

The statistical mechanics and machine learning of the α -Rényi ensemble

Andrew Jreissaty^{1,2,3,*} and Juan Carrasquilla^{1,2,3}

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

²*Vector Institute, MaRS Centre, Toronto, Ontario, M5G 1M1, Canada*

³*Department of Physics and Astronomy, University of Waterloo, Ontario, N2L 3G1, Canada*
(Dated: April 18, 2024)

We study the statistical physics of the classical Ising model in the so-called α -Rényi ensemble, a finite-temperature thermal state approximation that minimizes a modified free energy based on the α -Rényi entropy. We begin by characterizing its critical behavior in mean-field theory in different regimes of the Rényi index α . Next, we re-introduce correlations and consider the model in one and two dimensions, presenting an exact analysis of the former and devising an unconventional Monte Carlo approach to the study of the latter. Remarkably, we find that while mean-field predicts a continuous phase transition below a threshold index value of $\alpha \sim 1.303$ and a first-order transition above it, the Monte Carlo results in two dimensions point to a continuous transition at all α . We conclude by performing a variational minimization of the α -Rényi free energy using a recurrent neural network (RNN) ansatz where we find that the RNN performs well in two dimensions when compared to the Monte Carlo simulations. Our work highlights the potential opportunities and limitations associated with the use of the α -Rényi ensemble formalism in probing the thermodynamic equilibrium properties of classical and quantum systems.

I. INTRODUCTION

Simulating finite-temperature states, both in equilibrium and out-of-equilibrium, remains a significant challenge in the study of quantum many-body systems. Quantum Monte Carlo approaches, long considered state-of-the-art for the simulation of equilibrium states in quantum many-body systems, are plagued by fundamental sign problem issues in fermionic and frustrated quantum spin systems [1–7]. More recently, a large number of approaches for Gibbs state simulation involving the imaginary time evolution of a purified mixed state to produce thermal pure quantum states (TPQS) have been introduced [8–13]. Other pure state approaches such as minimally entangled typical thermal states (METTS) have also been proposed, leveraging matrix product state (MPS) algorithms along the way [14, 15]. In time, many in the community have turned to the variational method, proposing TPQS and density matrix ansätze parameterized by a set of parameters that are tuned to approximate the Liouvillian dynamics of mixed states coupled to Markovian baths using the time-dependent variational principle (TDVP) [16, 17]. Those behind the vast majority of these approaches have recognized a common issue: simulating the Gibbs state variationally by minimizing the Gibbs free energy at finite temperature is challenging due to the issues associated with computing the von Neumann entropy of a parameterized quantum density matrix. As such, thermal state approximations have started to emerge. One such approximation involves the minimization of a modified free energy known as the 2-Rényi free energy, where the von Neumann entropy is replaced by the second Rényi entropy [18]. In this way, the 2-Rényi ensemble, which

minimizes the 2-Rényi free energy, has provided a fresh breeding ground for quantum simulation of finite temperature states, in particular using MPS and neural network quantum state (NNQS) approaches [19, 20].

In the last few years, NNQS models ranging from restricted Boltzmann machines (RBMs) to convolutional neural networks (CNNs) have exploded onto the scene, providing highly expressive variational ansätze for the efficient simulation of ground state wavefunctions, the detection of continuous phase transitions and the reconstruction of quantum states [21–25]. The continued development of NNQS has since resulted in the emergence of a highly efficient autoregressive model based on recurrent neural networks (RNNs) which has been used for ground state wavefunction optimization in both frustrated and unfrustrated spin systems [26–28]. Some studies have sought to enhance RNN ground state optimizations by leveraging quantum simulation and Monte Carlo sampling data in the process [29, 30], demonstrating the flexibility of the overall NNQS approach.

Although the work in Refs. [18–20] has focused on studying finite-temperature properties of quantum systems through the Rényi ensemble, here we take a step back and examine whether the Rényi ensemble provides an accurate approximation of the Gibbs state at the classical level. We focus on the Ising model in one and two dimensions [31], which, in light of its analytical and numerical tractability, provides an ideal playground for understanding to what extent and in which regimes the α -Rényi ensemble reproduces the physics of the Gibbs state. We first consider a mean-field solution of the model within the ensemble, followed by a detailed exploration of the model in the presence of fluctuations through the development of a Markov-chain Monte Carlo technique specifically designed to target the Rényi ensemble. In the latter case, sampling via Monte Carlo presents a challenge as the distribution itself depends on the average en-

* ajreissaty@phys.ethz.ch

ergy, which we estimate via an iterative procedure. Beyond our Monte Carlo method, we consider variational approximations to the Rényi ensemble using recurrent neural networks and assess their quality by comparing their output to Monte Carlo and exact approaches.

Putting it all together, the paper is broken down as follows. In Sec. II, we introduce the Rényi ensemble machinery that is the foundation of this entire article. Next, we vet the Rényi ensemble approximation by applying it to the mean-field study of the Ising model in Sec. III, followed by an analytical treatment of the one-dimensional (1D) Ising model in this ensemble in Sec. IV. We reintroduce correlations in Sec. V, presenting Monte Carlo results for the two-dimensional (2D) Ising model in the Rényi ensemble, and in Sec. VI we compare those results with the RNN predictions. In Sec. VII, we conclude and discuss the future outlook of our work and thermal state approximations more broadly.

II. THE RÉNYI ENSEMBLE

We consider the α -Rényi free energy given by

$$F_\alpha = \text{Tr}(\rho H) - T \frac{1}{1-\alpha} \log [\text{Tr}(\rho^\alpha)]. \quad (1)$$

Here, α represents the Rényi index satisfying $\alpha \geq 1$, and ρ is the density matrix of the system. The α -Rényi ensemble is defined as the density matrix $\rho^{(\alpha)}$ that minimizes F_α . It is expressed as $\rho^{(\alpha)} = \sum_k p_k^{(\alpha)} |E_k\rangle \langle E_k|$, as previously derived in Ref. [18] and extensively explored in Refs. [19, 20]. The eigenstates of the Hamiltonian $\{|E_k\rangle\}$ have corresponding k^{th} energy levels E_k with degeneracies N_k . The probabilities $p_k^{(\alpha)} \in [0, 1]$ are given by

$$p_k^{(\alpha)} = \frac{[1 - \beta \frac{\alpha-1}{\alpha} (E_k - \bar{E})]^{\frac{1}{\alpha-1}}}{Z_\alpha}, \quad (2)$$

$$Z_\alpha = \sum_{k=0}^{n_\beta} N_k \left[1 - \beta \frac{\alpha-1}{\alpha} (E_k - \bar{E}) \right]^{\frac{1}{\alpha-1}}, \quad (3)$$

$$E_k \leq \frac{\alpha}{\beta(\alpha-1)} + \bar{E}, \quad (4)$$

where the temperature of the ensemble is T and its inverse is $\beta = 1/T$. The probabilities $p_k^{(\alpha)}$ satisfy the constraint $\text{Tr}(\rho^{(\alpha)} H) \equiv \bar{E}$, where $\bar{E} = \langle H \rangle$ is the average energy of the system, and Z_α is the partition function of the generalized ensemble. The condition in Eq. (4) must be satisfied in order to ensure the positive semi-definiteness of $\rho^{(\alpha)}$. It is possible to show that in the limit $\alpha \rightarrow 1$, the Rényi ensemble tends exactly to the Gibbs state [32]. The sum in Eq. (3) is over all n_β eigenstates that satisfy Eq. (4), a number that depends on the inverse temperature β .

The average energy \bar{E} is computed by solving the fixed

point equation

$$\text{Tr}(\rho^{(\alpha)} H) = \bar{E} \longrightarrow \sum_{k=0}^{n_\beta} N_k E_k p_k^{(\alpha)} = \bar{E}. \quad (5)$$

As discussed in App. B, we observe that the $\text{Tr}(\rho^{(\alpha)} H) = \bar{E}$ fixed point is attractive for all $\alpha > 1$ for the 1D and 2D Ising models with no external field. This feature proves especially useful in numerical simulations as it enables the possibility to find \bar{E} iteratively, which we use for both exact and Monte Carlo simulations of the Ising model.

III. ISING MODEL: MEAN-FIELD

We first consider the α -Rényi ensemble within mean-field theory, focusing on the classical Ising model $H = -J \sum_{\langle i,j \rangle} \sigma_i \sigma_j$ with $J > 0$. Our mean-field calculation follows the approach in Ref. [33], which is based on a factorized density matrix

$$\rho = \bigotimes_{i=1}^N \rho_i \equiv \bigotimes_{i=1}^N \begin{bmatrix} \frac{1+m}{2} & 0 \\ 0 & \frac{1-m}{2} \end{bmatrix}. \quad (6)$$

We minimize the resulting α -Rényi free energy with respect to the variational parameter m . Restricting m to the interval $[-1, 1]$ allows for the interpretation of ρ_i as a classical probability distribution over the binary spin values $\{+1, -1\}$ such that the average spin value is m . This product state approach is equivalent to other mean-field formulations and can be shown to recover the mean-field equation for the magnetization $m = \tanh[m(qJ)/T]$ of the Ising model in the Gibbs state, with $q = 2d$ the coordination number associated with a hypercubic lattice in D dimensions. Applying Eq. (6) to the α -Rényi free energy leads to a free energy per spin of

$$f_\alpha(m) = -\frac{1}{2}(qJ)m^2 - T \frac{1}{1-\alpha} \log \left[\left(\frac{1+m}{2} \right)^\alpha + \left(\frac{1-m}{2} \right)^\alpha \right]. \quad (7)$$

Remarkably, while the mean-field free energy in the Rényi index interval $1 \leq \alpha \lesssim 1.3$ predicts a continuous phase transition for the Ising model, for $\alpha \gtrsim 1.3$ the transition is first-order. This can be seen in Fig. A.1 (see App. A) and Fig. 2, where the hallmarks of continuous and first-order transitions emerge for different values of α . This stands in contrast with the well-known case of the Ising model in the Gibbs state ($\alpha \rightarrow 1$), where the mean-field transition is continuous with a critical temperature $T_c = qJ$. The appearance of a first-order transition at larger α arises because as T increases, higher energy states can be “suddenly” turned on and made accessible to the system due to the nature of the Rényi constraint (see RHS of Eq. (4)), which hints at the possibility of a discontinuous jump in the value of the mean-field order parameter at some transition temperature T^* . On

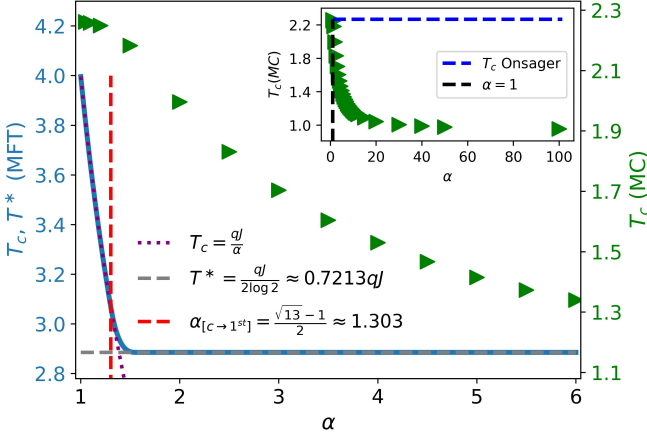


FIG. 1. Mean-field critical temperature T_c and transition temperature T^* (solid blue), as well as T_c as extracted from Monte Carlo simulations (green), as a function of the Rényi index α , for the 2D Ising model ($q = 4$). The mean-field data is extracted numerically from Eq. (7), while the Monte Carlo T_c values are computed by data collapse (Sec. V). The Monte Carlo error bars are smaller than the size of the data points. We show the exact functional form for T_c in the continuous regime (purple dotted curve for $T_c = qJ/\alpha$), the constant value of $qJ/(2 \log 2)$ that T^* takes for all $\alpha > \bar{\alpha} \sim 1.56$ (grey dashed horizontal line) and the threshold $\alpha_{[c \rightarrow 1^{\text{st}}]}$ (red dashed vertical line) that separates the continuous and first-order regimes. J has been set to 1. The limit $\lim_{\alpha \rightarrow 1} T_c = qJ$ recovers the mean-field result for the Ising model in the Gibbs state. The inset displays Monte Carlo T_c data for large Rényi index, where T_c is seen to approach $T_c \sim 1$ as $\alpha \rightarrow \infty$.

the other hand, values of α closer to 1 produce a continuous transition since the Rényi ensemble tends to the Gibbs state as $\alpha \rightarrow 1$. We can see from Eq. (4) that as α approaches 1, more and more higher energy states are rendered accessible to the system at any given temperature, making discontinuous jumps in m less likely at the mean-field level.

We now derive expressions for the critical temperature T_c in the continuous regime and the transition temperature T^* in the first-order regime. In particular, we focus on their dependence on the Rényi index α . In between, we also derive the value of α that exactly separates the two regimes, which we denote $\alpha_{[c \rightarrow 1^{\text{st}}]}$. Let us assume that α is such that the mean field α -Rényi free energy in Eq. (7) describes a continuous transition. Then T_c is the temperature at which the nature of the extremum at $m = 0$ changes from a local maximum to the global minimum. To derive it, we compute $\frac{\partial^2 f_\alpha}{\partial m^2} \Big|_{m=0}$, set it to 0 and solve for T_c . We find

$$\frac{\partial^2 f_\alpha}{\partial m^2} \Big|_{m=0, T=T_c} = -qJ + \alpha T_c = 0 \longrightarrow T_c = \frac{qJ}{\alpha}. \quad (8)$$

This expression recovers the Gibbs state mean-field limit $\lim_{\alpha \rightarrow 1} T_c = \lim_{\alpha \rightarrow 1} qJ/\alpha = qJ$.

Eq. (8) is valid for $\alpha \in [1, \alpha_{[c \rightarrow 1^{\text{st}}]}]$, i.e. the continuous regime of α values. It is possible to evaluate $\alpha_{[c \rightarrow 1^{\text{st}}]}$ exactly. The procedure involves computing the Taylor

expansion of f_α about $m = 0$ to 6th order in m , which we denote as $f_\alpha^{(6)}$, extremizing the result, and subsequently identifying the regime of α values for which $f_\alpha^{(6)}$ allows for the possibility of five real extrema depending on the temperature T , which is a hallmark of a first-order transition. The reason we conduct this analysis to only $\mathcal{O}(m^6)$ and not greater is that $f_\alpha^{(6)}$ captures the macroscopic "extremal shape" of the true free energy f_α when f_α has five extrema. In other words, whenever f_α has five extrema in the interval $m \in [-1, 1]$, $f_\alpha^{(6)}$ also has five extrema, although for the latter, the interval may have to be widened to observe them all depending on the specific Rényi index under consideration. As such, a higher order analysis is not needed. The details of the procedure are laid out in App. A. It finds

$$\alpha_{[c \rightarrow 1^{\text{st}}]} = \frac{\sqrt{13} - 1}{2} \approx 1.303, \quad (9)$$

which, as opposed to T_c , is independent of the dimensionality of the system. In summary, mean-field theory predicts a continuous symmetry-breaking phase transition for the Ising model in the α -Rényi ensemble for $\alpha \in \left[1, \frac{\sqrt{13}-1}{2}\right]$ and a first-order transition for $\alpha \in \left(\frac{\sqrt{13}-1}{2}, \infty\right)$.

We now focus on the dependence of the first-order transition temperature T^* on α . By examining the dependence of f_α on m at various values of α in the first-order regime, we find that for α greater than or equal to some value $\bar{\alpha}$, f_α is globally minimized in the interval $m \in [-1, 1]$ at $m = \pm 1$ or $m = 0$, depending on the temperature, meaning that the jump in magnetization as T^* is crossed is exactly $m_{\text{gap}} = 1$ for all $\alpha \geq \bar{\alpha}$. Thus, the transition temperature for all $\alpha \geq \bar{\alpha}$ can be derived by setting $f_\alpha(m = \pm 1) = f_\alpha(m = 0)$, which leads to

$$-\frac{1}{2}(qJ) = -T \frac{1}{1-\alpha} \log(2^{1-\alpha}) \longrightarrow T^* = \frac{qJ}{2 \log 2}. \quad (10)$$

Using a simple numerical approximation, we find $\bar{\alpha} \sim 1.56$. In Fig. 1, we collect all the above results and plot the mean-field critical temperature T_c and the first-order transition temperature T^* as a function of α for the two-dimensional Ising model in the α -Rényi ensemble. The plot also includes results from our Monte Carlo data collapse for comparison, described in detail in Sec. V below. In Fig. 2, we plot the absolute value of the mean-field magnetization per spin \bar{m} (i.e. the value of m that minimizes $f_\alpha(m)$) as a function of temperature for various values of α .

We now explore the critical exponents of the continuous phase transition regime predicted by the α -Rényi ensemble, i.e., for $1 \leq \alpha \leq \frac{\sqrt{13}-1}{2}$. Since the α -Rényi mean-field free energy can be written analytically in terms of m for $T \sim T_c$ (where we have $|\bar{m}| \sim 0$) by performing a Taylor approximation of the logarithm (see Eq. (A1) for the $\mathcal{O}(m^6)$ expression), there must exist critical exponents that describe the behavior of the magnetization

and the divergences of thermodynamic quantities such as the specific heat and magnetic susceptibility within mean-field theory. Specifically, our goal is to derive the dependence of these exponents on α , if any. To that end, we only require the free energy to 4th order in m , $f_\alpha^{(4)}(m)$, equivalent to Eq. (A1) less the 6th order term. We find that the mean-field critical exponents β , α_{c_v} , γ and δ (we denote the specific heat critical exponent as α_{c_v}) take the exact same values for the Ising model in the Rényi ensemble as they do in the Gibbs state: $\beta = 1/2$, $\alpha_{c_v} = 0$, $\gamma = 1$ and $\delta = 3$.

While the coefficients modulating the divergences of some of the quantities of interest (e.g. A_+ and A_- , as in $\chi \sim A_+|t|^{-\gamma}$ for $t > 0$, where $t \equiv |T - T_c|/T_c$ is the reduced temperature and χ is the susceptibility) do indeed depend on the Rényi index α , in mean-field theory, we find that the critical exponents listed above do not—a remarkable result. The last of the relevant critical exponents for this discussion is ν , which is the critical exponent describing the divergence of the correlation length ξ according to $\xi \sim |t|^{-\nu}$. For the Ising model in the Gibbs state, its derivation involves re-introducing fluctuations into the partition function with the derivation heavily reliant on the presence of the Gibbs state exponentials [34]. An attempt at following an analogous argument for the α -Rényi ensemble presents us with the challenge of evaluating a partition function whose number of terms depends on the temperature-dependent Rényi constraint in Eq. (4), and whose terms depend on the average energy, which appears daunting to solve for analytically for most values of α . However, mean-field theory predicts that β , α_{c_v} , γ and δ not only all exist for $\alpha > 1$ but are also α -independent, and since the divergence of thermodynamic quantities is understood in statistical physics to stem from the divergence of the correlation length and the resulting scale invariance, we claim that ν takes on the mean-field Gibbs state value of $1/2$ for all α .

IV. 1D ISING MODEL: ANALYTICAL ARGUMENTS

Extremal Cases—We now focus on the 1D Ising model given by $H = -J \sum_{i=1}^N \sigma_i \sigma_{i+1}$ with periodic boundary conditions $\sigma_{N+1} = \sigma_1$. It is known that the classical Ising model in 1D exhibits no spontaneous symmetry breaking at finite temperature in the Gibbs state. We investigate if long-range order at finite T is possible in the generalized α -Rényi ensemble and, if so, at what values of α . We start with the limit $\beta \rightarrow \infty$. The constraint in Eq. (4) becomes

$$E_j \leq \lim_{\beta \rightarrow \infty} \left[\frac{\alpha}{\beta(\alpha - 1)} + \bar{E} \right] = \bar{E},$$

which means all allowed microstates have energies lower than the average energy. Since $\bar{E} = \sum_{j=0}^{n_\beta} N_j E_j p_j^{(\alpha)}$ is a convex combination of the allowed E_j values, then for




FIG. 2. Absolute value of the mean-field magnetization per spin \bar{m} that minimizes $f_\alpha(m)$ as a function of T for various values of α . For all $\alpha > \bar{\alpha} \sim 1.56$, the magnetization jump m_{gap} is 1, while for $\alpha \in \left(\frac{\sqrt{13}-1}{2}, \sim 1.56\right)$, $m_{\text{gap}} < 1$. The red curve depicts the magnetization at exactly $\alpha_{[c \rightarrow 1st]} = \frac{\sqrt{13}-1}{2}$. Below this value, the curves are characteristic of a continuous transition. As $\alpha \rightarrow 1$, the critical temperature tends to the Gibbs ensemble prediction $T_c = qJ$. We have set $qJ = 1$ for simplicity.

a finite system, the only way to realize $E_j \leq \bar{E}$ for all allowed states is to ensure that only one state is allowed: the ground state, with all spins aligned. Thus, $\bar{E} = E_0$ where E_0 is the ground state energy. In the thermodynamic limit, higher energy fixed points are possible at $T = 0$ (see analysis in finite temperature section below which applies here as well), but the fixed point that globally minimizes the free energy must be the ground state fixed point $\bar{E} = E_0$ because at $T = 0$, the free energy is simply \bar{E} . The ground state is thus the only allowed state at zero temperature in the α -Rényi ensemble in the thermodynamic limit. In addition, while a finite system occupies the all-up and all-down ground states equally, if we assume local fluctuations only, an infinite system must choose one or the other, because in the limit $N \rightarrow \infty$, a jump from one symmetry-broken regime to the other cannot be considered local. Therefore, at zero temperature, the symmetry is broken in this limit.

In the opposite limit, $\beta = 0$, we can see from the Rényi constraint that all microstates become accessible, for all $\alpha \geq 1$. The partition function evaluates to 2^N , all probabilities equalize as $p_j^{(\alpha)} = 1/2^N$, and the magnetization vanishes. As expected, there is no long-range order at infinite temperature.

Finite Temperature—Let us now consider finite temperature. In the case of the Ising model in the Gibbs state, the solution is found by evaluating the partition function analytically and using the result to derive the magnetization. Instead, we follow a different strategy and make a Peierls argument [35]. If the system starts in one of the two symmetry-broken ground states at $T = 0$,

and T is then increased, if there is enough thermal energy to excite the system into flipping a single spin, then the minority droplet of flipped spins can grow and move until all states with two broken bonds become accessible with equal probability via local thermal fluctuations. The system can now reach the other symmetry-broken regime, making both ground states equally probable, and the magnetization vanishes. The first excited state in 1D with energy $E_1 = -JN + 4J$ has two broken bonds. Once E_1 is "turned on", any long-range order is destroyed. Our goal now is to solve the fixed point equation Eq. (5) at finite β , and ultimately determine if higher energy fixed points $\bar{E} > E_0$ are allowed at any finite T .

Firstly, we note that at a given finite temperature, there may be multiple fixed points \bar{E}_{fp} that solve Eq. (5). To demonstrate this, let us assume that the system is in a state such that it can only access one of the two E_0 configurations. If $\bar{E} = E_0$ is a fixed point, then E_1 must violate the constraint, i.e.

$$\begin{aligned} E_1 &> \frac{\alpha}{\beta(\alpha-1)} + E_0, \\ -JN + 4J &> \frac{\alpha T}{(\alpha-1)} - JN, \end{aligned}$$

which produces

$$T < \frac{\alpha-1}{\alpha} 4J. \quad (11)$$

In other words, $\bar{E} = E_0$ is a fixed point for all $T \in [0, \frac{\alpha-1}{\alpha} 4J)$ ($T \in [0, 2J)$ when $\alpha = 2$), a result valid in all dimensions and in the limit $N \rightarrow \infty$. However, for $T \in [0, \frac{\alpha-1}{\alpha} 4J)$, higher energy fixed points than $\bar{E} = E_0$ also exist in the thermodynamic limit. We can show this as follows. Firstly, in 1D, the degeneracy of the j^{th} energy level (with $2j$ broken bonds) is given by $N_j = 2 \times \binom{N}{2j}$ which is an $\mathcal{O}(N^{2j})$ number. We now ask whether, given some temperature β , we can find a valid solution to Eq. (5) for \bar{E} in the limit $N \rightarrow \infty$ that is a convex combination of E_0 and an arbitrary number of excited state energies. As an example, if we assume E_0 and E_1 are the only two accessible energies, then we have

$$\frac{\sum_{j=0}^1 N_j E_j [1 - \beta^{\frac{\alpha-1}{\alpha}} (E_j - \bar{E})]^{\frac{1}{\alpha-1}}}{\sum_{j=0}^1 N_j [1 - \beta^{\frac{\alpha-1}{\alpha}} (E_j - \bar{E})]^{\frac{1}{\alpha-1}}} = \bar{E}.$$

With $N_0 = 2 \times \binom{N}{0} = 2$ and $N_1 = 2 \times \binom{N}{2} = N(N-1)$, in the limit $N \rightarrow \infty$, the $\mathcal{O}(N^2)$ term dominates in both the numerator and denominator of the left-hand side, producing $\bar{E} = E_1$. Since E_2 was assumed a priori to be the lowest forbidden energy, then $E_2 > \frac{\alpha}{\beta(\alpha-1)} + \bar{E} = \frac{\alpha}{\beta(\alpha-1)} + E_1$, and with $E_2 = -JN + 8J$ and $E_1 = -JN + 4J$, we find $T < \frac{\alpha-1}{\alpha} 4J$ once again. Thus, if E_0 and E_1 are the only allowed energies, $\bar{E} = E_1$ is a thermodynamic limit fixed point for all $T \in [0, \frac{\alpha-1}{\alpha} 4J)$.

So far, that makes two fixed points in the limit $N \rightarrow \infty$ for any $T \in [0, \frac{\alpha-1}{\alpha} 4J)$: $\bar{E} = E_0$ and $\bar{E} = E_1$. We can continue with this line of thinking by introducing the next energy E_2 as an allowed energy a priori (with E_3 being the lowest forbidden energy), noting that N_2 is an $\mathcal{O}(N^4)$ term and that it dominates in both numerator and denominator of the left-hand side of the fixed point equation as $N \rightarrow \infty$, giving us $\bar{E} = E_2$ as another mathematically valid solution for all $T \in [0, \frac{\alpha-1}{\alpha} 4J)$.

In this way, higher energy fixed points for any $T \in [0, \frac{\alpha-1}{\alpha} 4J)$ can be found by continuing to introduce higher energies E_j as accessible states, until energies with degeneracies that have similar N -scaling to the maximum degeneracy of $2^{\binom{N}{N/2}} \sim \mathcal{O}(2^N/\sqrt{N})$ [36] are reached and multiple terms begin to survive in the fixed point equation in the limit $N \rightarrow \infty$ as opposed to the single dominant terms we have seen in the simple examples above. As a result, at each temperature, there is a maximum energy fixed point that can be found in the thermodynamic limit, and this maximum grows with increasing T . It can be shown analytically that in the limit $N \rightarrow \infty$, this fixed point globally minimizes the free energy in any finite dimension D , because higher energy microstates, which are not exponentially suppressed in the Rényi ensemble, have increasing degeneracies that significantly boost the entropy. This leads to the approach we use for the attractive fixed point search in our 2D Monte Carlo simulations, the results of which we present in Sec. V. In 1D, since the maximum energy fixed point at all $T > 0$ satisfies $\bar{E} > E_0$ in the thermodynamic limit, we conclude that there is no spontaneous symmetry-breaking at finite T in the 1D Ising model in the α -Rényi ensemble.

V. 2D ISING MODEL: MONTE CARLO

Let us now consider the case of the two-dimensional classical Ising model in the α -Rényi ensemble. We are interested in studying the critical behavior of the true, correlated model. A key goal of this study is to shed light onto the extent to which the Rényi ensemble reproduces the Gibbs state in light of the claims made in Refs. [19, 20] that these two ensembles reproduce each other for local observables in the thermodynamic limit. Similarly, we want to know if any phase transition that emerges coincides with the mean-field prediction that there is a "threshold" α separating continuous and first-order regimes. Unlike the Onsager result for the 2D Ising model in the Gibbs state [31], the challenges associated with evaluating the Rényi ensemble partition function in Eq. (3) make an exact solution difficult to derive, rendering the model ripe for numerical exploration.

We use the Monte Carlo (MC) method with the Metropolis algorithm to simulate the 2D Ising model in an equilibrium defined by the α -Rényi ensemble probabilities in Eq. (2), choosing single-spin flip dynamics for

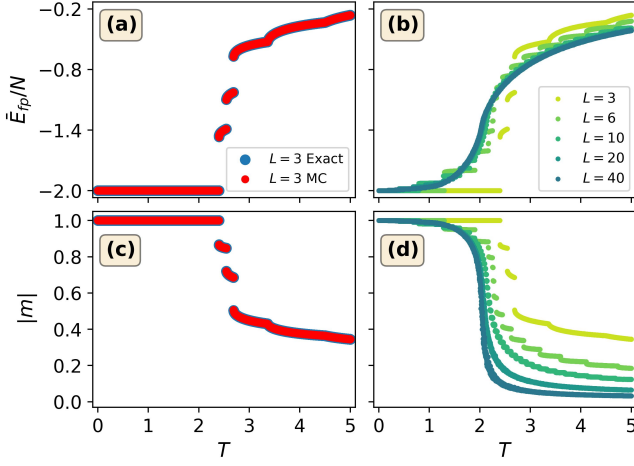


FIG. 3. (a)-(b) Average energy per spin \bar{E}_{fp}/N (with $N = L^2$) computed using the fixed point search technique of Sec. V and (c)-(d) absolute value of the magnetization per spin $|m|$ as a function of temperature T for the 2D Ising model in the 2-Rényi ensemble. In (a) and (c), the Monte Carlo (MC) results for a 3×3 system are compared with exact results. In (b) and (d), the effects of changing system size on the MC results are shown. The small system results exhibit discontinuous jumps in \bar{E}_{fp}/N and $|m|$ with changing T , hinting at the possibility of a first-order transition in the thermodynamic limit, but as L increases, the curves begin to display a more continuous character, providing evidence of a continuous phase transition. At all T , the MC error bars corresponding to the "minimum errors" discussed in Sec. V are smaller than the size of the data points.

simplicity. We customize the original Metropolis algorithm [37] and define

$$A_\alpha(\mu \rightarrow \nu) = \begin{cases} \frac{[1 - \beta^{\frac{\alpha-1}{\alpha}}(E_\nu - \bar{E})]^{\frac{1}{\alpha-1}}}{[1 - \beta^{\frac{\alpha-1}{\alpha}}(E_\mu - \bar{E})]^{\frac{1}{\alpha-1}}}, & E_\nu > E_\mu \parallel E_\nu \leq \frac{\alpha}{\beta(\alpha-1)} + \bar{E} \\ 1, & E_\nu \leq E_\mu \parallel E_\nu \leq \frac{\alpha}{\beta(\alpha-1)} + \bar{E} \\ 0, & E_\nu > \frac{\alpha}{\beta(\alpha-1)} + \bar{E}, \end{cases} \quad (12)$$

where $A_\alpha(\mu \rightarrow \nu)$ represents the acceptance ratio associated with a transition from the current state μ to a proposed state ν (parametrized by the Rényi index α) and where it is assumed that the system is already in a state μ that satisfies the Rényi constraint prior to the update.

The algorithm satisfies detailed balance, but for a certain range of lower temperature values, it does not necessarily satisfy ergodicity. For example, if T is very low, to reach one of the two ground states starting from the other using local dynamics may require accessing states that are strictly forbidden by the Rényi constraint. This becomes less important as α approaches 1 and the Rényi ensemble tends to the Gibbs state, but for all $\alpha > 1$, it is a factor nonetheless for at least a small range of nonzero temperatures. We argue that this does not affect our analysis for the observables we consider in our simula-

tions. In the Gibbs state at low temperature, Monte Carlo simulations of the 2D Ising model for large system sizes that result in excellent approximations of the critical temperature and critical exponents can be conducted in such a way that only one of the symmetry-broken regimes ends up being explored in the typical amount of Monte Carlo time for which such simulations are usually performed. Thus, we do not expect the lack of ergodicity in specific (T, α, N) parameter regimes to affect the study of the critical behavior of the model.

A critical issue that must be resolved in order to simulate the Rényi ensemble is the presence of the average energy \bar{E} in the corresponding probabilities and by association the acceptance ratio in Eq. (12). At each temperature β , we must solve for \bar{E} by solving the fixed point equation Eq. (5). In App. B, we argue that the 2-Rényi ensemble has an attractive fixed point for the Ising model with no external field, and we expect this to remain true for all $\alpha \geq 1$. We leverage this result in our Monte Carlo simulations as follows. For each temperature T and Renyi index α , we start by pre-selecting an initial value of \bar{E} , defined as $\bar{E}^{(0)}$, which allows the Rényi acceptance ratio Eq. (12) to be fully characterized. We use this ratio to perform a full Monte Carlo simulation of the 2D Ising model and extract a new estimate of the average energy $\bar{E}^{(1)}$ using importance sampling and the binning technique [38, 39]. The attractive nature of the fixed point means that, unless $\bar{E}^{(0)}$ happens to be the true fixed point \bar{E}_{fp} , $\bar{E}^{(1)}$ should be closer to \bar{E}_{fp} than $\bar{E}^{(0)}$, barring Monte Carlo errors in the estimation of the average energies. Next, we take $\bar{E}^{(1)}$, plug it into Eq. (12) to form a new acceptance ratio, and repeat the process to extract $\bar{E}^{(2)}$ at the new equilibrium. We continue in this vein until we have found some $\bar{E}^{(k)} \approx \bar{E}_{\text{fp}}$ after k Monte Carlo simulations. The fixed point search is defined by the recursion

$$\begin{aligned} g_\alpha(\bar{E}^{(i-1)}) &\approx \bar{E}^i \\ g_\alpha(\bar{E}^{(i)}) &\approx \bar{E}^{i+1} \\ &\dots \end{aligned} \quad (13)$$

where $g_\alpha \equiv \text{Tr}(\rho^{(\alpha)} H)$. Here, there is Monte Carlo error involved in the estimation of $g_\alpha(\bar{E}^{(i)})$ for every i in the iteration. This noise is propagated through the recursion as the simulation searches for the fixed point, but we find that, at each step in the recursion, if the Monte Carlo time is large enough and an accurate estimation process based on the binning technique is used, this noise has little effect when it comes to moving in the general direction of the fixed point and ultimately extracting a reasonable estimate for \bar{E}_{fp} . To identify the fixed point, we choose to define a new hyperparameter N_{osc} that counts the number of times the fixed point search "oscillates". In other words, once the general vicinity of the fixed point has been approximately found, its attractive nature means that continuing the recursion should make

the Monte Carlo estimate for \bar{E} oscillate about some average value that is very close to the true \bar{E}_{fp} , and we quantify this oscillation by counting the number of times $(\bar{E}^{(i+1)} - \bar{E}^{(i)})$ changes sign from one iteration to the next, defining N_{osc} as precisely this number. In practice, we find that as long as N_{osc} is large enough, changing its value does not significantly affect the final results for the Monte Carlo averages and data collapse.

At each temperature T , there may be more than one fixed point. In Sec. IV, we showed that in the 1D model, the Rényi ensemble can generate a large number of fixed points at each temperature in the thermodynamic limit. While the analysis to prove this in the 2D Ising model is more involved, our Monte Carlo simulations provide strong evidence for the existence of multiple fixed points at most T , for the finite system sizes that we choose to study. Following on from the discussion in Sec. IV, we make the explicit decision to follow the “maximum energy” approach at each temperature. We start at $T = 0.001$, where we expect the maximum energy fixed point to be either exactly the ground state energy of $E_0 = -2JN$ or just above it (because $\bar{E} = E_0$ is the only fixed point at $T = 0$ when the system is finite). To be certain of finding it, we start the search from above, setting the initial average energy in our recursion equal to $\bar{E}^{(0)} = E_0 + \Delta$ where Δ is any number that is large enough to ensure $\bar{E}^{(0)}$ is greater than the expected value of \bar{E}_{fp} . We terminate the fixed point search once $N_{\text{osc}} = 30$ oscillations have been detected for $\bar{E}^{(i)}$. Defining k as the number of MC simulations required to reach $N_{\text{osc}} = 30$, and $\bar{E}^{(k)}$ as the energy of the k^{th} run, we take this $\bar{E}^{(k)}$ value and run one final simulation with $\bar{E} = \bar{E}^{(k)}$ in Eq. (12), this time significantly increasing the Monte Carlo time, and compute the final average energy estimate which we define as our maximum energy fixed point $\bar{E}_{\text{fp}} \approx \bar{E}^{(k+1)}$. It is this last Monte Carlo simulation that we also use for the final estimate of the magnetization $|m|$ and the corresponding error. We note that in all the above simulations, we choose the all-down ground state as the initial configuration.

With the $T = 0.001$ simulation now complete, we seek results for $T \in [0.001, 5.000]$ in increments of $dT = 0.001$. We increment T as $T \rightarrow T + dT$, and, at every subsequent temperature, we embark on an annealing strategy for the fixed point search defined by

$$\begin{aligned} \bar{E}^{(0)}(T + dT) &= \bar{E}_{\text{fp}}(T) + \Delta_E \\ &\approx \bar{E}^{(k+1)}(T) + \Delta_E. \end{aligned} \quad (14)$$

In other words, for each temperature $T + dT$ we set the initial average energy used in the search for the elusive fixed point equal to the fixed point estimate from the previous temperature T plus some Δ_E that must be large enough to ensure we are conducting the next search from above. We note that k , the number of simulations required to reach $N_{\text{osc}} = 30$, is temperature-dependent.

In Fig. 3, we plot Monte Carlo results at $\alpha = 2$ for the average energy per spin \bar{E}_{fp}/N and absolute value




FIG. 4. Absolute value of the Monte Carlo magnetization per spin $|m|$ as a function of temperature T for the 2D Ising model in the α -Rényi ensemble at $L = 30$, for a variety of values of α . For each curve, the corresponding estimate of T_c as extracted by collapsing the MC data using Eq. (15) is shown. The curves are relatively continuous in character, hinting at a continuous phase transition for all α . The shape of the curves through the transition appears to be independent of α , but the estimates of T_c exhibit a strong α -dependence. The error bars are smaller than the size of the data points, but these are “minimum errors” as previously described.

of magnetization $|m|$ as a function of T for various system sizes of interest. The near-perfect overlap between Monte Carlo and exact results in the 3×3 case highlights the strength of the MC approach and the accuracy of the annealing method we use for the fixed point search. We note that the exact results were computed to high precision by leveraging the attractive nature of the fixed point as well, with each search starting from above at exactly $\bar{E}^{(0)} = 0$. For the 3×3 results, \bar{E}_{fp}/N and $|m|$ exhibit discontinuous jumps at various temperatures, hinting at the possibility of a first-order transition in the thermodynamic limit. However, that possibility is quashed by the results at larger N , which provide stronger evidence for the presence of a continuous transition at $\alpha = 2$, with the curves becoming more and more continuous with increasing N , and tending to the shapes we would expect in the $\alpha \rightarrow 1$ case for the 2D classical Ising model.

Turning to Monte Carlo error, the error bars in Fig. 3 are nominally smaller than the size of the data points, but these errors must be termed “minimum errors”, because the fixed point $\bar{E}^{(k+1)} \approx \bar{E}_{\text{fp}}$ found at each temperature after the final Monte Carlo simulation is an approximation and not exact, with $|\bar{E}^{(k+1)} - \bar{E}^{(k)}|$ small but nonzero. We do not attempt to quantify the error beyond computing the errors of the final averages as per the procedure outlined in Ref. [39].

So what happens when we vary the Rényi index α ? In Fig. 4, we plot the Monte Carlo results for the magnetization $|m|$ as a function of T for an $L = 30$ Ising model in 2D at various values of α . We find that for each α probed, the shape of the curve remains effectively

the same, generally continuous in character, indicating a continuous phase transition in the thermodynamic limit. What does change, however, is the position of the apparent critical temperature T_c for each curve, which we estimate using the data collapse technique described below. The continuous nature of the curves becomes even more apparent as L is increased beyond 30 (not shown).

To extract T_c , we take inspiration from renormalization group (RG) scaling theory for continuous phase transitions in classical models in the Gibbs state, elucidated in full detail in Refs. [40–44]. Specifically, we make use of the following RG-derived scaling function:

$$|m|L^{\beta/\nu} = f(tL^{1/\nu}). \quad (15)$$

Here, β is the critical exponent governing the behavior of the magnetization m as $T \rightarrow T_c^-$ such that $m \sim (T_c - T)^\beta$, $t \equiv (T - T_c)/T_c$ is the reduced temperature, and ν is the critical exponent that characterizes the divergence of the correlation length as the critical point is approached, i.e. $\xi \sim |t|^{-\nu}$. Eq. (15) tells us that given access to high quality Gibbs state ($\alpha \rightarrow 1$) data for $|m|$ vs T for all values of L , all the data points collapse onto a single $|m|L^{\beta/\nu}$ vs $tL^{1/\nu}$ curve.

Eq. (15) is only valid in the vicinity of the critical point (i.e. near $t = 0$), because the analysis that produced it was a single RG step performed under the assumption that $|t| = |(T - T_c)/T_c| \ll 1$ [41, 42]. This means that in theory, only data points corresponding to temperatures near $T = T_c$ should form part of this collapse. Although we know with certainty that Eq. (15) applies to the Ising model in the Gibbs state, we note that the $|m|$ vs T curves in Fig. 4 do not seem to change shape significantly as α departs from 1. This suggests that the collapse may also apply to all $\alpha > 1$ so long as the system is large enough and the curves are “continuous enough” through the transition. To obtain a collapse, we must tune T_c , β and ν . We recall that the Gibbs state values for the 2D Ising model as derived by Onsager are given by $T_c = 2J / [\log(1 + \sqrt{2})] \sim 2.269J$, $\beta = 1/8$ and $\nu = 1$.

In Fig. 5, the results for the data collapse of the magnetization data associated with four relatively large values of L are shown, for four different values of α . The collapse is performed by fixing the critical exponents to the Gibbs state values $\beta = 1/8$ and $\nu = 1$, i.e., we only tune T_c through the following procedure. At each value of α , we select a range of T_c values to test within an interval that contains the approximate critical temperature as estimated from the raw magnetization data, and perform a grid search for the value of T_c in this interval that minimizes the distance between the $(tL^{1/\nu}, mL^{\beta/\nu})$ data points and a polynomial fit that includes only a certain percentage of the data points either side of $t = 0$, since the collapse is only supposed to apply in the vicinity of $T \sim T_c$. A simple average for T_c and rudimentary error bars are computed by varying the polynomial degree (testing degrees of 15, 20 and 25) and the specific percentage of data points above and below $t = 0$ that are




FIG. 5. Results for the collapse of the Monte Carlo magnetization data for the 2D Ising model, for four different values of α . The data collapses quite well for all α in the vicinity of $t \sim 0$. The collapse was performed as follows: β and ν are fixed to the Gibbs state ($\alpha \rightarrow 1$) critical exponent values for the 2D Ising model, and T_c is tuned using the polynomial fit technique described in Sec. V, with a 25th order polynomial used to fit 5% of the data either side of $t = 0$.

used in the fit (5%, 7.5% and 10%). We keep the percentages relatively low to focus on collapsing the data in the vicinity of $t = 0$ only.

Returning to Fig. 1, the results of this procedure showcase T_c as a function of α as compared to both the mean-field results in 2D and the Onsager critical temperature ($\alpha \rightarrow 1$). As α tends to 1, T_c increases and approaches $\sim 2.269J$ at a decreasing rate. In the other direction, T_c initially decreases at a decreasing rate as α grows beyond 1, but the data points quickly revert to a decrease at an increasing rate at larger values of α . It appears as if T_c might tend to an asymptote near $T_c \sim 1$ in the limit $\alpha \rightarrow \infty$, but we have so far been unable to find an analytical argument supporting this claim. The T_c values we extract seem to be a good fit when plotted against the large- L magnetization data, as shown in Fig. 4.

Our Monte Carlo results imply that there is a finite-temperature continuous phase transition that spontaneously breaks the symmetry in the 2D Ising model in the α -Rényi ensemble for all $\alpha \geq 1$, with critical exponents independent of α . The latter claim is backed by the mean-field prediction for the exponents, the apparent α -independence of the qualitative shape of the $|m|$ vs T curves through the transition and the quality of the data collapse. However, from the strong α -dependence of m (a local observable) and T_c , we cannot conclude that the Rényi ensemble locally reproduces the Gibbs state in the limit $N \rightarrow \infty$ near the critical point.

It is in a way remarkable that in mean-field theory, the predicted transition is only continuous below a threshold value of $\alpha_{[c \rightarrow 1^{\text{st}}]} = \frac{\sqrt{13}-1}{2} \sim 1.303$. In the Rényi ensemble, the constraint in Eq. (4) is such that at a given value of α , as temperatures are increased, more and more higher energy states are made accessible to the sys-




FIG. 6. Comparison between RNN, Monte Carlo and exact results for the absolute value of the magnetization per spin $|m|$ of the 2D Ising model in the α -Rényi ensemble as a function of temperature T , for three combinations of system size and Rényi index. The exact and Monte Carlo results are based on the maximum energy fixed point approach of Sec. IV. A single-layer RNN was used with 50 memory units [26], 500 samples for gradient estimation, 2×10^4 training steps and 10^6 (2×10^5) samples for the final observable estimation for the 4×4 (6×6) system. Overall, the RNN performs strongly for system sizes up to 6×6 , despite some minor discrepancies near certain discontinuities when the RNN appears to find lower energy fixed points as per the analysis in Sec. VI. At $\alpha = 1.05$ in panel (c), $|m|$ becomes “Gibbs”-like in its continuous nature when compared to the $\alpha = 2$ results in panel (b).

tem discontinuously, each suddenly “turning on” at some temperature T . Now each state has a different magnetization m , and in the true model, consecutive eigenstates can be separated by a single spin flip, producing a small gap in energy and magnetization between such states. In the thermodynamic limit, this gap vanishes when considering quantities on a “per spin” basis, and so, as new states become accessible with increasing T , their emergence into phase space occurs continuously in this “per spin” context. This is essentially what takes place in our Monte Carlo simulations for all α as N increases. The difference in mean-field is that the field at each site takes on the same value—that of the mean-field order parameter, m . Thus, the per-spin gaps in energy and magnetization between consecutive eigenstates do not vanish as $N \rightarrow \infty$, implying first-order behavior at a prospective transition, assuming α is large enough. If α is small, the right-hand side of Eq. (4) is such that most states become accessible at all T , making continuous changes more likely. This is why mean-field theory predicts the existence of both continuous and first-order regimes.

VI. 2D ISING MODEL: RECURRENT NEURAL NETWORKS

Having shown that Monte Carlo methods can successfully simulate the α -Rényi ensemble, we now turn to variational Monte Carlo (VMC). We take inspiration from Ref. [19, 20] where the authors developed tensor network and RBM ansätze to variationally simulate quantum spin models in the the 2-Rényi ensemble at finite temperature. Specifically, we leverage the positive recurrent neural network (RNN) approach of Refs. [26, 28, 45] and apply it to the study of the 2D classical Ising model in our ensemble of interest. The cost function we wish to minimize is the

Rényi free energy Eq. (1), which in the case of a classical spin model can be rewritten as

$$F_\alpha = \sum_{\{\sigma\}} P(\sigma) E(\sigma) - \frac{T}{1-\alpha} \log \left[\sum_{\{\sigma\}} P(\sigma) P(\sigma)^{\alpha-1} \right], \quad (16)$$

where $P(\sigma)$ is the probability associated with a configuration σ and $E(\sigma)$ is the corresponding energy. If $P(\sigma)$ is parameterized by the variational parameters $\{\lambda\}$ as $P(\sigma) \rightarrow P_\lambda(\sigma)$ then the gradients of Eq. (16) are given by

$$\begin{aligned} \partial_\lambda F_\alpha = & \sum_{\{\sigma\}} P_\lambda(\sigma) [\partial_\lambda \log P_\lambda(\sigma)] E(\sigma) \\ & - \frac{\alpha T / (1-\alpha)}{\sum_{\{\sigma\}} P_\lambda(\sigma) P_\lambda(\sigma)^{\alpha-1}} \sum_{\{\sigma\}} \left[P_\lambda(\sigma) [\partial_\lambda \log P_\lambda(\sigma)] \right. \\ & \left. P_\lambda(\sigma)^{\alpha-1} \right]. \end{aligned} \quad (17)$$

For large systems, the sum $\sum_{\{\sigma\}}$ cannot be performed exactly—instead it must be evaluated by drawing independent samples $\{\sigma^{(i)}\}$ from $P_\lambda(\sigma)$, rendering the evaluation of the gradients stochastic. The presence of the generalized purity in the denominator of the second term of Eq. (17), a quantity that tends to zero rapidly with increasing temperature and increasing system size and one that must be approximated stochastically, complicates the VMC process due to the resulting large variance of the gradient estimate. One way around this is to use the variance reduction technique proposed by Refs. [46–48],




FIG. 7. Comparison between the exact 2-Rényi ensemble probability distribution (corresponding to the maximum energy fixed point) and the distribution of the trained RNN for a 4×4 system corresponding to the 2D Ising model results in Fig. 6(a) at (a) $T = 1.0$, (b) $T = 2.4$ and (c) $T = 4.0$. At all temperatures, the RNN finds a mixture of the two symmetry-broken regimes. The exact results are computed in such a way to also display this mixture. At $T = 2.4$, the RNN finds a fixed point that is both lower in energy and free energy compared to the maximum \bar{E}_{fp} , showing that the maximum energy approach does not always result in the global free energy minimum for small systems near energy discontinuities, as per the discussion in Sec. VI.

which modifies the gradients as

$$\begin{aligned} \partial_\lambda F_\alpha = & \sum_{\{\sigma\}} P_\lambda(\sigma) [\partial_\lambda \log P_\lambda(\sigma)] [E(\sigma) - E] \\ & - \frac{\alpha T / (1 - \alpha)}{\sum_{\{\sigma\}} P_\lambda(\sigma) P_\lambda(\sigma)^{\alpha-1}} \sum_{\{\sigma\}} \left[P_\lambda(\sigma) [\partial_\lambda \log P_\lambda(\sigma)] \right. \\ & \left. \left(P_\lambda(\sigma)^{\alpha-1} - \sum_{\{\sigma\}} P_\lambda(\sigma) P_\lambda(\sigma)^{\alpha-1} \right) \right]. \end{aligned} \quad (18)$$

It can be shown that the new terms in Eq. (18) do not bias the gradient estimates [26]. The base parameterization we select for $P_\lambda(\sigma)$ is a positive RNN with the two-dimensional tensorized gated recurrent unit cell (2D GRU) of Ref. [45], which we couple to the periodic RNN structure introduced in Ref. [28] with a two-dimensional sampling path. We use the Adam optimizer of Ref. [49] to update the parameters.

In Fig. 6, the RNN magnetization results for three different combinations of the Rényi index α and system size are compared with relevant exact and Monte Carlo results, both computed using the maximum energy fixed point approach of Sec. IV. To produce the RNN results, at each temperature, a single-layer RNN was used with selected hyperparameters (see Fig. 6 caption). We anneal from $T = 6$, where an RNN initialized with weights drawn from a Gaussian distribution is optimized, after which we decrement T by $dT = 0.01$ and initialize the RNN at each subsequent temperature using the trained RNN from the previous temperature. The RNN performs strongly for the system sizes shown, with the 6×6 results in particular giving cause for optimism.

At $\alpha = 2$, minor discrepancies emerge at some intermediate temperatures. To see this clearly, in Fig. 7 we plot the exact 2-Rényi ensemble probability distribution associated with the maximum energy fixed point for the 4×4 2D Ising model of Fig. 6(a) and compare it with the

corresponding RNN prediction in three different temperature regimes. The results are a near-perfect match in the high and low temperature regimes, but at the specific intermediate temperature shown ($T = 2.4$), the RNN finds a lower energy, less entropic fixed point, one that is shown to be lower in free energy. The conclusion from Sec. IV that, in the thermodynamic limit, the maximum energy fixed point produces the global free energy minimum remains valid, but for smaller systems, the free energy may be minimized by a lower energy fixed point in the vicinity of discontinuities. This might explain the slight discrepancies between the 6×6 RNN and Monte Carlo results near intermediate temperature discontinuities in Fig. 6(b), despite the curves overlapping well overall. We observe that as system size increases, this effect becomes less pronounced due to the increasingly continuous nature of the curves, until at some large enough L , the free energy minimum is produced by the maximum energy fixed point at all T , justifying our approach to the Monte Carlo fixed point search. The RNN's success in finding the equilibrium fixed point at $T = 2.4$ in Fig. 6(a) speaks to its expressive power as a variational ansatz.

As $\alpha \rightarrow 1$, the magnetization becomes more continuous and the RNN and Monte Carlo results at 6×6 (see Fig. 6(c)) generate near perfect overlap. As Fig. 7 shows, the RNN tends to find mixtures of the two symmetry-broken regimes after the optimization, which contrasts with some of our low temperature Monte Carlo simulations that are confined to one or the other.

All in all, the RNN produces promising results, but our approach is susceptible to difficulties at larger system sizes, when we expect the issue of having to estimate the generalized purity in the denominator of the gradient to have an effect on convergence. And while this approach works well for system sizes up to 6×6 (and possibly slightly larger) in the 2D classical Ising model, we expect further challenges to arise when applying this

approach to quantum spin models. This emphasizes the importance of the work of Ref. [20], where a method for gradient estimation that avoids vanishing denominators is introduced and applied to the quantum Ising model.

VII. CONCLUSION & OUTLOOK

We developed various techniques to study the generalized α -Rényi ensemble thermal state approximation of the classical Ising model. First we analyzed the model at the mean-field level, and found that there is a threshold value of the Rényi index $\alpha_{c \rightarrow 1^{\text{st}}} = (\sqrt{13} - 1)/2 \sim 1.303$ separating continuous and first-order phase transition regimes. We proceeded to present an analytical argument as to why there is no finite-temperature symmetry-breaking phase transition in 1D for all values of the Rényi index α . For 2D, we developed a Monte Carlo technique that targets the Rényi ensemble distribution by leveraging an attractive fixed point, and concluded that the true phase transition of the fully correlated model is continuous for all $\alpha \geq 1$. We argued that the Monte Carlo results, combined with the mean-field predictions, provide strong evidence that the critical exponents associated with this transition are independent of α . However, the predicted critical temperature as extracted from a data collapse of the magnetization curves is strongly α -dependent. While our numerical simulations away from the critical point support the arguments in Refs. [19, 20] that the Gibbs state and the Rényi ensemble become locally indistinguishable in the thermodynamic limit, our results near the critical point suggest that the Rényi ensemble predictions can diverge from the Gibbs ensemble even for local observables.

Turning to recurrent neural networks (RNNs), we presented variational Monte Carlo results for the simulation of the 2D classical Ising model in the Rényi ensemble, finding that the RNN of Refs. [26, 28, 45] performs strongly as a variational ansatz for system sizes up to 6×6 . For larger systems, our approach to the optimization of the α -Rényi free energy suffers from difficulties involved in estimating a vanishing purity in the denominator of the entropy gradient term. For this, we pay tribute to the work of Ref. [20], which found a way to use an RBM ansatz to simulate the 2D quantum Ising model in the 2-Rényi ensemble while avoiding vanishing denominators in the gradient.

We anticipate that the combination of the Ref. [20] approach with highly expressive models such as the RNN could prove promising for larger systems in the context of variational studies of the Rényi ensemble in various classical and quantum spin models of interest. The apparent independence of the critical exponents of the 2D Ising transition on the Rényi index makes the α -Rényi ensemble a potentially rich playground for the extraction of universal properties and critical exponents of finite-temperature phase transitions in classical and quantum systems through variational techniques.

OPEN-SOURCE CODE

Our Monte Carlo code, including code for the fixed point search, is made publicly available at "https://github.com/andrewjreissaty91/ising_renyi_ensemble_MonteCarlo", while details of the RNN implementation of Sec. VI can be found at "https://github.com/andrewjreissaty91/ising_renyi_ensemble_RNN".

ACKNOWLEDGEMENTS

We thank R. Wiersema, M. Duschenes, M. S. Moss, A. Orfi, M. Hibat-Allah and A. Ijaz for their wisdom, insights and valuable discussion. We also thank R. Brekelmans, R. G. Melko, L. Hayward, M. Reh, S. Czischek and F. Oyedemi for their expert guidance and support throughout this project. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC). JC 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 www.vectorinstitute.ai/#partners.

Appendix A: Mean-Field Details

As detailed in the main text, the mean-field product-state technique encapsulated by Eq. (6) produces an expression for the Rényi free energy per spin $f_\alpha(m)$ detailed in Eq. (7). The expression f_α as a function of m showcases the existence of a "threshold" value of the Rényi index $\alpha_{[c \rightarrow 1^{\text{st}}]} = \frac{\sqrt{13}-1}{2} \sim 1.303$ below which the Ising model in the α -Rényi ensemble exhibits a continuous phase transition ($\alpha \in [1, \alpha_{[c \rightarrow 1^{\text{st}}]]$), and above which it produces a first-order transition ($\alpha > \alpha_{[c \rightarrow 1^{\text{st}}]}$). This can be seen in Fig. A.1, where f_α is plotted in the continuous regime above and below the critical temperature T_c , as well in the first-order regime above and below the transition temperature T^* .

Let us now derive the threshold value $\alpha_{[c \rightarrow 1^{\text{st}}]}$. To that end, we need an analytic expression for f_α as a function of the order parameter m , in classic Landau tradition. We return to Eq. (7), Taylor expand the logarithm about

$m = 0$ to 6th order, and obtain

$$\begin{aligned}
 f_\alpha^{(6)} \approx & -T \log 2 + \frac{1}{2}(\alpha T - qJ)m^2 \\
 & + \frac{\alpha T}{24} [(\alpha - 2)(\alpha - 3) - 3\alpha(\alpha - 1)]m^4 \\
 & + \frac{\alpha T}{720} [30\alpha^2(\alpha - 1)^2 - 15\alpha(\alpha - 1)(\alpha - 2)(\alpha - 3) \\
 & + (\alpha - 2)(\alpha - 3)(\alpha - 4)(\alpha - 5)]m^6,
 \end{aligned} \tag{A1}$$

where $f_\alpha^{(6)}$ has been defined as the $\mathcal{O}(m^6)$ approximation to f_α . We can show that only even orders survive the Taylor expansion, fulfilling the Landau theory vision of having an analytic free energy that captures the symmetries of the Hamiltonian, in this case the \mathbb{Z}_2 spin-flip symmetry of the Ising model. We justify ignoring higher orders than $\mathcal{O}(m^6)$ because in the vicinity of $\alpha_{[c \rightarrow 1^{\text{st}}]}$, the transition is either continuous or "nearly continuous", and so $m \sim 0$ for all T near T_c or T^* . We also choose to restrict ourselves to 6th order specifically because free energy expressions that capture first-order transitions tend to have five extrema (three minima and two maxima) near the transition temperature, and to produce this number of extrema, at minimum an $\mathcal{O}(m^6)$ expression is needed. By plotting f_α as a function of m for values of α just above the numerically deduced $\alpha_{[c \rightarrow 1^{\text{st}}]} \sim 1.3$, we confirm that five extrema emerge in the interval $m \in [-1, 1]$ when $T \sim T^*$, as can be seen in Fig. A.1.




FIG. A.1. Mean-field α -Rényi free energy per spin (Eq. (7)) as a function of magnetization m for a value of α for which the mean-field phase transition is continuous ($\alpha = 1.2$) and a value of α for which the transition is first-order ($\alpha = 1.5$). The dashed and solid curves respectively showcase the free energy just above and below the given transition. At each temperature, the phase is determined by the value of m that minimizes f_α globally, \bar{m} . Below the transition, the symmetry is broken ($\bar{m} \neq 0$) while above it, it is maintained ($\bar{m} = 0$).

We find that the $f_\alpha^{(6)}$ approximation captures the ex-

act value of α at which five extrema begin to emerge for f_α near the transition. When f_α has five extrema, $f_\alpha^{(6)}$ also has five extrema, although in the case of the latter, the extrema may extend beyond the interval $m \in [-1, 1]$ depending on α . In short, the $\mathcal{O}(m^6)$ analysis is sufficient for our goal. We now turn to Eq. (A1) to extract $\alpha_{[c \rightarrow 1^{\text{st}}]}$. Setting $\frac{\partial f_\alpha^{(6)}}{\partial m} = 0$ and pulling out a factor m which produces an extremum at $m = 0$, the remaining extrema are the solutions of

$$Am^4 + Bm^2 + C = 0, \tag{A2}$$

where we have defined

$$\begin{aligned}
 A &\equiv \frac{\alpha T}{720} [30\alpha^2(\alpha - 1)^2 - 15\alpha(\alpha - 1)(\alpha - 2)(\alpha - 3) \\
 &\quad + (\alpha - 2)(\alpha - 3)(\alpha - 4)(\alpha - 5)], \\
 B &\equiv \frac{\alpha T}{24} [(\alpha - 2)(\alpha - 3) - 3\alpha(\alpha - 1)], \\
 C &\equiv \frac{1}{2}(\alpha T - qJ).
 \end{aligned}$$

The solution to Eq. (A2) is

$$m^2 = \frac{-B \pm \sqrt{B^2 - 4AC}}{2A}. \tag{A3}$$

In order for Eq. (A3) to be able to generate four real extrema (on top of the $m = 0$ extremum discussed above), we must have $-B > 0$, while m^2 and the argument of the square root must also be positive. We find that for values of α just above the approximate value of $\alpha_{[c \rightarrow 1^{\text{st}}]} \sim 1.3$, the sign of $B^2 - 4AC$ depends on the temperature T . Thus, the emergence of five extrema is determined first by the sign of $-B$, and then by the choice of T . As such, setting $-B > 0$ ends up being a sufficient condition for the derivation of the exact $\alpha_{[c \rightarrow 1^{\text{st}}]}$. We have

$$\begin{aligned}
 -[(\alpha - 2)(\alpha - 3) - 3\alpha(\alpha - 1)] &> 0, \\
 \alpha^2 + \alpha - 3 &> 0, \\
 \left(\alpha + \frac{1}{2}\right)^2 &> \frac{13}{4} \rightarrow \alpha > \frac{\sqrt{13} - 1}{2}.
 \end{aligned}$$

In other words, in order to have the possibility of five real extrema and thus a first-order transition, we must have $\alpha > \frac{\sqrt{13} - 1}{2}$, which recovers equation Eq. (9) for $\alpha_{[c \rightarrow 1^{\text{st}}]}$.

Computing the exact f_α for any α above and below this exact threshold α confirms that $\alpha_{[c \rightarrow 1^{\text{st}}]} = \frac{\sqrt{13} - 1}{2}$ separates the continuous and first-order regimes exactly. In addition, a related analysis using the $\mathcal{O}(m^4)$ Taylor expansion of f_α , one that we do not detail here, can be performed, and we find that it also produces the same exact result for $\alpha_{[c \rightarrow 1^{\text{st}}]}$.

Appendix B: The α -Rényi Ensemble Fixed Point

Let us now consider the Rényi ensemble at $\alpha = 2$. We wish to study the nature of its fixed point. We define the

function $f_2(\bar{E})$ as

$$f_2(\bar{E}) = \frac{\sum_{j=0}^{n_\beta} N_j E_j [1 - \frac{1}{2}\beta(E_j - \bar{E})]}{\sum_{j=0}^{n_\beta} N_j [1 - \frac{1}{2}\beta(E_j - \bar{E})]}, \quad (\text{B1})$$

which corresponds to the left-hand side of Eq. (5) with $\alpha = 2$. At the fixed point $\bar{E} = \bar{E}_{\text{fp}}$, we have $f_2(\bar{E}_{\text{fp}}) = \bar{E}_{\text{fp}}$. Now, in order to prove that the fixed point is attractive, we would have to show that

$$|f_2(\bar{E}_{\text{fp}} + d\bar{E}) - f_2(\bar{E}_{\text{fp}})| < |d\bar{E}|, \quad (\text{B2})$$

where we have defined $d\bar{E}$ as an infinitesimal perturbation away from the fixed point. Given Eq. (B1), we can write

$$\begin{aligned} f_2(\bar{E}_{\text{fp}} + d\bar{E}) &= \frac{\sum_{j=0}^{n_\beta} N_j E_j [1 - \frac{1}{2}\beta(E_j - \bar{E}_{\text{fp}})] + \frac{1}{2}\beta d\bar{E} \sum_{j=0}^{n_\beta} N_j E_j}{\sum_{j=0}^{n_\beta} N_j [1 - \frac{1}{2}\beta(E_j - \bar{E}_{\text{fp}})] + \frac{1}{2}\beta d\bar{E} \sum_{j=0}^{n_\beta} N_j} \\ &\equiv \frac{X + X'}{Y + Y'}, \end{aligned} \quad (\text{B3})$$

$$\equiv \frac{X + X'}{Y + Y'}, \quad (\text{B4})$$

where the variables $X \equiv \sum_{j=0}^{n_\beta} N_j E_j [1 - \frac{1}{2}\beta(E_j - \bar{E}_{\text{fp}})]$,

$X' \equiv \frac{1}{2}\beta d\bar{E} \sum_{j=0}^{n_\beta} N_j E_j$, $Y \equiv \sum_{j=0}^{n_\beta} N_j [1 - \frac{1}{2}\beta(E_j - \bar{E}_{\text{fp}})]$

and $Y' \equiv \frac{1}{2}\beta d\bar{E} \sum_{j=0}^{n_\beta} N_j$ have been defined to supplement the subsequent analysis. In all the above equations, we have assumed that a perturbation of the fixed point \bar{E}_{fp} by an infinitesimal quantity $d\bar{E}$ does not change the set of $n_\beta + 1$ energies that satisfy the constraint in Eq. (4) when $\bar{E} = \bar{E}_{\text{fp}}$.

Next, we consider Eq. (B2). In the language of X and Y , if the fixed point is attractive, it must mean that

$$\left| \frac{X + X'}{Y + Y'} - \frac{X}{Y} \right| < |d\bar{E}|. \quad (\text{B5})$$

Some algebraic work yields

$$\begin{aligned} \left| \frac{X + X'}{Y + Y'} - \frac{X}{Y} \right| &= \left| \frac{\frac{1}{2}\beta \sum_{j=0}^{n_\beta} N_j (Y E_j - X) d\bar{E}}{Y^2 + Y \frac{1}{2}\beta \sum_{j=0}^{n_\beta} N_j d\bar{E}} \right| \\ &\approx \underbrace{\left| \frac{\frac{1}{2}\beta \sum_{j=0}^{n_\beta} N_j (Y E_j - X)}{Y^2} \right|}_{\equiv A} |d\bar{E}|, \end{aligned} \quad (\text{B6})$$

where the fraction A , which we can write as

$$\begin{aligned} A &= \left| \frac{\frac{1}{2}\beta \sum_{j=0}^{n_\beta} \sum_{k=0}^{n_\beta} N_j N_k [1 - \frac{1}{2}\beta(E_j - \bar{E}_{\text{fp}})] (E_k - E_j)}{\sum_{j=0}^{n_\beta} \sum_{k=0}^{n_\beta} N_j N_k [1 - \frac{1}{2}\beta(E_j - \bar{E}_{\text{fp}})] [1 - \frac{1}{2}\beta(E_k - \bar{E}_{\text{fp}})]} \right| \\ &= \left| \frac{\frac{1}{2}\beta \sum_{k=0}^{n_\beta} N_k (E_k - \bar{E}_{\text{fp}})}{\sum_{k=0}^{n_\beta} N_k [1 - \frac{1}{2}\beta(E_k - \bar{E}_{\text{fp}})]} \right| \\ &= \left| \frac{\frac{1}{2}\beta \sum_{k=0}^{n_\beta} N_k (E_k - \bar{E}_{\text{fp}})}{\sum_{k=0}^{n_\beta} N_k - \frac{1}{2}\beta \sum_{k=0}^{n_\beta} N_k (E_k - \bar{E}_{\text{fp}})} \right|, \end{aligned} \quad (\text{B7})$$

must be less than 1 in order for the fixed point to be attractive. In the high-temperature limit, β tends to zero and we have $\sum_{k=0}^{n_\beta} N_k \gg \mathcal{O}(\beta)$, producing $A < 1$ and thus an attractive fixed point. In the opposite limit $T \rightarrow 0$, the ground state fixed point $\bar{E}_{\text{fp}} = E_0$ must be attractive as well since the sums in Eq. (B7) have only one term $E_k - \bar{E}_{\text{fp}} = E_0 - E_0 = 0$, so at low temperatures, the ground state fixed point, which our 2D Monte Carlo simulations of the Ising model indeed find, is attractive as well with $A < 1$. We also showed in Sec. IV that in the 1D Ising model in the generalized α -Rényi ensemble, there are an infinite number of lower energy fixed points $E_{\text{fp}} = E_0, E_1, E_2, \dots$ that can be found in the thermodynamic limit for any $T \in [0, \frac{\alpha-1}{\alpha} 4J]$ such that $E_j \leq E_{\text{fp}}$ for all allowed energies $\{E_j\}$. For each of those those fixed points, we have $\sum_{k=0}^{n_\beta} N_k (E_k - \bar{E}_{\text{fp}}) \leq 0$, producing $A < 1$ in Eq. (B7) and thus a large set of attractive fixed points.

In our exact and Monte Carlo simulations of the 1D and 2D Ising models with no external field, we observe that the maximum energy fixed point, which we explicitly target, is also attractive for all values of the Rényi index $\alpha > 1$ that we test.

-
- [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, Congjun Wu, Yi Li, Shiwei Zhang, and T. Xiang, “Majorana positivity and the fermion sign problem of quantum Monte Carlo simulations,” *Phys. Rev. Lett.* **116**, 250601 (2016).
- [8] S. Sugiura and A. Shimizu, “Thermal pure quantum states at finite temperature,” *Phys. Rev. Lett.* **108**, 240401 (2012).
- [9] S. Sugiura and A. Shimizu, “Canonical thermal pure quantum state,” *Phys. Rev. Lett.* **111**, 010401 (2013).
- [10] K. Takai, K. Ido, T. Misawa, Y. Yamaji, and M. Imada, “Finite-temperature variational Monte Carlo method for strongly correlated electron systems,” *Journal of the Physical Society of Japan* **85**, 034601 (2016).
- [11] A. Iwaki, A. Shimizu, and C. Hotta, “Thermal pure quantum matrix product states recovering a volume law entanglement,” *Phys. Rev. Research* **3**, L022015 (2021).
- [12] Y. Nomura, N. Yoshioka, and F. Nori, “Purifying deep boltzmann machines for thermal quantum states,” *Phys. Rev. Lett.* **127**, 060601 (2021).
- [13] N. Irikura and H. Saito, “Neural-network quantum states at finite temperature,” *Phys. Rev. Research* **2**, 013284 (2020).
- [14] S. R. White, “Minimally entangled typical quantum states at finite temperature,” *Phys. Rev. Lett.* **102**, 190601 (2009).
- [15] E. M. Stoudenmire and S. R. White, “Minimally entangled typical thermal state algorithms,” *New J. Phys.* **12**, 055026 (2010).
- [16] J. Nys, Z. Denis, and G. Carleo, “Real-time quantum dynamics of thermal states with neural thermofields,” (2023), [arXiv:2309.07063 \[quant-ph\]](#).
- [17] F. Vicentini, R. Rossi, and G. Carleo, “Positive-definite parametrization of mixed quantum states with deep neural networks,” (2022), [arXiv:2206.13488 \[quant-ph\]](#).
- [18] A. G. Bashkurov, “Maximum Rényi entropy principle for systems with power-law Hamiltonians,” *Phys. Rev. Lett.* **93**, 130601 (2004).
- [19] G. Giudice, A. Çakan, J. I. Cirac, and M. C. Bañuls, “Rényi free energy and variational approximations to thermal states,” *Phys. Rev. B* **103**, 205128 (2021).
- [20] S. Lu, G. Giudice, and J. I. Cirac, “Variational neural and tensor network approximations of thermal states,” (2024), [arXiv:2401.14243 \[quant-ph\]](#).
- [21] G. Carleo and M. Troyer, “Solving the quantum many-body problem with artificial neural networks,” *Science* **355**, 602–606 (2017).
- [22] J. Carrasquilla and R. G. Melko, “Machine learning phases of matter,” *Nature Physics* **13**, 431–434 (2017).
- [23] 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).
- [24] R. G. Melko, G. Carleo, J. Carrasquilla, and J. I. Cirac, “Restricted boltzmann machines in quantum physics,” *Nature Physics* **15**, 887–892 (2019).
- [25] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. G. Melko, and G. Carleo, “Neural-network quantum state tomography,” *Nature Physics* **14**, 447–450 (2018).
- [26] M. Hibat-Allah, M. Ganahl, L. E. Hayward, R. G. Melko, and J. Carrasquilla, “Recurrent neural network wave functions,” *Phys. Rev. Research* **2**, 023358 (2020).
- [27] M. Hibat-Allah, E. M. Inack, R. Wiersema, R. G. Melko, and J. Carrasquilla, “Variational neural annealing,” *Nature Machine Intelligence* **3**, 952–961 (2021).
- [28] M. Hibat-Allah, R. G. Melko, and J. Carrasquilla, “Investigating topological order using recurrent neural networks,” *Phys. Rev. B* **108**, 075152 (2023).
- [29] S. Czischek, M. S. Moss, M. Radzihovsky, E. Merali, and R. G. Melko, “Data-enhanced variational Monte Carlo simulations for Rydberg atom arrays,” *Phys. Rev. B* **105**, 205108 (2022).
- [30] M. S. Moss, S. Ebadi, T. T. Wang, G. Semeghini, A. Bohrdt, M. D. Lukin, and R. G. Melko, “Enhancing variational monte carlo simulations using a programmable quantum simulator,” *Phys. Rev. A* **109**, 032410 (2024).
- [31] L. Onsager, “Crystal Statistics. I. A two-dimensional model with an order-disorder transition,” *Phys. Rev.* **65**, 117 (1944).
- [32] Alfred Rényi, “On measures of entropy and information,” in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics* (University of California Press, 1961) pp. 547–561.
- [33] D. Arovas, *Lecture Notes on Thermodynamics and Statistical Mechanics* (CreateSpace Independent Publishing Platform, 2014).
- [34] M. Kardar, *Statistical Physics of Fields* (Cambridge University Press, 2007).
- [35] R. Peierls and M. Born, “On Ising’s model of ferromagnetism,” *Mathematical Proceedings of the Cambridge Philosophical Society* **32**, 477 (1936).
- [36] “The asymptotic growth of $n \text{ choose } \text{floor}(n/2)$,” (Stack Overflow 2014).
- [37] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, “Equation of state calculations by fast computing machines,” *Journal of Chemical Physics* **21**, 1087 (1953).
- [38] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, 1999).

- [39] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, 2017).
- [40] R. Shankar, *Quantum Field Theory and Condensed Matter: An Introduction* (Cambridge University Press, 2017).
- [41] N. Goldenfeld, *Lectures on Phase Transitions and the Renormalization Group* (CRC Press, 2018).
- [42] H. Nishimori and G. Ortiz, *Elements of Phase Transitions and Critical Phenomena* (Oxford University Press, 2010).
- [43] M. E. J. Newman and G. T. Barkema, “Monte Carlo study of the random-field Ising model,” *Phys. Rev. E* **53**, 393 (1996).
- [44] H. Rieger and A. P. Young, “Critical exponents of the three-dimensional random field Ising model,” *J. Phys. A* **26**, 5279 (1993).
- [45] M. Hibat-Allah, R. G. Melko, and J. Carrasquilla, “Supplementing recurrent neural network wave functions with symmetry and annealing to improve accuracy,” in *Machine Learning and the Physical Sciences Workshop* (NeurIPS 2021).
- [46] D. Wu, L. Wang, and P. Zhang, “Solving statistical mechanics using variational autoregressive networks,” *Phys. Rev. Lett.* **122**, 080602 (2019).
- [47] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016).
- [48] Andriy Mnih and Karol Gregor, “Neural variational inference and learning in belief networks,” in *Proceedings of the 31st International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 32 (PMLR, Beijing, China, 2014) pp. 1791–1799.
- [49] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” (2018), [arXiv:1810.04805 \[cs.CL\]](https://arxiv.org/abs/1810.04805).