f_{\text{FL}} = \mathbb{1}$ in the FFL, in the presence of the final layer softmax and a circulant attention scheme, leads to a reduction in objective entanglement entropy values, but only a weak one, as shown in Fig. 17. More importantly, it has no clear impact on the entanglement scaling behavior (see Table III in Appendix H), which appears to

TABLE III. Values of a , ν , and b corresponding to the $S_{\text{fit}}(L) = aL^\nu + b \ln L + c$ curves of best fit for the ATF entanglement scaling data of Figs. 8 and 17. Values are reported to an accuracy of $\mathcal{O}(10^{-4})$.

$(\mathcal{A}, f_{\text{FL}}, g)$	σ	a	ν	b
(Circ, ReLU, SM)	0.1	0.0000	0.0000	0.0000
(Circ, ReLU, SM)	0.2	0.0033	1.3705	0.0000
(Circ, ReLU, SM)	0.4	1.3256	0.3937	0.0000
(Circ, ReLU, SM)	0.5	0.2049	0.7133	0.0000
(Circ, ReLU, SM)	0.8	0.0090	1.0893	0.0366
(SM, ReLU, SM)	0.1	0.0654	0.0000	0.0000
(SM, ReLU, SM)	0.2	0.0296	0.1251	0.0000
(SM, ReLU, SM)	0.4	0.0566	0.6316	0.0000
(SM, ReLU, SM)	0.5	0.0018	1.2883	0.0823
(SM, ReLU, SM)	0.8	0.0023	0.9533	0.0207
(SM, ReLU, $\mathbb{1}$)	0.1	0.0868	0.0021	0.0064
(SM, ReLU, $\mathbb{1}$)	0.2	0.1097	0.0000	0.0356
(SM, ReLU, $\mathbb{1}$)	0.4	0.1919	0.4002	0.0000
(SM, ReLU, $\mathbb{1}$)	0.5	0.1332	0.5837	0.0000
(SM, ReLU, $\mathbb{1}$)	0.8	0.8890	0.3484	0.0000
(Circ, $\mathbb{1}$, SM)	0.1	0.0000	0.0000	0.0000
(Circ, $\mathbb{1}$, SM)	0.2	0.0172	1.1428	0.0000
(Circ, $\mathbb{1}$, SM)	0.4	0.1838	0.7134	0.0000
(Circ, $\mathbb{1}$, SM)	0.5	0.0379	0.9447	0.0138
(Circ, $\mathbb{1}$, SM)	0.8	0.0045	1.0991	0.0240

remain power law at all σ , with possible saturation in the thermodynamic limit.

APPENDIX G: ATF CORRELATIONS

In Fig. 18, we plot the average connected correlation function $\ln |\langle \hat{\sigma}_i^z \hat{\sigma}_j^z \rangle_c|$ as a function of $|i - j|$ for random $L = 20$ ATF wavefunction ensembles with circulant attention and a final layer softmax. We investigate the impact of varying σ , d_{emb} , and the FFL activation function f_{FL} on the correlation behavior. Just like in the case of the RNN, there appears to be a possible transition in decay behavior between the $\sigma = 0.1$ curve and all curves $\sigma > 0.1$, especially at larger d_{emb} , though this effect is not as clear as in the case of the RNN in

Appendix E. Compared to the nonlinearity in the fully connected layer of the RNN (f), f_{FL} appears to have very little impact on the nature of the decay. Many of the correlation functions on display are essentially constant over large distances, similar to the eigenstates of the random Heisenberg Hamiltonian in the delocalized phase [72]. We suspect that varying the attention mechanism will have a stronger impact on decay behavior than f_{FL} , but we leave this investigation for a future study.

APPENDIX H: PARAMETER VALUES

In this Appendix, the values of the fitted exponents in Figs. 4, 8, and 17 are shown in Tables II (RNN) and III (ATF).

-
- [1] E. Y. Loh, Jr., J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, Sign problem in the numerical simulation of many-electron systems, *Phys. Rev. B* **41**, 9301 (1990).
 - [2] A. W. Sandvik and J. Kurkijärvi, Quantum Monte Carlo simulation method for spin systems, *Phys. Rev. B* **43**, 5950 (1991).
 - [3] P. Henelius and A. W. Sandvik, Sign problem in Monte Carlo simulations of frustrated quantum spin systems, *Phys. Rev. B* **62**, 1102 (2000).
 - [4] M. Troyer and U.-J. Wiese, Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations, *Phys. Rev. Lett.* **94**, 170201 (2005).
 - [5] C. Wu and S.-C. Zhang, Sufficient condition for absence of the sign problem in the fermionic quantum Monte Carlo algorithm, *Phys. Rev. B* **71**, 155115 (2005).
 - [6] Z.-X. Li, Y.-F. Jiang, and H. Yao, Solving the fermion sign problem in quantum Monte Carlo simulations by Majorana representation, *Phys. Rev. B* **91**, 241117(R) (2015).
 - [7] Z. C. Wei, C. Wu, Y. Li, S. Zhang, and T. Xiang, Majorana positivity and the fermion sign problem of quantum Monte Carlo simulations, *Phys. Rev. Lett.* **116**, 250601 (2016).
 - [8] R. Mondaini, S. Tarat, and R. T. Scalettar, Quantum critical points and the sign problem, *Science* **375**, 418-424 (2022).
 - [9] P. Broecker and S. Trebst, Entanglement and the fermion sign problem in auxiliary field quantum Monte Carlo simulations, *Phys. Rev. B* **94**, 075144 (2016).
 - [10] C. H. Mak and D. Chandler, Solving the sign problem in quantum Monte Carlo dynamics, *Phys. Rev. A* **41**, 5709 (1990).
 - [11] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
 - [12] F. Verstraete and J. I. Cirac, Matrix product states represent ground states faithfully, *Phys. Rev. B* **73**, 094423 (2006).
 - [13] F. Verstraete, V. Murg, and J. I. Cirac, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, *Adv. Phys.* **57**, 143 (2008).
 - [14] L. Tagliacozzo, T. R. de Oliveira, S. Iblisdir, and J. I. Latorre, Scaling of entanglement support for matrix product states, *Phys. Rev. B* **78**, 024410 (2008).
 - [15] X. Yu, D. Pekker, and B. K. Clark, Finding matrix product state representations of highly excited eigenstates of many-body localized Hamiltonians, *Phys. Rev. Lett.* **118**, 017201 (2017).
 - [16] J. Tindall, A. Mello, M. Fishman, E. M. Stoudenmire, and D. Sels, Dynamics of disordered quantum systems with two- and three-dimensional tensor networks, [arXiv:2503.05693](https://arxiv.org/abs/2503.05693).
 - [17] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
 - [18] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborova, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
 - [19] J. Carrasquilla, Machine learning for quantum matter, *Adv. Phys.: X* **5**, 1797528 (2020).
 - [20] H. Lange, A. Van de Walle, A. Abedinnia, and A. Bohrdt, From architectures to applications: A review of neural quantum states, *Quantum Sci. Technol.* **9**, 040501 (2024).
 - [21] A. Chen and M. Heyl, Empowering deep neural quantum states through efficient optimization, *Nat. Phys.* **20**, 1476 (2024).
 - [22] M. Hibat-Allah, M. Ganahl, L. E. Hayward, R. G. Melko, and J. Carrasquilla, Recurrent neural network wave functions, *Phys. Rev. Res.* **2**, 023358 (2020).
 - [23] M. Hibat-Allah, R. G. Melko, and J. Carrasquilla, Investigating topological order using recurrent neural networks, *Phys. Rev. B* **108**, 075152 (2023).
 - [24] Z. Denis, A. Sinibaldi, and G. Carleo, Comment on “can neural quantum states learn volume-law ground states?” *Phys. Rev. Lett.* **134**, 079701 (2025).
 - [25] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, *Nat. Phys.* **13**, 431 (2017).
 - [26] F. Schindler, N. Regnault, and T. Neupert, Probing many-body localization with neural networks, *Phys. Rev. B* **95**, 245134 (2017).
 - [27] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. G. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
 - [28] M. Schmitt and M. Heyl, Quantum many-body dynamics in two dimensions with artificial neural networks, *Phys. Rev. Lett.* **125**, 100503 (2020).
 - [29] J. Nys, Z. Denis, and G. Carleo, Real-time quantum dynamics of thermal states with neural thermofields, *Phys. Rev. B* **109**, 235120 (2024).
 - [30] K. Donatella, Z. Denis, A. Le Boité, and C. Ciuti, Dynamics with autoregressive neural quantum states: Application to critical quench dynamics, *Phys. Rev. A* **108**, 022210 (2023).

- [31] L. Gravina, V. Savona, and F. Vicentini, Neural projected quantum dynamics: A systematic study, *Quantum* **9**, 1803 (2025).
- [32] A. Van de Walle, M. Schmitt, and A. Bohrdt, Many-body dynamics with explicitly time-dependent neural quantum states, *Mach. Learn.: Sci. Technol.* **6**, 045011 (2025).
- [33] L. Mauron and G. Carleo, Challenging the quantum advantage frontier with large-scale classical simulations of annealing dynamics, *arXiv:2503.08247*.
- [34] O. Sharir, Y. Levine, N. Wies, G. Carleo, and A. Shashua, Deep autoregressive models for the efficient variational simulation of many-body quantum systems, *Phys. Rev. Lett.* **124**, 020503 (2020).
- [35] D. Wu, L. Wang, and P. Zhang, Solving statistical mechanics using variational autoregressive networks, *Phys. Rev. Lett.* **122**, 080602 (2019).
- [36] T. D. Barrett, A. Malyshev, and A. I. Lvovsky, Autoregressive neural-network wavefunctions for *ab initio* quantum chemistry, *Nat. Mach. Intell.* **4**, 351 (2022).
- [37] C. Roth, Iterative retraining of quantum spin models using recurrent neural networks, *arXiv:2003.06228*.
- [38] M. Hibat-Allah, R. G. Melko, and J. Carrasquilla, Supplementing recurrent neural network wave functions with symmetry and annealing to improve accuracy, *arXiv:2207.14314*.
- [39] M. Hibat-Allah, E. M. Inack, R. Wiersema, R. G. Melko, and J. Carrasquilla, Variational neural annealing, *Nat. Mach. Intell.* **3**, 952 (2021).
- [40] M. S. Moss, R. Wiersema, M. Hibat-Allah, J. Carrasquilla, and R. G. Melko, Leveraging recurrence in neural network wavefunctions for large-scale simulations of Heisenberg antiferromagnets on the square lattice, *Phys. Rev. B* **112**, 134450 (2025).
- [41] M. S. Moss, R. Wiersema, M. Hibat-Allah, J. Carrasquilla, and R. G. Melko, Leveraging recurrence in neural network wavefunctions for large-scale simulations of Heisenberg antiferromagnets on the triangular lattice, *Phys. Rev. B* **112**, 134449 (2025).
- [42] E. Ibarra-García-Padilla, H. Lange, R. G. Melko, R. T. Scalettar, J. Carrasquilla, A. Bohrdt, and E. Khatami, Autoregressive neural quantum states of Fermi Hubbard models, *Phys. Rev. Res.* **7**, 013122 (2025).
- [43] K. Sprague and S. Czischek, Variational Monte Carlo with large patched transformers, *Commun. Phys.* **7**, 90 (2024).
- [44] D. Luo, Z. Chen, J. Carrasquilla, and B. K. Clark, Autoregressive neural network for simulating open quantum systems via a probabilistic formulation, *Phys. Rev. Lett.* **128**, 090501 (2022).
- [45] D. Luo, Z. Chen, K. Hu, Z. Zhao, V. M. Hur, and B. K. Clark, Gauge-invariant and anyonic-symmetric autoregressive neural network for quantum lattice models, *Phys. Rev. Res.* **5**, 013216 (2023).
- [46] Y.-H. Zhang and M. Di Ventura, Transformer quantum state: A multipurpose model for quantum many-body problems, *Phys. Rev. B* **107**, 075147 (2023).
- [47] H. Lange, G. Bornet, G. Emperauger, C. Chen, T. Lahaye, S. Kienle, A. Browaeys, and A. Bohrdt, Transformer neural networks and quantum simulators: A hybrid approach for simulating strongly correlated systems, *Quantum* **9**, 1675 (2025).
- [48] D. Fitzek, Y. H. Teoh, H. P. Fung, G. A. Dagnew, E. Merali, M. S. Moss, B. MacLellan, and R. G. Melko, RydbergGPT, *Mach. Learn.: Sci. Technol.* **6**, 045057 (2025).
- [49] M. Bortone, Y. Rath, and G. H. Booth, Impact of conditional modelling for a universal autoregressive quantum state, *Quantum* **8**, 1245 (2024).
- [50] A. D. King *et al.*, Beyond-classical computation in quantum simulation, *Science* **388**, 199 (2025).
- [51] D.-L. Deng, X. Li, and S. Das Sarma, Quantum entanglement in neural network states, *Phys. Rev. X* **7**, 021021 (2017).
- [52] X.-Q. Sun, T. Nebabu, X. Han, M. O. Flynn, and X.-L. Qi, Entanglement features of random neural network quantum states, *Phys. Rev. B* **106**, 115138 (2022).
- [53] S. D. Geraedts, R. Nandkishore, and N. Regnault, Many-body localization and thermalization: Insights from the entanglement spectrum, *Phys. Rev. B* **93**, 174202 (2016).
- [54] S. Bravyi, D. DiVincenzo, R. Oliveira, and B. Terhal, The complexity of stoquastic local Hamiltonian problems, *Quantum Inf. Comput.* **8**, 361 (2006).
- [55] K. Hornik, M. Stinchcombe, and H. White, Multilayer feedforward networks are universal approximators, *Neural Networks* **2**, 359 (1989).
- [56] M. Raghu, B. Poole, J. Kleinberg, S. Ganguli, and J. Sohl-Dickstein, On the expressive power of deep neural networks, in *Proceedings of the 34th International Conference on Machine Learning* (PMLR, Cambridge, MA, 2017), Vol. 70, pp. 2847–2854.
- [57] S. Roca-Jerat, M. Gallego, F. Luis, J. Carrete, and D. Zueco, Transformer wave function for quantum long-range models, *Phys. Rev. B* **110**, 205147 (2024).
- [58] L. L. Viteritti, R. Rende, and F. Becca, Transformer variational wave functions for frustrated quantum spin systems, *Phys. Rev. Lett.* **130**, 236401 (2023).
- [59] R. Rende and L. L. Viteritti, Are queries and keys always relevant? A case study on transformer wave functions, *Mach. Learn.: Sci. Technol.* **6**, 010501 (2025).
- [60] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, Attention is all you need, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Red Hook, NY, 2017), Vol. 30.
- [61] G. Ke, D. He, and T.-Y. Liu, Rethinking positional encoding in language pre-training, in *9th International Conference on Learning Representations, Vienna, Austria* (OpenReview.net, 2021).
- [62] U. Wennberg and G. E. Henter, The case for translation-invariant self-attention in transformer-based language models, in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing* (Volume 2 Short Papers) (Association for Computational Linguistics, 2021), pp. 130–140.
- [63] P. Shaw, J. Uszkoreit, and A. Vaswani, Self-attention with relative position representations, in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics Human Language Technologies* (Volume 2 Short Papers) (Association for Computational Linguistics, New Orleans, Louisiana, 2018), pp. 464–468.
- [64] J. L. Ba, J. R. Kiros, and G. E. Hinton, Layer normalization, *arXiv:1607.06450*.

- [65] R. M. Neal, Connectionist learning of belief networks, *Artif. Intell.* **56**, 71 (1992).
- [66] K. He, X. Zhang, S. Ren, and J. Sun, Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, in *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)* (IEEE Computer Society, Los Alamitos, CA, 2015), pp. 1026–1034.
- [67] A. Daniely, R. Frostig, and Y. Singer, Toward deeper understanding of neural networks: The power of initialization and a dual view on expressivity, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Red Hook, NY, 2016), Vol. 29, pp. 2261–2269.
- [68] M. B. Hastings, I. González, A. B. Kallin, and R. G. Melko, Measuring Rényi entanglement entropy in quantum Monte Carlo simulations, *Phys. Rev. Lett.* **104**, 157201 (2010).
- [69] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, Oxford, UK, 1999).
- [70] T. Grover and M. P. A. Fisher, Entanglement and the sign structure of quantum states, *Phys. Rev. A* **92**, 042308 (2015).
- [71] T. Grover and M. P. A. Fisher, Quantum disentangled liquids, *J. Stat. Mech.* (2014) P10010.
- [72] A. Pal and D. A. Huse, Many-body localization phase transition, *Phys. Rev. B* **82**, 174411 (2010).
- [73] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, Colloquium: Many-body localization, thermalization, and entanglement, *Rev. Mod. Phys.* **91**, 021001 (2019).
- [74] D. Wu, R. Rossi, F. Vicentini, and G. Carleo, From tensor-network quantum states to tensorial recurrent neural networks, *Phys. Rev. Res.* **5**, L032001 (2023).
- [75] Y. Levine, O. Sharir, and A. Shashua, Benefits of depth for long-term memory of recurrent networks, in *6th International Conference on Learning Representations, Vancouver, BC, Canada* (OpenReview.net, 2018).
- [76] Y. Levine and O. Sharir, N. Cohen, and A. Shashua, Quantum entanglement in deep learning architectures, *Phys. Rev. Lett.* **122**, 065301 (2019).
- [77] T.-H. Yang, M. Soleimanifar, T. Bergamaschi, and J. Preskill, When can classical neural networks represent quantum states? [arXiv:2410.23152](https://arxiv.org/abs/2410.23152).
- [78] H.-Q. Zhou, J.-H. Zhao, and B. Li, Fidelity approach to quantum phase transitions: Finite-size scaling for the quantum ising model in a transverse field, *J. Phys. A: Math. Theor.* **41**, 492002 (2008).
- [79] M. Campostrini, J. Nespolo, A. Pelissetto, and E. Vicari, Finite-size scaling at first-order quantum transitions, *Phys. Rev. Lett.* **113**, 070402 (2014).
- [80] H. Tian, T. He, and X. Wu, Finite-size scaling of half-chain entanglement entropy in the one-dimensional transverse field ising model and the xx model, *Phys. Rev. B* **111**, 104437 (2025).
- [81] Y. Y. Atas, E. Bogomolny, O. Giraud, and G. Roux, Distribution of the ratio of consecutive level spacings in random matrix ensembles, *Phys. Rev. Lett.* **110**, 084101 (2013).
- [82] V. A. Marchenko and L. A. Pastur, Distribution of eigenvalues for some sets of random matrices, *Math. USSR-Sb.* **1**, 457 (1967).
- [83] M. V. Berry and M. Tabor, Level clustering in the regular spectrum, *Proc. R. Soc. London, Ser. A* **356**, 375 (1977).
- [84] O. Bohigas, M. J. Giannoni, and C. Schmit, Characterization of chaotic quantum spectra and universality of level fluctuation laws, *Phys. Rev. Lett.* **52**, 1 (1984).
- [85] T. A. Brody, J. Flores, J. B. French, P. A. Mello, A. Pandey, and S. S. M. Wong, Random-matrix physics: Spectrum and strength fluctuations, *Rev. Mod. Phys.* **53**, 385 (1981).
- [86] T. Guhr, A. Müller-Groeling, and H. A. Weidenmüller, Random-matrix theories in quantum physics: Common concepts, *Phys. Rep.* **299**, 189 (1998).
- [87] P. J. Forrester, Log-gases and random matrices, *London Mathematical Society Monographs* (Princeton University Press, Princeton, NJ, 2010), Vol. 34.
- [88] *The Oxford handbook of random matrix theory*, edited by G. Akemann, J. Baik, and P. Di Francesco (Oxford University Press, Oxford, UK, 2011).
- [89] M. Rigol, V. Dunjko, V. Yurovsky, and M. Olshanii, Relaxation in a completely integrable many-body quantum system: An *ab initio* study of the dynamics of the highly excited states of 1D lattice hard-core bosons, *Phys. Rev. Lett.* **98**, 050405 (2007).
- [90] L. Vidmar and M. Rigol, Generalized Gibbs ensemble in integrable lattice models, *J. Stat. Mech.* (2016) 064007.
- [91] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, On the properties of neural machine translation: Encoder–decoder approaches, in *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation* (Association for Computational Linguistics, Doha, Qatar, 2014), pp. 103–111.
- [92] X. Glorot and Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, edited by Y. W. Teh and M. Titterton (PMLR, Chia Laguna Resort, Sardinia, Italy, 2010), Vol. 9, pp. 249–256.
- [93] Z.-C. Yang, C. Chamon, A. Hamma, and E. R. Mucciolo, Two-component structure in the entanglement spectrum of highly excited states, *Phys. Rev. Lett.* **115**, 267206 (2015).
- [94] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [95] K. P. Murphy, *Machine Learning: A Probabilistic Perspective* (MIT Press, Cambridge, MA, 2012).
- [96] P. Zhang, H. Shen, and H. Zhai, Machine learning topological invariants with neural networks, *Phys. Rev. Lett.* **120**, 066401 (2018).
- [97] W. T. Lou, H. Sutterud, G. Cassella, W. M. C. Foulkes, J. Knolle, D. Pfau, and J. S. Spencer, Neural wave functions for superfluids, *Phys. Rev. X* **14**, 021030 (2024).
- [98] Z. Qin, W. Sun, H. Deng, D. Li, Y. Wei, B. Lv, J. Yan, L. Kong, and Y. Zhong, cosFormer: Rethinking softmax in attention, in *The Tenth International Conference on Learning Representations* (OpenReview.net, 2022).
- [99] J. Zhu, X. Chen, K. He, Y. LeCun, and Z. Liu, Transformers without normalization, [arXiv:2503.10622](https://arxiv.org/abs/2503.10622).
- [100] P. Calabrese and J. Cardy, Entanglement entropy and quantum field theory, *J. Stat. Mech.* (2004) P06002.
- [101] F. D. M. Haldane, Exact Jastrow-Gutzwiller resonating-valence-bond ground state of the spin- $\frac{1}{2}$ antiferromagnetic Heisenberg chain with $1/r^2$ exchange, *Phys. Rev. Lett.* **60**, 635 (1988).

- [102] B. S. Shastry, Exact solution of an $s = \frac{1}{2}$ Heisenberg antiferromagnetic chain with long-ranged interactions, *Phys. Rev. Lett.* **60**, 639 (1988).
- [103] D. Han, Y. Pu, Z. Xia, Y. Han, X. Pan, X. Li, J. Lu, S. Song, and G. Huang, Bridging the divide: Reconsidering softmax and linear attention, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Red Hook, NY, 2024), Vol. 37, pp. 79221–79245.
- [104] S. Fort and A. Scherlis, The goldilocks zone: Towards better understanding of neural network loss landscapes, in *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)* (AAAI Press, 2019), pp. 3574–3581.
- [105] A. Vysogorets, A. Dawid, and J. Kempe, Deconstructing the goldilocks zone of neural network initialization, in *Proceedings of the 41st International Conference on Machine Learning (ICML)* (PMLR, Vienna, Austria, 2024), Vol. 235, pp. 49717–49732.
- [106] S. Sorella and L. Capriotti, Green function Monte Carlo with stochastic reconfiguration: An effective remedy for the sign problem, *Phys. Rev. B* **61**, 2599 (2000).
- [107] S. Sorella, M. Casula, and D. Rocca, Weak binding between two aromatic rings: Feeling the van der Waals attraction by quantum Monte Carlo methods, *J. Chem. Phys.* **127**, 014105 (2007),.
- [108] G. Carleo, K. Choo, D. Hofmann, J. ET Smith, T. Westerhout, F. Alet, E. J. Davis, S. Efthymiou, I. Glasser, S.-H. Lin, M. Mauri, G. Mazzola, C. B. Pereira, and F. Vicentini, Netket: A machine learning toolkit for many-body quantum systems, *SoftwareX* **10**, 100311 (2019).
- [109] F. Vicentini, D. Hofmann, A. Szabó, D. Wu, C. Roth, C. Giuliani, G. Pescia, J. Nys, V. Vargas-Calderón, N. Astrakhantsev, and G. Carleo, NetKet 3: Machine learning toolbox for many-body quantum systems, *SciPost Phys. Codebases* **7** (2022).
- [110] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of Python+NumPy programs (2018), <https://github.com/google/jax>.
- [111] J. Heek, A. Levskaia, A. Oliver, M. Ritter, B. Rondepierre, A. Steiner, and M. van Zee, Flax: A neural network library and ecosystem for JAX (2024), <https://github.com/google/flax>.
- [112] https://github.com/andrewjreissaty91/RandomRNN_Entanglement_VMC.
- [113] www.vectorinstitute.ai/#partners.
- [114] https://github.com/andrewjreissaty91/RandomATF_Entanglement_VMC.
- [115] J. V. Uspensky, *Introduction to Mathematical Probability* (McGraw-Hill, New York, 1937).
- [116] H. Fischer, A history of the central limit theorem: From classical to modern probability theory, *Sources and Studies in the History of Mathematics and Physical Sciences* (Springer New York, New York, NY, 2011), 1st ed.