PhD program in QUANTUM TECHNOLOGIES

XXXVIII cycle

Neural quantum Monte Carlo simulations for disordered systems and quantum simulators

---

PHYS-04/A - Fisica teorica della materia, modelli, metodi matematici e applicazioni

Candidate:
Luca Brodoloni

Supervisor:
Prof. Sebastiano Pilati

a.y. 2024-2025

UNIONE EUROPEA
Fondo Sociale Europeo
Fondo Europeo di Sviluppo Regionale

Ministero dell'Istruzione, dell'Università e della Ricerca

PON RICERCA E INNOVAZIONE 2014 - 2020

Quello che facciamo qui non è nulla, rispetto a quello che sogniamo di fare.

What we do here is nothing compared to what we dream of doing.

Marchese de Sade

Abstract

Neural quantum Monte Carlo simulations for disordered systems and quantum simulators

This thesis extends the applicability of the self-learning Projector Quantum Monte Carlo (PQMC) framework, a method that intrinsically employs a neural quantum state to enhance its performance, to a range of complex quantum systems beyond testbed models. We begin by detailing the theoretical foundations of Projective Quantum Monte Carlo (PQMC) and its enhancement through Neural Quantum States. The self-learning PQMC scheme is described, where a Restricted Boltzmann Machine is iteratively trained on the walker population to approximate the ground state. This method is rigorously benchmarked, showing exact convergence for the 1D Transverse Field Ising Model. The power of this neural-guided approach is then demonstrated on a variety of challenging problems involving disorder and frustration. We enable the high-precision study of quantum phase transitions in the 2D Edwards-Anderson spin glass and in amorphous systems of Rydberg atoms, providing key insights into their critical properties and glassy phases. Furthermore, the framework is applied to unravel the influence of geometry and interaction range on phase transitions in frustrated triangular lattices, specifically revealing exotic phases and the emergence of clock phases via an order-by-disorder mechanism for both short and power-law interactions. The research also include a hybrid quantum-classical pipeline in which experimental data from a Rydberg atom quantum computer are used to guide the PQMC algorithm, leading to a substantial performance increase validated against more standard benchmarks. This work suggests a general protocol for leveraging current quantum devices to enhance classical simulations. Finally, we provide an investigation of the energy gap distributions of spin glasses. We develop a high-performance, GPU-accelerated exact diagonalization code to study finite-size scaling in the Sherrington-Kirkpatrick and Edwards-Anderson models. This foundational study is then extended to larger scales via a novel PQMC-based gap estimator, suggesting the emergence of heavy-tailed gap distributions in two-dimensional spin glass models. This investigation is useful to understand the behaviors and limitations in the context of adiabatic quantum computing.

Contents

Introduction

The classical simulation of quantum many-body systems has been a central challenge in theoretical and computational physics for more than half a century. From the outset of quantum computational methods, it has been evident that the exponential expansion of the Hilbert space with system size makes exact solutions intractable for all but the smallest systems. Over the decades, a rich variety of approximate and stochastic methods have been developed to overcome this limitation. Among them, Quantum Monte Carlo (QMC) techniques have played a key role, providing access to ground-state and finite-temperature properties of strongly correlated systems with remarkable accuracy whenever the so-called sign problem is absent [1, 2, 3].

Early variants of QMC, such as Diffusion Monte Carlo and Path-Integral Monte Carlo, were primarily designed for bosonic systems [4, 5], but soon extended to spin and fermionic models through imaginary-time projection techniques [1, 6]. In particular, Projection Quantum Monte Carlo (PQMC) algorithms emerged as a powerful framework to extract ground-state properties by evolving a trial wave function in imaginary time, effectively filtering out excited-state contributions [7]. Despite their success, these methods remained limited by systematic biases, such as population-control errors and the need for a good guiding wave function to reduce statistical fluctuations [7, 8, 9, 10].

The last decade has witnessed a huge methodological revolution brought by machine learning. Neural-network-based representations of quantum states, known as Neural Quantum States (NQS), have provided a new variational paradigm for simulating many-body quantum systems [11]. By encoding complex correlations into a compact set of trainable parameters, neural networks such as Restricted Boltzmann Machines (RBMs) [12], convolutional architectures, or autoregressive flows have demonstrated an unprecedented ability to approximate ground states beyond the reach of traditional variational ansätze [13, 14, 15]. When integrated into Monte Carlo frameworks, these neural networks can serve either as variational trial states or as adaptive

guiding functions in PQMC, drastically improving the efficiency of stochastic sampling and the accuracy of energy estimations [16, 17].

In parallel, the rapid progress in quantum simulation platforms, including trapped ions, superconducting qubits, and neutral Rydberg-atom arrays, has opened new frontiers in the exploration of quantum matter [18, 19, 20]. These devices now realize programmable many-body Hamiltonians with tunable interactions and disorder, offering an experimental window into problems that are otherwise classically intractable [21, 22, 23]. A prominent application of these programmable systems is, for example, quantum annealing and the quantum approximate optimization algorithm (QAOA), where they are used to solve complex combinatorial problems by adiabatically evolving a quantum system to its ground state [24, 25, 26]. The success of such protocols is critically governed by the system's physical properties, most notably the size of the minimum energy gap, which dictates the adiabatic evolution time and the susceptibility to thermal excitations [27, 28, 29]. In this context, classical simulations provide a benchmarks for validating quantum experiments, calibrating quantum hardware, and studying fundamental properties like gap distributions and critical points that determine the performance of quantum optimizers. Therefore, efficient and unbiased classical algorithms such as the neural-augmented PQMC presented in this thesis represent a tool for both the characterization of quantum matter and the benchmarking of modern quantum simulators and annealers. This is particularly crucial in the study of frustrated, disordered and spin-glass systems, where the interplay between quantum fluctuations, randomness, and competing interactions gives rise to rich phenomenology, such as glassy dynamics, many-body localization, and topological magnetic order, that is notoriously difficult to simulate with conventional methods. The complex energy landscape of these systems pose a significant challenge for both classical and quantum optimizers, making them an ideal testbed for our simulation techniques.

This thesis explores the integration of neural-network wave functions into Projection Quantum Monte Carlo simulations, with the aim of extending their applicability to complex disordered quantum systems and to benchmarking data from Rydberg-atom quantum platforms. In this work, we adopt and extend an existing self-learning

PQMC framework, originally proposed for simple benchmark models, in which a Restricted Boltzmann Machine adapts dynamically to the walker population [16]. The algorithm is here validated and applied to increasingly complex systems, such as disordered quantum models and spin glasses [30, 31, 32], demonstrating its robustness and versatility. The use of this technique enables unbiased investigations of the ground-state properties of these complex systems and, beyond that, allows for the estimation of their energy gaps, which is a key quantity for benchmarking modern quantum annealing platforms [33, 34], where the size of the minimum gap critically determines the efficiency of solving optimization problems.

Structure of the thesis

The thesis is organized as follows:

- **Chapter 1**: Introduces the theoretical foundations of Projection Quantum Monte Carlo algorithms. The imaginary-time formalism, branching dynamics, importance sampling, and continuous-time formulation are described in detail, together with unbiased estimators of observables via forward walking.
- **Chapter 2**: Presents the concept of Neural Quantum States and the implementation of Restricted Boltzmann Machines as guiding wave functions within PQMC. The self-learning PQMC scheme is introduced, where the neural network is trained iteratively on the walker population to approximate the ground-state probability distribution. The algorithm is benchmarked on the one-dimensional Transverse Field Ising Model (TFIM), demonstrating convergence to the exact Jordan-Wigner solution.
- **Chapter 3**: Extends the method to two-dimensional disordered quantum spin glasses, specifically the Edwards-Anderson model in a transverse field. Here, the neural-guided PQMC allows one to study zero-temperature quantum spin-glass transitions, providing insights into replica overlap statistics and the scaling behavior of the spin-glass susceptibility.
- **Chapter 4**: Applies the developed framework to amorphous two-dimensional arrays of Rydberg atoms, where positional disorder gives rise to spin-glass phases.

These classical simulations provide valuable insight into regimes that can be experimentally explored with programmable Rydberg quantum simulators.

- **Chapter 5**: Explores quantum phase transitions in triangular lattices, both with nearest-neighbor and powe-law Rydberg-type interactions, highlighting how frustration and geometry influence the critical properties of the system. The phase transition between the ordered clock phase and the disordered phase is investigated for both power-law and nearest-neighbor models.

- **Chapter 6**: A hybrid quantum-classical benchmarking framework is developed, in which experimental data from a Rydberg-atom quantum platform is used to guide the Projector Quantum Monte Carlo algorithm. This is enabled by a precise mapping between the Rydberg Hamiltonian and the Ising model. Rigorous benchmarking against classical methods like Variational Monte Carlo (VMC) demonstrates a substantial performance improvement, validating the efficacy of this data-enhanced approach.

- **Chapter 7**: Focuses on the implementation and benchmarking of a high-performance exact diagonalization code for sparse Hamiltonians, enabling the study of systems of up to 26 spins via GPU-accelerated Lanczos methods. This computational tool is used to establish finite-size scaling trends for energy gaps in the Sherrington-Kirkpatrick (SK) and Edwards-Anderson (EA) models at the critical point. These exact results for small systems provide a critical benchmark and foundation for the subsequent chapter, where this analysis is extended to larger scales using PQMC to reveal the emergence of heavy-tailed gap distributions.

- **Chapter 8**: Leverages a newly developed PQMC energy gap estimator, validated against exact diagonalization and Jordan-Wigner theory, to study the gap scaling of spin-glass models for larger system sizes. The analysis characterizes the heavy-tailed gap distributions and extracts critical exponents, extending the finite-size scaling analysis presented in Chapter 7.

- **Conclusions and outlook**: Summarizes the main results presented in the thesis. The chapter concludes by proposing promising avenues for future research.

Finally, it provides a comprehensive list of the scientific publications and open-source datasets produced during the doctoral research, which are the foundation for several chapters of this dissertation.

Chapter 1

Projection quantum Monte Carlo algorithms

In this first chapter we describe projection Quantum Monte Carlo (PQMC) algorithms [31, 4]. These are powerful computational techniques that allow one to extract the ground-state properties of many-body quantum systems whose Hamiltonian is not affected by the negative sign problem [2]. In the first part of the chapter we provide the basic ideas behind the PQMC technique, we then describe the details of the algorithm focusing on discrete quantum spin models. Furthermore, we analyze the common source of systematic error, namely the finite number of random walkers, and present a way to mitigate it by introducing the so-called importance sampling technique. It consists in the implementation of a guiding wave function that drives the walker population towards the exact solution, removing the bias mentioned before. In our case the guiding wave function is a Restricted Boltzmann Machine (RBM), one of the fundamental network architecture used in machine learning. RBMs have gained strong interest in recent years as they were among the first architectures employed to represent quantum many-body states under the name of Neural Quantum States (NQS), which have become a state-of-the-art numerical technique for simulating various physical systems [11]. Next, we discuss the continuous-time approximation of the algorithm, leading to the so-called continuous-time PQMC formulation, which eliminates time discretization errors. Finally, we describe the forward walking technique, a method that allows one to compute unbiased estimators of physical observables that are diagonal in the chosen computational basis. This provides a complete and accurate characterization of the ground-state properties.

In the next chapter, we benchmark the algorithm using exactly solvable models, focusing in particular on the one-dimensional Transverse Field Ising Model (1D-TFIM), which can be solved via the Jordan-Wigner transformation that maps the Ising Hamiltonian into a free-fermion one.

1.1 Introduction to PQMC

Let's consider the imaginary-time Schrödinger equation, where we apply the transformation $\tau = it/\hbar$:

$$-\frac{\partial}{\partial\tau}|\Psi(\tau)\rangle = (\hat{H} - E_{\text{ref}})|\Psi(\tau)\rangle. \tag{1.1}$$

Here, the reduced Planck constant $\hbar$ is set to one. $|\Psi(\tau)\rangle$ is the state of a generic quantum system at imaginary time τ. $\hat{H}$ is the quantum many-body Hamiltonian with eigenvectors $\{|\Psi_n\rangle\}$ and eigenvalues $\{E_n\}$. The index $n = 0, 1, 2, ...$ labels the eigenstates. The reference energy E_{ref} aims to normalize the wave function when it is set equal to the average energy $\langle\hat{H}\rangle$, which denotes the expectation value of the Hamiltonian in the ground state. In practice, it is kept close to this value in order to ensure the stability of the simulations, as will be explained later.

At the initial imaginary time $\tau = 0$, the system is initialized in a state $|\psi_{\text{ini}}\rangle$. Expanding this state in the eigenbasis of the Hamiltonian $\hat{H}$ gives

$$|\psi_{\text{ini}}\rangle = \sum_n c_n|\Psi_n\rangle. \tag{1.2}$$

The initial state can be arbitrary, with the constrain that it has a non-vanishing overlap with the ground state of the system, i.e. $\langle\Psi_0|\psi_{\text{ini}}\rangle \neq 0$. In the imaginary propagation time τ, the state evolves as

$$|\Psi(\tau)\rangle = e^{-(\hat{H}-E_{\text{ref}})\tau}|\psi_{\text{ini}}\rangle = \sum_n c_n e^{-(E_n-E_{\text{ref}})\tau}|\Psi_n\rangle. \tag{1.3}$$

The key idea of the algorithm emerges at this point: as τ increases, excited-state contributions are exponentially suppressed provided that $E_{\text{ref}} \simeq E_0$. Consequently, the

projection onto the exact ground state is obtained in the long time limit

$$|\Psi(\tau)\rangle \xrightarrow[\tau\to\infty]{} c_0 e^{-(E_0-E_{\text{ref}})\tau}|\Psi_0\rangle. \tag{1.4}$$

From Eq. (1.3) one can notice that if the energy difference $\Delta = E_1 - E_0$ between the first excited state E_1 and the ground state E_0 is small, the algorithm can struggle to suppress the contribution of the excited state and hence, to converge to the ground state. There are additional factors that may affect the efficiency of PQMC, such as the choice of the initial state, the number of walkers evolving stochastically during the simulation, and the possible use of a guiding wave function, the so-called importance sampling technique, which is well know to improve the performance of the algorithm [35, 36, 16]. All of these aspects will be discussed in the following. In the next section, we provide a guide to the implementation of a vanilla projection technique.

1.2 Simple PQMC implementation

The simple version of the Projection Quantum Monte Carlo (PQMC) algorithm consists in evolving the Schrödinger equation Eq. (1.1) in imaginary time. It is referred to as simple because it does not employ the importance sampling technique which, as mentioned previously, is widely used to enhance the efficiency of PQMC simulations. In this section, we introduce the basic formulation of the method, while the role of importance sampling will be addressed in the next chapter. The imaginary-time Schrödinger equation can be solved iteratively by applying the short-time propagator:

$$\Psi(\mathbf{x}, \tau + \Delta\tau) = \sum_{x'} G(\mathbf{x}, \mathbf{x}', \Delta\tau)\Psi(\mathbf{x}', \tau), \tag{1.5}$$

where $\Psi(\mathbf{x}, \tau) = \langle \mathbf{x}|\Psi(\tau)\rangle$ is the wave function at the time τ and

$$G(\mathbf{x}, \mathbf{x}', \Delta\tau) = \langle \mathbf{x}|e^{-\Delta\tau(\hat{H}-E_{\text{ref}})}|\mathbf{x}'\rangle \tag{1.6}$$

is the Green's function for a small time step $\Delta\tau$. By iteratively applying this short-time propagator, one can reach long imaginary times, so that in the limit $\tau \to \infty$, configurations $\mathbf{x}$ are sampled with probability proportional to the ground-state wave function $\Psi_0(\mathbf{x})$. Notice that, provided the Hamiltonian is free from sign problem, $\Psi_0(\mathbf{x})$ can be

chosen real and non-negative. The main issue here is that, in general, $G(\mathbf{x}, \mathbf{x}', \Delta\tau)$ is not a stochastic matrix and therefore cannot directly be used as the transition probability in a Markov chain Monte Carlo process. One can overcome this problem by rewriting

$$G(\mathbf{x}, \mathbf{x}', \Delta\tau) = G_T(\mathbf{x}, \mathbf{x}', \Delta\tau) b_{\mathbf{x}'}, \tag{1.7}$$

where $G_T(\mathbf{x}, \mathbf{x}', \Delta\tau)$ is a normalized stochastic transition matrix defining transition probabilities, while

$$b_{\mathbf{x}'} = \sum_{\mathbf{x}} G(\mathbf{x}, \mathbf{x}', \Delta\tau) \tag{1.8}$$

represents the weight associated with the configuration $\mathbf{x}'$. The algorithm is then implemented by evolving a population of walkers, namely independent copies of the system. Each walker, initially assigned weight $w_n = 1$, undergoes configuration updates according to G_T of Eq. (1.7), while its weight is updated with the rule $w_n \to w_n b_{\mathbf{x}'_n}$. However, as the imaginary time increases, the variance in walker weights grows rapidly, leading to a situation where only a few walkers dominate the signal, while most provide negligible contributions. To control this instability, one can introduce the so-called branching (or birth–death) technique. At each step, a walker is either replicated or killed depending on the value of its weight. Specifically, for each walker, the number of descendants is chosen as

$$n_d = \mathrm{int}[w_n + \eta] \tag{1.9}$$

where $\eta \in [0, 1]$ is a random number sampled from an uniform distribution. Walkers with $n_d > 1$ generate identical copies, while those with $n_d = 0$ are killed. After branching, all walkers are reset to weight $w_n = 1$. The total walker population fluctuates during the simulation, but it can be stabilized around a target value N_w by tuning the reference energy E_{ref}. A common update rule reads [37]

$$E_{\mathrm{ref}} = E + \mu \log\left(\frac{N_w}{N_{\mathrm{cur}}}\right). \tag{1.10}$$

Here, E is the average energy of the current walkers population of size N_{cur}, N_w is the chosen population target value and $\mu > 0$ is a simulation parameter which controls the walkers fluctuations. Although branching makes the algorithm feasible, it introduces

correlations among walkers sharing the same parents. This may bias the results if the walker population is not sufficiently large, but the effect vanishes in the limit of $N_w \to \infty$ [31, 38, 8, 39]. This is the so-called population bias and it is one of the main systematic errors in PQMC simulations. Another common source of error is due to the finite choice of the time step $\Delta\tau$. In this thesis, however, we employ the continuous-time formulation of the projection Quantum Monte Carlo method, which eliminates the bias introduced by a finite time discretization. In the next sections, we will discuss in detail the importance sampling scheme and the continuous-time approximation, as both play a crucial role in overcoming the limitations mentioned above.

1.3 Importance sampling

One of the key features of the simple PQMC is that no prior information of the groud state wave function is required. Nevertheless, the presence of rapidly fluctuating potential energies, can severely destabilize the branching procedure, sometimes to the extent of making the simulation impractical. A well-established way to mitigate this issue is the use of importance sampling.This technique, widely recognized for improving the efficiency of PQMC calculations (see, e.g, Ref. [1]) allows one to significantly reduce the number of walkers needed to achieve a target level of accuracy [31]. In this method the modified function

$$f(x,\tau) = \Psi(\mathbf{x},\tau)\psi_g(\mathbf{x},\tau) \tag{1.11}$$

evolves through an adapted imaginary-time Schrödinger equation. Here, $\psi_g(\mathbf{x},\tau)$ denotes the trial or guiding function, that is chosen to accurately approximate the ground-state wave function. Its role is to favor the sampling of configurations with high probability amplitude. In conventional PQMC implementations with importance sampling, the guiding function is typically assumed to be independent of time. For this reason, from now on, we consider a time-independent guiding wave function, $\psi_g(\mathbf{x},\tau) \equiv \psi_g(\mathbf{x})$. Thus, Eq. (1.11) reads

$$f(\mathbf{x},\tau) = \Psi(\mathbf{x},\tau)\psi_g(\mathbf{x}). \tag{1.12}$$

It is worth mentioning that there are situations in which the guiding wave function acquires an explicit time dependence. A clear example is simulated quantum annealing, where the kinetic component of the Hamiltonian evolves in time together with its spectrum. As a consequence, it is natural to expect that the guiding function should dynamically adjust to the instantaneous Hamiltonian.

To understand how the imaginary-time Schrödinger equation changes in this framework, let us multiply both sides of Eq. (1.5) by $\psi_g(\mathbf{x})$. We obtain

$$\Psi(\mathbf{x}, \tau + \Delta\tau)\psi_g(\mathbf{x}) = \sum_{\mathbf{x}'} G(\mathbf{x}, \mathbf{x}', \Delta\tau)\frac{\psi_g(\mathbf{x})}{\psi_g(\mathbf{x}')}\Psi(\mathbf{x}', \tau)\psi_g(\mathbf{x}'). \tag{1.13}$$

At this point, by introducing the modified Green's function with importance sampling as

$$\tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau) = G(\mathbf{x}, \mathbf{x}', \Delta\tau)\frac{\psi_g(\mathbf{x})}{\psi_g(\mathbf{x}')} \tag{1.14}$$

we get that the adapted wave function of Eq. (1.12), satisfies the same evolution of Eq. (1.5), with $\tilde{G}$, playing the role of G:

$$f(\mathbf{x}, \tau + \Delta\tau) = \sum_{\mathbf{x}'} \tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau) f(\mathbf{x}', \tau). \tag{1.15}$$

From Eq. (1.12) it is clear that in the limit of long imaginary time, the walkers sample spin configurations according to a probability distribution proportional to $f(\mathbf{x}, \tau \to \infty) = \Psi_0(\mathrm{x})\psi_g(\mathrm{x})$. Hence, when the guiding wave function is a good approximation of the ground state, this distribution closely represents the true probability of finding the configuration $\mathbf{x}$. The expectation value of the ground-state energy within PQMC is evaluated using the mixed estimator:

$$\begin{aligned}
\langle H \rangle_{\text{mixed}} &= \frac{\langle \Psi_0 | \hat{H} | \psi_g \rangle}{\langle \Psi_0 | \psi_g \rangle} \\
&= \frac{\sum_{\mathbf{x}} \langle \Psi_0 | \mathbf{x} \rangle \langle \mathbf{x} | \hat{H} | \psi_T \rangle}{\sum_x \langle \Psi_0 | \mathbf{x} \rangle \langle \mathbf{x} | \psi_g \rangle} \\
&= \frac{\sum_{\mathbf{x}} \Psi_0(\mathbf{x}) \, \frac{\langle \mathbf{x} | \hat{H} | \psi_g \rangle}{\langle \mathbf{x} | \psi_g \rangle} \, \langle \mathbf{x} | \psi_g \rangle}{\sum_{\mathbf{x}} \Psi_0(\mathbf{x}) \, \psi_T(x)} \\
&= \frac{\sum_{\mathbf{x}} \Psi_0(\mathbf{x}) \, E_{\text{loc}}(\mathbf{x}) \, \psi_g(\mathbf{x})}{\sum_{\mathbf{x}} \langle \Psi_0(\mathbf{x}) | \, \psi_g(\mathbf{x}) \rangle},
\end{aligned} \tag{1.16}$$

where $E_{\text{loc}}(\mathbf{x})$ is the *local energy* defined as

$$E_{\text{loc}}(\mathbf{x}) = \frac{\langle \mathbf{x}|\hat{H}|\psi_g\rangle}{\langle \mathbf{x}|\psi_g\rangle}. \tag{1.17}$$

The local energy plays a key role in stabilizing the simulation. Strong fluctuations of E_{loc} translate into large variations in the walker population. If the guiding function were exact, i.e. $\psi_g(\mathbf{x}) = \Psi_0(\mathbf{x})$, then the local energy would reduce to a constant, $E_{\text{loc}} = E_0$. In that case, the fluctuation of the number of walkers would vanish completely, a direct consequence of the zero-variance property [31], thus eliminating the bias associated with a finite number of walkers N_w. In summary, the purpose of the trial wave function is to guide the walkers in exploring the Hilbert space of the system. If this guiding function were exactly equal to the ground state, then even a single walker would be sufficient to get the exact convergence of the PQMC. However, even when $\psi_g(\mathbf{x})$ is not exact but still a good approximation to the ground state, the variance of the walker population is significantly reduced compared to the simple PQMC algorithm. This leads to a faster convergence toward the $N_w \to \infty$ limit and thus lowers the computational cost, particularly in large-scale simulations. The ground state energy is evaluated as

$$E = \lim_{N_s\to\infty} = \frac{1}{N_s}\sum_{i=1}^{N_s} E_{\text{loc}}(\mathbf{x}_i) \tag{1.18}$$

where N_s denotes the number of uncorrelated samples $\mathbf{x}_i$ of the probaiblity distribution

$$p(\mathbf{x}) = \frac{\Psi_0(\mathbf{x})\psi_g(\mathbf{x})}{\sum_{\mathbf{x}} \Psi_0(\mathbf{x})\psi_g(\mathbf{x})}. \tag{1.19}$$

In the special case of the Hamiltonian, as well as for any other operator that commutes with it, the mixed estimator defined above yields exact results (within statistical uncertainty) for the ground-state expectation value. More generally, the accuracy of the mixed estimator depends on the quality of the trial wave function $\psi_g(\mathbf{x})$. For all the other operators that do not commute with the Hamiltonian, alternative approaches such as pure estimators can be employed.

1.4 The continuous time limit

Let us suppose that we want simulate the generic quantum Ising Hamiltonian with uniform transverse field, it reads

$$\hat{H} = -\sum_{i,j} J_{ij}\hat{\sigma_i^z}\hat{\sigma_j^z} - \Gamma\sum_i \hat{\sigma_i^x}. \tag{1.20}$$

It can be decomposed as

$$\hat{H} = \hat{H}_{\text{pot}} + \hat{H}_{\text{kin}}, \tag{1.21}$$

where H_{pot} is the potential energy of the Hamiltonian representing the classical part of the physics of the system and $\hat{H}_{\text{kin}}$ is the kinetic part that drives the quantum fluctuations. The evolution under the modified imaginary-time Schrödinger equation can be represented as a Markov process, defined by Eq. (1.15). To obtain a practical approximation of the modified Green's function of Eq. (1.14), the time interval $\Delta\tau$ is split into M smaller steps of size $\delta\tau = \Delta\tau/M$. When $\delta\tau$ is sufficiently small, we can write the first-order approximation of the Green's function as

$$\tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau) \approx [\tilde{g}(\mathbf{x}, \mathbf{x}', \delta\tau)]^M. \tag{1.22}$$

with

$$\tilde{g}(\mathbf{x}, \mathbf{x}', \delta\tau) = p_{\mathbf{x},\mathbf{x}'}\tilde{g}(\mathbf{x}, \mathbf{x}', \delta\tau'), \tag{1.23}$$

where $p_{\mathbf{x},\mathbf{x}'}$ defines the stochastic transition probabilities of the Markov process:

$$p_{\mathbf{x},\mathbf{x}'} = \frac{(\delta_{\mathbf{x},\mathbf{x}'} - \delta\tau(H_{\mathbf{x},\mathbf{x}'} - E_{\text{ref}})\delta_{\mathbf{x},\mathbf{x}'})\,\frac{\psi_g(\mathbf{x})}{\psi_g(\mathbf{x}')}}{1 - \delta\tau\,(E_{\text{loc}}(\mathbf{x}') - E_{\text{ref}})}. \tag{1.24}$$

Clearly, convergence to the exact ground state is ensured as long as the time step $\delta\tau$ remains below a finite threshold, specifically small enough to guarantee that all the matrix elements of $\tilde{g}(\mathbf{x}, \mathbf{x}', \delta\tau)$ are non-negative [40, 36]. However, as the dimension of the Hilbert space grows, it becomes necessary to further reduce the time step $\delta\tau$. This leads to inefficient simulations, since most proposed transitions $\mathbf{x}' \to \mathbf{x}$ are rejected due to the dominance of the identity operator $p_{\mathbf{x},\mathbf{x}'} \approx \delta_{\mathbf{x},\mathbf{x}'}$, resulting in very long autocorrelation times. This issue can be overcome by implementing the continuous-time

Green's function Monte Carlo (CTGFMC) algorithm, originally introduced in [40, 38] (see also [31] for a detailed implementation). The main idea is to formally take the limit $M \to \infty$ and stochastically sample the time interval $\delta\tau'$ that elapses before the next configuration update. To maintain consistency with the PQMC framework, one can track the remaining time $\delta\tau_t$ needed to complete a total time step $\Delta\tau$, ensuring that each iteration of the algorithm corresponds to a fixed interval $\Delta\tau$. The random time step $\delta\tau'$ is sampled according to

$$\delta\tau' = \min\left(\delta\tau_t, \frac{\ln(1-\xi)}{E_{\text{loc}}(\mathbf{x}') - E_{\text{pot}}(\mathbf{x}')}\right), \tag{1.25}$$

where $\xi \in (0,1)$ is a uniform random variable. For the generic Hamiltonian in Eq. (1.21), one can rewrite the expression for the local energy in Eq. (1.17) as

$$\begin{aligned} E_{\text{loc}}(\mathbf{x}) &= \frac{\langle \mathbf{x}|\hat{H}|\psi_g\rangle}{\langle \mathbf{x}|\psi_g\rangle} \\ &= \frac{\langle \mathbf{x}|\hat{H}_{\text{pot}}|\psi_g\rangle}{\langle \mathbf{x}|\psi_g\rangle} + \frac{\langle \mathbf{x}|\hat{H}_{\text{kin}}|\psi_g\rangle}{\langle \mathbf{x}|\psi_g\rangle} \\ &= E_{\text{pot}}(\mathbf{x}) + \frac{-\Gamma\left(\langle \tilde{\mathbf{x}}_1| + \langle \tilde{\mathbf{x}}_2| + \cdots + \langle \tilde{\mathbf{x}}_N|\right)|\psi_g\rangle}{\langle \mathbf{x}|\psi_g\rangle} \\ &= E_{\text{pot}}(\mathbf{x}) - \Gamma\frac{\sum_{i=1}^{N}\psi_g(\tilde{\mathbf{x}}_i)}{\psi_g(\mathbf{x})}. \end{aligned} \tag{1.26}$$

Here, $|\tilde{\mathbf{x}}_\mathbf{i}\rangle$ denotes the spin configuration obtained from $|\mathbf{x}\rangle$ by flipping the i-th spin.

In this framework, the configuration update $\mathbf{x}' \to \mathbf{x}$, with $\mathbf{x}' \neq \mathbf{x}$, is ruled by the probability distribution

$$t_{\mathrm{x},\mathrm{x}'} = \frac{p_{\mathrm{x},\mathrm{x}'}}{\sum_{\mathbf{x}\neq\mathbf{x}'} p_{\mathrm{x},\mathrm{x}'}}. \tag{1.27}$$

$t_{\mathrm{x},\mathrm{x}'}$ is always zero, except in the cases where $\mathbf{x}'$ and $\mathbf{x}$ differ by a single spin flip. In that case, it reads

$$t_{\mathrm{x},\mathrm{x}'} = \frac{\psi_g(\mathbf{x}_i)}{\sum_{i=0}^{N-1}\psi_g(\tilde{\mathbf{x}}_i)}. \tag{1.28}$$

From the latter equation it is clear that for a constant guiding wave function the transition probability takes the value $1/N$ if the configurations differ from a spin-flip and 0 otherwise. The continuous-time approach is sketched below.

Algorithm 1 Continuous-Time Green's Function Monte Carlo (CTGFMC) Algorithm

Require: Number of walkers N_w, total time step $\Delta\tau$, trial wave function $\psi_T(x)$

Ensure: Ground-state properties using CTGFMC

for each walker $n = 1, \ldots, N_w$ **do**
  Initialize the walker configuration $\mathbf{x}_n$
  Initialize remaining time $\delta\tau_t \leftarrow \Delta\tau$
  Initialize weight $w_n \leftarrow 1$
  while $\delta\tau_t > 0$ **do**
    Sample the stochastic time interval $\delta\tau'$ using Eq. (1.25)
    if $\delta\tau' < \delta\tau_t$ **then**
      Propose a configuration update $\mathbf{x}_n \to \mathbf{x}'_n$ with probability $t_{\mathbf{x}'_n,\mathbf{x}_n}$
    else
      Set $\delta\tau' \leftarrow \delta\tau_t$
    end if
    Update the weight: $w_n \leftarrow w_n\, b_{\mathbf{x'}_n}$
    Reduce remaining time: $\delta\tau_t \leftarrow \delta\tau_t - \delta\tau'$
  end while
end for
Perform branching of walkers according to the accumulated weights w_n

1.5 Pure estimation of quantum observables

In PQMC simulations, the ground-state expectation value of a generical local operator $\hat{O}$ is computed with the probability distribution $p(\mathbf{x})$ defined in Eq. (1.19). Let us generalize the definition of the mixed estimator in Eq. (1.16) as

$$\langle \hat{O} \rangle_{\text{mixed}} = \frac{\langle \Psi_0 | \hat{O} | \psi_T \rangle}{\langle \Psi_0 | \psi_T \rangle}. \tag{1.29}$$

The analogous *pure estimator*, obtained from the probability distribution reads:

$$\langle \hat{O} \rangle_{\text{pure}} = \frac{\langle \Psi_0 | \hat{O} | \Psi_0 \rangle}{\langle \Psi_0 | \Psi_0 \rangle}. \tag{1.30}$$

When $\hat{O}$ commutes with the Hamiltonian, the mixed estimator, which is the quantity directly accessible in PQMC, coincides (up to statistical fluctuations) with the pure estimator, i.e., $\langle \hat{O} \rangle_{\text{pure}} \approx \langle \hat{O} \rangle_{\text{mixed}}$. However, if $\hat{O}$ does not commute with the Hamiltonian, the mixed estimator does not provide the same value as the pure one. In standard PQMC, pure estimators are usually approximated using the Ceperley estimator [41]:

$$\langle \hat{O} \rangle_{\text{pure}} \approx 2\langle \hat{O} \rangle_{\text{mixed}} - \frac{\langle \psi_T | \hat{O} | \psi_T \rangle}{\langle \psi_T | \psi_T \rangle}, \tag{1.31}$$

where the last term is obtained from a separate Variational Monte Carlo calculation. The accuracy of this approximation strongly depends on how close the trial wave function ψ_T is to the exact ground-state wave function. In this work, we approximate pure estimators directly within PQMC simulations, in a way closely related to the forward walking technique [42]. Assuming that $\hat{O}$ is diagonal in the computational basis, its pure estimator can be written as follows:

$$\langle \hat{O} \rangle_{\text{pure}} = \frac{\langle \Psi_0 | \hat{O} | \Psi_0 \rangle}{\langle \Psi_0 | \Psi_0 \rangle} = \frac{\sum_x \Psi_0(x)\, O(x)\, \Psi_0(x)}{\sum_x |\Psi_0(x)|^2}. \tag{1.32}$$

By multiplying numerator and the denominator by the trial wave function $\psi_T(x)$, one obtains

$$\langle \hat{O} \rangle_{\text{pure}} = \frac{\sum_x \Psi_0(x)\, O(x)\, \frac{\Psi_0(x)}{\psi_T(x)}\, \psi_T(x)}{\sum_x \Psi_0^*(x)\, \frac{\Psi_0(x)}{\psi_T(x)}\, \psi_T(x)}. \tag{1.33}$$

Finally, we get

$$\left\langle \hat{O} \right\rangle_{\text{pure}} = \frac{\left\langle O(x)\, \frac{\Psi_0(x)}{\psi_T(x)} \right\rangle_{\text{mixed}}}{\left\langle \frac{\Psi_0(x)}{\psi_T(x)} \right\rangle_{\text{mixed}}}, \tag{1.34}$$

where the averages $\langle \cdot \rangle_{\text{mixed}}$ are taken with respect to the mixed distribution sampled in PQMC. In the standard forward walking technique, the ratio $\Psi_0(x)/\psi_T(x)$ is estimated by keeping track of the long-time offspring of each walker configuration x (see Ref. [31] for details). An alternative approach, consists in exploiting directly the instantaneous walker population evolved during the PQMC with a fixed guiding wave function $\psi_T(x)$.

In this case, the amplitude of the ground-state wave function on a given configuration x can be approximated by the fraction of walkers occupying that configuration,

namely

$$\Psi_0(x) \approx \frac{N_{\text{cur}}(x)}{N_{\text{cur}}}, \tag{1.35}$$

where $N_{\text{cur}}(x)$ denotes the number of walkers in configuration x at a given Monte Carlo sweep, and N_{cur} is the total walker population.

The corresponding pure estimator of a local operator $\hat{O}$ can then be expressed as

$$\langle \hat{O} \rangle_{\text{pure}} \approx \frac{\sum_x O(x)\, \Psi_0^2(x)}{\sum_x \Psi_0^2(x)}. \tag{1.36}$$

In practice, two possible procedures can be employed to estimate $\Psi_0^2(x)$. A simple (though not fully statistically correct) option consists in squaring the estimate of Eq. (1.35), i.e.

$$\Psi_0^2(x) \approx \left(\frac{N_{\text{cur}}(x)}{N_{\text{cur}}} \right)^2. \tag{1.37}$$

A more robust and statistically correct strategy is obtained by evolving two independent PQMC populations (replicas A and B) in parallel, and defining

$$\Psi_0^2(x) \approx \frac{N_{\text{cur}}^A(x)}{N_{\text{cur}}} \times \frac{N_{\text{cur}}^B(x)}{N_{\text{cur}}}. \tag{1.38}$$

This two-replica formulation ensures that only configurations contributing significantly to the ground-state wave function are consistently sampled in both populations, thereby yielding an unbiased estimate.

It should be emphasized that, as in forward walking, the accuracy of this method relies on sufficiently large walker populations: the probability of overlap between the configurations sampled by the two replicas decreases with increasing system size, which may result in an exponential growth of the walker number required. Nevertheless, this scheme represents a practical alternative to forward walking for Hamiltonians independent of imaginary time, and can be readily implemented within standard PQMC simulations.

Chapter 2

Restricted Boltzmann Machines as a guiding wave functions

2.1 Basics of Neural Quantum States

A Neural Quantum State (NQS) is a representation of a many-body quantum state in the form of an artificial neural network. Since the dimension of the Hilbert space of a quantum ising model grows exponentially as 2^N (with N being the dimension of the system), a neural network can capture the complex correlations in a compressed representation that depends on a relatively small number of parameters. In other words, a Neural Quantum State provides a compact encoding of a quantum state that would otherwise be computationally intractable [11]. Let's consider a quantum system with N degrees of freedom in a computational basis $\{|\boldsymbol{\sigma}\rangle\}$, a generic state can be expressed as

$$|\psi\rangle = \sum_{\boldsymbol{\sigma}} \psi(\boldsymbol{\sigma})|\boldsymbol{\sigma}\rangle, \tag{2.1}$$

where $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ denotes the system configuration, namely a spin states $\sigma_i = \pm 1$ for spin-1/2 systems (or occupation numbers $\sigma_i = 0, 1, \ldots, n_{\max}$ for bosonic systems). In the standard NQS approach, the wave function coefficients $\psi(\boldsymbol{\sigma})$ are not provided explicitly, instead they are parametrized by a neural network. However, in explicit architectures such as autoregressive neural quantum states, the network directly outputs normalized probability for each configuration, giving access to the corresponding amplitudes. In this thesis, we focus on implicit architectures. The generic

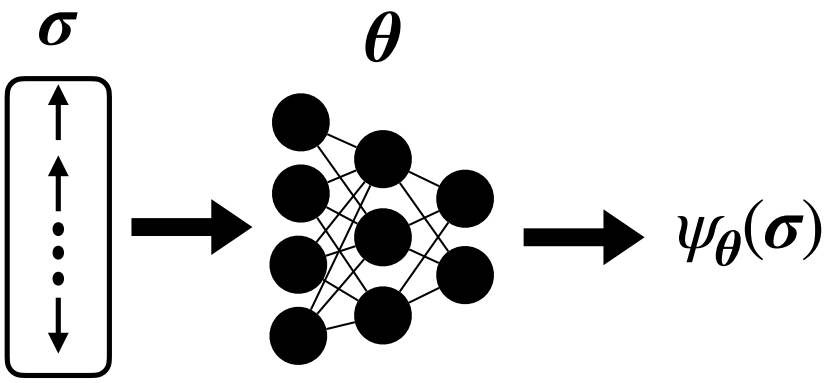

FIGURE 2.1: Example of a neural quantum state. The input configuration is represented by the physical system denoted by $\boldsymbol{\sigma}$ in this case, a spin chain, although it could also correspond to a continuous system of coordinates or other observables. The network adjusts its parameters $\boldsymbol{\theta}$, which in the most general case may be complex, in order to represent the wave function of the physical system provided as input. Once trained, given an input configuration $\boldsymbol{\sigma}$, the network outputs the corresponding value $\psi(\boldsymbol{\sigma})$.

output of the neural network is

$$\psi_{\boldsymbol{\theta}}(\boldsymbol{\sigma}) = \sqrt{p_{\boldsymbol{\theta}}(\boldsymbol{\sigma})}\, e^{i\phi_{\boldsymbol{\theta}}(\boldsymbol{\sigma})}, \tag{2.2}$$

where $\boldsymbol{\theta}$ denotes the set of variational parameters, namely weights and biases of the network. Here, $p_{\boldsymbol{\theta}}(\boldsymbol{\sigma}) = |\psi_{\boldsymbol{\theta}}(\boldsymbol{\sigma})|^2$ and $\phi_{\boldsymbol{\theta}}(\boldsymbol{\sigma})$ represent the amplitude and the phase respectively. A non-rigorous example of NQS is provided in Fig. 2.1. Sticking with this general case, the amplitude and the phase can be modeled with different technique, e.g. using two separates neural networks, two output nodes of the same network or a single network with complex coefficients. We emphasize that the systems considered in this thesis are characterized by real and positive wave functions. Consequently, we can restrict ourselves to the case in which Eq. (2.2) reduces to the simpler form

$$\psi_{\boldsymbol{\theta}}(\boldsymbol{\sigma}) = \sqrt{p_{\boldsymbol{\theta}}(\boldsymbol{\sigma})}. \tag{2.3}$$

The expressive power of neural networks, guaranteed by universal approximation theorems [43], ensures that NQS can, in principle, approximate highly complex quantum states, including those with non-trivial entanglement patterns. This makes them competitive with more conventional variational approaches, such as tensor network states (MPS, PEPS), while often being more flexible in higher dimensions. Nevertheless, the

training of NQS involves navigating a non-convex optimization landscape, which can make the convergence toward the true ground state challenging.

2.1.1 Variational estimation of observables

In original works on neural quantum states, wave functions in the form of Eq. (2.3) were employed as variational ansätze in variational Monte Carlo (VMC) simulations. This is a self-consistent technique, where the network parameters are iteratively updated based on the stochastic measurements of the algorithm until a good approximation of the ground state of the system is represented. Consequently, unlike traditional machine learning approaches, no external data are required to train the neural network. However, in recent years neural quantum states have also proven to be powerful tools when combined with real data, as will be discussed in the following chapters [44, 45, 46]. Following the same recipe of Eq. (1.16), physical observables can be estimated in the standard variational Monte Carlo fashion, hence evaluating their expectation values on a set of configurations $\{\boldsymbol{\sigma}\}$ sampled from the NQS. In particular, for a generic observable $\hat{O}$ we can write

$$\begin{aligned} \frac{\langle \psi_{\boldsymbol{\theta}} | \hat{O} | \psi_{\boldsymbol{\theta}} \rangle}{\langle \psi_{\boldsymbol{\theta}} | \psi_{\boldsymbol{\theta}} \rangle} &= \frac{\sum_{\boldsymbol{\sigma},\boldsymbol{\sigma}'} \langle \psi_{\boldsymbol{\theta}} | \boldsymbol{\sigma} \rangle \langle \boldsymbol{\sigma} | \hat{O} | \boldsymbol{\sigma}' \rangle \langle \boldsymbol{\sigma}' | \psi_{\boldsymbol{\theta}} \rangle}{\sum_{\boldsymbol{\sigma}} \langle \psi_{\boldsymbol{\theta}} | x \rangle \langle \boldsymbol{\sigma} | \psi_{\boldsymbol{\theta}} \rangle} \\ &= \frac{\sum_{\boldsymbol{\sigma},\boldsymbol{\sigma}'} \psi^*_{\boldsymbol{\theta}}(\boldsymbol{\sigma}) \, \hat{O}_{\boldsymbol{\sigma}\boldsymbol{\sigma}'} \, \psi_{\boldsymbol{\theta}}(\boldsymbol{\sigma}')}{\sum_{\boldsymbol{\sigma}} |\psi(\boldsymbol{\sigma})|^2} \end{aligned} \tag{2.4}$$

Notice that, to seek generality, we are considering complex wave functions. However, the simplest case of real coefficients is straightforward. Starting from the latter equation we can distinguish two cases:

- The operator $\hat{O}$ is diagonal in the computational basis, such as $\hat{O}_{\boldsymbol{\sigma}\boldsymbol{\sigma}'} = \delta_{\boldsymbol{\sigma}\boldsymbol{\sigma}'} O(\boldsymbol{\sigma})$, then it's easy to see that the expectation value of Eq. (2.4) becomes

$$\frac{\langle \psi_{\boldsymbol{\theta}} | \hat{O} | \psi_{\boldsymbol{\theta}} \rangle}{\langle \psi_{\boldsymbol{\theta}} | \psi_{\boldsymbol{\theta}} \rangle} = \frac{\sum_{\boldsymbol{\sigma}} |\psi_{\boldsymbol{\theta}}|^2 \hat{O}(\boldsymbol{\sigma})}{\sum_{\boldsymbol{\sigma}} |\psi_{\boldsymbol{\theta}}|^2} = \langle \hat{O} \rangle, \tag{2.5}$$

  where, $\langle \hat{O} \rangle$ is the expectation value over the probability distribution $|\psi_{\boldsymbol{\theta}}|^2$. This means that these expectation values can be easily computed by sample configurations of the Hilbert space according with the square modulus of the wave

function and calculating the statistical average.

- The operator $\hat{O}$ is off-diagonal in the computational basis. Recalling Eq. (1.17), we can define the local operator as

$$\begin{aligned}\hat{O}_{\text{loc}}(\boldsymbol{\sigma}) &= \frac{\langle \boldsymbol{\sigma}|\hat{O}|\psi_{\boldsymbol{\theta}}\rangle}{\langle \boldsymbol{\sigma}|\psi_{\boldsymbol{\theta}}\rangle} \\ &= \frac{\langle \boldsymbol{\sigma}|\hat{O}\Big(\sum_{\boldsymbol{\sigma}'}|\boldsymbol{\sigma}'\rangle\langle\boldsymbol{\sigma}'|\Big)|\psi_{\boldsymbol{\theta}}\rangle}{\langle \boldsymbol{\sigma}|\psi_{\boldsymbol{\theta}}\rangle} \\ &= \frac{\sum_{\boldsymbol{\sigma}'}\langle \boldsymbol{\sigma}|\hat{O}|\boldsymbol{\sigma}'\rangle\langle\boldsymbol{\sigma}'|\psi_{\boldsymbol{\theta}}\rangle}{\langle \boldsymbol{\sigma}|\psi_{\boldsymbol{\theta}}\rangle} \\ &= \frac{\sum_{\boldsymbol{\sigma}'} O_{\boldsymbol{\sigma}\boldsymbol{\sigma}'}\,\psi_{\boldsymbol{\theta}}(\boldsymbol{\sigma}')}{\psi_{\boldsymbol{\theta}}(\boldsymbol{\sigma})} \\ &= \sum_{\boldsymbol{\sigma}'} O_{\boldsymbol{\sigma}\boldsymbol{\sigma}'}\frac{\psi_{\boldsymbol{\theta}}(\boldsymbol{\sigma}')}{\psi_{\boldsymbol{\theta}}(\boldsymbol{\sigma})}.\end{aligned} \tag{2.6}$$

  Following the procedure in Eq. (1.16) it's easy to prove that the generic expectation value of Eq. (2.4) can be written as

$$\frac{\langle\psi_{\boldsymbol{\theta}}|\hat{O}|\psi_{\boldsymbol{\theta}}\rangle}{\langle\psi_{\boldsymbol{\theta}}|\psi_{\boldsymbol{\theta}}\rangle} = \frac{\sum_{\boldsymbol{\sigma}}|\psi_{\boldsymbol{\theta}}|^2\hat{O}_{\text{loc}}(\boldsymbol{\sigma})}{\sum_{\boldsymbol{\sigma}}|\psi_{\boldsymbol{\theta}}|^2} = \langle\hat{O}_{\text{loc}}\rangle. \tag{2.7}$$

  In this way, we can compute the expectation values of any operators as the statistical average of Hilbert space configuration sampled by the neural network wave function.

At this point, the reason behind the introduction of the concept of local operator in the context of the Hamiltonian in Chap. 1, shoul be clear. In general, this operator is not diagonal in the computational basis. The most simple case is the Transverse Field Ising Hamiltonian (TFIM) which will be extensively discussed in the following sections.

In recent years, different NQS architectures have been explored (a detailed review is provided in Ref. [47]). Among these, one of the most prominent is the Restricted Boltzmann Machine (RBM), which combines analytical tractability with the flexibility to represent real quantum systems. In the following, we will describe the model, the training procedure, and how it can be combined with the PQMC described in Chap. 1 to study unbiased ground state property of quantum many body systems.

2.2 Architecture

A Restricted Boltzmann Machine is a parametrized generative neural network that aims to learn the probability distribution from a set of data. Given samples from the train dataset, training an RBM means to adjust the parameters such that the probability distribution described by the model fits the probability distribution of the train samples. Graphically, the structure of an RBM is that of an undirected graph composed of two layers, called the *visible layer* and the *hidden layer*, each of which is made up of a fixed number of neurons. The neurons composing the visible layer are called *visible units* and represent the components of the sample (e.g., one visible unit for each pixel of an image, or for each particle site in a physical system). The neurons composing the hidden layer are called *hidden neurons* and represent the relations between visible variables. To align with the standard notation in machine learning, we employ the binary convention $\{0,1\}$. In computational physics, the $\{-1,1\}$ convention is more common; however, the two formulations are equivalent, and the mathematical mapping between them is straightforward. Let us call $\mathbf{v} = (v_1, v_2, ..., v_M)$, with $v_i = \{0,1\}$ the visible variables and $\mathbf{h} = (h_1, h_2, ..., h_N)$, with $h_i = \{0,1\}$ the hidden variables, where M and N are the number of visible and hidden units respectively. RBMs are also called energy-based model. The idea behind these models is to associate a certain energy with the configurations sampled by the network. This scalar value serves as a measure of compatibility between the generated samples and the training data. The so-called energy function of the RBM reads

$$H_{\text{RBM}}(\mathbf{v}, \mathbf{h}) = -\sum_{i,j} W_{ij} h_i v_j - \sum_i a_j v_j - \sum_j b_i h_i. \tag{2.8}$$

The parameters a_j and b_i are called *visible bias* and *hidden bias* respectively, the W_{ij} matrix are the couplings between the visible and the hidden units. The whole set of parameters can be denoted as $\boldsymbol{\theta} = \{W_{ij}, a_j, b_i\}$. In general, we are interested in probability of sample a visible configuration; this is given by the marginal distribution over all possible hidden-spin configurations $\mathbf{h}$:

$$p(\mathbf{v}) \equiv \sum_h p(\mathbf{v}, \mathbf{h}) = \frac{1}{Z} \sum_h \exp\left[-H_{\text{RBM}}(\mathbf{v}, \mathbf{h})\right], \tag{2.9}$$

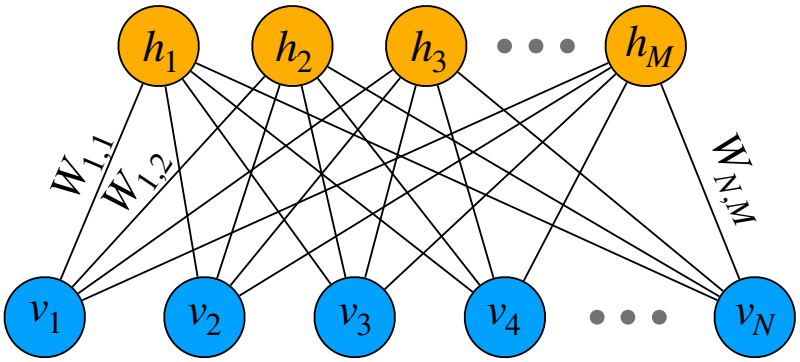

FIGURE 2.2: The graphical representation of an RBM. The two sets of neurons are connected by the coupling matrix W_{ij}, which is allowed only between neurons from different layers and not between neurons within the same layer. Furthermore, the structure is that of an undirected graph. Notice that here, the visible and hidden biases are not shown. However they act in each visible and hidden variables respectively.

where $Z = \sum_{v,h} \exp\left[-E_{\mathrm{RBM}}(\mathbf{v},\mathbf{h})\right]$ is the standard partition function. Notice from Fig. 2.2 and Eq. (2.8) that visible-visible and hidden-hidden connection are not allowed, hence the suffix restricted. This restriction allows us to compute explicitly the probability by tracing out the hidden-spin configurations as follows:

$$\begin{aligned} p(\mathbf{v}) &= \sum_h \exp\left[\sum_{i,j} W_{ij} h_i v_j + \sum_i a_j v_j + \sum_j b_i h_i\right] \\ &= e^{\sum_j v_j a_j} \times \sum_h \exp\left[\sum_j W_{ij} v_j h_i + b_i h_i\right] \\ &= e^{\sum_j v_j a_j} \times \prod_i \left(\exp\left[\sum_j W_{ij} v_j + b_i\right] + \left[-\sum_j W_{ij} v_j - b_i\right]\right) \\ &= e^{\sum_j v_j a_j} \times \prod_i 2\cosh\left(\sum_i W_{ij} v_j + b_i\right). \end{aligned} \tag{2.10}$$

2.3 Training

Training an RBM involves finding the optimal parameters such that the network can best represent the probability distribution of the training data. To do this one have to maximize the log-likelihood

$$\ln\left(\mathcal{L}(\mathbf{\Theta}, S)\right) = \ln\left(\prod_{k=1}^{N_s} p(\mathbf{v_k})\right) = \sum_{k=1}^{N_s} \ln p(\mathbf{v}_k), \tag{2.11}$$

where $S \equiv (\mathbf{v}_1, \mathbf{v}_2, ..., \mathbf{v}_{N_s})$ is the dataset and N_s the total number of samples. This is equivalent to minimizing the distance between the target distribution q describing S and the distribution p of the RBM, in terms of Kullback-Leibler divergence (KL-divergence):

$$\mathrm{KL}(q|p) \equiv \sum_{\mathbf{v}} q(\mathbf{v}) \ln q(\mathbf{v}) - \sum_{\mathbf{v}} q(\mathbf{v}) \ln p(\mathbf{v}) = \sum_{\mathbf{v}} q(\mathbf{v}) \ln\left[\frac{q(\mathbf{v})}{p(\mathbf{v})}\right]. \tag{2.12}$$

Since the term with only $q(\mathbf{v})$ does not depend on the RBM, there is no symmetry with respect to exchange $q \leftrightarrow p$.

2.3.1 Optimization: stochastic gradient ascent

The training process consists in the optimization of the parameters in order to maximize (or minimize) the loss function. In order to optimize the weights at step t, starting from an arbitrary point in the parameter space, the following update rule is applied until convergence:

$$\theta_j \leftarrow \theta_j + \eta \frac{\partial \mathcal{L}(\mathbf{\Theta})}{\partial \theta_j}. \tag{2.13}$$

The parameter η is the so-called *learning rate* and it can be adjusted according to needs. This algorithm is known as *gradient ascent* and is one of the standard process to modify the parameters of an artificial neural network in order to maximize the loss function.
In practice, a slightly modified version of this algorithm is used: instead of using all training instances ($N_s \sim 10^4, 10^5$), a small mini-batch indexed by N_b of 10, 100 data points is randomly chosen. This way, it is possible to calculate the "average direction"

in the parameter landscape, with statistical noise that allows escaping from local minima. Eq. (2.13) becomes:

$$\theta^{t+1} = \theta^t + \eta \frac{\partial}{\partial \theta} \left(\sum_{l=1}^{N_b} \mathcal{L}(\theta^t, \mathbf{v}_l) \right). \tag{2.14}$$

This modified version of the gradient ascent is the so called *stochastic gradient ascent*. In order to compute the second term of Eq. (2.14) we have to compute the derivative of the loss function. It results:

$$\begin{aligned} \frac{\partial \ln \mathcal{L}(\mathbf{v}, \boldsymbol{\Theta})}{\partial \theta} &= \frac{\partial}{\partial \theta} \left[\ln \sum_{\mathbf{h}} \exp\left(-H(\mathbf{v}, \mathbf{h})\right) \right] - \frac{\partial}{\partial \theta} \left[\ln \sum_{\mathbf{v},\mathbf{h}} \exp\left(-H(\mathbf{v}, \mathbf{h})\right) \right] \\ &= -\sum_{\mathbf{h}} p(\mathbf{h}|\mathbf{v}) \frac{\partial H(\mathbf{v}, \mathbf{h})}{\partial \theta} + \sum_{\mathbf{v},\mathbf{h}} p(\mathbf{v}, \mathbf{h}) \frac{\partial H(\mathbf{v}, \mathbf{h})}{\partial \theta}, \end{aligned} \tag{2.15}$$

where the conditional probability $p(\mathbf{h}, \mathbf{v})$ reads

$$p(\mathbf{h}|\mathbf{v}) = \frac{p(\mathbf{v}, \mathbf{h})}{p(\mathbf{v})} = \frac{\frac{1}{Z} \exp\left(-H(\mathbf{v}, \mathbf{h})\right)}{\frac{1}{Z} \sum_{\mathbf{h}} \exp\left(-H(\mathbf{v}, \mathbf{h})\right)} = \frac{\exp\left(-H(\mathbf{v}, \mathbf{h})\right)}{\sum_{\mathbf{h}} \exp\left(-H(\mathbf{v}, \mathbf{h})\right)}. \tag{2.16}$$

2.3.2 Gibbs sampling and K-step contrastive divergence

The restrictive structure of the Boltzmann Machine comes into play in the calculation of these conditional probabilities. Since the hidden variables are conditionally independent given fixed visible values and vice versa, the conditional probability distributions are separable:

$$\begin{aligned} p(\mathbf{h}|\mathbf{v}) &= \prod_{j=1}^{M} p\left(h_j|\mathbf{v}\right) \\ p(\mathbf{v}|\mathbf{h}) &= \prod_{i=1}^{N} p\left(v_i|\mathbf{h}\right). \end{aligned} \tag{2.17}$$

Let us consider the derivative of the log likelihood with respect to the inter-layer parameters matrix W_{ij}:

$$\frac{\partial \mathcal{L}(\mathbf{v}, \boldsymbol{\Theta})}{\partial W_{ij}} = \sum_{\mathbf{h}} p(\mathbf{h}|\mathbf{v}) h_i v_j - \sum_{\mathbf{h},\mathbf{v}} p(\mathbf{h}, \mathbf{v}) h_i v_j \tag{2.18}$$

For a single batch S the latter equation is often written in the form:

$$\sum_{\mathbf{v}\in S} \frac{\partial \ln \mathcal{L}(\boldsymbol{\Theta}, \mathbf{v})}{\partial W_{ij}} \propto \langle v_j h_i \rangle_{data} - \langle v_j h_i \rangle_{model}. \tag{2.19}$$

While the positive term can be computed directly from the data, the computation of the second expectation value involves the sum over all possible configuration of hidden and visible variables under the model distribution. Since the number of configurations grows exponentially with the number of units, this summation can become impractical.

The workaround is to employ a stochastic estimate, obtained using the so-called Gibbs sampling algorithm. One starts from a visible configuration from the train dataset $\mathbf{v}^{(0)}$. At every step $s = 1, \ldots, k$, one samples $\mathbf{h}^{(s)}$ from the conditional $p(\mathbf{h}|\mathbf{v}^{(s)})$ and then $\mathbf{v}^{(s+1)}$ from $p\left(\mathbf{v}|\mathbf{h}^{(s)}\right)$. In practice, one can show that the conditionals in Eq. (2.17) can be written in the following form:

$$\begin{aligned} p(h_j = 1|\mathbf{v}) &= \sigma\left(\sum_{j=1}^{N_v} W_{ij} v_j + b_i\right) \\ p(v_i = 1|\mathbf{h}) &= \sigma\left(\sum_{i=1}^{N_h} W_{ij} h_i + a_j\right), \end{aligned} \tag{2.20}$$

where $\sigma(x) = 1/[1 + \exp(-x)]$ is the *sigmoid* function. In this way, the values of hidden and visible neurons are sampled in parallel (once the other layer has fixed values). The sampling procedure for the visible variables is shown in Alg. 2:

Algorithm 2 Sampling visible variables

```
for j ← 1 to N do                                   ▷ For each visible variable v_j
    r ← UniformRandomNumber()  ▷ Sample a uniform random variable in (0, 1)
    if r < p(v_j = 1, h) then
        v_j ← 1                                                  ▷ Set v_j to 1
    else
        v_j ← 0                                                  ▷ Set v_j to 0
    end if
end for
```

The whole k-step contrastive divergence for $v_j \in \{0, 1\}$ and $h_i \in \{0, 1\}$ is sketched in Alg. 3.

Algorithm 3 k-step contrastive divergence.

Input: mini-batch $\mathbf{v}_l$ with $l = 1, ..., N_b$
Output: partial derivatives $\mathcal{L}_{W_{ij}}$, $\mathcal{L}_{a_j}$, $\mathcal{L}_{b_i}$

for $l = 1, ..., N_b$ **do**
  $\mathbf{v}^{(0)} \leftarrow \mathbf{v}_l$
  for $t = 0, ..., k-1$ **do**
    set $h_i^t = 1$ with $p(h_i = 1|\mathbf{v}^t)$, else $h_i^t = 0$
    set $v_j^{t+1} = 1$ with $p(v_j = 1|\mathbf{h}^t)$, else $v_j^{t+1} = 0$
  end for
  for $i = 1, ..., N_h$ and $j = 1, ..., N_v$ **do**
    $\mathcal{L}_{W_{ij}} \leftarrow \mathcal{L}_{W_{ij}} + p(h_i = 1|\mathbf{v}^{(0)})v_j^{(0)} - p(h_i = 1|\mathbf{v}^{(k)})v_j^{(k)}$
    $\mathcal{L}_{a_j} \leftarrow \mathcal{L}_{a_j} + v_j^0 - v_j^k$
    $\mathcal{L}_{b_i} \leftarrow \mathcal{L}_{b_i} + p(h_i = 1|\mathbf{v}^{(0)}) - p(h_i = 1|\mathbf{v}^{(k)})$
  end for
end for

In the next section, we show how this statistical model, extensively employed in the framework of Variational Monte Carlo, can be employed as a guiding wave function for the PQMC algorithm.

2.4 Self-Learning PQMC

As explained in Chap. 1, the convergence of PQMC is guaranteed provided that a sufficiently large number of random walkers explores the Hilbert space, so in the limit $N_w \to \infty$, and that an appropriate guiding wave function ψ_g is employed. It is worth emphasizing that if the guiding wave function is not adequate, for instance in the trivial case $\psi_g(\boldsymbol{\sigma}) = 1 \ \forall \boldsymbol{\sigma}$ where $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, ..., \sigma_N)$ are the configurations of the system, the number of walkers required to estimate the ground-state energy within a target accuracy grows exponentially with the system size [48]. This leads to a computational cost that grows exponentially. On the other hand, if the guiding wave function is exact, i.e., $\psi_g(\boldsymbol{\theta}) = \psi_0(\boldsymbol{\theta})$, with ψ_0 denoting the exact ground state wave function, the local energy $E_{\text{loc}}(\boldsymbol{\theta})$ is a constant function. In this case, the annihilation and replication of random walkers in the branching process are completely suppressed, thereby eliminating any source of bias. More generally, when $\psi_g(\boldsymbol{\theta})$ is a reasonable approximation to $\psi_0(\boldsymbol{\theta})$, the fluctuations in the walker population are significantly reduced

compared to the case of the simple PQMC algorithm with $\psi_g(\boldsymbol{\theta}) = 1$, resulting in a much faster convergence towards the exact $N_w \to \infty$ limit. In Ref. [36], an unrestricted Boltzmann machine, a slightly different model than the one of our interest, was employed as a parametrized guiding wave function for PQMC simulations of a 1D TFIM. It was shown that, when the network parameters are properly optimized, the exponential scaling of the number of walkers with system size, characteristic of plain PQMC, is reduced to a polynomial scaling.In that study, the optimal set of parameters was first determined through an independent variational simulation. Subsequently, the resulting network, optimized to faithfully represent the ground state of the system, was implemented into the PQMC simulation, where it served as the guiding wave function.

In this section we describe the self-learning PQMC which consists in the implementation of an adaptive RBM in the projection algorithm. It was first propsed in [16] and together with a few variants, it constitutes the algorithm on which all the original results of this thesis are based. The key idea behind the self-learning PQMC is to train the RBM on-the-fly during the simulation via the k-step contrastive divergence of Alg. 3, thereby eliminating the need of a separate variational procedure. As a consequence, this method provide an accurate ansatz for the ground-state wave function, obtained by minimizing the Kullback–Leibler divergence in Eq. (2.12) with respect to the PQMC samples, rather than the energy expectation value as in conventional variational approaches. As in the case of PQMC guided by unrestricted Boltzmann machines, the effectiveness of this self-learning PQMC scheme is demonstrated on the paradigmatic ferromagnetic quantum Ising chain, where it achieves excellent agreement with both Jordan–Wigner theory.

In the adaptive scheme of self-learning PQMC, the unsupervised learinig algorithm is employed to train the RBM to represent the unnormalized probability distribution $f(\boldsymbol{\sigma}, \tau \to \infty)$, corresponding to the equilibrium population of random walkers generated by a PQMC simulation. As mentioned in the first part of this chapter, an RBM with a sufficiently large number of hidden units N_h can, in principle, approximate any discrete distribution and several studies have shown that the training algorithms outlined above are capable of determining RBM parameters that are optimal or nearly optimal. In this protocol, the training dataset for the RBM is provided directly by the random walkers generated during the PQMC simulation. At the initial stage,

the RBM is initialized with random parameters, producing a rather *dummy* guiding wave function. Nevertheless, even with such a poor ansatz, PQMC intrinsically tends to sample the most relevant regions of the Hilbert space. After the first stint, the walkers that survive the branching process are accumulated and used as the training dataset for the RBM. As a result, in the subsequent stint the RBM provides a better approximation of the ground-state wave function and thus serves as a more effective guide for the random walkers. This iterative refinement continues until convergence is reached. To be more specific, the scheme consists of a sequence of stints labeled by $s = 0, 1, 2, ..., N_{\text{stints}} - 1$. In each stint, a PQMC simulation with a guiding function $\psi_g^s(\boldsymbol{\sigma})$ is performed over an imaginary time τ_s, followed by a learning stage where the RBM is trained on the equilibrium walkers from the precedent stint, producing a new guiding function $\psi_g^{s+1}(\boldsymbol{\sigma})$. The total imaginary time is $\tau = \tau_s N_{\text{stints}}$. Notice that the RBM parameters at each stage are initialized with the optimized values obtained in the previous stint.

For the initial stint $s = 0$, the guiding function can be extremely simple for example, a constant function or the square root of an RBM distribution with random weights, $\psi_g^{(0)}(\boldsymbol{\sigma}) \propto \sqrt{p(\boldsymbol{\sigma})}$. After the s-th stint, the trained RBM distribution satisfies $p_s(\boldsymbol{\sigma}) \propto \psi_g^{(s)}(\boldsymbol{\sigma})\psi_0(\boldsymbol{\sigma})$, provided that N_w and τ_s are sufficiently large. With the iterative update rule $\psi_g^{s+1}(\boldsymbol{\sigma}) = \sqrt{p_s(\boldsymbol{\sigma})}$, the accuracy of the guiding function systematically improves. Under idealized conditions—large enough N_h, successful unsupervised optimization, and sufficiently large N_w and τ_s the walkers in stint s sample the (unnormalized) distribution

$$f(\boldsymbol{\sigma}, \tau \to \infty) = \psi_g^{(0)}(\boldsymbol{\sigma})^{1/2^s} \psi_0(\boldsymbol{\sigma})^{2-1/2^s}, \tag{2.21}$$

which converges rapidly to the exact ground-state wave function, $\psi_g^{(\infty)}(\boldsymbol{\sigma}) \to \psi_0(\boldsymbol{\sigma})$. Consequently, even under sub-optimal conditions, the bias due to a finite number of walkers is expected to vanish after only relative small number of stints. A qualitative sketch of the algorithm is provided in Fig. 2.3. A detailed description of the algorithm can be found in the original paper [16].

In the following subsection, we demonstrate the effectiveness of the method for a paradigmatic system, namely the one-dimensional TFIM. In particular, we show how the expressivity of the network depends on the number of hidden neurons N_h,

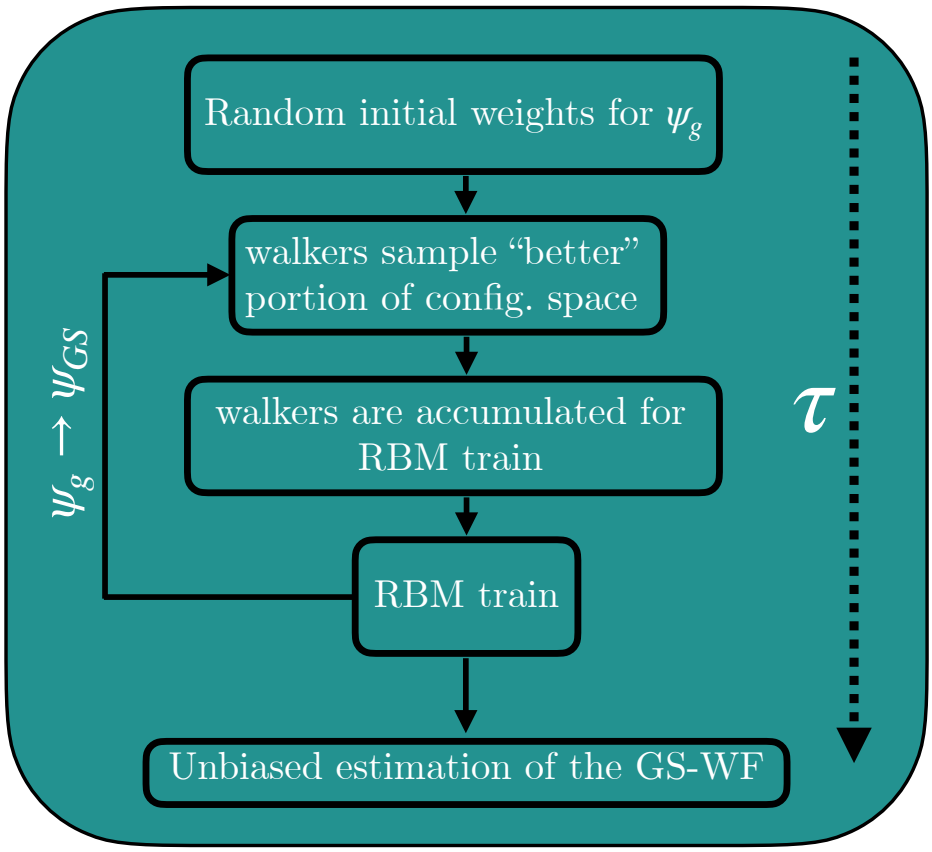

FIGURE 2.3: Diagram of the Self-learining algorithm. Starting from a random RBM, PQMC generates a population of random walkers sampling the most relevant regions of the Hilbert space. The surviving walkers after each stint are bookkeped and used as the training dataset for the RBM, whose parameters are updated to approximate an improved guiding wave function. This refined wave function then guides the next PQMC stint, leading progressively to a more accurate approximations of the ground state until convergence is reached.

and how this algorithm allows unbiased simulations of the ground-state properties of quantum systems of sizes that are otherwise intractable with exact methods.

2.4.1 Test case: The Transverse Field Ising Model (TFIM)

To demonstrate the effectiveness of the method, we consider a paradigmatic case in statistical physics: the Ising chain in the presence of a transverse magnetic field. This model serves as a useful benchmark, since it can be directly mapped onto a free-fermion Hamiltonian via the Jordan–Wigner transformation. As a result, it allows for exact solutions even for system sizes that would be inaccessible to standard exact diagonalization techniques. The Hamiltonian reads

$$\hat{H} = -J \sum_{i=1}^{N} \sigma_i^z \sigma_{i+1}^z - \Gamma \sum_{i=1}^{N} \sigma_i^x. \tag{2.22}$$

Here, the interaction constant is set to $J = 1$ determining the ferromagnetic TFIM. Hence, all the energy scales are expressed in units of J. We consider open boundary conditions (OBC). Recalling Eq. (1.21), it is worth emphasizing that the interaction term of the Ising Hamiltonian $\hat{H}_{\text{pot}} = -J \sum_{i=1}^{N-1} \sigma_i^z \sigma_{i+1}^z$ corresponds to the so-called potential energy, as it is diagonal in the computational basis. The kinetic term, on the other hand, is represented by $\hat{H}_{\text{kin}} = -\Gamma \sum_{i=1}^{N-1} \sigma_i^x$, which is off-diagonal. This contribution induces the quantum fluctuations of the system, which are responsible for driving the phase transitions. It is well known that, at zero temperature, this model undergoes a quantum phase transition from an ordered ferromagnetic phase ($\Gamma < J$) to a disordered paramagnetic phase ($\Gamma > J$). Due to critical correlations, Monte Carlo simulations become particularly challenging at the critical point. For this reason, we consider a chain of $N = 60$ spins and place the system precisely at the transition point $\Gamma_c = 1.0J$. The purpose of the simulations is to demonstrate that, even in this critical regime, increasing the number of hidden neurons N_h in the RBM guiding the simulation allows us to overcome the bias introduced by the finite walker population, which in this case is fixed at $N_w = 4000$.

The results shown in Fig. 2.4 indicate that, provided the number of hidden neurons is not too small, unbiased ground-state properties of the system can be extracted. An important point to stress is that PQMC is not a variational method since it projects the

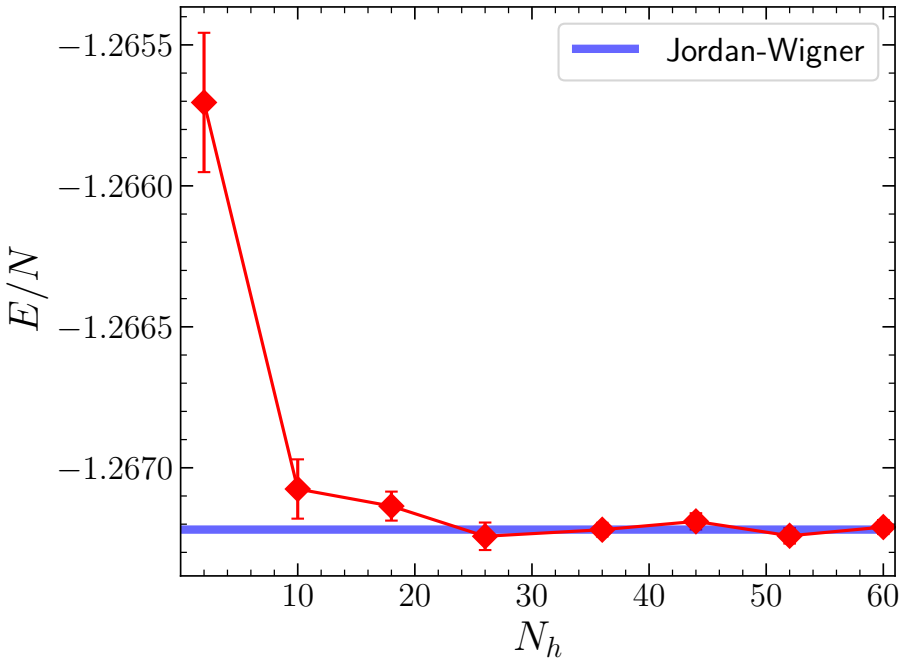

FIGURE 2.4: Energy per spin as a function of the number of hidden neurons N_h of the RBM guiding wave function. The blue solid line is the exact energy obtained via the Jordan–Wigner transformation. Simulations are performed for a spin chain of size $N = 60$ at the quantum critical point between the paramagnetic and ferromagnetic phases, $\Gamma_c = 1.0J$. The number of walkers used in the PQMC is $N_w = 4000$.

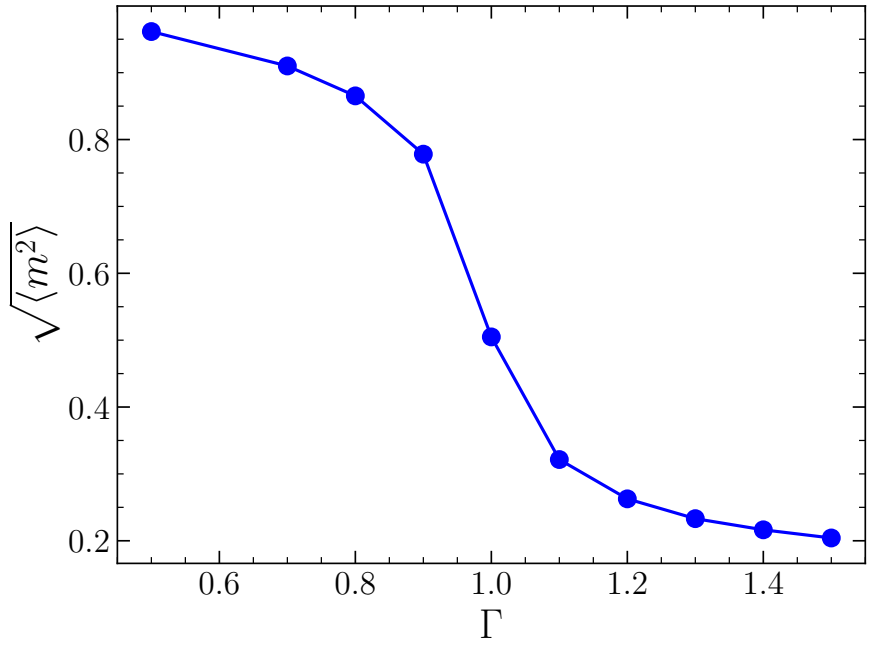

FIGURE 2.5: Pure estimation of the square root of the average squared magnetization, $\sqrt{\langle m^2 \rangle}$, as a function of the transverse field Γ for a system of $N = 60$ spins. A rather sharp transition is clearly visible at $\Gamma = \Gamma_c = 1.0J$, marking the ferromagnetic–paramagnetic quantum phase transition. This result shows that the self-learning PQMC algorithm can successfully identify quantum phase transitions by sampling ground-state configurations according to the unbiased ground-state wave function. Error bars are smaller than the marker size.

state in imaginary time providing an approximation of the ground state wave function. In the limit $\tau \to \infty$, the results is fully unbiased.

Throughout this section, we have shown that the continuous time PQMC, with an appropriate choice parameters such as the number of walkers N_w and the number of hidden neurons N_h in the RBM, allows us to simulate the correct state of the system. Once convergence is achieved, in principle, any physical observable can be computed. In Fig. 2.5 we report the square root of the average of the squared magnetization along the z-direction, defined as

$$\sqrt{\langle m^2 \rangle} = \sqrt{\left\langle \left(\frac{1}{N} \sum_{i=1}^{N} \sigma_i^z \right)^2 \right\rangle}. \tag{2.23}$$

In systems with a $\mathbb{Z}_2$ symmetry, such as the ferromagnetic Ising model, sufficiently

long simulations typically restore the symmetry between the two degenerate ground states (most spins up and most spins down). As a result, the direct average magnetization $\langle m \rangle$ tends to vanish even in the ordered phase. For this reason, the quantity $\sqrt{\langle m^2 \rangle}$ is used here as an effective order parameter to characterize the transition from the ferromagnetic to the paramagnetic phase. It remains finite in the ordered phase and approaches zero in the disordered one for large τ, thus providing a clear signal of the quantum phase transition. In the following chapters, we will consider cases where the continuous-time self-learning PQMC can be used to simulate the ground state not only of paradigmatic systems such as the TFIM discussed in this section, but also of much more complex and still mostly unexplored systems, namely quantum spin glasses.

Chapter 3

Spin glass quantum phase transition of two-dimensional Edwards-Anderson model via neural quantum Monte Carlo simulations

In the following chapter, we apply the continuous-time self-learning Projection Quantum Monte Carlo method, introduced in the Chap. 1 and Chap. 2, to simulate a paradigmatic disordered quantum system: the two-dimensional Edwards-Anderson (2D EA) spin-glass Hamiltonian in a transverse field, with Gaussian nearest-neighbor couplings. This study is a step forward from the benchmark case of the TFIM, extending the approach to a far more complex and less explored class of systems. We show that guiding wave functions based on self-learned neural networks are able to suppress the population-control bias down to small statistical uncertainties, at least up to system sizes of the order of one hundred spins. To study the spin-glass phase, we modify the self-learning PQMC projecting a two-fold replicated Hamiltonian and compute the spin glass order parameter, namely the spin overlap. We perform a finite-size scaling analysis to estimate the critical transverse field at which the spin-glass transition occurs, together with the critical exponents of the correlation length and the spin-glass susceptibility. For the latter two, our results are in good agreement with recent findings reported in the literature for different types of random couplings. Finally, we analyze the spin-overlap distribution within the spin-glass phase. For the system sizes accessible to our simulations, we observe a non-trivial double-peak structure with significant weight at zero overlap, providing further evidence of the complex nature of

the spin-glass ground state. The results of this chapter have been published in Ref. [32].

3.1 Quick overview on quantum spin glasses

Despite decades of research, Ising spin glasses continue to raise fundamental open questions, including the fate of replica symmetry breaking in finite dimensional lattices, i.e., beyond the mean-field Sherrington-Kirkpatrick (SK) model with all-to-all couplings, and the possible existence of the Almeida-Thouless line in the presence of a longitudinal field [30, 49]. In recent years, quantum random Ising models have attracted growing interest, as understanding their zero-temperature properties is crucial for assessing the potential efficiency of quantum annealers [50, 26]. Nevertheless, the accurate simulation of the ground state of quantum spin systems with frustrated interactions remains a formidable computational challenge [51].

Most computational investigations on quantum spin glasses have focused on finite-temperature properties, typically employing path-integral Monte Carlo (PIMC) techniques. The finite-temperature spin-glass transition of the quantum SK model in a weak transverse field was studied in Refs. [52, 34]. Some properties of the quantum phase transition were extracted approaching the zero-temperature limit, considering two- and three-dimensional lattices [53, 54] (with finite imaginary-time step models), the (quantum) SK model [55], and the Bethe lattice [56]. The quantum SK model with longitudinal field has been investigated also via a mapping to a one-dimensional model in the thermodynamic limit [57] and using a continuous-time PIMC algorithm in presence of full replica symmetry breaking [58].

However, PIMC simulations become increasingly demanding in the zero-temperature limit. Only very recently, extreme-scale PIMC studies have succeeded in accurately determining the system-size scaling of the first excitation gap, which determines the computational complexity of quantum-annealing optimization [27]. Analyzing the zero-temperature quantum critical properties is possible [33, 27], but it requires effectively reducing a two-dimensional scaling function to a one-dimensional form. For few-spin models, exact diagonalization methods represent a suitable alternative to study ground state properties [59, 34]. Alternatively, variational ansatzes for the ground-state wave function have been proposed, specifically, employing generalized coherent states in the case of quantum SK model [60].

As we have seen in previous chapters, Projection Quantum Monte Carlo (PQMC) algorithms provide an effective framework for simulating the ground states of general quantum many-body systems [31]. These methods are potentially exact, provided that the negative sign problem is absent. However, the need to control the random walk population during the stochastic evolution may sometimes introduce systematic biases [48, 39, 8, 10, 9]. As discussed in Chapter 1, population-control bias in PQMC can be efficiently mitigated by using accurate guiding wave functions. Neural-network-based ansätze, known as Neural Quantum States, have proven particularly particularly effective [11, 61]. Such representation have been successfully integrated into PQMC algorithms, including self-learning variants that update the neural guiding function during the simulation [16, 62]. This strategy circumvents the need for explicit variational parameter optimization, which can sometimes be problematic. However, so far the self-learning PQMC approach has been applied only to one-dimensional nearest-neighbour ferromagnetic and random Ising models, where frustration effects are absent.

In this chapter, we present continuous-time PQMC simulations of the two-dimensional quantum Edwards-Anderson (EA) Hamiltonian with Gaussian random couplings. A self-learned Restricted Boltzmann machine (RBM) is employed as neural guiding wave function. Our numerical results demonstrate that an RBM with a manageable number of hidden neurons effectively suppresses the systematic biases possibly affecting the energy and the replica spin overlap below sufficiently small statistical uncertainties, at least up to a hundred spins. The spin-overlap estimator is implemented through a single imaginary-time evolution of a two-fold replicated Hamiltonian, and the pure estimator is obtained via the forward-walking technique. A finite-size scaling analysis of the mean-squared EA order parameter and the corresponding Binder cumulant is performed. Finite-size corrections to the scaling ansatz turn out to be negligible for the accessible system sizes, enabling a reliable determination of the critical transverse field at which the quantum spin-glass transition occurs, together with the critical exponents of the correlation length and spin-glass susceptibility. For both quantities, we find good agreement with results from recent extreme-scale PIMC simulations [33, 27]of models belonging to the same universality class but featuring different random couplings. We further investigate the spin-overlap distribution within the spin glass phase. We observe that, at a transverse field approximately 20% lower than the critical

value, for the feasible sizes it displays a non-trivial structure with two (symmetry related) peaks and also substantial weight at zero overlap. If this feature persists in the thermodynamic limit, it would indicate the presence of replica symmetry breaking, consistent with the behavior of classical spin-glass. [63, 64, 65, 66, 67, 68]. Finally, to facilitate future numerical studies of short-range quantum spin glasses, we provide in the repository of Ref. [69] a dataset of ground-state energies for 50 independent realizations of random couplings. This dataset serves as a benchmark reslusrce for testing and validating new computational approaches to frustrated quantum spin systems.

The rest of the chapter is organized as follows: Section 3.2 defines the spin-glass Hamiltonian we study and it provides the details of the modified PQMC algorithm to simulate the double-fold Hamiltonian. The computation of the replica overlap is highlighted. The elimination of the population control bias as a function of the random-walker population size and of the number of hidden neurons in the RBM is analyzed. The quantum phase transition from the paramagnetic to the spin-glass phase is analyzed in Section 3.3. Finally, in Section 3.4 we summarize the main findings of this chapter and we discuss some future perspectives.

3.2 Double-fold RBM for Edwards-Anderson model

3.2.1 Two-dimensional Edwards-Anderson Hamiltonian with transverse field

Our goal is to study the ground-state properties of the two-dimensional quantum EA model. For simulations performed with a single replica of the system, the Hamiltonian reads:

$$\hat{H} = -\sum_{\langle i,j \rangle} J_{ij}\sigma_i^z\sigma_j^z - \Gamma\sum_{i=1}^{N}\sigma_i^x. \tag{3.1}$$

σ_i^x and σ_i^z are conventional Pauli matrices at the lattice sites $i = 1, \ldots, N$, and $N = L^2$ is the number of spins (in the single replica), with L the (adimensional) linear system size. The triangular brackets $\langle i, j \rangle$ indicate that the summation is performed over the nearest-neighbor nodes of a square lattice. The couplings J_{ij} are randomly sampled from a Gaussian distribution with zero mean and unit variance, denoted as $\mathcal{N}(0, 1)$. Γ

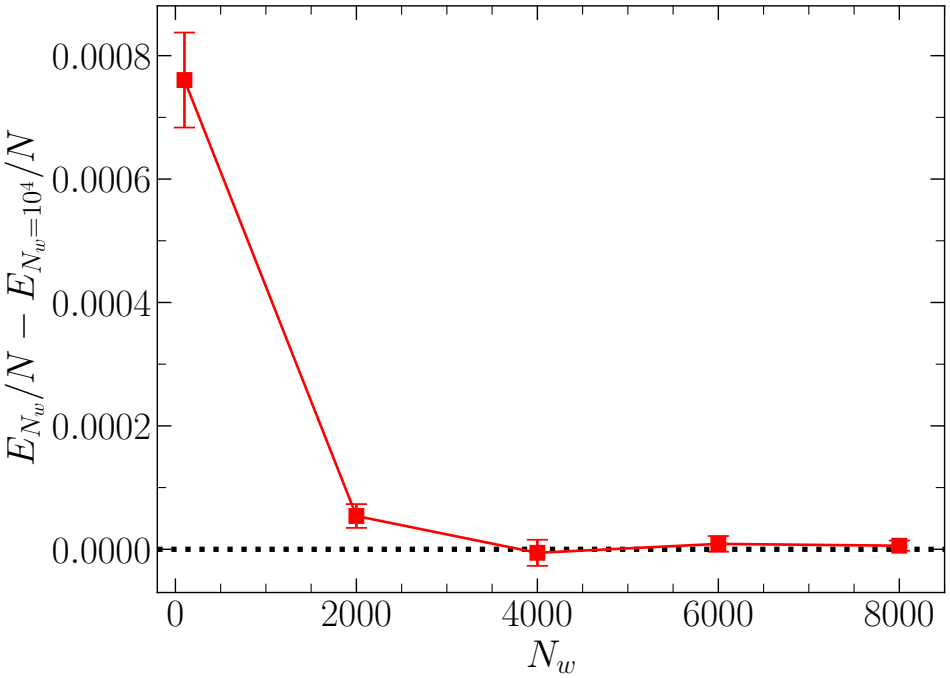

FIGURE 3.1: Bias in the energy per spin E_{N_w}/N as a function of the average number of random walkers N_w, averaged over 5 instances of the 2D Gaussian EA model of size $L = 10$, at transverse field $\Gamma = 1.8$. The reference energy is the results with $N_w = 10^4$ walkers. The RBM guiding wave-function features $N_h = 128$ hidden neurons.

is the intensity of the uniform transverse magnetic field. Hereafter, the eigenstates of the Pauli matrix σ_i^z with eigenvalues $x_i = \pm 1$ are denoted as $|x_i\rangle$. A convenient computational basis is formed by states of N spins $|x\rangle = |x_1 ... x_N\rangle$, with $x = (x_1, \ldots, x_N)$. With $|\psi\rangle$ we denote the state corresponding to the wave-function $\langle x|\ \psi\rangle = \psi(x)$.

3.2.2 PQMC algorithm guided by neural network states

For the sake of completeness, we provide a brief recap of the PQMC algorithm used to simulate the ground-state properties of the Hamiltonian (3.1). For an exhaustive description, see Sec. 2.4. A similar projection method was previously used to simulate the annealing dynamics in Ref. [70]. Another zero-temperature method, but based on the stochastic series expansion algorithm, has recently been applied to the random-field quantum Ising model with ferromagnetic couplings [71].

PQMC algorithms stochastically simulate the imaginary time Schrödinger equation. This allows projecting out the ground-state wave function $\psi_0(x)$, starting from

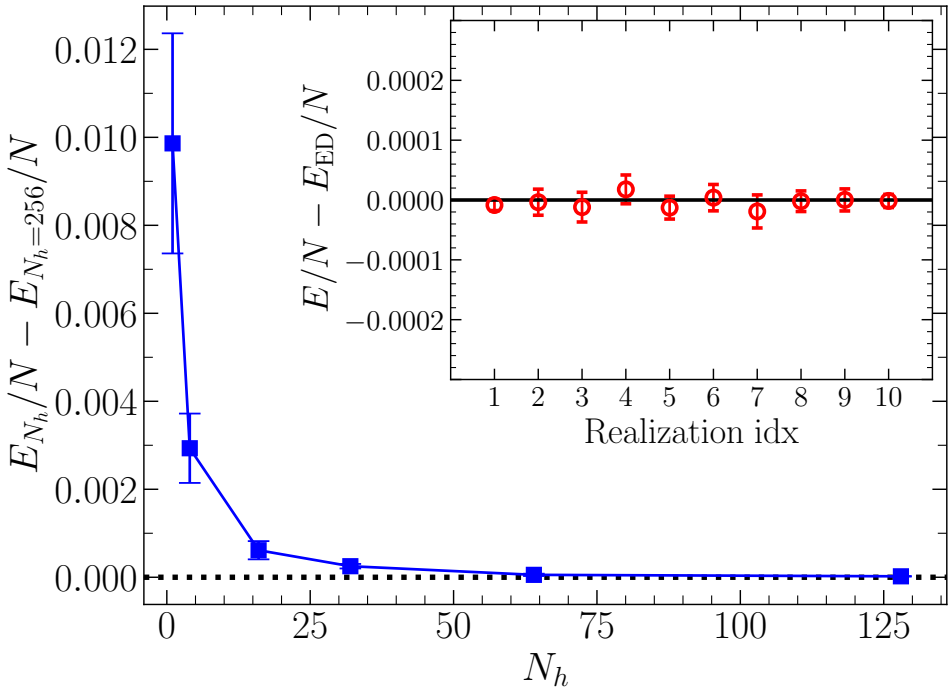

FIGURE 3.2: Main panel: Bias in energy per spin E_{N_h}/N as a function of the number of hidden neurons N_h in the RBM guiding-wave function, averaged over 5 instances as in Fig. 3.1. The reference energy is the results with $N_h = 256$ neurons. The average number of random walkers is $N_w = 4000$. Inset: difference between PQMC predictions E/N and exact-diagonalization results E_{ED}/N for 10 realizations of size $L = 5$, labeled by the index $\mathrm{idx} = 1, 2, \ldots, 10$.

an initial state $\psi(x, 0)$ at imaginary time $\tau = 0$ (assumed not orthogonal to $\psi_0(x)$), as follows:

$$\psi_0(x) = \lim_{\tau\to\infty} \exp\left[-\tau\left(\hat{H} - E_{\mathrm{ref}}\right)\right]\psi(x, 0), \tag{3.2}$$

where E_{ref} is a reference energy introduced to stabilize the numerics, as explained below. To improve efficiency, one introduces a guiding wave function $\psi_g(x)$, namely, a suitable approximation for the ground state. The aim is to sample computational basis states proportionally to the (unnormalized) distribution $f(x, \tau) = \psi(x, \tau)\psi_g(x)$. Long imaginary times are reached iterating many short time-steps $\Delta\tau$, according to the update rule:

$$f(x, \tau + \Delta\tau) = \sum_{x'} \tilde{G}(x, x', \Delta\tau) f(x', \tau), \tag{3.3}$$

where $\tilde{G}(x, x', \Delta\tau) = G(x, x', \Delta\tau)\frac{\psi_g(x)}{\psi_g(x')}$, and

$$G(x, x', \Delta\tau) = \left\langle x \left| \exp\left[-\Delta\tau(\hat{H} - E_{\text{ref}})\right] \right| x' \right\rangle \tag{3.4}$$

is the imaginary-time Green's function. In this work, we adopt the continuous-time PQMC algorithm detailed in Refs. [40, 31, 36], allowing us to avoid finite time-step errors. Notice that continuous-time approaches are possible also for finite temperature simulations; see, e.g., the continuous Wolff algorithm of Ref. [72]. The imaginary-time steps described by Eq. (3.3) are implemented by evolving a population of random walkers, which undergo stochastic updates in the computational basis and a replication process, in jargon called branching, which accounts for the normalization of the update rule. The population of random walkers is controlled to match a target size N_w by dynamically adjusting the reference energy E_{ref}. After a sufficiently long imaginary-time evolution, and provided that N_w is large enough, the walkers sample the computational basis according to $f(\mathbf{x}, \tau) \xrightarrow{\tau\to\infty} \psi_0(\mathbf{x})\psi_g(\mathbf{x})$, where $f(\mathbf{x}, \tau)$ is assumed to be non-negative in the absence of a sign problem. As discussed in Chap. 1, for finite N_w the sampling may exhibit a population-control bias, which can be systematically reduced by improving the guiding wave function $\psi_g(\mathbf{x})$. In the ideal case $\psi_g = \psi_0$, the algorithm becomes exact even with a single walker.

Expectation values of observables are computed via Monte Carlo integration using the mixed estimator,

$$\langle \hat{O} \rangle_{\text{mixed}} = \frac{\langle \psi_0 | \hat{O} | \psi_g \rangle}{\langle \psi_0 | \psi_g \rangle},$$

which coincides with the ground-state value for commuting operators, such as the Hamiltonian [1]. For non-commuting operators, residual bias due to the imperfect guiding function can be further mitigated using the standard forward-walking technique [42, 6]. In this work, we employ neural-network wave functions as flexible guiding ansätze, systematically improving $\psi_g(\mathbf{x})$ by increasing the number of hidden units in the RBM. Again, an accurate guiding function shortens the required forward propagation time, allowing obtaining for large N_w the unbiased ground-state expectation value.

Neural network states can be formed using energy-based generative neural networks in the form of RBMs [11]. These models define an unnormalized probability

distribution over the spin variables x, as follows:

$$P_{\mathrm{RBM}}(x) \propto \exp\left(\sum_i a_i x_i\right) \prod_m F_m(x), \tag{3.5}$$

where

$$F_m(x) = 2\cosh\left[b_m + \sum_{i=1}^{N} w_{im} x_i\right]. \tag{3.6}$$

Here, $m = 1, \ldots, N_h$ labels the N_h neurons in the (unique) hidden layer of the RBM, while the weights w_{im} and the biases a_i and b_m are determined by training the RBM. By setting $a_i = b_m = 0$ we enforce the $\mathbb{Z}_2$ symmetry. Here, we adopt an unsupervised learning protocol, following Ref. [16]. The RBM is trained by minimizing the Kullback-Leibler divergence in Eq. (2.12) with respect to the walkers distribution. For this, the standard $k-$step contrastive divergence algorithm (see Alg. 3) is employed [12]. A sequence of PQMC simulations is performed. In each simulation, a large dataset of random walkers is accumulated, and this dataset is used to train an RBM. In every simulation, the guiding wave function $\psi_g(x) = \sqrt{P_{\mathrm{RBM}}(x)}$ based on the RBM trained after the previous run is used. Notice that this ansatz is legitimate since the ground-state wave function can be assumed to be real and nonnegative. Indeed, the Hamiltonian (3.1) features non-positive off-diagonal elements in the chosen computational basis, such models are sometimes referred to as stoquastic Hamiltonians [73], allowing one to apply the Perron-Frobenius theorem. The resulting self-learning protocol allows us to systematically improve the guiding wave function, thus reducing the population control bias that might occur for finite N_w.

For the typical simulation performed for this work, each RBM is trained on datasets featuring $\simeq 1.2 \times 10^5$ walker configurations. The one-step contrastive divergence algorithm is iterated for 3000 steps using mini-batches of 40 configurations. We iterate PQMC runs followed by RBM training $\simeq 5$ times. A long PQMC simulation is then performed with the optimal RBM. The computer times required for the whole process, considering, e.g., the results discussed in Fig. 3.1, range from $\simeq 0.25$ to $\simeq 15$ hours using, e.g., 6 cores of an Intel(R) Xeon(R) Gold 6154 CPU processor, depending on N_w.

The roles of the walker-population size N_w and of the hidden neuron number N_h are analyzed in Fig. 3.1 and in Fig. 3.2, respectively. The average energy per particle $E/N = \langle \hat{H} \rangle / N$ of five representative instances of the Hamiltonian (3.1) is shown as

a function of N_w and N_h. The convergence to the unbiased asymptotic limits, which we identify with the cases $N_w = 10^4$ (at $N_h = 128$) and $N_h = 256$ (at $N_w = 4000$), is rapid. This demonstrates that the population control bias can be reduced well below the (modest) statistical uncertainties. To further benchmark the neural PQMC algorithm, we make comparison against an alternative computational approach, namely, the exact diagonalization (ED) algorithm. This is practical up to system sizes $N \approx 30$. In the inset of Fig. 3.2, the discrepancy between the PQMC and the ED predictions for the ground-state energy per spin is shown, considering 10 Hamiltonian realizations of size $L = 5$. We find vanishing values within the (small) statistical uncertainties, confirming the absence of bias in the PQMC predictions.

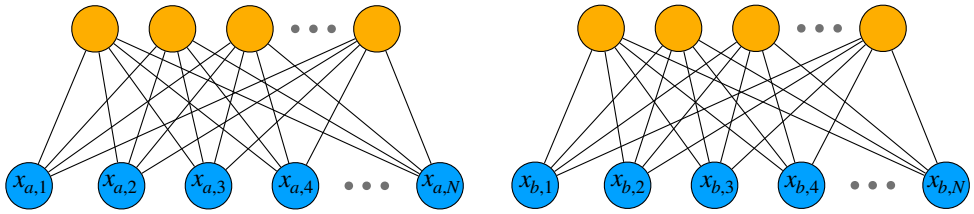

FIGURE 3.3: Representation of the connectivity of the RBM adopted to simulate the two-fold replicated EA Hamiltonian (3.7). The two sets of spins $x_{a,1}, \ldots, x_{a,N}$ and $x_{b,1}, \ldots, x_{b,N}$ (blue circles in the bottom) are connected to two corresponding sets of hidden variables (orange circles in the top). The two replicas feature the same random couplings J_{ij}.

3.2.3 Replica overlap in PQMC simulations

To discern the spin-glass phase from the paramagnetic phase, we analyze the replica spin overlap. To determine this quantity in a PQMC simulation, we simulate a two-fold replicated system formed by the union of two identical copies of the EA model defined in Eq. (3.1). This replication leads to the following Hamiltonian of $2N$ spins:

$$\hat{H}_2 = \sum_{\alpha=a,b} \left[-\sum_{\langle i,j \rangle} J_{ij} \sigma^z_{\alpha,i} \sigma^z_{\alpha,j} - \Gamma \sum_i \sigma^x_{\alpha,i} \right] ; \tag{3.7}$$

here, $\sigma^z_{\alpha,i}$ and $\sigma^x_{\alpha,i}$ indicate standard Pauli matrices acting on spin i of replica $\alpha = a, b$. It is worth emphasizing that the two replicas feature the same random couplings J_{ij} and transverse field Γ. Furthermore, the whole Hamiltonian is stochastically evolved in imaginary time using a unique random walker population. Due to the absence

of inter-replica interactions, the ground-state wave function of $\hat{H}_2$ must be separable: $\Psi_0(x_a, x_b) = \psi_0(x_a)\psi_0(x_b)$. A separable guiding wave function is conveniently implemented via an RBM featuring separable inter-layer connectivity, as shown in Fig. 4.6.

In the modified PQMC algorithm used to simulate replicated systems, we consider a double RBM, i.e., two independent copies of the same physical system, each with identical couplings. For this purpose, we define the RBM weight matrix as

$$W \in \mathbb{R}^{N_h \times N},$$

where N_h and N denote, respectively, the total number of hidden and visible spins across both replicas. For example, to simulate a single system with 25 spins, we set $N = 50$. The same holds for N_h. By construction, the matrix W is block diagonal:

$$W = \begin{pmatrix} W^{(a)} & 0 \\ 0 & W^{(b)} \end{pmatrix},$$

where $W^{(1)}$ and $W^{(2)}$ describe the connections within the two RBMs, while the off-diagonal blocks are null-matrices since the two replicas are not coupled to each other. The local connectivity of the physical spins is encoded in the adjacency matrix

$$R_{\text{near}} \in \mathbb{Z}^{N \times M_{\max}},$$

where $M_{\max}$ is the maximum number of nearest-neighbor connections of a spin in the lattice. For a two-dimensional system with nearest-neighbor interactions, we have $M_{\max} = 2$, neglecting double counting of symmetric connections. The explicit structure of R_{near}, which specifies the indices of the spins coupled to each visible spin, is described below

- For the first replica ($i = 0, 1, \ldots, N_{\text{half}} - 1$):

$$R_{\text{near}}[i, 0] = \begin{cases} i + 1, & (i+1) \bmod L \neq 0, \\ i + 1 - L, & (i+1) \bmod L = 0, \end{cases}$$

$$R_{\text{near}}[i, 1] = \begin{cases} i + L, & i + L < N_{\text{half}}, \\ i + L - N_{\text{half}}, & i + L \geq N_{\text{half}}, \end{cases}$$

where $N_{\text{half}} = N/2$ and L is the linear size of the system.

- For the second replica ($i = N_{\text{half}}, \ldots, N-1$), the same rules apply with shifted indices:

$$R_{\text{near}}[i,0] = \begin{cases} i+1, & (i+1) \bmod L \neq 0, \\ i+1-L, & (i+1) \bmod L = 0, \end{cases}$$

$$R_{\text{near}}[i,1] = \begin{cases} i+L, & i+L < N, \\ i+L-N_{\text{half}}, & i+L \geq N. \end{cases}$$

The spin-overlap operator is defined as:

$$\hat{q} = \frac{1}{N} \sum_{i=1}^{N} \sigma_{a,i}^{z} \sigma_{b,i}^{z}. \tag{3.8}$$

This operator is diagonal in the chosen computational basis. Thus, we are able to estimate via forward walking (see Sec. 1.5 of Chap. 1) the pure ground-state expectation value of the squared overlap $\langle q^2 \rangle$, namely, the mean-squared EA order parameter. This quantity allows discerning the spin glass from the paramagnetic phase. The disorder average, obtained as the average over a large ensemble of N_r realizations of the random couplings J_{ij}, will be indicated with square parenthesis:

$$\left[\left\langle q^2 \right\rangle \right] = \frac{\chi_{\text{sg}}}{N}, \tag{3.9}$$

where χ_{sg} indicates the spin-glass susceptibility. It is important to inspect whether the RBM wave function allows us to suppress the population control bias. In Fig. 3.4, we show $\langle q^2 \rangle$ for several representative instances of the Hamiltonian (3.1) for $L = 10$, obtained with different numbers N_h of hidden neurons. The inset displays the average over an ensemble of 256 disorder realizations. One notices that the results for $N_h \geq 32$ are well within the statistical uncertainties, both for the ensemble average and, more stringently, for individual realizations. This indicates that possible systematic biases are under control.

In the study of the quantum critical point, it is convenient to consider the Binder ratio [53, 55]:

$$R = \left[\frac{\langle q^4 \rangle}{\langle q^2 \rangle^2} \right]. \tag{3.10}$$

Notice that the Monte Carlo estimate of R is possibly biased, since one evaluates a square and a ratio of expectation values approximated via Monte Carlo integration performed over a finite number of samples, say, N_{MC}. Still, as discussed in Ref. [33], this bias is expected to vanish as N_{MC}^{-1}, i.e., faster than the statistical uncertainty, which scales as $N_{\mathrm{MC}}^{-1/2}$. A reduced bias [33] is expected for the following alternative definition [74]:

$$R' = \frac{[\langle q^4 \rangle]}{[\langle q^2 \rangle]^2}. \tag{3.11}$$

As pointed out previously [34], we also find that the latter definition leads to consistent results as with the definition of Eq. (3.10), but with enlarged fluctuations due to disorder averaging. See also the discussion in Ref. [71]. Therefore, the results presented in the following Section have been computed with Eq. (3.10). Following the common convention, they are cast in the form of the Binder cumulant, defined as

$$B = (3 - R)/2. \tag{3.12}$$

To inspect the nature of the spin-glass phase, it is instrumental to determine the spin-overlap distribution, defined, for $q \in [-1, 1]$ and for a given choice of the random couplings, as:

$$P_J(q) = \langle \delta(q - \hat{q}) \rangle \,. \tag{3.13}$$

In practice, we determine a discrete histogram with N bins, corrisponding to the number of spins. The ensemble-average spin-overlap distribution $P(q) = [P_J(q)]$ is the mean over N_r realizations of the couplings.

3.3 Results: quantum spin glass transition

The 2D EA Hamiltonian with transverse field is expected to host a quantum phase transition from a paramagnetic to a spin-glass phase when the transverse field is reduced below a critical value. It is worth mentioning here that, in the classical EA

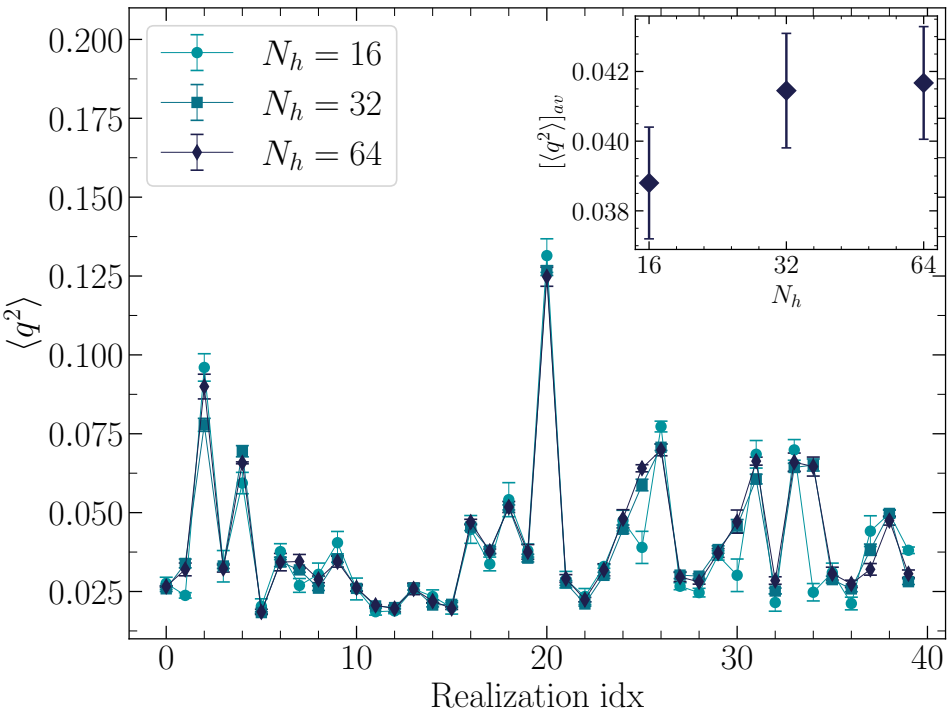

FIGURE 3.4: Mean-squared replica overlap $\langle q^2 \rangle$ for 40 instances of the 2D Gaussian EA model of size $L = 10$, at transverse field $\Gamma = 1.8$. The three datasets correspond to RBM guiding wave-functions with different numbers of hidden neurons N_h. The inset shows the average over 256 realizations as a function of N_h.

model, the spin-glass phase only occurs at zero temperature [75]. This is in contrast with the corresponding 3D model or the SK Hamiltonian, which host a classical spin-glass transition at a finite critical temperature. For the quantum SK model, the occurrence of replica symmetry breaking at sufficiently small transverse field was recently proven [76]. Hereafter, we analyze the quantum phase transition in the 2D quantum EA Hamiltonian with Gaussian random couplings. Very recent extreme-scale PIMC simulations investigated the analogous quantum phase transition considering binary random couplings [33, 27].

First, we analyze the finite-size scaling of the (realization averaged) Binder cumulant B of the mean-squared EA order parameter, defined in Eq. (3.12). This quantity is expected to be independent of the system size at the critical point, apart from corrections to the universal scaling ansatz. The latter is expected to hold when approaching the thermodynamic limit, but finite-size corrections might be sizable for small L. This

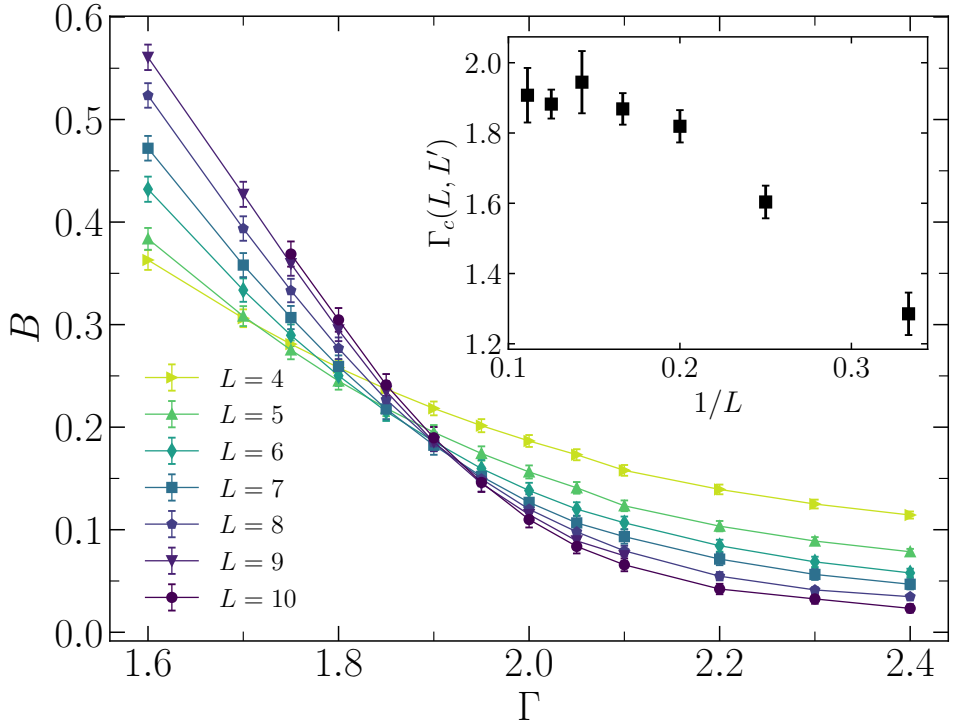

FIGURE 3.5: Binder cumulant B as a function of the transverse field Γ. Different curves correspond to different system sizes L. The inset shows the crossing point $\Gamma_c(L, L')$ of consecutive system sizes L and L', as a function of the (smaller) system size L.

behaviour is indeed what we observe; see results in Fig. 3.5. These data are averaged over a number of realizations of the random couplings ranging from $N_r = 256$ to $N_r = 1536$. Due to the finite-size effects, the crossing points of consecutive system sizes drift to larger transverse fields. However, for $L \geq 6$, the crossing points saturate within statistical uncertainties. In the large L regime, the crossings occur around the transverse field $\Gamma_c \approx 1.9$. This represents a first approximate estimate of the quantum critical point.

The Binder cumulant is expected to scale as $B = f_B\left((\Gamma - \Gamma_c)L^{1/\nu}\right)$, where ν is the critical exponent of the correlation length, Γ_c is the critical transverse field, and $f_B(x)$ is a universal scaling curve. Γ_c and ν are obtained from a best-fit analysis. In this analysis, we use the software provided in Ref. [77] to extract the best fit parameters. The collapse of the data corresponding to different system sizes is indeed verified in Fig. 3.6, considering the sizes $L \geq 6$.

To pinpoint the critical transverse field more precisely, we consider the (realization averaged) mean-squared EA order parameter $\left[\langle q^2 \rangle\right]$. Notice that this coincides with

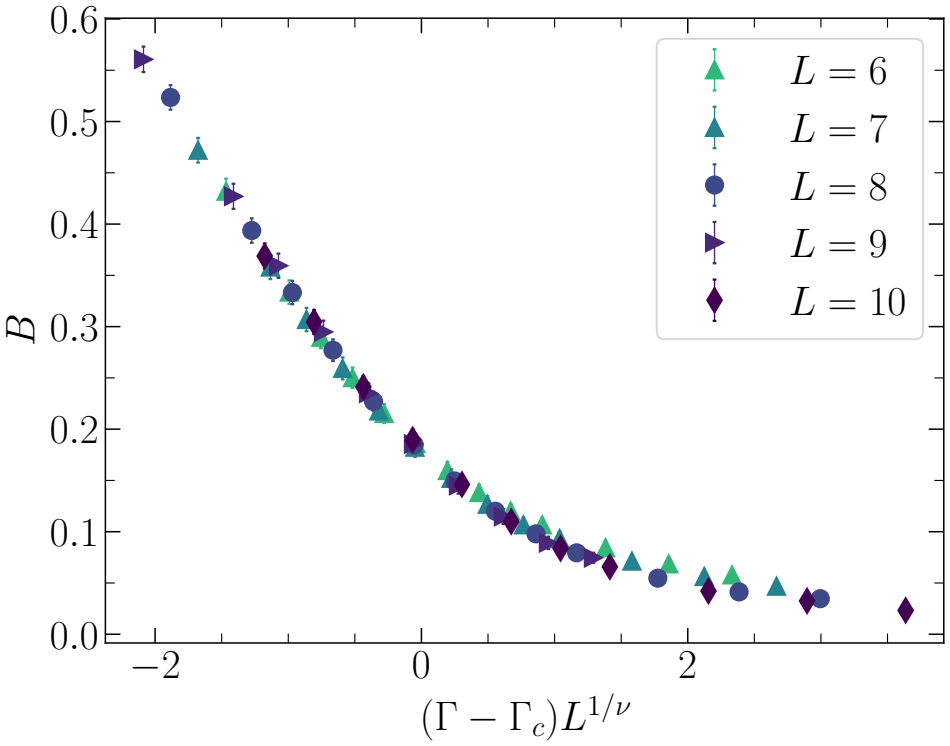

FIGURE 3.6: Collapse of the Binder cumulant B as a function of the rescaled transverse field $(\Gamma - \Gamma_c)L^{1/\nu}$. ν denotes the critical exponent of the correlation length. Different symbols correspond to different system sizes L.

$\chi_{\rm sg}/N$. Its finite-size scaling ansatz is

$$\left[\left\langle q^2 \right\rangle\right] L^{b_q} = f_q\left((\Gamma - \Gamma_c)L^{1/\nu}\right), \tag{3.14}$$

where b_q is the corresponding critical exponent. Accordingly, the rescaled quantity $\left[\left\langle q^2 \right\rangle\right] L^{b_q}$ should be system size independent at the critical point Γ_c. As shown in Fig. 3.7, data corresponding to different system sizes cross essentially at the same value of the transverse field, even for system sizes as small as $L \simeq 3$. This suggests that the corrections to the scaling ansatz are less pronounced here than in the case of the Binder cumulant B, allowing us to better estimate Γ_c. Still, to avoid finite size biases, in the data collapse we consider sizes $L \geq 6$. The best-fit analysis is performed using the software available from Ref. [77]. The data collapse is verified in Fig. 3.8. Importantly, the best-fit analysis provides estimates of ν, of b_q, and of the critical point Γ_c. For the latter, we find $\Gamma = 1.98(7)$, consistently with the findings based on B obtained in the large L regime (see inset of Fig. 3.5). The two critical exponents are expected to be

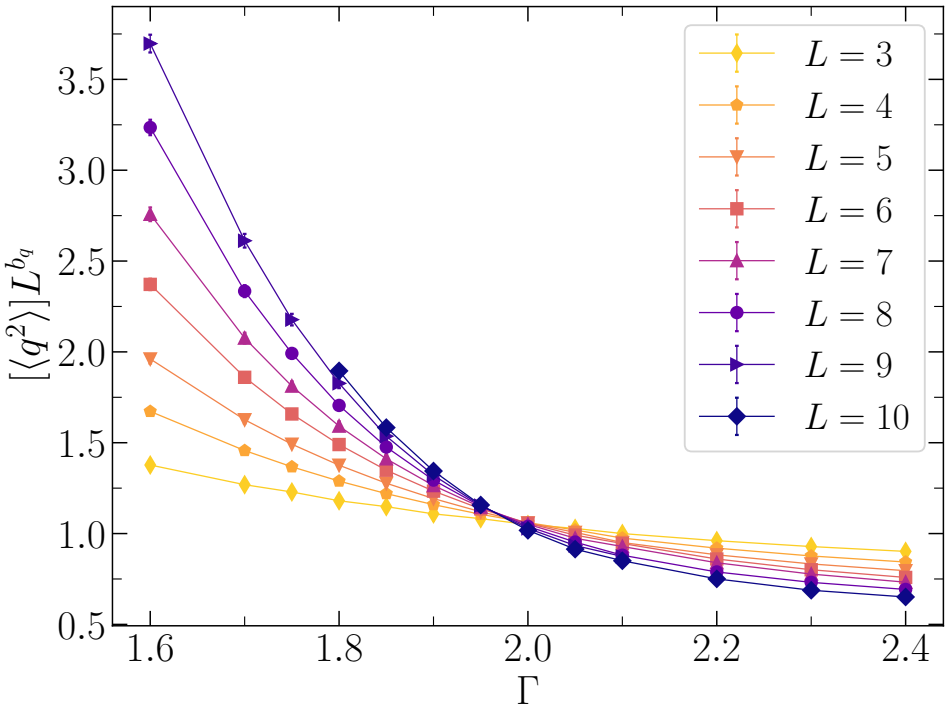

FIGURE 3.7: Rescaled (disordered averaged) EA order parameter $[\langle q^2 \rangle] L^{b_q}$ as a function of the transverse field Γ. b_q is the critical exponent associated with $\langle q^2 \rangle$. Different datasets correspond to different system sizes.

universal. In fact, we find good agreement with recent estimates, obtained via PIMC simulations in the zero-temperature limit, based on different choices of the random couplings. For the correlation-length critical exponent, we find $1/\nu = 1.11(22)$. Within the reported error bar, this results is consistent with the bound $\nu \geq 2/D$, where D is the dimensionality [78].

Our results, as well as the corresponding data from the recent literature, are summarized in Table 4.1. To determine the statistical uncertainties, we repeat the fitting process 55 times starting from randomized initial guesses. This allows computing the means of the standard deviations. Systematic errors due to scaling corrections are estimated considering the fluctuations obtained excluding the $L = 6$ size and then also the case $L = 7$. The error bars we report in parenthesis correspond to the larger between the statistical and the systematic uncertainty.

In order to facilitate comparison with other results on the same equivalence class systems, we also compute the effective exponent $1/\nu$, obtained from the finite-size scaling collapse between pairs of adjacent lattice sizes, keeping Γ_c and b_q fixed. We

compare the values of $1/\nu$ extracted from our accessible system sizes, $L = 3, \ldots, 10$, with those reported for $L = 8, 12, 16, 20, 24$ in Tab. 2 of Ref. [27]. As shown in the plot in Fig. 3.9. The value of the critical exponent $1/\nu$ saturates, within statistical uncertainty, for our largest simulated sizes.

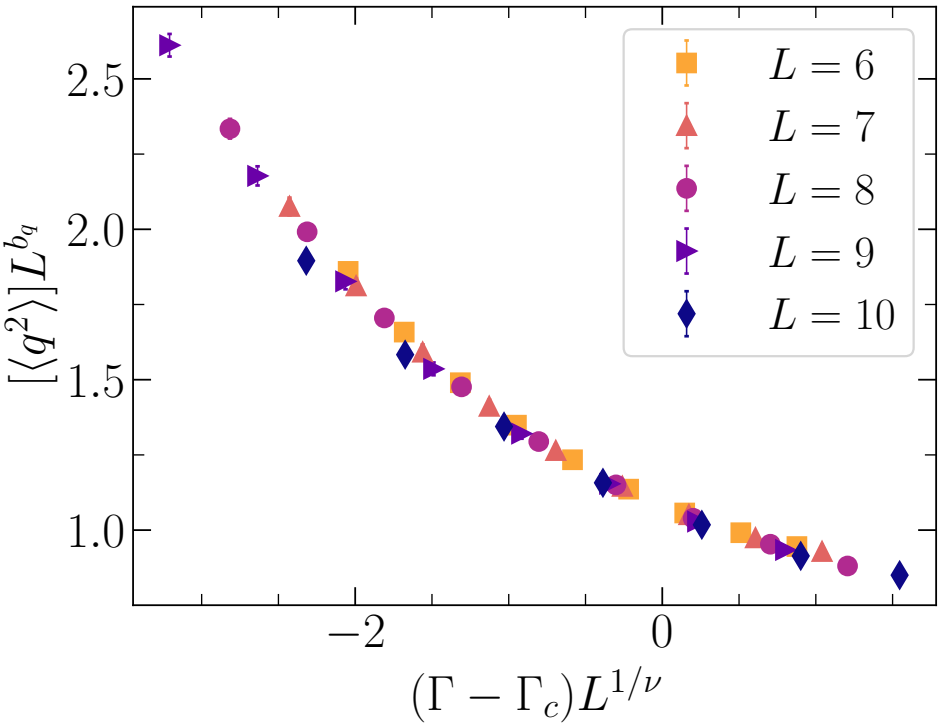

FIGURE 3.8: Collapse of the rescaled (disordered averaged) EA order parameter $[\langle q^2 \rangle] L^{b_q}$ as a function of the rescaled transverse field $(\Gamma - \Gamma_c) L^{1/\nu}$. b_q is the critical exponent of $\langle q^2 \rangle$.

Couplings	$1/\nu$	b_q	Γ_c	Ref.
50%. ± 1	1.02(16)	1.76(3)	2.11(1)	[33]
50% ± 1	0.71(24)(9)	1.73(8)(8)	2.18(1)	[27]
$\mathcal{N}(0,1)$	1.11(22)	1.68(8)	1.98(7)	This work

TABLE 3.1: Summary of the critical exponents of the correlation length ν and of the squared EA order parameter b_q, as well as of the critical transverse fields of the spin-glass transitions Γ_c, corresponding to different choices of the random couplings J_{ij} (reported in the first column). The last column reports the reference. When two errors are reported in parentheses, the first corresponds to the systematic error, the last to statistical uncertainty.

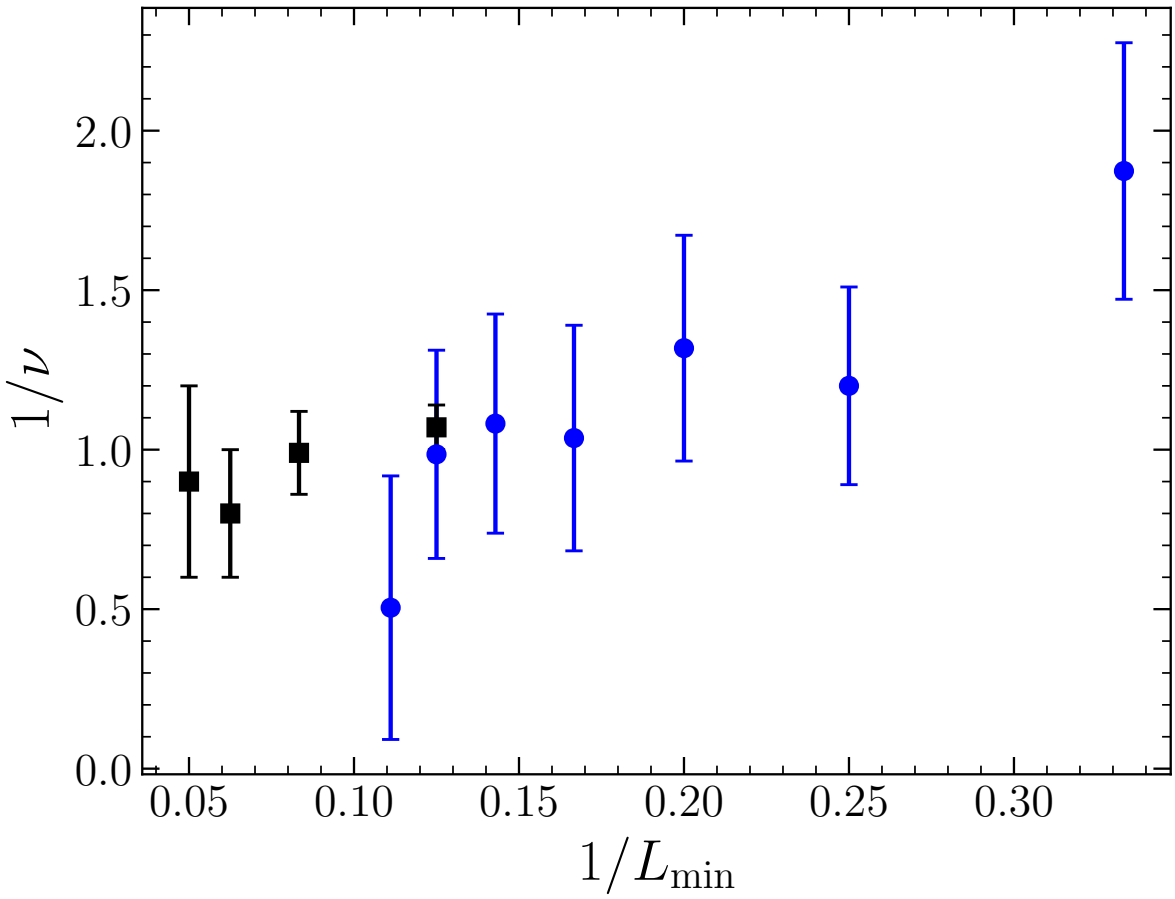

FIGURE 3.9: Critical exponent $1/\nu$ obtained from fits at fixed b_q and Γ_c, plotted as a function of the smallest lattice size $L_{\min}$ in each pair of sizes. Blue points correspond to fits of adjacent sizes considered in this work ($L = 3, \ldots, 10$), while black points represent the values from successive size reported in Tab. 2 of Ref. [27]. The critical exponent values $1/\nu$ saturates, within statistical error bars, for the largest system sizes analyzed in this work.

To inspect the nature of the spin-glass phase, we perform an exploratory analysis of the zero-temperature replica spin-overlap distribution $P(q)$ for $\Gamma < \Gamma_c$. Specifically, we set $\Gamma = 1.6$, and we consider the system sizes $L = 7, 8, 9$, and 10. The histograms, which are averaged over a number of realizations ranging from $N_r = 1186$ to $N_r = 1931$, are shown in Fig. 3.10. For relatively large L, the distributions display two sizable peaks, symmetrically located around $q \simeq \pm 0.25$. For the smallest size $L = 7$ the peaks are significantly less pronounced. At this small Γ, the ground state is well within the spin-glass phase. This implies that the Monte Carlo autocorrelation times increase compared to the regime $\Gamma \simeq \Gamma_c$ considered above. It is verified that for $L = 7$ essentially all disorder realizations satisfy the $\mathbb{Z}_2$ symmetry, corresponding to the invariance

$P_J(q) = P_J(-q)$. For the largest sizes, in particular for $L = 10$, many simulations randomly break the symmetry for the simulations times considered in this analysis. It is verified that, performing much longer simulations for representative disorder realizations, the symmetry is recovered. Also, when ensemble averaging is performed, the symmetry is fulfilled within statistical uncertainties. The latter are estimated from the fluctuations over a five-fold random splitting of the realization ensemble.

In the case of classical spin-glass models, special attention is devoted to the distribution weight at $q \simeq 0$. Two prominent theories of the spin-glass phase lead to different predictions. The droplet theory predicts vanishing values, namely $P(0) = 0$, in the thermodynamic limit, while in the framework of replica symmetry breaking a finite value is expected (see, e.g., Refs. [66, 68]). In the quantum model addressed here, large values are found for the workable system sizes, with a slow decrease as a function of L. Finite values $P(0)$ were also found in finite temperature PIMC simulations of the quantum SK model at weak transverse field [52]. However, it is worth pointing out that larger system sizes are needed to ascertain the correct asymptotic behaviour valid for $L \to \infty$. In fact, Refs. [79, 80] argue that this is the case, in particular, when the overlap distribution is evaluated very close to the phase transition. Notice that this argument has been debated [81]. Performing PQMC simulations even deeper in the spin-glass phase is in principle possible, but it comes at the cost of further increased correlation times, which require longer simulation times for ergodic sampling and, thus, problematic computational costs.

3.4 Summary of the chapter

In this chapter, we have investigated the spin-glass quantum phase transition in the 2D EA Hamiltonian with transverse field. We focused on the case of Gaussian couplings, complementing recent investigations that addressed the case of binary random couplings [33, 27]. Notably, we have shown that guiding wave functions in the form of self-learned neural network states allow removing the population control bias in PQMC simulations of frustrated quantum spin models. To compute the replica spin overlap, which gives us access to the mean-squared EA order parameter $\left[\langle q^2 \rangle\right]$, a two-fold replicated Hamiltonian has been evolved in imaginary time. A finite-size scaling

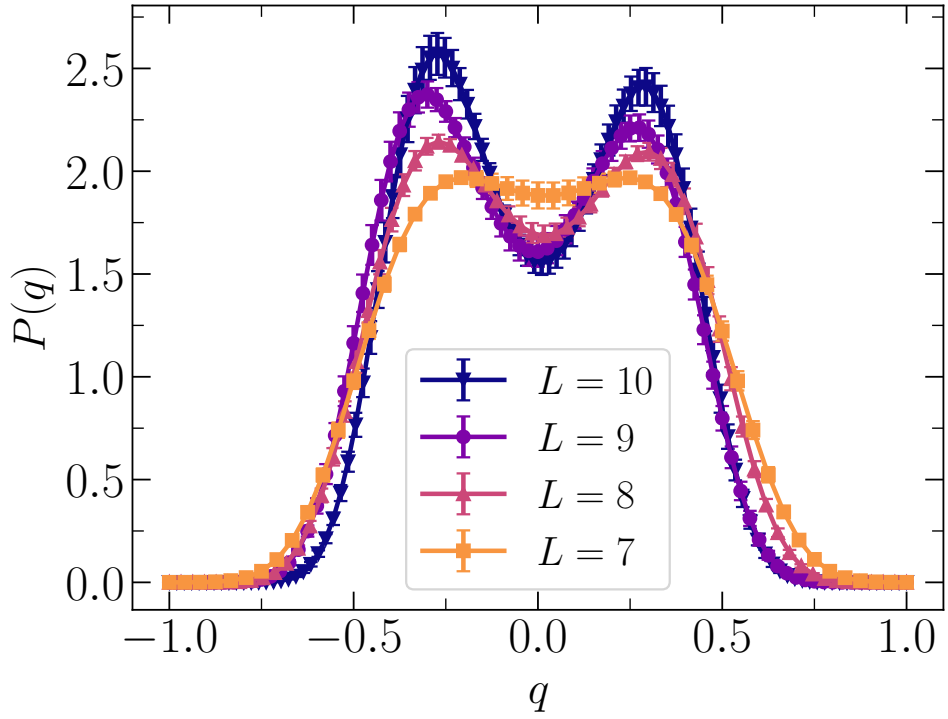

FIGURE 3.10: Disorder averaged replica spin overlap distribution $P(q)$. Different datasets correspond to different sizes L. The transverse field is $\Gamma = 1.6 < \Gamma_c$.

of $[\langle q^2 \rangle]$ allowed us to pinpoint the critical transverse field where the spin-glass quantum phase transition occurs, finding $\Gamma_c = 1.98(7)$. For the critical exponents, good agreement is found with the recent PIMC results for binary random couplings (see Table 4.1). We also performed an exploratory analysis of the overlap distribution $P(q)$ in zero-temperature quantum spin glasses. For the feasible system sizes $N \lesssim 100$ and for a transverse field approximately 20% lower than the critical value Γ_c, the function $P(q)$ displays a non-trivial shape. It features two (symmetry related) peaks, and substantial weight at $q \simeq 0$. The latter feature, when confirmed in the thermodynamic limit, is associated with replica symmetry breaking. Notably, to favour the development of novel computational methods for quantum Ising models, we provided at the repository of Ref. [69] a dataset of ground-state energies for an ensemble of disorder realizations. Specifically, the dataset includes couplings and energies of 50 Hamiltonian instances at $\Gamma = 1.8$ and system size $L = 10$. This dataset represent a useful benchmark also for quantum simulation platforms, e.g., for Rydberg-atom arrays.

The results presented in this chapter highlight the instrumental role of neural network states in simulating otherwise computationally overwhelming quantum many-body systems. Specifically, we have used them to implement an unbiased quantum Monte Carlo algorithm for the ground state of frustrated disordered quantum Ising models. This technique is complementary to, e.g., highly optimized PIMC algorithms [82] or stochastic series expansion methods [71]. Neural-network driven quantum Monte Carlo algorithms might also be useful to simulate the tunneling dynamics of quantum annealers [83, 29, 84, 28, 48, 7, 85]. Such simulations could help understanding how adiabatic quantum computers solve hard optimization problems. Future endeavours might focus on more experimentally relevant setups, e.g., on models for trapped-ion [86] or Rydberg-atom quantum computers [87, 88]. In the latter platform, frustrated lattices, such as the kagome geometries, have been implemented, and the possibility to observe glassy phases has been discussed [89, 90, 91, 92, 93]. In particular, positional disorder can be controlled essentially at will, allowing the implementation of, e.g., models for amorphous materials [94]. This paves the way for the investigation of the different localization properties of purely random noise, correlated disorder, and quasi-periodic patterns [95, 96, 97, 98, 99, 100].

Starting from the promising results shown in this chapter, in the next one we show the results obtained in a work conducted in collaboration with the PASQAL simulation group on spin-glass phase transition in modern Rydberg quantum platforms used as neutral atoms QPUs.

Chapter 4

Amorphous Rydberg-atom arrays and the emergence of spin-glass quantum phases

Experiments with neutral atoms trapped in optical tweezers and coherently coupled to Rydberg states allow quantum simulations of paradigmatic Hamiltonians for quantum magnetism. Previous studies have primarly investigated periodic tweezers arrangements, revealing a variety of spatially ordered magnetic phases. Here, we perform unbiased quantum Monte Carlo simulations of the ground state of quantum Ising models for amorphous arrays of Rydberg atoms. These models are constructed to exhibit well-controlled local structural properties while lacking long-range positional order. By computing the Edwards-Anderson order parameter, we provide evidence for a quantum phase transition from a paramagnetic to a spin-glass phase. The magnetic structure factor reveals short-range isotropic antiferromagnetic correlations. For accessible system sizes, the spin-overlap distribution display a nontrivial structure characterized two broad peaks and a sizable weight at zero overlap. The comparison against results for the clean kagome lattice, which features local structural properties similar to those of our amorphous arrays, highlights the important role of the absence of long-range structural order of the underlying array. Our findings indicate a route to experimentally implement the details of a Hamiltonian which hosts a quantum spin-glass phase. The results of this chapter have been published in Ref. [101].

4.1 Introduction to frustrated and amorphous solids

Frustrated random interactions in Ising models can lead to spin-glass phases, characterized by disordered magnetism, memory effects, and potentially the breakdown of ergodicity[30, 49]. These models were originally introduced to describe dilute magnetic alloys and play a central role in the study of combinatorial optimization problems [102]. In their quantum formulation, which is the focus of this chapter, spin-glass modes are also investigated as benchmark to assess the potential quantum advantage of quantum annealers [83, 50, 26]. Despite decades of research, several fundamental questions remain open, e.g., where replica symmetry breaking (RSB) occurs [103, 104, 105, 106, 107, 34, 76] or whether the spin-glass transition survives longitudinal fields [57, 108, 58, 60]. Another open problem is the relation between spin-glass and many-body localized phases [109, 110, 56, 111].

Quantum simulations of paradigmatic quantum Ising Hamiltonians are currently realized across a variety of experimental platforms, including ultracold atomic gases [22, 21], trapped ions [18], superconducting qubits [19], or neutral atoms precisely positioned using optical tweezers and then coupled to the Rydberg state [112, 20, 23]. For the latter platform, numerous magnetically ordered phases have been theoretically predicted [113, 114, 115, 116, 117, 118, 119, 120] and in some cases experimentally observed [121, 88, 122]. Most studies have focused on periodic tweezer arrays, including one-dimensional and two-dimensional geometries such as square, triangular, and kagome lattices.

Incommensurate phases that separate ordered and paramagnetic phases have been studied in single- and double-leg chains [123, 124, 125, 126, 127]. Interestingly, memory effects associated with glassy behavior have been predicted to occur in a periodic array at the boundary of differently ordered phases [90]. Yet, this finding has been questioned [128].

Frustration in these systems can arise either from the lattice geometry or from random mixtures of ferromagnetic and antiferromagnetic interactions. In Rydberg-atom platforms, the interatomic couplings are intrinsically antiferromagnetic, yet the atomic layout can be engineered almost arbitrary [129, 130, 131]. This flexibility has been exploited in experiments addressing combinatorial optimization [129, 132, 133]

and machine-learning problems [134, 135, 136]. In fact, by placing the optical tweezers in disordered configurations, one can effectively randomize the interaction couplings[98]. As recently shown [94], this feature allows one to create amorphous solids, namely configurations featuring well-defined short-range structural properties in the absence of long-range positional or orientational orders. However, the phase diagram of quantum Ising models with positional disorder remains largely unexplored, and it is still unclear whether purely antiferromagnetic interactions can sustain quantum spin-glass phases.

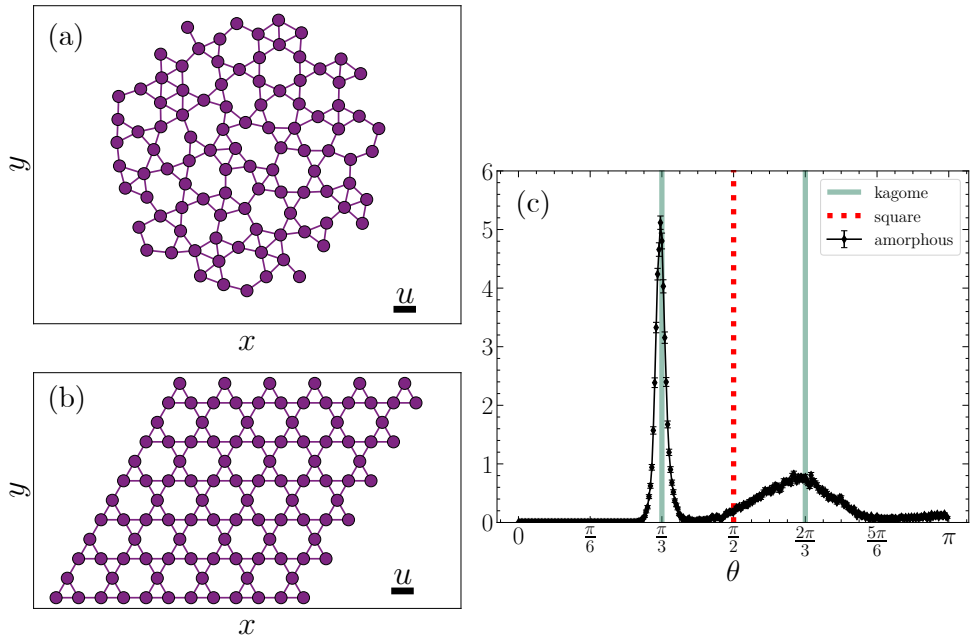

FIGURE 4.1: Panels (a) and (b): representation of a 2D amorphous array (a) and of the kagome lattice (b). Segments connect nearest-neighbor atoms. Panel (c): probability distribution of bond-angle θ in the amorphous arrays. The vertical lines denote the angles of the kagome lattice (continuous lines) and of the square lattice (dashed line), with periodic boundary conditions.

In this work, we perform unbiased quantum Monte Carlo (QMC) simulations to study the ground-state properties of quantum Ising models defined on amorphous arrays. The arrays are constructed to reproduce, on average, the local structural features of the periodic kagome lattice, such as the coordination number and preferred bond angles (see Fig. 4.1). The interactions are purely antiferromagnetic and decay with the sixth power of the interatomic distance, as realized in Rydberg-atom experiments. This setup combines the geometric frustration characteristic of the kagome lattice with

positional disorder, thereby providing a controlled framework to explore the interplay between the two effects. We determine the static magnetic structure factor and the replica-overlap Edwards-Anderson (EA) order parameter. The latter is suitable for identifying spin-glass phases, and we indeed find clear indications of a quantum phase transition from a paramagnetic to a spin-glass phase. The critical transverse field is determined via a finite-size scaling analysis. The magnetic structure factor indicates that, in the spin-glass regime, the antiferromagnetic correlations are short-ranged and isotropic, denoting the absence of spatial ordering. We also analyze the spin-overlap distribution for the feasible sizes. To highlight the role of the amorphous structure, we make a comparison with results corresponding to the periodic kagome lattice.

4.2 Ising-like Rydberg Hamiltonian

Specifically, we simulate a 2D transverse field Ising Hamiltonian, which reads

$$H = \sum_{i<j} J_{ij}\sigma_i^z\sigma_j^z - \Gamma\sum_{i=1}^{N}\sigma_i^x. \tag{4.1}$$

Here, σ_i^α, for $\alpha = x, z$, are usual Pauli operators acting on the spin $i = 1, \ldots, N$ at position $\mathbf{r}_i$. Γ is the uniform transverse magnetic field, and the couplings J_{ij} follow a power-low decay with respect to the distance: $J_{ij} = J_0/|\mathbf{r}_i - \mathbf{r}_j|^6$. J_0 is set as the energy unit. These parameters describe the Ising model that can be implemented in neutral atom quantum simulators. Notice that, to fulfill the Z_2 symmetry of the above Hamiltonian, local addressing of the detuning is required [137, 138, 139, 140].

The ground state of the Hamiltonian (4.1) is simulated using the continuous-time projection QMC algorithm detailed in Refs. [31, 36, 62] and in Sec. 2.4 of this thesis. It is verified that the possible systematic bias due to the population control [48, 9, 10] is below statistical uncertainties and the pure estimators are obtained via the standard forward-walking technique [42]. The guiding wave function is a neural network state in the form of a restricted Boltzmann machine [11], which is trained using the self-learning protocol of Ref. [16]. A relevant benefit of this QMC method is the absence of imaginary time discretization errors.

On the other hand, finite-temperature path-integral Monte Carlo simulations with discrete imaginary time allow one to reach larger system sizes than those considered

below [27, 82, 33]. In general, QMC algorithms allow simulating power-law interactions without the truncations usually implemented in relevant alternative computational methods based on tensor-network methods [141]. A viable QMC algorithm might also be the stochastic series expansion method [71, 142]. For the amorphous arrays, open boundary simulations are performed, while for the kagome lattice we implement periodic boundary conditions and account also for the first periodic images beyond the fundamental simulation cell, given that the couplings with further-apart images are negligible.

4.3 Amorphous solids

The amorphous configurations $(\mathbf{r}_1, \ldots, \mathbf{r}_N)$ are created starting from the algorithm described in Ref. [94]. An exemplary array is visualized in Fig. 4.1, panel (a). The full procedure we use to generate the amorphous solid is detailed in Ref. [94]; here we give a brief overview of the finer details relevant to the results reported in this chapter. In particular, when generating the amorphous array, we target the coordination number $C = 4$. C can be understood as the average number of nearest-neighbor (NN) atoms or, more precisely, those closer than the first minimum of the pair correlation function from a reference atom. Furthermore, in this work, we tune the hyper-parameters of the loss function given in Ref. [94] to favor amorphous structures composed of kagome-type plaquettes rather than square-type ones, both of which have $C = 4$. The core mechanism lies in the flexibility to control the NN separation, and the penalty for next-nearest neighbor (NNN) separations and beyond, both of which are governed by the choice of kernel function used in the design process. The NN separation is determined by the peak of the kernel; note, a sharp peak enforces uniform NN distances and hence low disorder, while a broader, flatter peak allows for greater disorder. Additionally, the rate at which the kernel decays away from the peak sets the penalty in the loss for atoms that are not NN. By tuning this decay rate, the system can be biased toward lattice types characterized by specific NN to NNN ratios; for example, $\sqrt{2}$ for square lattices or $\sqrt{3}$ for kagome lattices. To shows that the amorphous configurations under investigation are kagome-like, we compute the bond-angle distribution, which displays sizable broad peaks at the angles corresponding to the kagome lattice, namely, $\theta = \pi/3$ and $\theta = 2\pi/3$, and essentially no weight at the angle $\theta = \pi/2$ occurring in

the square lattice (see panel (c) of Fig 4.1). The average nearest-neighbor distance is $1.062(2)u$, where u is the length unit shown in Fig. 4.1.

4.3.1 Radial distribution function for amorphous systems

Defining nearest neighbors in an amorphous solid is more challenging than in a crystalline lattice, where the separation of the n-th order neighbors is fixed by geometry. To generalize the concept of n-th order NN we follow the formulation of Ref. [94], which provides a detailed discussion on how to characterize an amorphous system.

Let us introduce the radial distribution function $g(r)$, which characterizes how the atomic density varies as a function of the distance r from a reference atom. For a material with atomic density ρ, the radial distribution function is defined as

$$g(r) = \frac{1}{\rho}\Big\langle \sum_i \delta(r - r_i) \Big\rangle, \tag{4.2}$$

where r_i is the distance between the reference atom and the i-th atom, and the sum excludes the reference atom itself (in practice, a representative subset of atoms is usually considered). From $g(r)$, the nearest-neighbor coordination number can be computed as

$$n_1 = 2\pi \int_0^{R_1} r g(r) dr, \tag{4.3}$$

where R_1 corresponds to the position of the first minimum of $g(r)$. More generally, two atoms are defined as n-th neighbors if their separation r satisfies $R_{i-1} < r < R_i$, with $R_0 = 0$. Unlike crystalline systems, where neighbors of arbitrarily high order can be well defined, in amorphous solids the definition is practically restricted to the first few coordination shells. In crystalline lattices, the well-ordered structure gives rise to sharp peaks of $g(r)$ at fixed distances. Introducing positional disorder, as in a disordered lattice, broadens these peaks while retaining their overall structure. In Fig. 4.2 the radial distribution function of an amorphous solid is shown. The short-range order is still reflected by well-defined peaks of $g(r)$ at small r, whereas at large distances $g(r)$ gradually approaches unity, consistent with the absence of long-range order and the emergence of complete randomness.

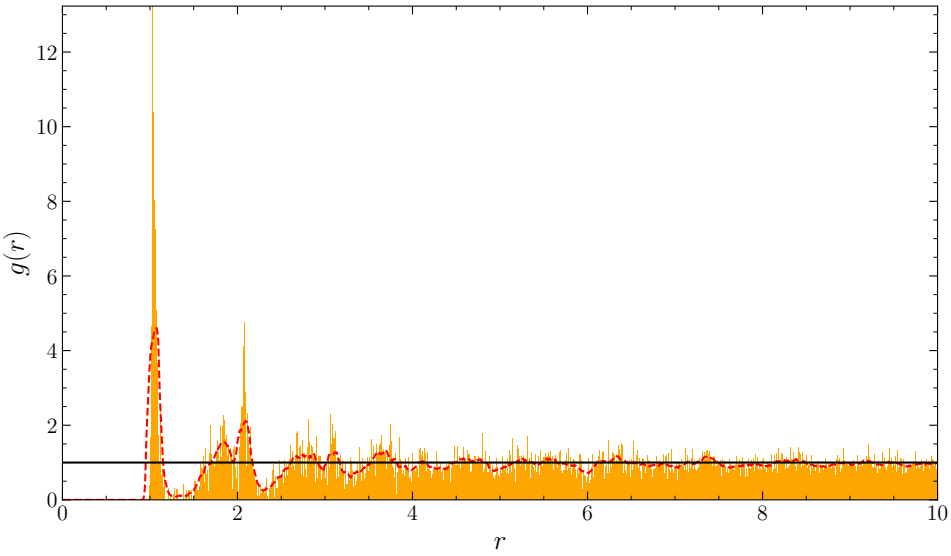

FIGURE 4.2: Radial distribution function $g(r)$ for an amorphous solid of size $N = 100$ simulated in this work. The solid is considered to be in the bulk of a bigger amorphous array. Notice that $g(r) \to 1$ (black solid line) for large values of r, indicating the lack of long-range positional order. Instead, for small value of r, the two sharp peaks of $g(r)$ are the clear sign of the local positional order. The red dashed line is the moving average of the $g(r)$ signal.

4.3.2 Geometrical structure factor

To further characterize the difference between an amorphous solid and an ordered lattice, we consider the static structure factor $S_g(\mathbf{k})$, defined as the Fourier transform of the atomic coordinates:

$$S_g(\mathbf{k}) = \frac{1}{N} \sum_{j=1}^{N} \sum_{k=1}^{N} e^{-i\mathbf{k}\cdot(\mathbf{R}_j - \mathbf{R}_k)}, \tag{4.4}$$

where $\mathbf{R}_j - \mathbf{R}_k$ is the relative distance between the j-th and k-th atoms. This quantity provides a direct probe of the spatial organization of the system, allowing one to distinguish between ordered and disordered structures.

In Fig. 4.3 we report the static geometric structure factor for both the clean kagome lattice and the kagome-like amorphous system. In the amorphous case, panel (a), the structure factor is averaged over 80 independent realizations. The absence of sharp Bragg peaks and the emergence of continuous circular features at fixed $|\mathbf{k}|$ clearly signal the presence of long-range disorder, consistent with the lack of translational invariance. Conversely, in the clean kagome lattice, panel (b), the reciprocal-space structure

is dominated by sharp peaks located at the reciprocal lattice vectors, a hallmark of long-range periodic order.

This comparison highlights the capability of $S_g(\mathbf{k})$ to discriminate between ordered crystalline systems and amorphous structures, and it will be used in the following as a diagnostic tool to analyze structural properties of the systems under study.

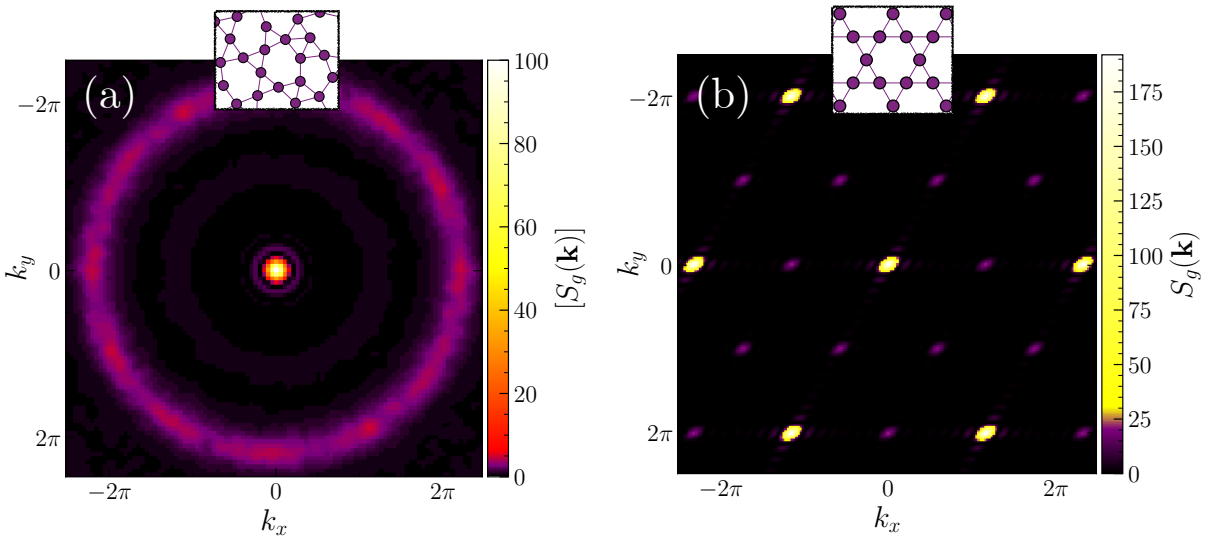

FIGURE 4.3: Static geometric structure factor $S_g(\mathbf{k})$ for the amorphous system and the clean kagome lattice with periodic boundary conditions (PBC). (a) The static structure factor averaged over 80 independent realizations of the amorphous system. The presence of continuous circular features at fixed $|\mathbf{k}|$ is a clear fingerprint of long-range disorder. (b) The static geometric structure factor of the ordered kagome lattice, where sharp Bragg peaks of the reciprocal lattice are visible, indicating the presence of long-range periodic order.

4.3.3 Patch extraction

As we seen in previous chapter, in the study of spin glasses, all results have to be averaged over copious ensembles of different realizations of the random couplings. To optimize this generation process, we first create a few large amorphous arrays featuring $N \approx 1100$ atoms. We then extract from each array several well-separated patches, each containing $N \in [40 : 100]$ atoms. To avoid overlap among different patches originating from the same array, we apply K-means clustering on the atomic coordinates. The cluster centroids serve as reference points for the patch extraction, where we select the N atoms closest to each centroid. For each large array, we perform K-means

clustering with five clusters. This number is chosen to ensure that each patch is independent, avoiding overlaps as well as boundary atoms, even for the largest patch size $N = 100$. A sketch of the extraction process is provided in Fig. 4.4. As discussed in the main text, the distribution of bond angles of the extracted patches exhibits broad peaks centered at the angles corresponding to the kagome lattice (with periodic boundary conditions), namely, $\theta = \pi/3$ and $\theta = 2\pi/3$. A minimal weight is also observed at $\theta \approx \pi$, which we attribute to boundary atoms. On the other hand, the square lattice, which also corresponds to $C = 4$, would lead to bond angles only at $\theta = \pi/2$. The actual coordination number, averaged over bulk atoms of large arrays of size $N = 1100$, is $C \simeq 4.01(3)$. Averaging over all atoms leads to a lower average $C \simeq 3.870(5)$ due to boundary effects. When computed on the extracted patches, boundary atoms contribute more to the average; if these are taken into account, the coordination number is $C \simeq 3.57(2)$ for $N = 100$.

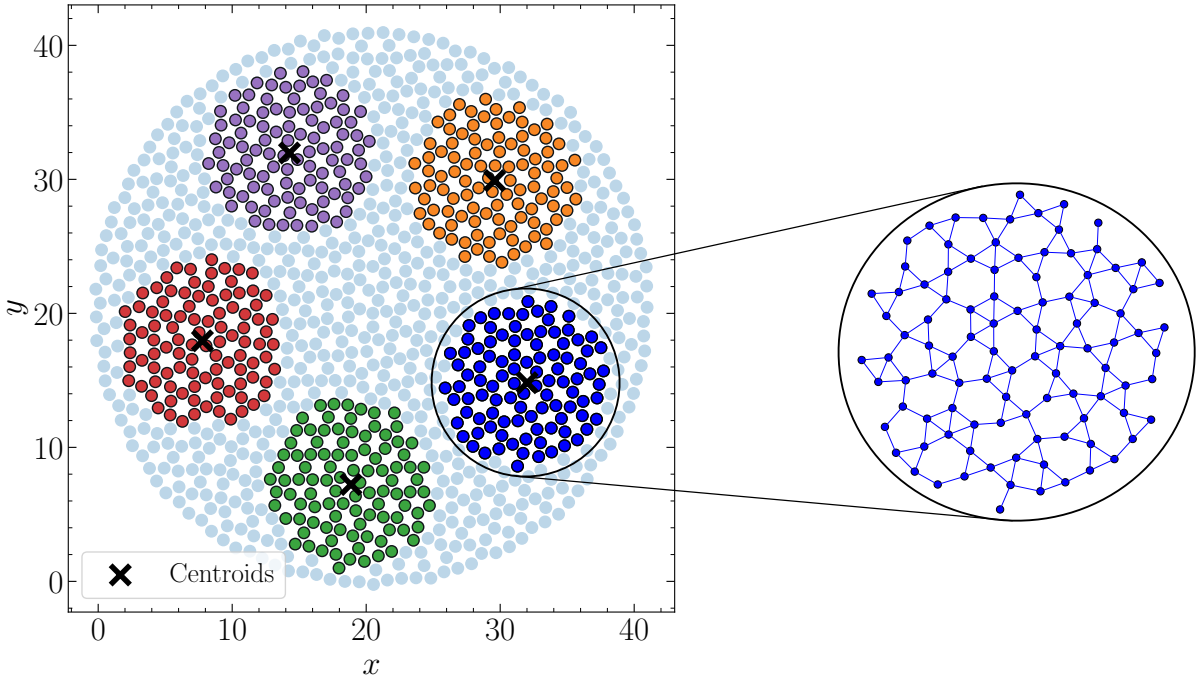

FIGURE 4.4: Example of patch extraction. Points in light blue represent the entire amorphous array created with the algorithm described in Ref. [94]. The black crosses are the centroids identified by K-means clustering performed on the coordinates. The colored clusters are the five patches of $N = 100$ atoms each. The zoom shows the detail of the blue patch.

To favor future studies, the amorphous arrays employed in this work are provided via the public repository at Ref. [143]. To highlight the role of the lack of long-range

structural order, below we also discuss the QMC results for the clean kagome lattice with periodic boundary conditions and nearest-neighbor distance a.

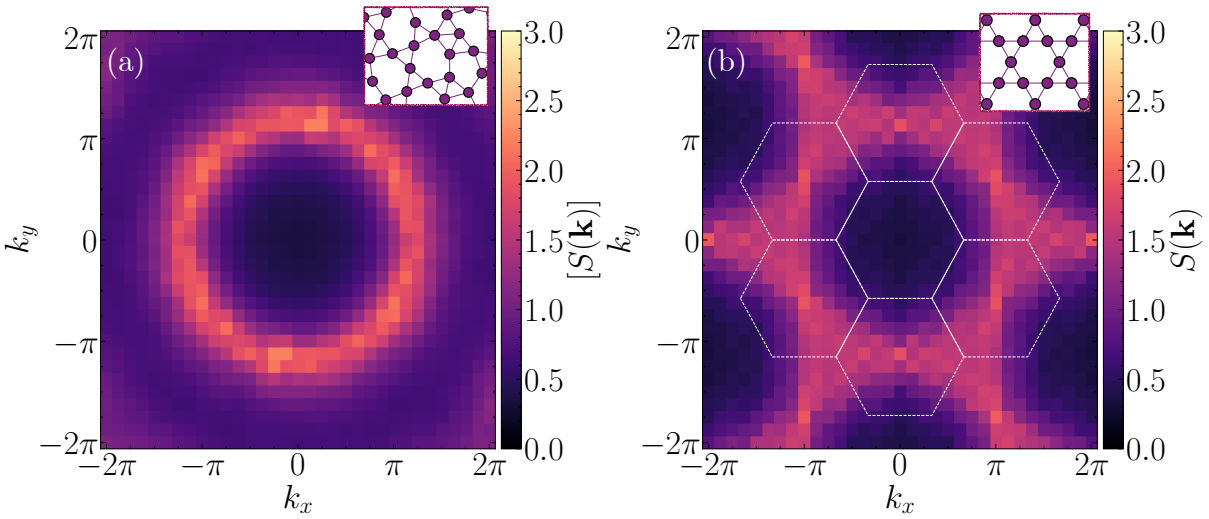

FIGURE 4.5: Panels (a) and (b): magnetic structure factor $S(\mathbf{k})$ versus k_x and k_y, with $\mathbf{k} = (k_x, k_y)$, for (ensemble-averaged) amorphous arrays (a) and for the periodic kagome lattice (b). In panel (b), the (white) hexagons represent the Brillouin zones.

4.4 Magnetic structure factor

The first observable we analyze is the static magnetic structure factor, defined as

$$S(\mathbf{k}) = \langle \mathcal{S}^z(\mathbf{k}) \mathcal{S}^z(-\mathbf{k}) \rangle \,, \tag{4.5}$$

where

$$\mathcal{S}^z(\mathbf{k}) = \frac{1}{\sqrt{N}} \sum_i \sigma_i^z \exp\left(i\mathbf{k} \cdot \mathbf{r}_i\right). \tag{4.6}$$

For the amorphous arrays, the ensemble-average, denoted with $[S(\mathbf{k})]$, is determined considering $N_r = 30$ realizations. The results at transverse field $\Gamma = 0.737$ display a broad circular hill, corresponding to isotropic magnetic correlations (see Fig. 4.5). The peak height does not scale with the system size, meaning that these correlations decay with distance. To show that, we compute the magnetic structure factor $S(\mathbf{k})$ for amorphous patches of sizes $N = [100, 80, 60]$. The ensemble-averaged structure factor, is shown in Fig. 4.7(a). It features very broad circular peaks, whose height does not significantly depend on the system size N. This indicates the lack of long-rage

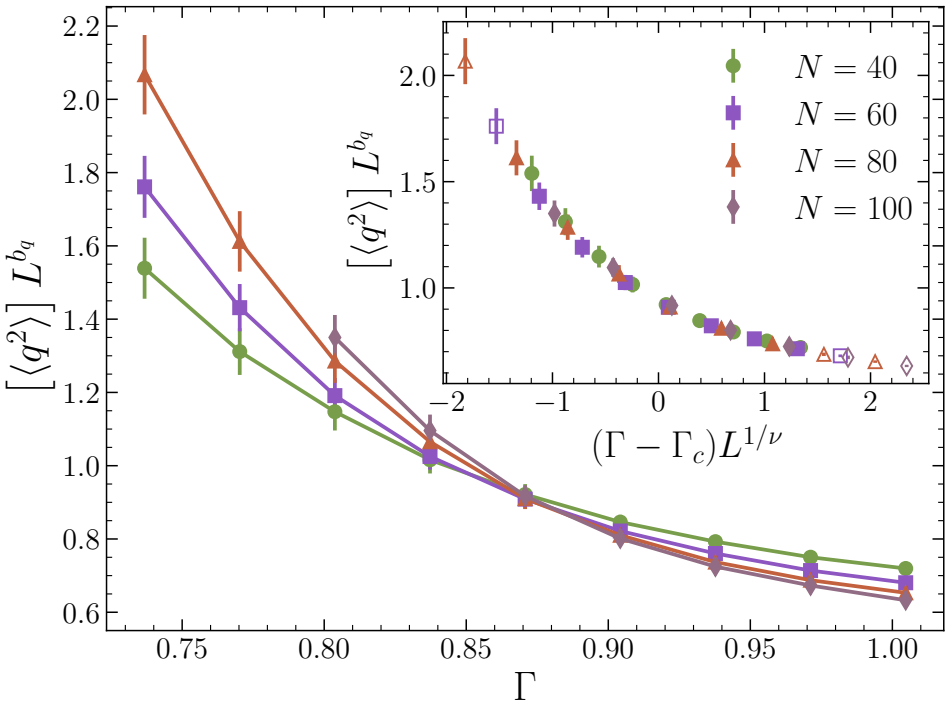

FIGURE 4.6: Main panel: rescaled (disordered averaged) EA order parameter $[\langle q^2 \rangle]\, L^{b_q}$ as a function of the transverse field Γ. Different datasets correspond to different system sizes. Inset: universal collapse of $[\langle q^2 \rangle]\, L^{b_q}$ plotted as a function of the transverse field $(\Gamma - \Gamma_c)\, L^{1/\nu}$. The critical parameters Γ_c, $1/\nu$, and b_q are obtained by optimizing the data collapse [77], excluding the data shown as empty symbols.

magnetic ordering. Additionally, we compare this result with the structure factor of the periodic kagome lattice, shown in Fig. 4.7(b). Also in this case, there are no sharp peaks, indicating short-range correlations. However, it is important to note that the pattern reveals the hexagonal structure of the Brillouin zones. The results shown in Fig. 4.7 are obtained for a transverse field in the spin-glass phase of the amorphous system, namely $\Gamma = 0.737$, for both the amorphous and periodic kagome systems.

Going back to Fig. 4.5, in the periodic kagome lattice, we again find a broad continuum, but this follows the hexagonal structure of the Brillouin zones. Notice that also in this case the correlations have a short-range character, pointing to a paramagnetic ground state, as found in previous studies addressing kagome Ising models with nearest-neighbor antiferromagnetic interactions [144, 145]. As discussed below, the system with $\Gamma = 0.737$ is in the spin-glass regime.

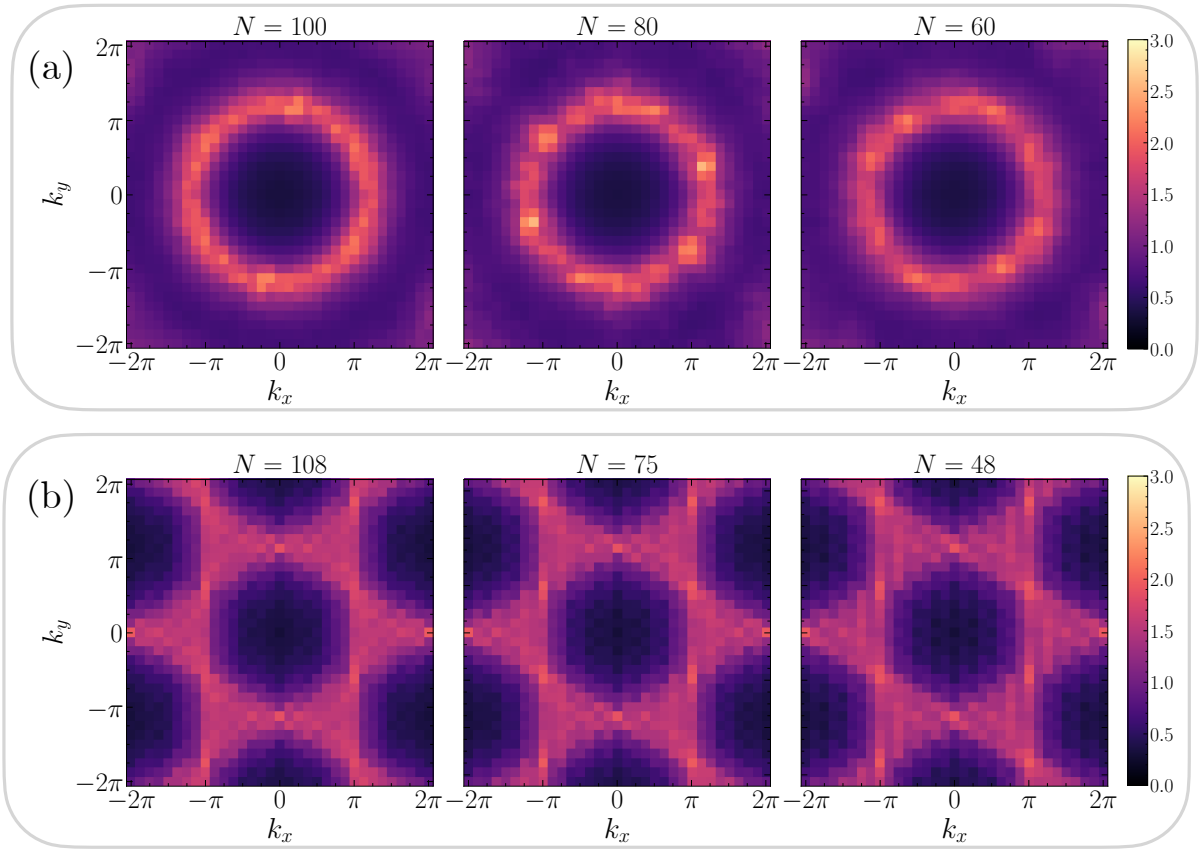

FIGURE 4.7: Panel (a): Magnetic structure factor $S(\mathbf{k})$ for the amorphous patches of system sizes $N = [100, 80, 60]$ (from left to right). The results are averaged over $N_r = 30$ realizations for $N = 100$ and $N_r = 10$ realizations for the other two sizes. Panel (b): Magnetic structure factor $S(\mathbf{k})$ for the clean kagome lattice for the system sizes $N = [108, 75, 48]$ (from left to right).

4.5 Spin-glass phase transition

Once again, to identify the spin-glass phase, we analyze the (mean squared) EA order parameter $\langle \hat{q}^2 \rangle$, which can be evaluated without systematic bias considering the overlap operator $\hat{q}$ between two identical replicas of the system [146]; $\hat{q}$ is defined as

$$\hat{q} = \frac{1}{N} \sum_{i=1}^{N} \sigma_{a,i}^z \sigma_{b,i}^z. \tag{4.7}$$

Here, the subscripts a and b denote the two replicas, which are simulated within the projection QMC algorithm by evolving a two-fold replicated Hamiltonian $H_T = H_a + H_b$ [147], where $H_{a/b}$ are defined as in Eq. (4.1) with Pauli matrices $\sigma_{a/b,i}^{\alpha}$. The ensemble

average $\left[\left\langle \hat{q}^2 \right\rangle\right]$ is performed over $N_r \in [100 : 200]$ amorphous arrays. Close to the critical point Γ_c, one assumes the following scaling Ansatz with a critical exponent b_q [33]:

$$\left[\left\langle \hat{q}^2 \right\rangle\right] L^{b_q} = f\left(x\right), \tag{4.8}$$

where $f(x)$ is a universal function of the reduced transverse field $x = (\Gamma - \Gamma_c)L^{1/\nu}$ and ν is the correlation-length critical exponent. The linear size is defined as $L = \sqrt{N}$. The critical parameters are obtained by optimizing the data collapse using the software available from Refs. [77, 148], which minimizes deviations from a self-consistently derived master curve [149]. Restricting the scaling analysis to $|x| \leq 1.34$, we obtain the following estimates: $\Gamma_c = 0.864(33)$, $b_q = 1.63(12)$, and $1/\nu = 1.22(45)$. Smaller fitting windows provide comparable estimates. Interestingly, universal critical parameters b_q and ν are consistent with recent estimates for the 2D EA Hamiltonian on the square lattice with binary or Gaussian random couplings [33, 27, 147], within the statistical errorbars. The comparison is detailed in Table 4.1. It is worth mentioning that the estimate of $1/\nu$ is expected to slightly decrease when very small sizes are excluded from the data collapse, due to corrections to the universal scaling ansatz [27]. This effect cannot be discerned here due to the limited system sizes. In fact, large-scale QMC simulations performed well within the spin-glass phase are computationally demanding due to the slow stochastic dynamics. Hence, we obtain a relatively large uncertainty on $1/\nu$, while the comparison of the parameter b_q is more stringent. As already noted [27], the values obtained for b_q denote a very slow divergence of the spin-glass susceptibility

$$\chi_{\mathrm{SG}} = \left[\left\langle \hat{q}^2 \right\rangle\right] L^2 \tag{4.9}$$

at the critical point.

4.6 Spin-overlap distribution

To shed some light on the nature of the spin-glass phase of the amorphous arrays, we analyze the distribution of the ensemble-averaged replica spin-overlap $P(q) = [P_J(q)]$ at $\Gamma = 0.67$; for a single realization, this distribution is computed as: $P_J(q) = \langle \delta(q - \hat{q}) \rangle$. Simulations of $P(q)$ are often used to assess the applicability of any of the most prominent theories of the spin-glass phase, in particular the droplet-picture or

Couplings	$1/\nu$	b_q	Γ_c	Ref.
50% ± 1	1.02(16)	1.76(3)	2.11(1)	[33]
50% ± 1	0.71(24)(9)	1.73(8)(8)	2.18(1)	[27]
$\mathcal{N}(0,1)$	1.11(22)	1.68(8)	1.98(7)	[147]
$\propto 1/r^6$	1.22(45)	1.63(12)	0.864(33)	This [101]

TABLE 4.1: Critical exponents ν and b_q and transition point Γ_c for different choices of random couplings J_{ij}, namely, binary random values (50% ± 1) and Gaussian values $\mathcal{N}(0,1)$ on the square lattice, and the positional disorder decaying as $\propto 1/r^6$ considered in this Article. Numbers in parentheses denote the errorbar. Where two numbers are given, the first denotes the statistical uncertainty, the second the systematic effects.

the RSB scenario [66, 68]. While the first theory predicts two δ-function peaks in the thermodynamic limit, the latter predicts a nontrivial distribution with sizable weight at $q = 0$. Our QMC results at $\Gamma = 0.67$ are shown in Fig. 4.8. For the feasible sizes, we find a distribution with two (symmetry related) broad peaks and sizable values of $P(q = 0)$. Yet, due to the limited system sizes and the statistical uncertainties, it is not possible to perform a reliable extrapolation of $P(q = 0)$ to $L \to \infty$ [79, 80, 81]. Furthermore, while for the smaller sizes essentially all realizations satisfy the expected symmetry $P_J(q) = P_J(-q)$, with the largest size $L = 80$ some realizations break the symmetry for the feasible simulation time, leading to a small asymmetry of $P(q)$. It is worth comparing the replica overlap distribution of amorphous arrays with the one corresponding to the periodic kagome lattice at the same transverse field. In the latter case, we find a narrow single-peak distribution that shrinks in the thermodynamic limit.

4.7 Scaling of the spin-glass susceptibility

Transitions to spin-glass phases can be signaled by a divergence of the spin-glass susceptibility χ_{SG} with the system size N. In Fig. 4.9, χ_{SG} is plotted as a function of the transverse field Γ, both for (ensemble-averaged) amorphous arrays and for the periodic kagome lattice, for different system sizes. The difference between the two configurations is noticeable. In amorphous systems, χ_{SG} clearly diverges for small Γ. On the other hand, the ordered kagome lattice exhibits an essentially flat signal in the range $\Gamma \in [0.74, 1.005]$, and no sizable dependence on the system size. This highlights the

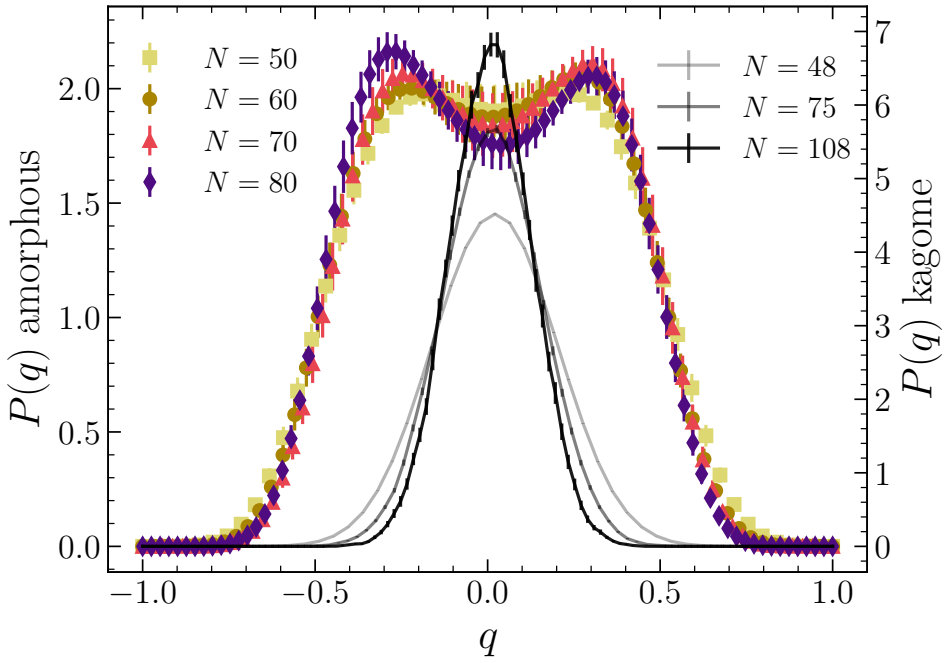

FIGURE 4.8: Distribution of the replica spin-overlap distribution $P(q)$ for the (ensemble averaged) amorphous arrays (left vertical axis) and for the kagome lattice (right vertical axis). Different datasets correspond to different sizes L. The transverse field is $\Gamma = 0.67 < \Gamma_c$.

important role of the absence of long-range structural order of the optical tweezers, compared with a periodic lattice with similar local structural properties.

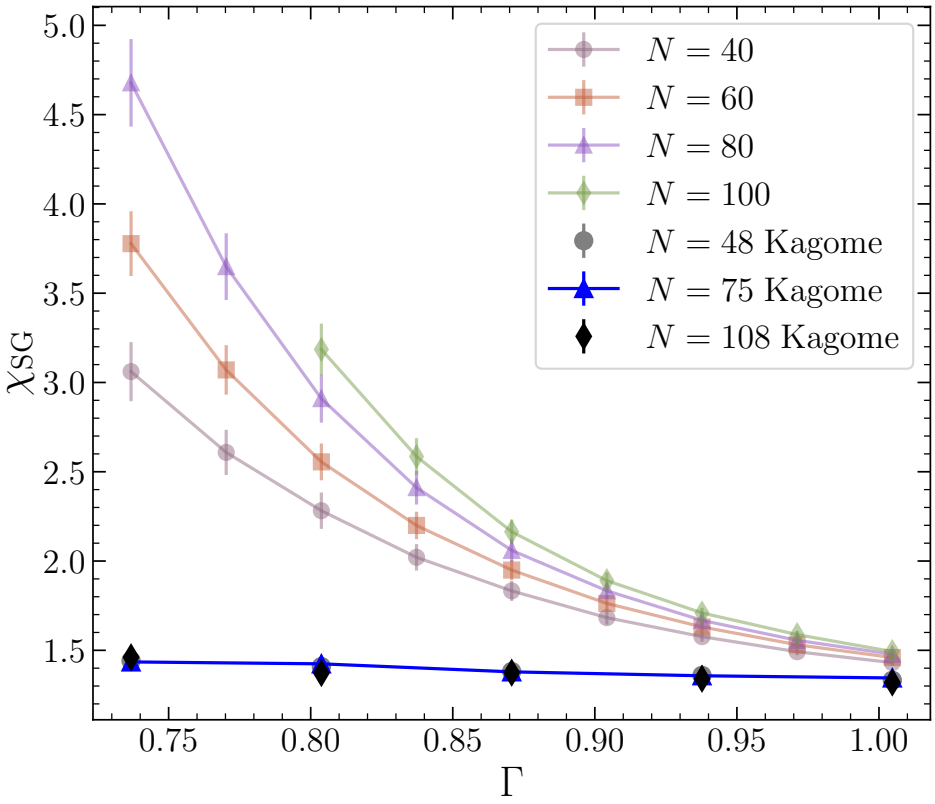

FIGURE 4.9: Spin glass susceptibility χ_{SG} for different system sizes N, for amorphous arrays (empty markers) and for the clean kagome lattice (filled markers)

4.8 Summary of the chapter

We have investigated the ground-state properties of quantum Ising models with positional disorder using unbiased neural QMC simulations. The Hamiltonian was designed to describe Rydberg atoms arranged in 2D amorphous arrays, whose local structural properties mimic those of kagome lattices. Notably, the finite-size scaling analysis of the EA order parameter revealed the occurrence of a quantum phase transition from a paramagnetic to a spin-glass phase. The critical exponents turned out to be consistent, within the statistical uncertainties, with those corresponding to the 2D EA model. By contrast, the results corresponding to the periodic kagome lattice display no hint of glassy behavior. This highlights the important role of the aperiodic structure and suggests that frustration effects in clean systems are not sufficient to induce

a spin-glass phase. Both the amorphous and the kagome setups host short-range antiferromagnetic correlations. However, these are isotropic in the former case, whereas they display a six-fold orientational symmetry in the latter. Notably, the ground state of the kagome lattice is paramagnetic in the explored parameter regime.

Our findings show that positionally distorted antiferromagnetic interactions can give rise to a spin-glass quantum phase transition, paving the way to the experimental observation of this phenomenon in Rydberg-atom platforms. As next steps, it would be interesting to explore if and how spin-glass behavior occurs in Hamiltonians without local detuning compensations and how the spin-glass phase could be prepared using quantum annealing protocols. Temperature effects are also worth studying. It is known that, in the case of power-law decaying random interactions on periodic lattices, the classical spin-glass phase transition might occur at finite or zero temperature depending on the power [150, 67, 151]; on the other hand, the case of finite-temperature amorphous arrays has not been studied. Anyway, the 2D quantum phase transition has strong experimentally observable effects in the critical dynamics [33]. Paradigmatic spin-glass models can also be implemented in quantum annealers built with superconducting-flux qubits [33], but these devices do not allow readout at finite transverse fields, except indirectly via anneal-pause protocols [145]. Trapped ions are promising alternative platforms. Many-body localization has already been observed in a chain of ten ions [86]. Yet, ion traps are still limited to one-dimensional geometries and fewer spins, although small 2D crystals have recently been implemented [152]. On the other hand, 2D arrays of several thousands of neutral atoms have already been realized [153, 154]. Future studies could address other types of positional disorder, other forms of correlated aperiodic couplings, e.g., quasi-crystals featuring long-range orientational order, randomly distorted frustrated lattices, or longer-range interactions relevant for trapped ions [155]. To favor future numerical and experimental investigations, the coordinates of the amorphous configurations employed in this work are made available at the repository in Ref. [143].

Chapter 5

Quantum phase transition in triangular lattices

As we have seen in the previous chapter, frustrated quantum magnets are paradigmatic systems in which the geometry of the lattice prevents the simultaneous satisfaction of the coupling conditions. In general, this leads to degenerate ground states and the emergence of exotic phase transitions. In addition to the specific case of amorphous solids we have already explored, the triangular lattice Ising antiferromagnetic is the canonical example of such frustration. At zero temperature, quantum fluctuations induced by the transverse magnetic field give rise to the phenomenon known as order by disorder. In contrast to ordered systems where thermal or quantum fluctuations typically increase degeneracy and promote disorder the presence of the transverse field in a frustrated magnet helps the system lift its ground-state degeneracy by selecting a preferred configuration. The resulting ordered state is referred to as the clock phase, since the spins align along one of six symmetry-related directions, resembling the hands of a clock. This ordered phase persists within a finite window of transverse field, beyond which it melts into a disordered paramagnetic phase. In this chapter, we review the state of the art in the study of triangular-lattice Ising models with nearest-neighbor and long-range interactions, and present Projection quantum Monte Carlo results on quantum phase transitions at zero temperature. We characterize the ordered phases through Binder cumulant analysis and magnetic structure factors, with particular focus on the fate of clock order in the presence of long-range couplings decaying as $1/r^6$. Our findings highlight the rich interplay between frustration, quantum fluctuations, and interaction range, and provide benchmarks for current and future quantum simulation platforms.

5.1 Order-by-disorder mechanism

Geometrical frustration in magnetism occurs when the topology of the lattice prevents all spin interactions from being simultaneously satisfied [156, 157]. The triangular lattice is the simplest two-dimensional geometry where this phenomenon appears. For Ising spin systems with antiferromagnetic couplings, each elementary triangle cannot satisfy all pairwise interactions, leading to the degeneracy of the classical ground state. The introduction of thermal or quantum fluctuations can partially break the degeneracy, inducing ordered states through the so-called *order-by-disorder* mechanism [158]. In the quantum case, a transverse field Γ induces spin flips that lower the energy of certain configurations, leading to three sublattice clock order. For this peculiarity, the ordered phase emerging from fluctuations is called *clock phase* [159]. In this phase, the three sublattices identifiable in the global system can develop a ferromagnetic ordering, characterized by the tuple $(1, -1, 0)$, or an antiferromagnetic ordering, described by the tuple $\left(-\frac{1}{2}, -\frac{1}{2}, 1\right)$. Each of these patterns admits six distinct states, generated by cyclic permutations and the rotational symmetries of the triangular lattice. This multiplicity reflects the competition between geometrical frustration and quantum fluctuations, making the clock phase a paradigmatic example of order emerging from a highly degenerate manifold. From a symmetry perspective, the ordered states correspond to the spontaneous breaking of a discrete Z_6 symmetry, associated with the sixfold degeneracy of the clock configurations. The system thus selects one among six equivalent orientations in order parameter space, a hallmark of the order-by-disorder mechanism in frustrated quantum magnets.

While the nearest-neighbor TLIAF has been extensively studied, recent experimental platforms now realize effective spin Hamiltonians with long-range interactions. In Rydberg atom arrays, spin-spin couplings follow a van der Waals decay $J_{ij} \sim 1/r_{ij}^6$ [87]. In trapped-ion quantum simulators, spin-spin interactions are tuned via optical dipole forces mediated by phonon modes. In particular, the interaction follows the power law $J_{ij} \sim 1/r_{ij}^{\alpha}$ with $0 < \alpha < 3$ [160]. The recent emergence of these experimental platforms has stimulated considerable interest in exploring quantum phase transitions of frustrated models as a function of the interaction decay exponent α. While several studies have primarily focused on the regime of highly long-range couplings [113, 161], in this chapter we present results for both the nearest-neighbor model and

for interactions of van der Waals type. Since the latter are particularly relevant as benchmarks for implementations with neutral atom platforms.

In this chapter, we present results for both $\alpha \to \infty$ and $\alpha = 6$ determining the critical transverse field Γ_c that signs the phase transition between the clock phase and the paramagnetic phase, as well as the correlation length critical exponent $1/\nu$.

5.2 Ising-like triangular lattice

The Hamiltonian under investigation is again the transverse field ising Hamiltonian

$$H = \sum_{i<j} J_{ij}\sigma_i^z\sigma_j^z - \Gamma \sum_i \sigma_i^x, \tag{5.1}$$

with $\sigma^{x,z}$ Pauli matrices, Γ the transverse field, and $J_{ij} = 1/|r_i - r_j|^\alpha$ antiferromagnetic couplings.

Geometrically, the triangular lattice can be divided into three sublattices.

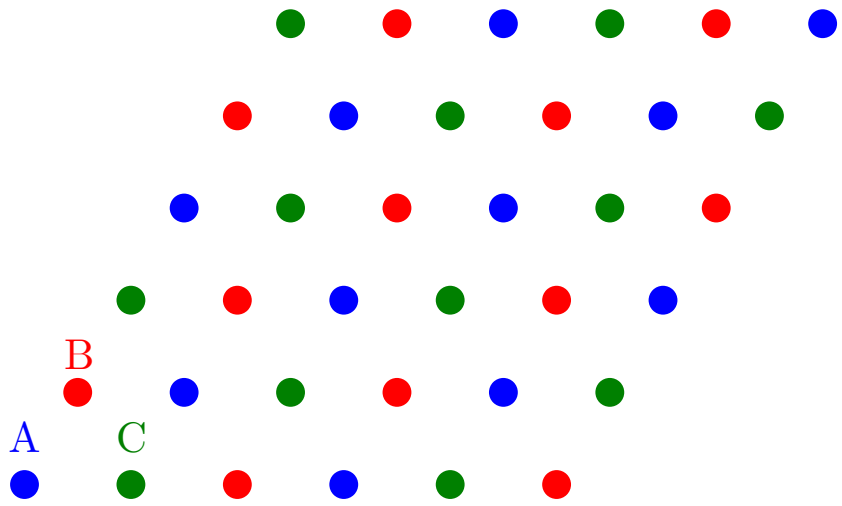

FIGURE 5.1: Triangular lattice structure. The blue (A), red (B) and the (C) sublattices are highlated.

Fig. 5.1 illustrates this decomposition, highlighting the A, B, and C sublattices. We consider periodic boundary conditions (PBC) with the method of the first images, as for the case of the amorphous Rydberg atoms array in Chap. 4.

5.3 Numerical observables

In order to characterize the phase transition due to the symmetry breaking in the triangular lattice Ising antiferromagnet under a transverse field, we introduce the XY complex order parameter that captures the structure of the three sublattices in the clock phase. Following the Landau-Ginzburg-Wilson construction (see Ref. [158]) one defines:

$$me^{i\theta} \equiv \frac{1}{C}(m_A + m_B e^{i4\pi/3} + m_C e^{-i4\pi/3}), \tag{5.2}$$

where m_A, m_B and m_C denote the sublattice magnetizations. The normalization factor C is chosen such that $|m| = 1$ in a perfectly ordered phase: in specifically, $C = \sqrt{3}$ for the $(1, -1, 0)$ pattern and $C = 3/2$ for the $(1, -\frac{1}{2}, -\frac{1}{2})$ pattern. This complex order parameter provides a natural description of the six degenerated clock states, with the phase θ identifying the preferred sublattice orientation. Its magnitude $|m|$ represents a sccalar measure of long-range order, while the angular distribution of θ signs the Z_6 symmetry breaking [113].

To investigate the critical properties and accurately locate phase transitions, we again rely on the Binder cumulant. Here, we write the generic form as reported in Ref. [113]. For an n-component order parameter (with $n = 2$ in our case, corresponding to the real and imaginary parts of $me^{i\theta}$), the appropriate definition is

$$U = \frac{n+2}{2}\left[1 - \frac{n}{n+2}\frac{\langle m^4 \rangle}{\langle m^2 \rangle^2}\right]. \tag{5.3}$$

This form ensures that U approaches a universal value at the critical point in the thermodynamic limit, while tending to zero in the disordered phase and to unity in the ordered state. As we have seen in Chap. 3. The scaling behavior of the Binder cumulant crossings across different system sizes provides a robust method for estimating the location and universality class of the quantum critical point.

Complementary information on the phase behavior can be extracted from the magnetic structure factor defined in Eq. 4.5. Let us recall it in an equivalent form:

$$S(\mathbf{q}) = \frac{1}{N}\sum_{i,j}\langle S_i^z S_j^z \rangle \, e^{i\mathbf{q}\cdot(\mathbf{r}_i - \mathbf{r}_j)}. \tag{5.4}$$

Here, N is the total number of spins, and the sum runs over all lattice sites i, j. Peaks in

$S(\mathbf{q})$ sign the presecte of long-range order at the corresponding wavevectors. In particular, as we have seen in Chap. 4, in the ordered phase, the heights of the peaks increase with the system size. In the case of the triangular lattice Ising antiferromagnet, the evidence of sharp peaks at the ordering wavevectors associated with the three-sublattice pattern provides clear evidence of the clock phase. Conversely, the suppression of these peaks at larger transverse fields indicates the melting of long-range order into the paramagnetic regime. Together, the complex order parameter, Binder cumulant, and structure factor form a good set of tools to characterize both the ordered and disordered phases, and to quantitatively determine the critical properties of the quantum phase transition.

5.4 The clock phase transition

In this section, we present results obtained using the self-learning projection quantum Monte Carlo method applied to this intrinsically frustrated systems. Our analysis focuses on two interaction regimes of particular relevance: the limit $\alpha \to \infty$, corresponding to the nearest-neighbor model, and the case $\alpha = 6$, representative of van der Waals interactions realized in neutral-atom quantum simulators.

For both models, we compute the Binder cumulant and perform finite-size scaling analyses, which enable us to accurately determine the critical point separating the clock-ordered phase from the paramagnetic phase, as well as the critical correlation-length exponent $1/\nu$. Furthermore, we investigate the structure factor $S(\mathbf{q})$ of the complex magnetization at different points in the phase diagram, illustrating how quantum fluctuations stabilize clock order at low transverse field strengths, while ultimately destroying it when fluctuations become too strong at larger fields.

5.4.1 Nearest-neighbour interaction

The results for the Binder cumulant B and the finite size scaling for the nearest-neighbor model is reported in Fig. (5.2). From its finite-size scaling behavior we obtain an estimate of the critical transverse field Γ_c as well as the critical exponent of the correlation length, $1/\nu$. Within the three system sizes that can be reliably accessed by our PQMC simulations, the extracted values are consistent with the results available in

the literature, $\Gamma_c = 1.65(5)$ and $1/\nu \approx 3/2$ (see, e.g., Ref. [113, 158]). The quantitative discrepancies in our preliminary estimates of both Γ_c and $1/\nu$ can plausibly be attributed to the restricted set of lattice sizes included in the scaling analysis, which may enhance finite-size corrections. With these limitations in mind, our best estimates are $\Gamma_c = 1.57(1)$ and $1/\nu = 1.44(7)$. Despite the fact that they are encouraging, these results should be regarded as preliminary, and further simulations at larger system sizes will be necessary to consolidate the scaling analysis and firmly establish the universality class.

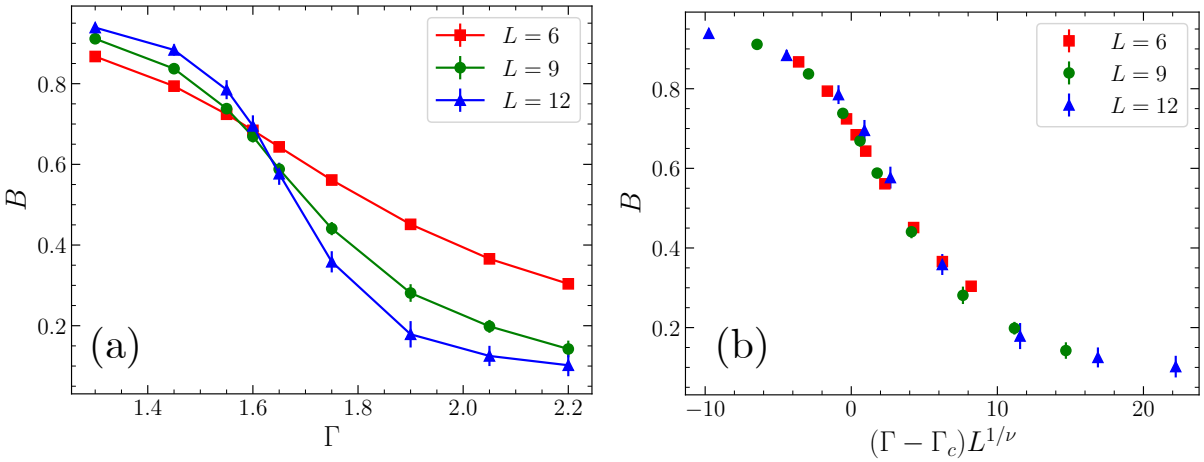

FIGURE 5.2: (a) Binder cumulant B for the nearest-neighbor triangular-lattice model as a function of the transverse field Γ. The curves corresponding to system sizes $L = 6, 9, 12$ exhibit a clear crossing, signaling the quantum phase transition between the clock-ordered phase and the paramagnetic phase. (b) Finite-size scaling collapse of the same data, where B is plotted against the reduced field $(\Gamma - \Gamma_c)L^{1/\nu}$. The optimal collapse is obtained for $\Gamma_c = 1.57(1)$ and $1/\nu = 1.44(7)$, providing estimates of the critical point and correlation-length exponent that are consistent with the universality class expected for this transition.

5.4.2 Rydberg-type interaction

For the model with power-law interactions characterized by $\alpha = 6$, we performed the same finite-size analysis as in the nearest-neighbor case. The Binder cumulant data, reported in Fig. (5.3), display a clear crossing that allows us to extract both the critical field and the correlation-length exponent. The estimates obtained are $\Gamma_c = 1.37(1)$ and $1/\nu = 1.26(8)$. Although subject to finite-size effects due to the restricted number of accessible system sizes, these values are consistent with the expected scaling scenario.

In particular, the observed shift of the critical field towards lower values compared to the nearest-neighbor limit is in agreement with theoretical expectations and with previous studies of long-range interacting frustrated models: as α decreases, longer-range couplings enhance quantum fluctuations and destabilize the ordered phase, thereby reducing Γ_c [113]. While further simulations on larger lattices will be necessary to firmly establish the critical behavior, our results confirm the qualitative trend reported in the literature and support the robustness of the order-by-disorder mechanism in the presence of van der Waals-type interactions relevant to Rydberg atom platforms.

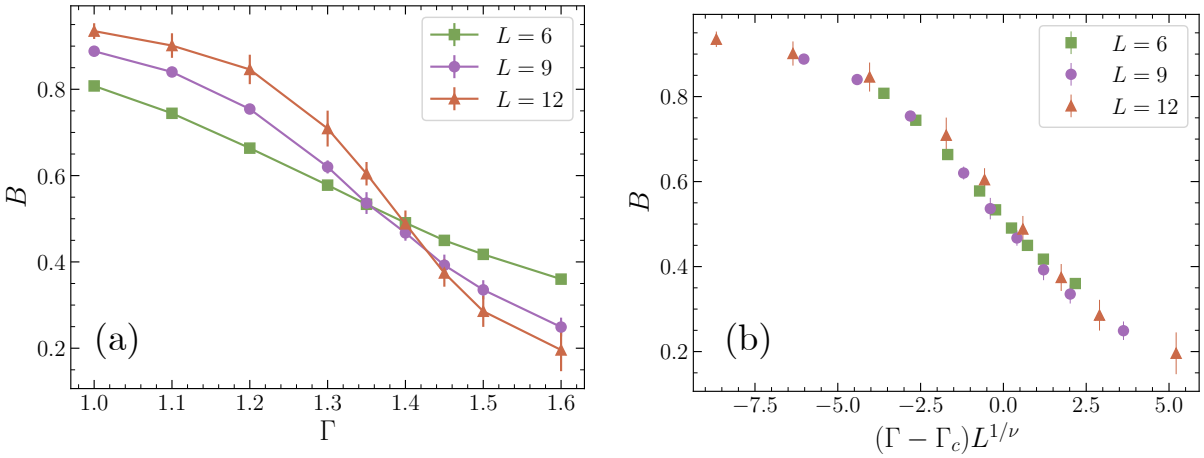

FIGURE 5.3: (a) Binder cumulant B for the triangular-lattice model with power-law interactions ($\alpha = 6$) as a function of the transverse field Γ. The curves for the accessible system sizes exhibit a crossing, indicating the quantum phase transition between the clock-ordered phase and the paramagnetic phase. (b) Finite-size scaling collapse of the Binder cumulant plotted against the reduced field $(\Gamma - \Gamma_c)L^{1/\nu}$. The best collapse is obtained for $\Gamma_c = 1.37(1)$ and $1/\nu = 1.26(8)$, confirming the downward shift of the critical field compared to the nearest-neighbor limit, consistent with theoretical expectations that longer-range interactions enhance quantum fluctuations and destabilize the ordered phase.

For this model, which is of particular interest due to its relevance to Rydberg atom platforms, we further analyze the magnetic structure factor $S(\mathbf{q})$ at three representative points in the phase diagram, namely $\Gamma = 1.0, 1.4, 1.6$, corresponding respectively to a field well inside the clock-ordered phase, near the critical point, and in the disordered paramagnetic phase. As discussed before, the magnetic structure factor is defined as the Fourier transform of the spin-spin correlations and provides a direct measure of the spatial ordering in the system. In particular, it allows one to identify the characteristic ordering wavevectors and the degree of long-range magnetic order. For the clock phase, which is stabilized by quantum fluctuations via the order-by-disorder

mechanism, $S(\mathbf{q})$ exhibits six sharp peaks corresponding to the three-sublattice pattern of the triangular lattice. These peaks serve as a clear fingerprint of the clock order and directly reflect the breaking of the discrete Z_6 symmetry. The results are shown in Fig. 5.4. At $\Gamma = 1.0 < \Gamma_c$, the six peaks are well developed, confirming that the system is in the ordered phase. As the transverse field is increased to $\Gamma = 1.4$, close to the critical point, the peaks start to broaden and decrease in intensity, signaling the progressive weakening of long-range correlations. Finally, at $\Gamma = 1.6 > \Gamma_c$, the peaks are largely replaced by broad, diffuse hills in $S(\mathbf{q})$, indicating that quantum fluctuations have effectively destroyed the long-range clock order, and the system has entered the disordered paramagnetic phase.

Analyzing $S(\mathbf{q})$ in this manner is crucial because it provides complementary evidence to the order parameter and Binder cumulant measurements. While the latter quantify the global degree of ordering, the structure factor reveals how the order manifests spatially, allowing one to track the evolution of correlations and directly visualize the mechanism through which quantum fluctuations destabilize the ordered phase. This makes $S(\mathbf{q})$ an essential diagnostic for both characterizing the clock phase and understanding the nature of the quantum phase transition.

5.5 Summary of the chapter

In this chapter, we have systematically investigated the zero-temperature quantum phase transitions of the triangular-lattice Ising antiferromagnet under a transverse field, focusing on both nearest-neighbor and long-range interactions with a power-law decay $\alpha = 6$. Our analysis highlights the central role of frustration and the order-by-disorder mechanism in stabilizing the clock-ordered phase, a paradigmatic example of emergent order from a highly degenerate classical manifold.

Using self-learning Projection Quantum Monte Carlo simulations, we have characterized the transition from the clock phase to the paramagnetic phase through the Binder cumulant and finite-size scaling analysis. For both interaction regimes, we extracted estimates of the critical transverse field Γ_c and the correlation-length exponent $1/\nu$. Notably, we observed a downward shift of Γ_c for the $\alpha = 6$ model compared to the nearest-neighbor limit, in agreement with theoretical expectations and previous

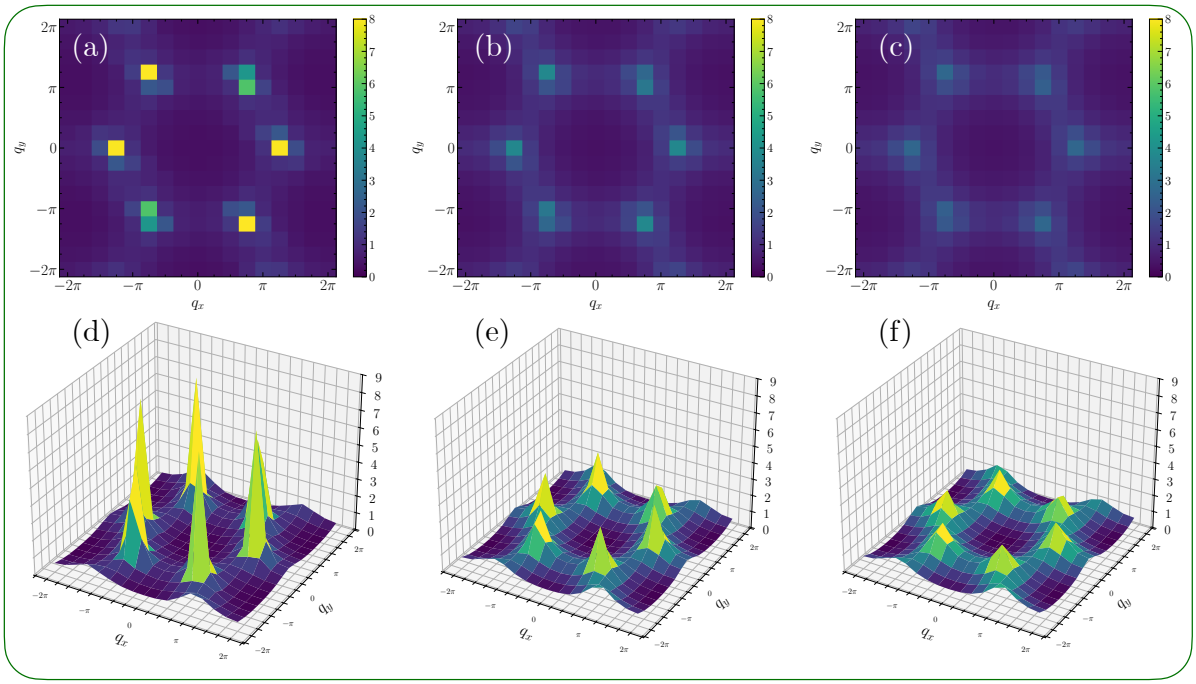

FIGURE 5.4: Magnetic structure factor $S(\mathbf{q})$ for the triangular-lattice model with power-law interactions $\alpha = 6$ at three representative transverse field strengths. Panels (a-c) show 2D intensity maps, while panels (d-f) display the corresponding 3D representations. Panels (a) and (d) correspond to $\Gamma = 1.0$, well inside the clock-ordered phase, where six sharp peaks indicate the presence of long-range three sublattice magnetic order. Panels (b) and (e), at $\Gamma = 1.4$, near the critical point, show a clear reduction and broadening of the peaks, reflecting the progressive weakening of the clock order due to increasing quantum fluctuations. Finally, panels (c) and (f), at $\Gamma = 1.6$, in the paramagnetic phase, reveal that the six peaks have largely disappeared, replaced by diffuse hills, signaling the destruction of long-range order. This figure illustrates how the magnetic structure factor provides a direct visualization of the evolution of spatial correlations across the quantum phase transition and complements the Binder cumulant and order-parameter analysis.

studies: longer-range interactions enhance quantum fluctuations, thereby destabilizing the ordered phase.

Complementary insight was provided by the magnetic structure factor $S(\mathbf{q})$, which allowed us to visualize the spatial manifestation of clock order. Our results clearly demonstrate the sixfold peak structure characteristic of the clock phase at low fields, its progressive weakening near the critical point, and the complete disappearance of long-range order in the paramagnetic regime. This confirms that the clock phase is robust against moderate long-range couplings but ultimately melt in a disordered phase due to strong quantum fluctuations.

Overall, this chapter underlines the rich interplay between geometric frustration,

quantum fluctuations, and interaction range in determining the phase behavior of triangular-lattice antiferromagnets. The results provide both a quantitative benchmark for future numerical studies and a direct connection to experimental implementations in Rydberg atom arrays and other quantum simulators. In particular, they establish a concrete foundation for exploring how long-range interactions modify the critical properties of fluctuation-induced ordered phases.

In the next chapter, we will explore how to simulate an effective Rydberg Hamiltonian using Projection Quantum Monte Carlo. We will detail the mapping of a Rydberg Hamiltonian onto an Ising-like model that is amenable to our projection simulation techniques. Using this approach, we will reproduce finite-size scaling results reported in the literature for the transition from a checkerboard-ordered phase to a disordered phase. This will provide a concrete framework for simulating experimentally relevant systems with Rydberg atoms. Finally, we will discuss how even non-optimal outputs from a quantum computer can be leveraged to enhance PQMC performance, thereby improving the efficiency of classical simulations.

Chapter 6

Enhancing PQMC with Rydberg atoms quantum platforms

In this chapter, we present a detailed study of effective Rydberg Hamiltonians mapped onto Ising-like models, suitable for simulation via Projection Quantum Monte Carlo (PQMC). We focus on systems with open boundary conditions, as these are most relevant for experimental realizations with Rydberg atom arrays. A checkerboard order parameter, defined as $\mathcal{F}_{\pi,\pi}$ is introduced to characterize the phase transition from an ordered checkerboard phase to a disordered phase. Finite-size scaling of this parameter is employed to accurately locate the critical point, allowing direct comparison with existing literature results.

We then extend our approach by incorporating experimental data obtained from quantum computers implementing Rydberg atom arrays. Using these outputs, we train a restricted Boltzmann machine (RBM) to learn the underlying probability distribution, providing an informed initial ansatz for the self-learning PQMC. This hybrid quantum-classical scheme demonstrates that initializing PQMC with data from real quantum devices enables simulations of system sizes that are otherwise intractable, surpassing the accuracy achieved in analogous studies based on variational optimization methods. Our results highlight the potential of combining experimental quantum data with classical stochastic simulation techniques to explore complex many-body systems beyond current computational limits [44, 46, 162, 45].

6.1 Neural Network representation of Rydberg many-body states

The accurate description of quantum many-body systems is a fundamental and long-standing challenge in modern physics. Classical numerical techniques, such as exact diagonalization or tensor network methods, face severe scalability issues due to the exponential growth of the Hilbert space with system size. Quantum Monte Carlo methods offer powerful alternatives but are often limited by the well-known sign problem, which renders many relevant models intractable. Against this limitation, programmable quantum simulators, especially Rydberg atom arrays, have recently emerged as a compelling platform to probe complex quantum phenomena experimentally, complementing and challenging theoretical methods [112, 87, 91, 139]. Rydberg atom arrays consist of neutral atoms trapped in optical tweezers and excited to highly energetic Rydberg states. Their coherent evolution is governed by the Rydberg Hamiltonian, which can be expressed as

$$\hat{H} = \frac{\Omega}{2} \sum_{i=1}^{N} \hat{\sigma}_i^x - \delta \sum_{i=1}^{N} \hat{n}_i + \sum_{i,j} V_{ij} \hat{n}_i \hat{n}_j, \tag{6.1}$$

where Ω is the Rabi frequency, δ the laser detuning, and $V_{ij} \sim 1/|r_i - r_j|^6$ the van der Waals interaction between atoms i and j. The operator $\hat{\sigma}_i^x$ induces coherent transitions between the ground $|g\rangle$ and Rydberg state $|r\rangle$, while $\hat{n}_i = |r\rangle_i\langle r|$ represents the number operator. Crucially, the Rydberg blockade mechanism prevents simultaneous excitations within a characteristic distance called Rydberg blockade radius R_b, giving rise to effective constraints that stabilize rich many-body phases. By tuning the parameters Ω, δ, and the array geometry, one can engineer a wide variety of quantum states and access quantum phase transitions between ordered and disordered regimes [115, 163, 164, 127]. The characterization of such states through projective measurements is complicated by the destructive nature of readout, which necessitates repeated state preparation. Current experimental platforms achieve only a few shots per second, leading to limited data availability compared to other architectures such as superconducting circuits or trapped ions. This bottleneck has motivated the search for efficient strategies to extract maximal information from scarce and noisy measurement data.

In recent years, neural-network quantum states have been introduced as versatile

variational ansatz for representing wave functions. Generative models such as restricted Boltzmann machines, exhaustively descripted in Chap. 2, Autoregressive Neural Networks (ANNs), and more recently Transformers, are particularly well suited to encode probability distributions associated with many-body states [46, 44, 45]. In this context, the wave function amplitudes can be mapped to the output of the network, e.g.

$$|\Psi(\sigma)|^2 = p_{\mathrm{ANN}}(\sigma;\boldsymbol{\theta}), \tag{6.2}$$

where p_{ANN} is the probability distribution generated by the neural network with parameters $\boldsymbol{\theta}$. For autoregressive models like Recurrent Neural Networks (RNNs), the joint distribution is factorized into conditional probabilities, guaranteeing normalization and allowing for efficient independent sampling without the need for Markov chains.

Two complementary training paradigms have been developed. In the data-driven setting, networks are optimized to reproduce the empirical distribution of projective measurement outcomes, typically by minimizing the Kullback-Leibler divergence between the network distribution and the data, as explained in Sec. 2.3 of Chap. 2. In the Hamiltonian-driven setting, variational Monte Carlo (VMC) is employed: the energy expectation value

$$E(\boldsymbol{\theta}) = \frac{\langle\Psi_{\boldsymbol{\theta}}|\hat{H}|\Psi_{\boldsymbol{\theta}}\rangle}{\langle\Psi_{\boldsymbol{\theta}}|\Psi_{\boldsymbol{\theta}}\rangle} = \frac{1}{N_s}\sum_{\sigma\sim p_{\mathrm{ANN}}} H_{\mathrm{loc}}(\sigma), \tag{6.3}$$

with local energies

$$H_{\mathrm{loc}}(\sigma) = \frac{\langle\sigma|\hat{H}|\Psi_{\boldsymbol{\theta}}\rangle}{\langle\sigma|\Psi_{\boldsymbol{\theta}}\rangle}, \tag{6.4}$$

is minimized with respect to the variational parameters. Both approaches have their strengths and limitations: purely data-driven training requires large data sets, while Hamiltonian-driven training suffers from slow convergence due to the rugged optimization landscape.

To overcome these limitations, data-enhanced variational Monte Carlo has been proposed. In this hybrid approach, networks are first pretrained on limited experimental data and then fine-tuned using Hamiltonian-driven optimization. This strategy has been shown to significantly accelerate convergence and, in some cases, improve accuracy beyond what can be achieved by either method alone [44, 45].

Nevertheless, RNN-based models struggle to capture long-range correlations in large two-dimensional systems due to their sequential architecture. To address this, transformer models have recently been adopted as wave-function ansatz. Transformers leverage self-attention mechanisms to account for all-to-all correlations across the system, overcoming the locality constraints of RNNs. Even more powerful are patched and large patched transformers, inspired by advances in computer vision [46, 15]. By processing blocks of qubits instead of single sites, these architectures reduce sequence length and computational cost while retaining the ability to accurately model correlations within and between patches. Benchmark studies have demonstrated that transformers can reach accuracies surpassing both RNNs and traditional QMC methods, particularly in the vicinity of quantum critical points [46].

Looking forward, the combination of experimental data with advanced neural-network architectures opens a path toward scalable, high-accuracy simulations of quantum matter. Hybrid strategies leveraging autoregressive sampling, transformer-based architectures, and efficient pretraining protocols promise not only to enhance the performance of variational methods but also to provide new insights into the physics of strongly correlated systems. This convergence of machine learning, quantum simulation, and experimental platforms represents a pivotal step toward unlocking the full potential of quantum many-body physics. In this chapter, we depart from the standard use of Variational Monte Carlo with autoregressive neural networks. The aim of this study is to leverage experimental data obtained from real Rydberg atom platforms [88] to enhance the performance of Projection Quantum Monte Carlo. Specifically, we train a neural quantum state in the form of a Restricted Boltzmann Machine on experimental data collected at different values of the detuning parameter δ. Together with the blockade radius R_b, this parameter determines the phase diagram of the model under consideration.

The experimental data originate from quantum simulations performed on Rydberg platforms and are expected to represent, as closely as possible, the ground state of the system. The trained RBM is then employed as a guiding wave function within PQMC. We demonstrate that this strategy allows PQMC to simulate systems that would otherwise be intractable for self-learning PQMC due to their size.

Finally, we benchmark this approach against other established methods, including conventional VMC and PQMC guided by neural networks. The results show that the

RBM-enhanced PQMC achieves accuracies comparable to, and in some regions of the phase diagram even exceeding, those obtained with existing techniques.

6.2 Experimental data

The experimental data employed in this work were obtained on a state-of-the-art Rydberg atom platform developed in previous studies [88]. In these experiments, a spatial light modulator was used to create large two-dimensional arrays of optical tweezers, which were then loaded with single Rb^{87} atoms from a magneto-optical trap. By means of an additional set of movable tweezers controlled through acousto-optic deflectors, the atoms were rearranged into defect-free geometries with filling fractions above 99%. This approach enabled the realization of programmable lattice structures such as square, honeycomb, and triangular arrays, typically containing several hundred atoms.

The qubits were encoded in the ground state $|g\rangle$ and the highly excited Rydberg state $|r\rangle$ ($n = 70$), coherently coupled via a two-photon transition driven by counter-propagating laser beams. The many-body dynamics of the system were governed by the interplay between coherent driving and strong van der Waals interactions, giving rise to the Rydberg blockade effect. By tuning the detuning parameter δ, the Rabi frequency Ω, and the lattice spacing a, the effective blockade radius R_b could be controlled, thereby giving access to a wide range of quantum phases.

State readout was performed through high-fidelity fluorescence imaging, which detected atoms in $|g\rangle$ directly while identifying Rydberg excitations as losses.

In this work, we focus specifically on the 16×16 Rydberg atom array, as this system size provides a sufficiently large Hilbert space to capture non-trivial correlations while still being experimentally accessible. Our primary interest lies in the quantum phase transition from the disordered regime to the ordered checkerboard phase, which represents one of the most studied and well-characterized phenomena on these platforms.

In the case of square arrays of size 16×16, the ordered checkerboard phase emerges when the ratio R_b/a, where a is the lattice spacing, approaches unity, which enforces the constraint that no two nearest-neighbor atoms can be simultaneously excited to

System size	16×16
Sweep rates	$15, 21, 30, 4260, 85, 120$ MHz$/\mu$s
Raby frequency Ω	4.24
Effective lattice spacing R_b/a	1.15
Detuning δ/Ω	-1.78 to 3.645 by 0.236
Critical detuning δ_c/Ω	$1.12(1)$

TABLE 6.1: Table of experimental values for the checkerboard phase transition [44, 88].

the Rydberg state. This condition stabilizes a $\mathbb{Z}_2$ symmetry-broken configuration corresponding to an antiferromagnetic ground state. To prepare this phase experimentally, the system is initialized at $R_b/a \approx 1.15$ (with lattice spacing $a = 6.7\,\mu$m and Rabi frequency $\Omega = 2\pi \times 4.3\,$MHz) with all atoms in the ground state $|g\rangle$. The detuning δ is then swept from negative to positive values while keeping Ω fixed, thereby driving the system quasi-adiabatically into the checkerboard phase. Following [44], we report all the experimental values for the checkerboard transition in Tab. 6.1.

6.2.1 Dataset

Our dataset consists of spin configurations sampled from the quantum platform at different sweep rates of 15, 21, 30, 42, 60, 85, 120 MHz$/\mu$s, and for several values of the normalized detuning δ/Ω. For each sweep rate and detuning value we have access to 500 measurement shots, except for the slowest sweep at 15 MHz$/\mu s$, where 1000 configurations were collected. To enlarge the dataset, we combine the configurations obtained at different sweep rates for the same detuning value. This procedure yields a total of 4000 configurations per detuning point. In order to further augment the dataset, we also include rotated versions of each configuration by angles of $\pi/2$, π, and $3\pi/2$. As a result, for each value of δ we obtain a final dataset of 16000 configurations, which is used to train the RBM which serves as guiding wave function for the PQMC.

Although real quantum simulators represent a promising tool for addressing many-body problems, they are also subject to technical limitations, most notably the presence of noise. As a result, data collected from these devices are often imperfect. In Fig. 6.1 we show a single snapshot of the system at two different values of the detuning, $\delta/\Omega = -1.78$ and $\delta/\Omega = 1.78$, corresponding to points deep inside the ferromagnetic and checkerboard phases, respectively. It is evident that, despite being well within these ordered regimes, the sampled configurations still display defects. The central

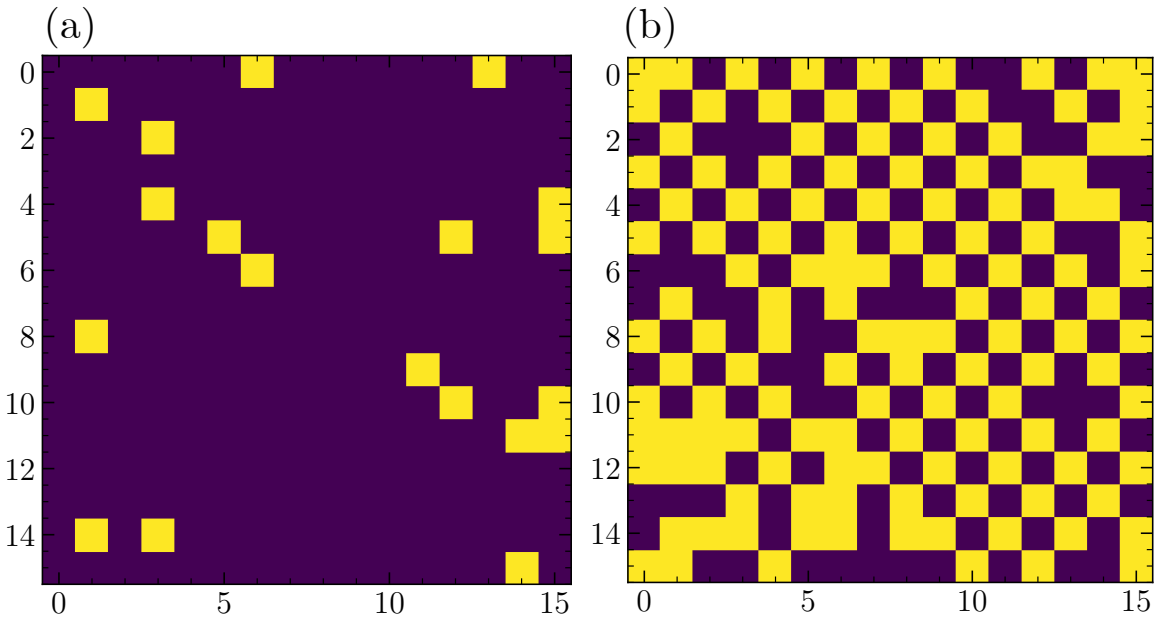

FIGURE 6.1: Single snapshot of experimental data for the 16×16 system taken at 21 MHz/μs sweep rate in the ferromagnetic phase $\delta/\Omega = -1.78$ (a) and deep in the checkerboard phase $\delta/\Omega = 3.645$ (b).

goal of this study is to show that, even in the presence of such noise, experimental data can play a crucial role in enhancing the performance of classical techniques such as Projector Quantum Monte Carlo (PQMC) for the simulation of systems that would otherwise be extremely challenging to study.

Nevertheless, these data clearly capture the physical properties of the system, most notably the phase transition. In the experimental literature it is common to use the average density of Rydberg excitations, $\bar{n}$, as an order parameter, plotted as a function of the detuning δ. In Fig. 6.2 we report the behavior of $\bar{n}$ extracted from the measurements of the quantum simulator as a function of δ/Ω. It is evident that for $\delta < \delta_c$, the excitation density vanishes, corresponding to the configuration in which all atoms remain in the ground state $|g\rangle$. Conversely, for $\delta > \delta_c$, the excitation density increases and saturates around 0.5, as expected for the checkerboard-ordered phase.

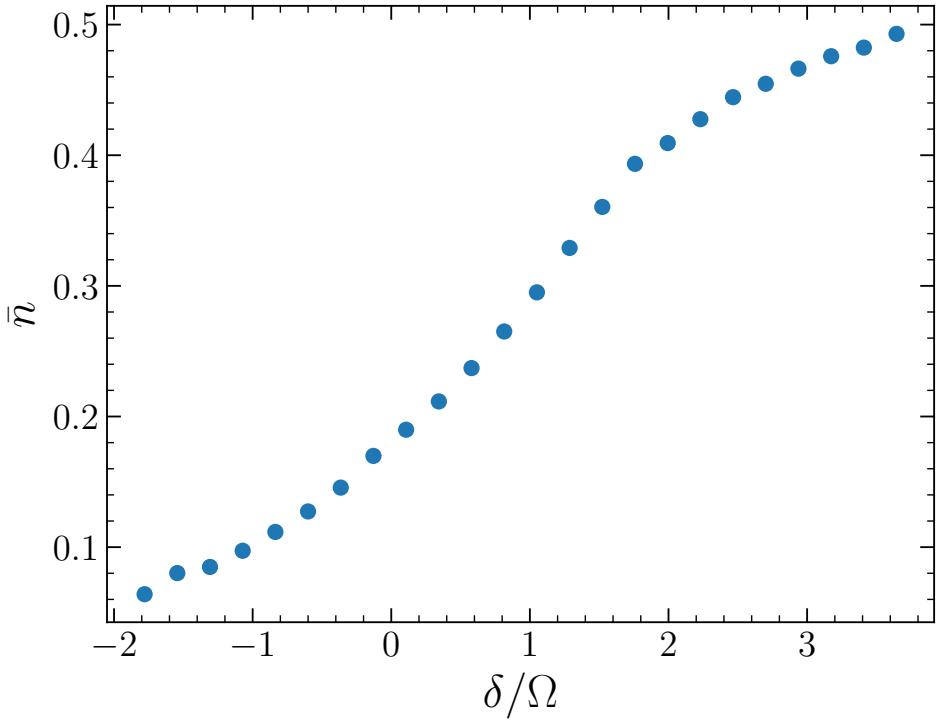

FIGURE 6.2: Mean Rydberg excitation density $\bar{n}$ as a function of detuning δ/Ω at sweep rate 21 MHz/μs signing the trasition between the ferromagnetic and the checkerboard phase.

6.3 Rydberg Hamiltonian in Ising form

First, in order to simulate the full Rydberg Hamiltonian, it is convenient to map it onto a more familiar framework suitable for PQMC. In particular, as long as we can identify the interaction term, the transverse field term, and, when relevant, as in the present case, the longitudinal field term, we are able to access and characterize the ground-state properties of these systems. The map of the Rydberg Hamiltonian to Ising Hamiltonian follows reference [115]. Let us start with the standard Rydberg Hamiltonian,

$$H = \Omega \sum_{i<j} \left(\frac{R_b}{R_{ij}} \right)^6 n_i n_j - \Delta \sum_i n_i + \frac{\Omega}{2} \sum_i \sigma_i^x. \tag{6.5}$$

where $n_i \equiv (\sigma_i^z + 1)/2$ measures the Rydberg excitation density at site i. Using this relation, we can write

$$H = \Omega \sum_{i<j} \left(\frac{R_b}{R_{ij}}\right)^6 \frac{(\sigma_i^z + 1)(\sigma_j^z + 1)}{4} - \Delta \sum_i \frac{(\sigma_i^z + 1)}{2} + \frac{\Omega}{2} \sum_i \sigma_i^x. \tag{6.6}$$

Thus,

$$\begin{aligned} H = &\frac{\Omega}{4} \sum_{i<j} \left(\frac{R_b}{R_{ij}}\right)^6 \sigma_i^z \sigma_j^z + \frac{\Omega}{4} \sum_{i<j} \left(\frac{R_b}{R_{ij}}\right)^6 \sigma_i^z + \frac{\Omega}{4} \sum_{i<j} \left(\frac{R_b}{R_{ij}}\right)^6 \sigma_j^z \\ &+ \frac{\Omega}{4} \sum_{i<j} \left(\frac{R_b}{R_{ij}}\right)^6 - \frac{\Delta}{2} \sum_i \sigma_i^z - \frac{\Delta}{2} N + \frac{\Omega}{2} \sum_i \sigma_i^x. \end{aligned} \tag{6.7}$$

Collecting the second and third terms, we get the following.

$$\begin{aligned} H = &\frac{\Omega}{4} \sum_{i<j} \left(\frac{R_b}{R_{ij}}\right)^6 \sigma_i^z \sigma_j^z + \sum_i I_i \sigma_i^z + \frac{\Omega}{4} \sum_{i<j} \left(\frac{R_b}{R_{ij}}\right)^6 - \frac{\Delta}{2} \sum_i \sigma_i^z \\ &- \frac{\Delta}{2} N + \frac{\Omega}{2} \sum_i \sigma_i^x, \end{aligned} \tag{6.8}$$

where $I_i = \frac{\Omega}{4} \sum_{j \neq i} \left(\frac{R_b}{R_{ij}}\right)^6$. Finally, placing the constant terms last, we conclude with

$$H = \frac{\Omega}{4} \sum_{i<j} \left(\frac{R_b}{R_{ij}}\right)^6 \sigma_i^z \sigma_j^z + \sum_i \left(I_i - \frac{\Delta}{2}\right) \sigma_i^z + \frac{\Omega}{2} \sum_i \sigma_i^x + \frac{\Omega}{4} \sum_{i<j} \left(\frac{R_b}{R_{ij}}\right)^6 - \frac{\Delta}{2} N. \tag{6.9}$$

The first, second, and third terms correspond to the interactions, the longitudinal field, and the transverse field contributions, respectively, with the last two being constants. To investigate the properties of the system's ground state, we can disregard the last two constants terms, as they are only necessary to obtain an exact value of the ground state energy.

6.3.1 Checkerboard phase in Rydberg square lattice

To demonstrate the effectiveness of the self-learning PQMC method, we present numerical results on a square lattice governed by a Rydberg Hamiltonian of the form

given in Eq. (6.9), with open boundary conditions (OBC). In order to characterize the phase transition from the ferromagnetic phase to the checkerboard phase, we use the Fourier transform of the magnetization evaluated at the specific wave vector $(k_x, k_y) = (\pi, \pi)$, which is defined as follows

$$\mathcal{F}_{\pi,\pi} = \frac{1}{N} \sum_i \sigma_i^z \exp\left[i(\pi, \pi)(x_i, y_i)\right], \tag{6.10}$$

where N is the number of atoms. An interested reader may refer to Ref. [115] for further details, where variants of this observable are employed to characterize different phase transitions, such as the striated and star-ordered phases.

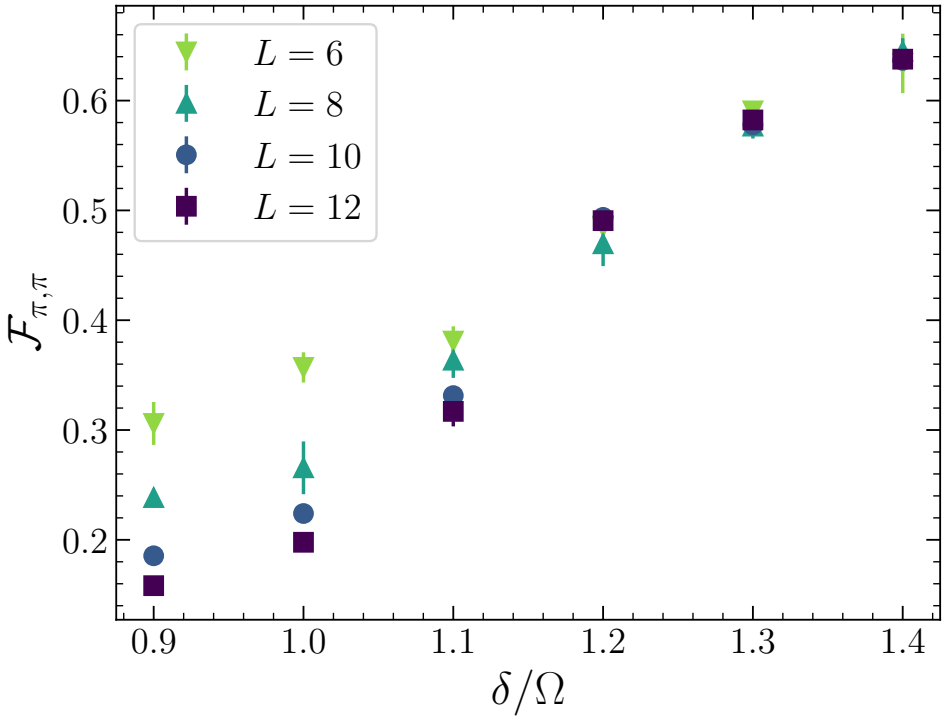

FIGURE 6.3: Checkerboard order parameter $\mathcal{F}_{\pi,\pi}$ as a function of detuning δ/Ω for $L = 6, 8, 10, 12$, where L is the side length of the square lattice.

In Fig. 6.3 we show the dependence of $F_{\pi,\pi}$ on the detuning, while Fig. 6.4 report the scaling collapse of the data across different system sizes. The simulations were performed with the Hamiltonian of Eq. 6.9, using parameters $R_b/a = 1.2$, $\Omega = 1$, and lattice spacing $a = 1$. From the scaling analysis, carried out with the numerical suite in Ref. [77], we obtain an estimation of the critical detuning $\delta_c/\Omega = 1.12(1)$ together

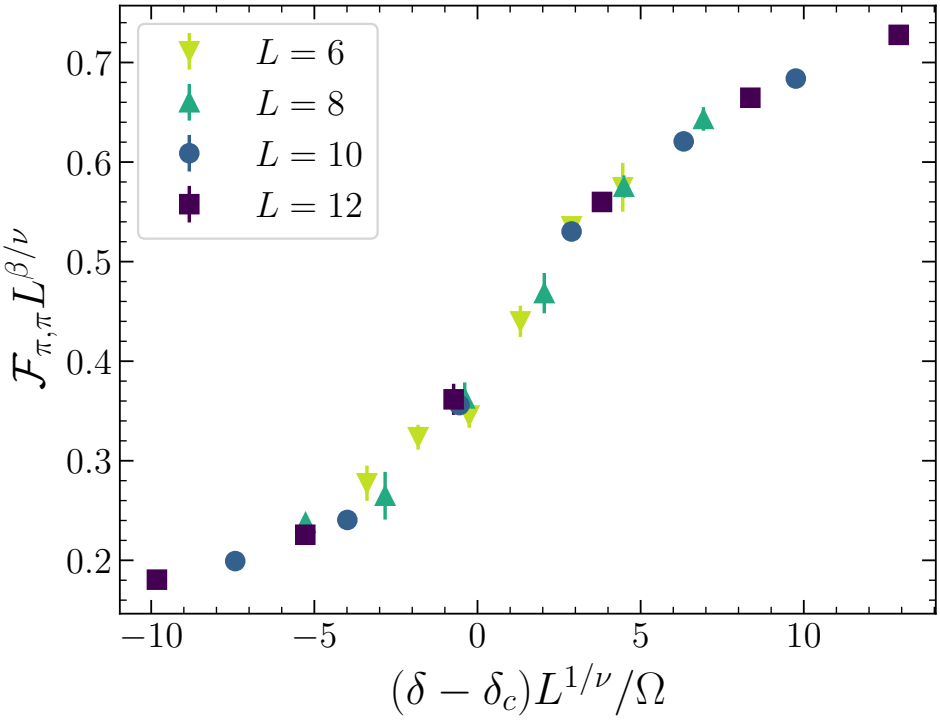

FIGURE 6.4: Finite size scaling of $\mathcal{F}_{\pi,\pi} L^{\beta/\nu}$ as a function of the reduced detuning $(\delta - \delta_c) L^{1/\nu}/\Omega$. The critical parameters δ_c/Ω, $1/\nu$ and β are obtained by optimizing the data collapse [77].

with the critical exponents. In particular, we extract the inverse correlation length exponent $1/\nu = 1.54(18)$ and a magnetization exponent $\beta/\nu = 0.332(40)$. These values are in fair agreement with previous results in the literature, who found $1/\nu \approx 1.582$ and $\beta \approx 0.29$ for the same set of Hamiltonian parameters. It is worth noting that the small discrepancies observed in the critical exponents may be attributed to the use of open boundary conditions, which are not ideal for extracting universal scaling behavior, particularly at the relatively modest system sizes accessible in our simulations. Nevertheless, we consider this benchmark highly relevant, since the focus of this study is on real quantum platforms, where open boundary conditions naturally arise.

6.4 Methods and results

In the previous section, we have shown that self-learning PQMC, combined with the mapping of the Rydberg Hamiltonian onto an effective Ising model, provides a promising approach for the simulation of these systems. Nevertheless, one of the main limitations of the method lies in its applicability to relatively small system sizes. As illustrated in Figs. 6.3 and 6.4, the largest lattice considered in our simulations is $L = 12$. However, to access the physics in the thermodynamic limit, it is essential to simulate larger systems, which is the regime that can be reached experimentally on quantum simulators.
The goal of this work is therefore to exploit experimental data described in the previous section to enhance classical simulation techniques like PQMC. In our approach, the platforms output are used to train a neural quantum state, which then serves as a guiding wave function for the PQMC simulation. We will show that this protocol, when combined with self-learning, provides a robust method that outperforms variational NQS approaches based on different training strategies. As a benchmark, we compare our results against those obtained from Stochastic Series Expansion quantum Monte Carlo (SSE QMC) simulations [165].

All the data presented in this study are based on the experimental work reported in Ref. [88], and therefore, we fix the blockade radius to $R_b = 1.15$ throughout this chapter. In this regime, the transition point between the ferromagnetic and checkerboard phases occurs at $\delta_c/\Omega = 1.12(1)$. For this reason, we restrict our analysis to a set of positive detuning values within the range listed in Tab. 6.1.

In order to provide a comprehensive benchmark, we compare the performance of different strategies, ranging from projector-based methods to purely variational approaches. This allows us to assess the effectiveness of incorporating experimental data, as well as to evaluate the role of different neural-network ansätze and optimization protocols. The benchmark can be structured as follows:

- Projector Quantum Monte Carlo (PQMC) methods
  - Self-learning PQMC: the guiding neural network is randomly initialized and subsequently trained on the fly using the walkers generated during the

PQMC simulation. This baseline approach, already discussed in earlier sections, highlights the intrinsic capability of PQMC to refine its own guiding wave function without prior knowledge.

- PQMC with experimental data: here the guiding wave function is provided by a Restricted Boltzmann Machine (RBM) trained on real experimental configurations. In this way, the PQMC starts from a physically meaningful approximation of the ground state, reflecting the information already encoded in the measurements. In Fig. 6.5 we show the workflow of the method.
- PQMC with variationally optimized RBM: in place of experimental data, the RBM is pretrained within the standard Variational Monte Carlo (VMC) framework using the NetKet library [61]. This method isolates the benefit of a well-optimized guiding state obtained purely from numerical simulations, without experimental input.
- Hybrid PQMC (self-learning + experimental data): in this combined scheme, the RBM is first initialized from experimental configurations, ensuring an informed starting point, and then refined dynamically during the PQMC run. This approach merges the advantages of experimental guidance and self-learning capability, and, as we will demonstrate, provides the best overall performance among the projector methods. Also in this case, the scheme of the algorithm is sketched in Fig. 6.5.

In the case of PQMC with experimental data, the Restricted Boltzmann Machine is trained to reproduce the probability distribution of the experimentally measured spin configurations, namely the marginal probability $p(\boldsymbol{\sigma})$ extracted from projective measurements on the Rydberg platform. Once this distribution has been learned, we construct a corresponding neural quantum state by taking the square root of the RBM probability,

$$\psi_{\mathrm{RBM}}(\boldsymbol{\sigma}) = \sqrt{p_{\mathrm{RBM}}(\boldsymbol{\sigma})}.$$

This choice is physically justified by the fact that the Rydberg Hamiltonian can be mapped onto an effective transverse-field Ising model whose ground state can be taken as real and positive. Under this assumption, the phase of the wave function

is irrelevant for the observables considered in this work, which are all diagonal in the computational basis. Therefore, the RBM trained on experimental data directly provides a consistent and physically meaningful guiding wave function for PQMC, faithfully encoding the probability amplitudes of the ground-state manifold.

The overall workflow of this approach is illustrated in Fig. 6.5.

- Variational Monte Carlo (VMC) methods: To complement the projector approaches, we also include purely variational benchmarks.
  - RBM ansatz: optimized using two different training protocols, namely Stochastic Gradient Descent (SGD), the standard baseline algorithm and Stochastic Reconfiguration (SR), a more advanced optimization scheme that better accounts for the geometry of parameter space. In both cases we use a density of hidden neurons $\alpha = N_h/N = 2$
  - Autoregressive neural network (ARNN) ansatz: likewise optimized with both SGD and SR. ARNNs have shown significant advantages in representing probability distributions of quantum states, especially due to their autoregressive structure that enforces normalization by construction. We use a single layer architecture with $N_h = 512$ neurons.

It is important to note that the VMC calculations performed with the NetKet library [61] for both the RBM and ARNN architectures were carried out using a fixed set of hyperparameters. Therefore, the results presented in this work should be regarded as representative of this specific choice, rather than optimal performance benchmarks. In particular, for the RBM we used between $N_{\text{iters}} \approx 5000$ optimization iterations, with Monte Carlo batches of $N_b = 2000$ configurations after discarding the first $N_{\text{therm}} = 100$ samples for thermalization. When using stochastic gradient descent (SGD), the learning rate was set to $\text{lr} = 10^{-4}$, while for the Adam optimizer combined with SGD we used a learning rate of $\text{lr} = 10^{-3}$. Similarly, for the ARNN ansatz, we performed approximately $N_{\text{iters}} \approx 10000$ optimization iterations with batches of $N_b = 2000$ configurations and the same learning-rate settings as for the RBM. Nevertheless, we emphasize that the performance of VMC methods can likely be improved through a more systematic hyperparameter tuning, as well as by adopting more expressive neural-network architectures beyond those explored here.

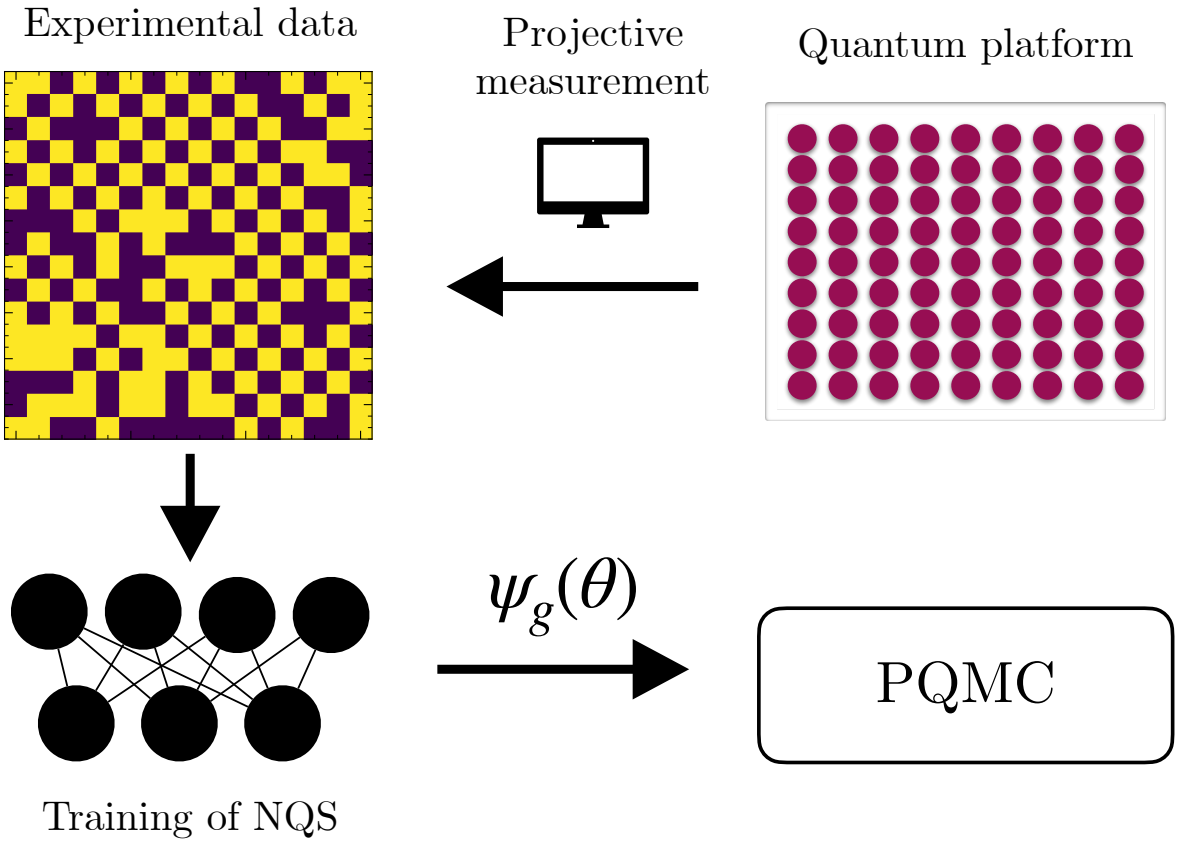

FIGURE 6.5: Workflow of the proposed method. Experimental data are obtained from a real Rydberg atom platform through projective measurements of the system state. The resulting spin configurations, sampled at different values of the detuning parameter δ, are used as a dataset to train a neural quantum state in the form of a restricted Boltzmann machine, which serves as the guiding wave function for the PQMC. The same workflow applies to the hybrid scheme, with the only difference that self-learning is additionally employed during the PQMC run.

We start by presenting the ground-state simulation results obtained with the different techniques at several values of δ/Ω. In Fig. 6.6 we show a qualitative plot, since the absolute energy scales are large, that provides at a glance a clear comparison of the relative performance of the various methods. As anticipated, self-learning PQMC with finite walker population struggles to accurately simulate large systems, despite the fact that in principle the algorithm converges to the exact result either when the guiding wave function is extremely accurate or in the limit of an infinite number of walkers.

We also observe that Variational Monte Carlo with an RBM ansatz optimized via vanilla stochastic gradient descent suffers from a systematic bias, which persists across

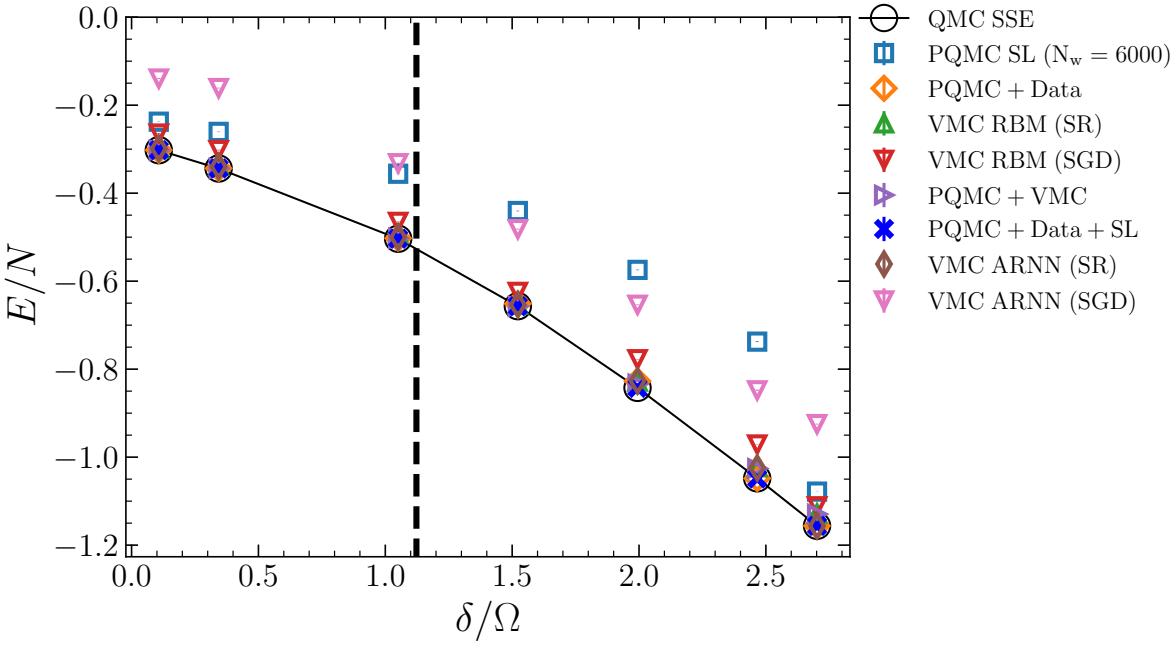

FIGURE 6.6: Energy per atom as a function of the detuning δ/Ω for different simulation techniques. The black circles indicate the reference results obtained with SSE QMC. Both self-learning PQMC with finite walkers population $N_w = 6000$ and all VMC-based methods display a systematic bias in the energy estimation, which becomes more pronounced as δ increases, i.e., within the checkerboard phase to the right of the dashed vertical line marking the critical point δ_c. By contrast, PQMC initialized with experimental data, either in the pure or hybrid self-learning version, shows convergence to the exact values across the phase diagram. Comparable accuracy is also achieved by PQMC initialized with a variationally optimized RBM obtained in a separate VMC simulation using the NetKet library [61].

all points of the phase diagram analyzed. In contrast, the other approaches display stable and reliable performance, producing acceptable results for every detuning value considered.

To be more precise, Fig. 6.7 shows the absolute difference between the energy expectation value denoted by $\langle H \rangle$ obtained with the various techniques and those from SSE QMC simulations, denoted with E_{SSE}. As already anticipated, both self-learning PQMC and VMC with RBM, optimized with either stochastic gradient descent or stochastic reconfiguration, perform suboptimally compared to the other approaches. The three protocols that deserve particular attention are: (i) PQMC guided by an RBM trained on experimental data, (ii) self-learning PQMC initialized with the experimentally trained RBM, and (iii) PQMC guided by a variationally optimized RBM using the stochastic reconfiguration algorithm.

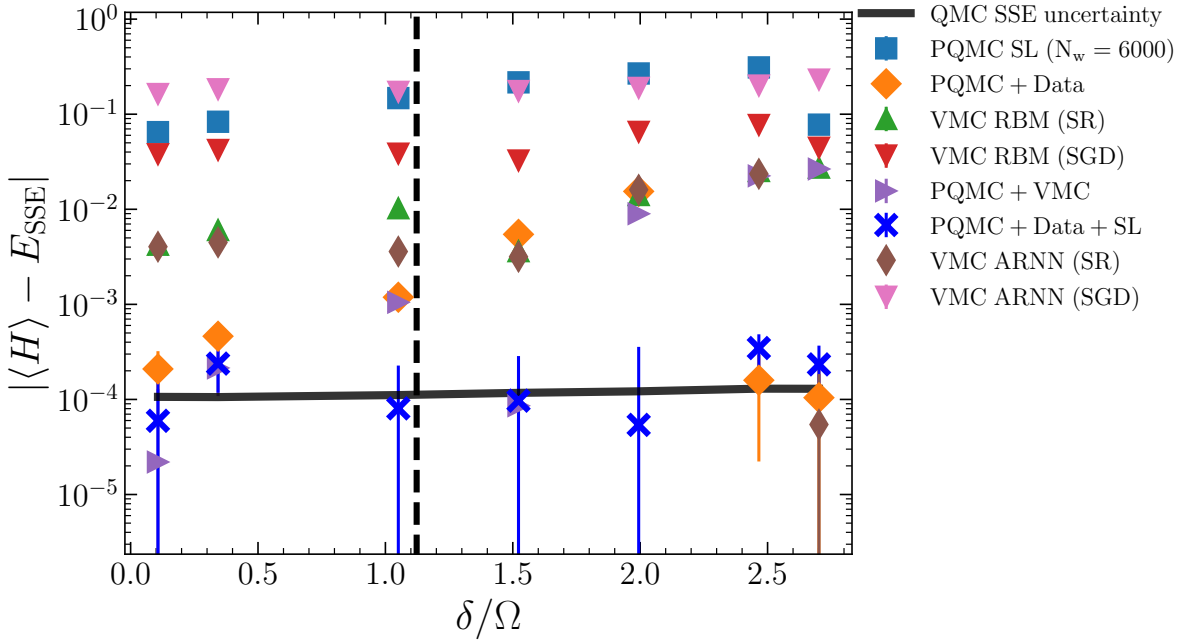

FIGURE 6.7: Absolute error in the energy estimation as a function of the detuning δ/Ω. The dashed vertical line indicates the critical detuning δ_c, while the solid horizontal line corresponds to the uncertainty of the SSE QMC results reported in Ref. [165]. Variational methods as well as the stand-alone self-learning PQMC display large errors, which systematically increase with δ, i.e., within the checkerboard phase. This dependence on δ essentially disappears when self-learning PQMC is initialized with experimental data, making this the only protocol whose accuracy does not degrade across the phase diagram. Remarkably, for the last two detuning values, the self-learning PQMC initialized with experimental data even achieves energies lower than the SSE QMC benchmark.

Initializing PQMC with a variationally optimized ansatz obtained in a separate VMC simulation is the standard practice. In this case, we employ the NetKet library to optimize an RBM within the VMC framework using the efficient stochastic reconfiguration method, which is well known to outperform stochastic gradient descent and similar schemes [61]. For small detuning values, corresponding to the ferromagnetic phase, the results are very accurate. However, as the critical point is approached the error increases, and the accuracy generally degrades within the checkerboard phase. A similar trend is observed for PQMC guided by an RBM trained on experimental data. Interestingly, the main difference appears deep inside the checkerboard phase, namely at $\delta/\Omega = 2.466$ and $\delta/\Omega = 2.702$ where the use of experimental data provides a more effective guiding wave function for PQMC than a variationally optimized ansatz.

Finally, the best results are obtained with the hybrid PQMC approach. In this protocol, the algorithm is initialized with a guiding wave function obtained from experimental data and subsequently refined on the fly using the walkers sampled during the PQMC run. Unlike the other methods, we find that the accuracy of the ground-state estimates is essentially independent of the detuning value. This is a highly promising result, as it suggests that the hybrid strategy is robust even in notoriously difficult regimes, such as near the critical point, where correlations grow exponentially and standard methods typically struggle. It is also worth discussing the last two points of the plot in Fig. 6.7. At first sight, the error associated with the hybrid protocol appears to be larger than that of the method relying solely on experimental data. However, this effect is in fact an artifact of the way the data are displayed. To clarify this point, in Fig. 6.8 we present the full energy history as a function of Monte Carlo steps for both approaches. This comparison reveals two important features. First, the hybrid protocol consistently exhibits a much smaller variance during the sampling process, a behavior that we observe throughout the phase diagram. Second, and perhaps more strikingly, the average energy obtained with the hybrid protocol lies slightly below the value reported by the SSE QMC. This bias may be attributed to the fact that Stochastic Series Expansion is intrinsically a finite-temperature technique, which only approximates the zero-temperature limit used in our PQMC simulations. Residual thermal contributions in the SSE data can therefore lead to slightly higher energy estimates compared to ground-state projection methods, probably due to the low-temperature approximation bias of SSE technique.

While further analysis would be necessary to confirm the generality of this observation, these findings provide strong evidence that the proposed hybrid strategy is a highly promising avenue for improving classical simulations of strongly correlated quantum systems.

Although the results discussed so far may suggest that self-learning PQMC performs poorly compared to data-assisted approaches, it is important to stress that this apparent inefficiency is not intrinsic to the method itself, but rather a consequence of the finite walker population used in practice. In the limit of an infinite number of walkers, PQMC is an unbiased ground-state projection algorithm and thus expected to converge exactly to the true ground-state energy.

To demonstrate this, in Fig. 6.9 we show the energy per atom as a function of the

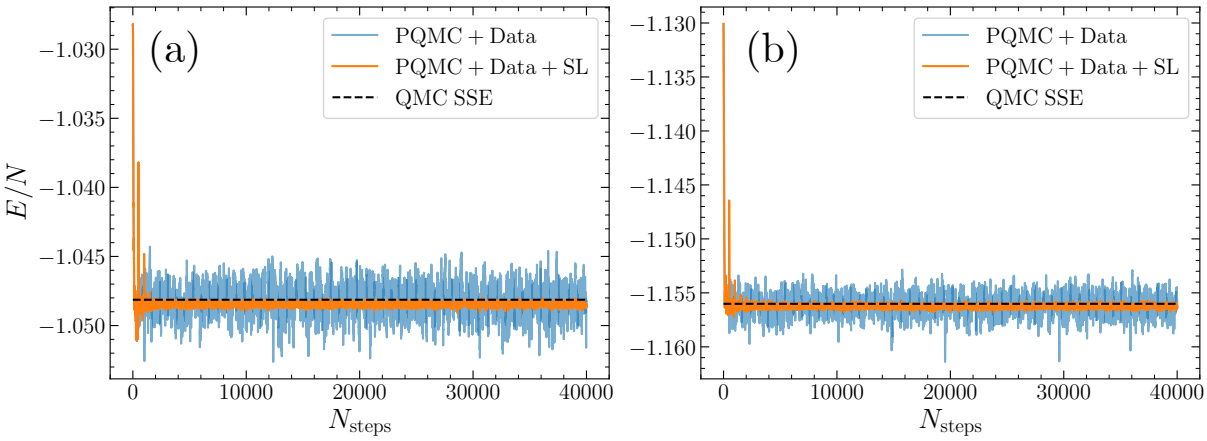

FIGURE 6.8: Monte Carlo energy as a function of PQMC steps N_{steps} for the last two detuning values, $\delta/\Omega = 2.466$ (a) and $\delta/\Omega = 2.702$ (b). The blue lines show the history of PQMC initialized with experimental data, while the orange lines correspond to self-learning PQMC initialized with the same data. The horizontal dashed black line indicates the energy obtained with SSE QMC. The PQMC initialized with experimental data reproduces the SSE QMC benchmark on average, albeit with larger variance. In contrast, the self-learning PQMC initialized with experimental data samples configurations with lower variance and, notably, achieves average energies below the SSE QMC reference.

walker population N_w for a representative value of the detuning, $\delta/\Omega = 1.0507$, near the critial point. In this case, the guiding wave function is fixed to a well-optimized ansatz previously obtained from a separate self-learning PQMC simulation. As N_w increases, the estimated energy systematically approaches the benchmark value obtained from Stochastic Series Expansion (SSE) QMC, confirming the asymptotic correctness of the method.

This analysis highlights that the limitations observed for self-learning PQMC at finite N_w arise solely from population-control effects and the imperfect guiding function used during on-the-fly training. When a sufficiently accurate guiding wave function is provided, or when the walker population is increased, PQMC fully recovers its unbiased character and achieves quantitative agreement with exact or benchmark results. Therefore, rather than being a flawed approach, PQMC constitutes a robust and systematically improvable framework, whose performance can be enhanced either by improving the guiding function (e.g., through experimental data) or by increasing computational resources.

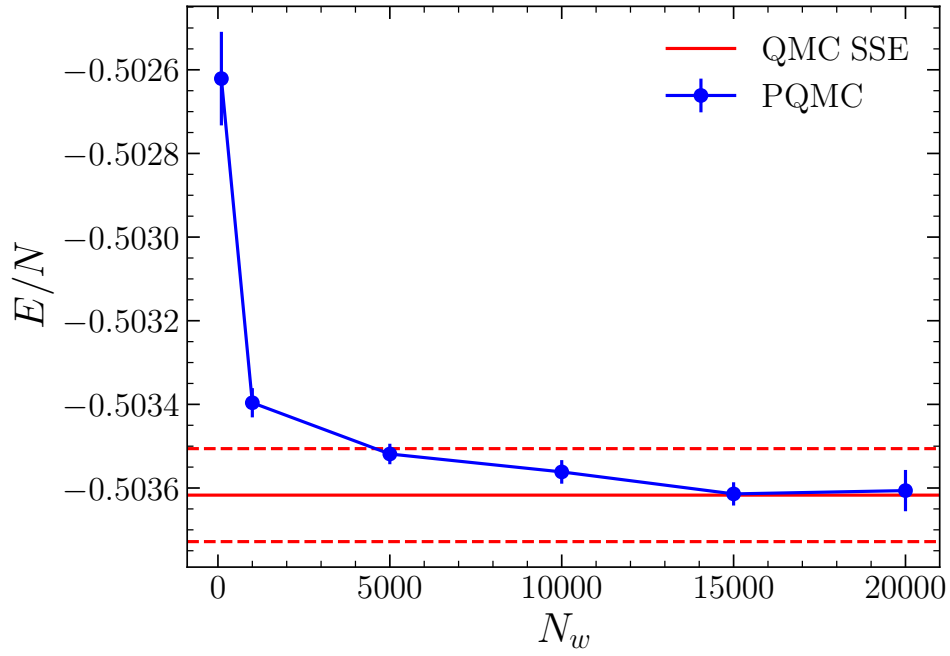

FIGURE 6.9: Energy per atom E/N as a function of the walker population N_w for PQMC simulations at $\delta/\Omega = 1.0507$. The solid red line indicates the reference energy obtained from SSE QMC and the red dashed lines represent the uncertainty. The PQMC results (blue points) converge systematically to the benchmark as N_w increases, confirming that the algorithm remains unbiased and exact in the limit of a large walker population.

Finally, we present the results for the staggered magnetization.

$$m_s = \frac{1}{L^2} \sum_{i=0}^{L-1} \sum_{j=0}^{L-1} (-1)^{i+j} \langle \sigma_{i,j}^z \rangle \tag{6.11}$$

This observable can be directly obtained from the Fourier transform of the magnetization in Eq. 6.10 by choosing the wave-vector values $(k_x, k_y) = (\pi, \pi)$.

In Fig. 6.10 we show the results. As in the previous case, we find that within the ferromagnetic phase all the methods considered yield satisfactory results. The most interesting regimes are, however, within the checkerboard phase and in the vicinity of the critical point. In these regions, simple variational methods fail as δ increases, leading to noticeable deviations from the expected behavior. In contrast, when experimental data are used to guide PQMC either alone or in combination with self-learning the results remain highly accurate and closely match those obtained with the stochastic series expansion approach. Furthermore, we note that the PQMC+VMC protool

does not yield accurate results at large values of δ (i.e., deep inside the ferromagnetic phase), while purely variational approaches remain systematically inaccurate for the chosen set of simulation parameters.

These findings further reinforce the central message of this study: incorporating experimental data into classical simulation protocols significantly improves both accuracy and robustness, especially in challenging regimes where purely variational approaches are prone to fail.

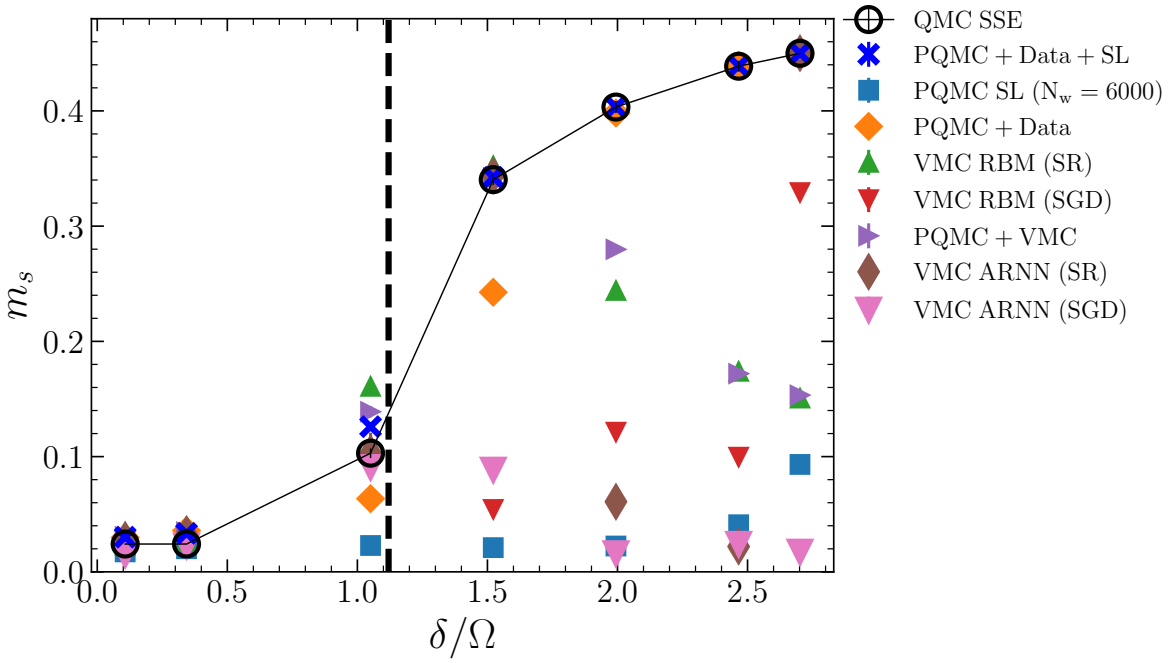

FIGURE 6.10: Staggered magnetization m_s as a function of the detuning δ/Ω for different simulation techniques. The dashed vertical line marks the critical detuning δ_c/Ω. In the ferromagnetic regime ($\delta < \delta_c$), all methods provide consistent results. However, within the checkerboard phase and close to the transition point, purely variational approaches deviate significantly as δ increases. In contrast, PQMC guided by experimental data, both in its standard and hybrid self-learning form, closely reproduces the SSE QMC benchmark, highlighting the enhanced accuracy and robustness of data-assisted methods.

6.5 Summary of the chapter

In this chapter, we have investigated the role of experimental data obtained from Rydberg atom platforms in enhancing classical simulation techniques, with particular emphasis on Projector Quantum Monte Carlo (PQMC). The Rydberg Hamiltonian is

mapped onto an effective Ising model, which provided a convenient framework for numerical analysis and for the comparison with established methods such as Stochastic Series Expansion (SSE) quantum Monte Carlo. This mapping allows us to characterize the checkerboard transition in square lattices and to extract critical exponents consistent with previous studies, thereby validating the numerical framework adopted throughout the work.

The benchmarking performed on different algorithms highlights both the strengths and limitations of the considered approaches. In particular, self-learning PQMC with fixed random walker popultation, while conceptually appealing due to its ability to improve the guiding wave function dynamically, suffers from scalability issues and becomes less reliable for larger system sizes. Variational Monte Carlo (VMC) based on restricted Boltzmann machines (RBMs) or autoregressive neural networks (ARNNs) provides acceptable results in some regimes, but exhibits systematic biases, particularly when the optimization is performed with stochastic gradient descent. More sophisticated protocols, such as stochastic reconfiguration, mitigate these issues but still show limitations when applied to phases characterized by strong correlations.

A central result of this study is that the inclusion of experimental data leads to a systematic improvement in performance. RBMs trained directly on spin configurations sampled from Rydberg platforms yield guiding states that significantly enhance the convergence of PQMC across the phase diagram. This outcome demonstrates that experimental configurations, despite being affected by noise and imperfections, capture essential correlations that are difficult to reproduce with purely classical optimization strategies.

Among the various protocols analyzed, the hybrid PQMC approach has shown the most robust performance. In this method, the simulation is initialized with an RBM trained on experimental data and subsequently refined during the PQMC run through self-learning. The numerical evidence indicates that this combined strategy maintains high accuracy throughout the different detuning regimes, including in the vicinity of the critical point where other techniques generally degrade. Furthermore, the reduced variance observed during the Monte Carlo sampling suggests a more stable and efficient convergence process. Interestingly, in some instances the average energy obtained with the hybrid protocol lies slightely below the reference values of SSE QMC, possibly due to residual fintie temperature biases in these simulations. Although such

results require further investigation to establish their generality, they represent a noteworthy indication of the potential of hybrid quantum-classical approaches.

The analysis of physical observables, such as the checkerboard order parameter and the staggered magnetization, further corroborates these findings. In the ferromagnetic regime, all considered methods yield comparable results. However, within the checkerboard phase and near the transition point, variational approaches alone are insufficient, while the inclusion of experimental data in PQMC ensures results in close agreement with the SSE QMC benchmarks. This reinforces the conclusion that the contribution of real experimental data is not merely auxiliary but constitutes a substantial enhancement to the reliability of classical simulations.

Overall, the results presented in this chapter demonstrate that the integration of experimental data from Rydberg atom platforms into PQMC provides a powerful tool for investigating strongly correlated quantum systems. By exploiting the physical information encoded in real configurations, one can overcome some of the limitations inherent in classical simulations, thereby enabling the study of larger systems and more complex regimes. The methodological framework developed here highlights a promising direction for future research, particularly in view of its possible extension to more expressive neural-network quantum states and to experimental platforms of increasing scale and tunability.

Chapter 7

Study of the energy gaps in disordered systems via Exact Diagonalization

In this chapter, we present a detailed numerical investigation of the low-energy spectrum of paradigmatic disordered quantum spin models in a transverse field, focusing on the fully-connected quantum Sherrington-Kirkpatrick (SK) model and the short-range two-dimensional quantum Edwards-Anderson (EA) model, already introduced in Chap. 3. Our goal here is to characterize the finite-size scaling and statistical properties of the even and odd parity energy gaps obtained via large-scale Exact Diagonalization (ED).

The chapter is organized around two main objectives. First, we introduce the numerical method in detail, an optimized ED algorithm exploiting Hamiltonian sparsity and GPU-accelerated Lanczos iterations. Second, we apply the ED methodology to the SK and EA models, to analyze the scaling of their characteristic gap scales and the probability distribution of the logarithm of even and odd gaps across disorder realizations. These results reveal strong finite-size fluctuations and broad log-normal–like distributions, highlighting the intrinsic variability of the low-energy spectrum in disordered quantum systems.

While the present chapter focuses entirely on exact diagonalization data, the next chapter will extend this analysis to significantly larger systems by employing the neural PQMC. There, we will compare our findings with recent results in the literature suggesting that the odd-parity gap of the quantum EA model may exhibit super-algebraic scaling, i.e., decay faster than any polynomial in the system size [27]. In addition to the average scaling behavior, we also compute the gap ratio statistic r_n as a diagnostic tool for level correlations and potential ergodicity breaking in the SK

model. This provides complementary information to the gap distributions, enabling a distinction between ergodic and nonergodic regimes within the disordered quantum phase. A key technical contribution of this work is the development of a dedicated, high-performance ED code that exploits the sparsity of the Hamiltonian matrix and leverages GPU acceleration to optimize the Lanczos algorithm. This implementation enables simulations of up to $N = 26$ spins for the fully connected SK model and $N = 25$ spins for the two-dimensional EA model on modern HPC clusters. We also document several design aspects of this code, which may serve as a useful reference for future large-scale numerical studies of disordered quantum systems.

7.1 Low-energy spectrum of quantum many-body systems

The low-energy spectrum of a quantum many-body system is a key quantity that governs its fundamental properties, from ground state characteristics to thermodynamic behavior. Of particular interest is the energy gap, which is the energy difference between the ground state and the first excited state. The behavior of this gap is crucial for understanding quantum phase transitions and the efficiency of quantum adiabatic optimization algorithms. For a system undergoing a second-order quantum phase transition, the energy gap closes at the critical point, which represents a major bottleneck for adiabatic computation [166]. Indeed, according to the adiabatic theorem, the annealing time required to remain in the ground state scales inversely with the square of the minimum gap,

$$T_{\text{anneal}} \propto \frac{1}{\Delta_{\text{min}}^2}. \tag{7.1}$$

A comprehensive understanding of gap scaling is therefore essential for evaluating the performance and potential speedup of modern quantum devices such as quantum annealers.

The Hamiltonian of the systems under investigation is of the general form:

$$H = -\sum_{i,j} J_{ij}\sigma_i^z\sigma_j^z - \Gamma\sum_i \sigma_i^x, \tag{7.2}$$

where J_{ij} are random interaction strengths and Γ is the external transverse field controlling quantum fluctuations. We investigate the following models:

- Quantum Sherrington-Kirkpatrick (SK) model: all-to-all connectivity, with couplings J_{ij} drawn from a normal distribution $\mathcal{N}(0, N^{-1/2})$.
- Quantum Edwards–Anderson (EA) model: short-range interactions on a two-dimensional square lattice, with couplings J_{ij} sampled from $\mathcal{N}(0, 1)$.

The SK model represents an infinite-range mean-field limit, while the EA model captures the effects of finite connectivity and spatial locality of interactions. Comparing their low-energy spectral properties allows us to identify universal versus model-dependent aspects of the gap scaling in disordered quantum systems. Both Hamiltonians exhibits a Z_2 symmetry, generated by the parity operator

$$P = \prod_i \sigma_i^x. \tag{7.3}$$

which acts globally by flipping all spins in the σ^z basis, i.e.

$$P|\uparrow\uparrow\downarrow\ldots\rangle = |\downarrow\downarrow\uparrow\ldots\rangle. \tag{7.4}$$

Since the parity operator P commutes with the Hamiltonian, $[P, H] = 0$, the eignestates of H are eigenstates of P. If we apply the parity operator P to a basis state $|\psi\rangle$, we get

$$P|\psi\rangle = \lambda|\psi\rangle, \tag{7.5}$$

with $\lambda = \pm 1$. The eigenvalue λ defines the parity sector of the state: states with $\lambda = +1$ belong to the even sector, while those with $\lambda = -1$ belong to the odd sector. This classification of eigenstates into even and odd parity sectors is a general consequence of Z_2 symmetry, and therefore applies not only to the transverse-field Ising model but to any Hamiltonian with this type of discrete symmetry.

In both SK and EA models, it is empirically observed that, for relatively small values of the transverse field Γ, the absolute ground state belongs to the even parity sector of the Hamiltonian. Moving up in the spectrum, one typically encounters a sequence of states in the odd sector, followed by the first even excited state. It is therefore of general interest to study both the even–even gap, defined as the energy difference between the ground state and the first excited state in the even sector, denoted as Δ_e, and the even-odd gap, which usually corresponds to the fundamental gap between

the ground state and the first globally excited state. We will refer to this latter as Δ_o. In the next section we provide the detailed explanation of the Exact Diagonalization algorithm that allow us to investigate the spectral properties of systems up to $N = 26$ spins in en-premise HPC cluster.

7.2 Methods: Exact Diagonalization of sparse Hamiltonian

The most reliable method for obtaining the energy spectrum of a quantum system is Exact Diagonalizaton (ED) of the associated Hamiltonian. However, the exponential scaling of the Hilbert space dimension $D = 2^N$ with the number of spins N is the major drawback of this technique. For instance, a system with $N = 26$ spins results in a Hilbert space of $2^{26} \approx 67 \times 10^6$ states. To overcome the overwhelming computational challenge posed by such large dimensions, both in therms of memory and processing time, we implement a highly optimized computational framework. This approach allow us to push the practical limit of ED to systems with up to $N = 26$ spins in modern HPC cluster, namely LEONARDO HPC Booster partition provided by CINECA. Our methodology is based on three main computational strategies.

7.2.1 Z_2 symmetry exploitation

The first and most effective step to reduce the problem size is to exploit the global Z_2 symmetry of the Hamiltonian, generated by the parity operator $P = \prod_i \sigma_i^x$. As we said before, since $[P, H] = 0$, the total Hilbert space can be divided into two orthogonal subspaces corresponding to the eigenvalues $+1$ that label the even-parity sector D_{even} and -1 that label the odd-parity sector D_{odd}. This decomposition halves the dimension of the matrix that must be diagonalized, as $D_{\text{even}} = D_{\text{odd}} = 2^{N-1}$. This block-diagonalization transforms the diagonalization of a $2^N \times 2^N$ matrix into two separate diagonalizations of two $2^{N-1} \times 2^{N-1}$ matricis, leading to a significant reduction in computational cost.

7.2.2 Sparse Matrix Representation

The Hamiltonian matrix H remains sparse, meaning most of its entries are zero. To handle these large matricese efficiently, we use the *QuSpin* library [167]. *QuSpin* is a Python library specifically designed for quantum dynamics and exact diagonalization of quantum many-body systems. It allows to efficiently constructs the Hamiltonian as a sparse matrix in the Compressed Sparse Column (CSC) format. The library is crucial for automatically applying the symmetry constraint; the `spin_basis_general` function, when called with the `pblock` argument, ensures that the resulting sparse matrix is already restricted to the desired parity sector ($H_{\text{even}}, H_{\text{odd}}$). This dramatically reduces the memory footprint compared to storing a dense matrix, which would quickly exceed the available GPU memory.

Since our focus is only on the low-energy spectrum, namely the ground state and the first few excited states, we do not need to calculate all 2^N eigenvalues. Instead, we use an iterative method called Lanczos , which aims to find the extreme lower and upper eigenvalues of the spectrum.

7.2.3 GPU-accelerated Lanczos iterative method

To find the lowest k eigenvalues and eigenvectors of the large, sparse Hamiltonians $H_{\text{even}}, H_{\text{odd}}$, we employ the Lanczos algorithm. This is an iterative projection technique that is optimal for finding the extreme (highest or lowest) eigenvalues of sparse, Hermitian matrices. The complexity of the Lanczos method is primarily determined by the cost of sparse matrix-vector multiplication, which scales as $O(N \cdot n_z)$, where n_z is the number of non-zero elements in the matrix, making it far more efficient than dense diagonalization.

To assess this task we leverage the power of modern GPUs by using *CuPy*, a NumPy-compatible library that exploits the parallel processing capabilities of NVIDIA GPUs [168]. CuPy implements arrays and sparse matrix structures directly on the GPU memory and provides GPU-accelerated versions of common linear algebra routines. Specifically, we use the function
`cupyx.scipy.sparse.linalg.eigsh`, a GPU implementation of the sparse eigenvalue

solver (Lanczos/Arpack). This process overcomes the most computationally demanding part of the ED, the iterative matrix multiplication, from the CPU to the GPU, resulting in an order of magnitude speedup essential for reaching large system sizes and for averaging over many instances of the random Hamiltonians.

For generic quantum spin systems without translational symmetry, the memory requirement ($O(2^N)$) is the primary constraint, often outweighing the time constraint. In the literature focused on one-dimensional systems with only local or short-range interactions, system sizes can sometimes reach $N \approx 30$ by exploiting additional symmetries and conservation laws (such as particle number or momentum conservation) and specialized memory architectures [167]. However, for fully connected or long-range disordered systems, which lack translational symmetry and have highly complex sparse matrix structures, the achievable limit is considerably lower. Our target size of $N = 26$ for the infinite-range SK model with parity simmetry is at the high end of what considered feasible for extensive statistical analysis in the context of many-body physics. Prior studies on the quantum SK model using ED generally focused on system sizes up to $N = 20$ or $N = 22$ for statistical reliability [169, 170]. Approximations like the Hartree-Fock or Configuration Interaction (CI) methods have been used to push the limit higher (e.g., up to $N = 90$), but these are approximated methods that sacrifice accuracy for size [169]. By contrast, achieving $N = 26$ using the exact Lanczos method provides highly reliable spectral data, which is crucial for distinguishing between competing theories of the gap closing behavior. This achievement highlights the efficacy of combining symmetry exploitation with GPU-accelerated iterative solvers on state-of-the-art HPC platforms.

7.2.4 Implementation details

The implementation of the ED algorithm for the quantum SK model is developed in Python, combining high-level libraries for quantum many-body physics with GPU-accelerated numerical routines.

The first step of the simulation consists in generating the interaction graph of the SK model. For each spin i, we construct the neighborhood structure that encodes all possible couplings to spins $j > i$. The random couplings J_{ij} are sampled from a Gaussian distribution with zero mean and variance $1/\sqrt{N}$, consistent with the infinite-range

scaling of the SK model. Once the couplings are generated, the Hilbert space is divided into two distinct sectors corresponding to the eigenvalues of the global Z_2 parity operator. This decomposition is handled automatically by the *QuSpin* library, which provides the function `spin_basis_general` to generate the spin basis restricted to a given symmetry block. In practice, two bases are constructed: one for the even-parity sector and one for the odd-parity sector. The Hamiltonian is then built separately in each basis, leading to two independent sparse matrices H_{even} and H_{odd}.

The Hamiltonian itself is composed of two contributions: the Ising interaction term, defined as $-\sum_{i<j} J_{ij}\sigma_i^z\sigma_j^z$, and the transverse field term, $-\Gamma\sum_i \sigma_i^x$. These terms are encoded as operator lists, which QuSpin uses to efficiently build the sparse Hamiltonian in Compressed Sparse Column (CSC) format. This representation is particularly suited for iterative solvers, as it minimizes both memory usage and the cost of sparse matrix vector multiplications.

The central computational task is the determination of the few lowest eigenvalues of H_{even} and H_{odd}. To this end, the code transfers the sparse matrices to the GPU, making use of the *CuPy* library, which provides GPU-compatible data structures and linear algebra routines. The eigenvalue problem is then solved using the Lanczos algorithm, implemented as
`cupyx.scipy.sparse.linalg.eigsh`, which allows us to compute only the lowest few eigenvalues in each sector. Typically, three eigenvalues are extracted per sector, which is sufficient to evaluate the ground state energy density and the lowest excitation gaps.

The output of the simulation consists of the ground-state energy in both parity sectors, E_0^{even} and E_0^{odd}, together with the first two excited levels. From these quantities, the code computes the parity gaps, defined as $\Delta_e = E_1^{\text{even}} - E_0^{\text{even}}$ and $\Delta_o = E_0^{\text{odd}} - E_0^{\text{even}}$, so the fundamental gap. For each disorder realization, specified by a random seed, the couplings J_{ij} and the computed spectra are stored in log files. This guarantees reproducibility and allows for large-scale statistical sampling with many instances of the disorder.

In summary, the implementation combines three essential ingredients: (i) exact construction of the disordered Hamiltonian with symmetry distinction, (ii) efficient sparse matrix representation, and (iii) GPU-accelerated iterative diagonalization. Together, these design choices enable exact diagonalization of the quantum SK model for

system sizes up to $N = 26$ in sufficiently short times to allow solving large ensambles of disordered realizations. This pushes the boundary of what is computationally feasible in this class of fully connected disordered quantum models.

7.3 Results

7.3.1 Energy gap distribution in random Ising chain

Before presenting the results obtained via Exact Diagonalization for the SK and EA models, it is instructive to briefly discuss a simpler disordered system that can be solved exactly using fermionization techniques. In this subsection, we present the distribution of odd energy gaps for the one-dimensional transverse field Ising model with disordered ferromagnetic couplings in open boundary conditions. This analysis follows the approach of Ref. [171], who reported similar results for the same system but with a different choice of coupling distribution. In our case, the Hamiltonian 7.2 includes nearest-neighbor interactions J_i drawn from a uniform distribution in the interval $[0.75, 1.25]$. This choice allows us to determine the critical transverse field in the thermodynamic limit using the analytical expression

$$\Gamma_c = \exp(\langle \log J \rangle)\,, \tag{7.6}$$

where the average is taken over the coupling distribution. For the present distribution, this formula yields a critical point at $\Gamma_c \approx 0.989$. In Fig. 7.1 we show the histograms of the logarithmic gaps of this model for different system sizes (up to $N = 200$) and 50000 couplings realizations at $\Gamma = 1.0$, close to the critical point. In addition, in this case, we see that the instances with low values of gaps increase with the system sizes and the distribution is broader as the thermodynamic limit is approached.

7.3.2 Sherrington-Kirkpatrick model

Here, we present the results obtained from the study of the energy gap in the quantum SK model using Exact Diagonalization. In the first part, we focus on the scaling of the even gap Δ_e and the fundamental gap Δ_o as a function of the system size. These data provide promising evidence for addressing the question of how the gap closes at the

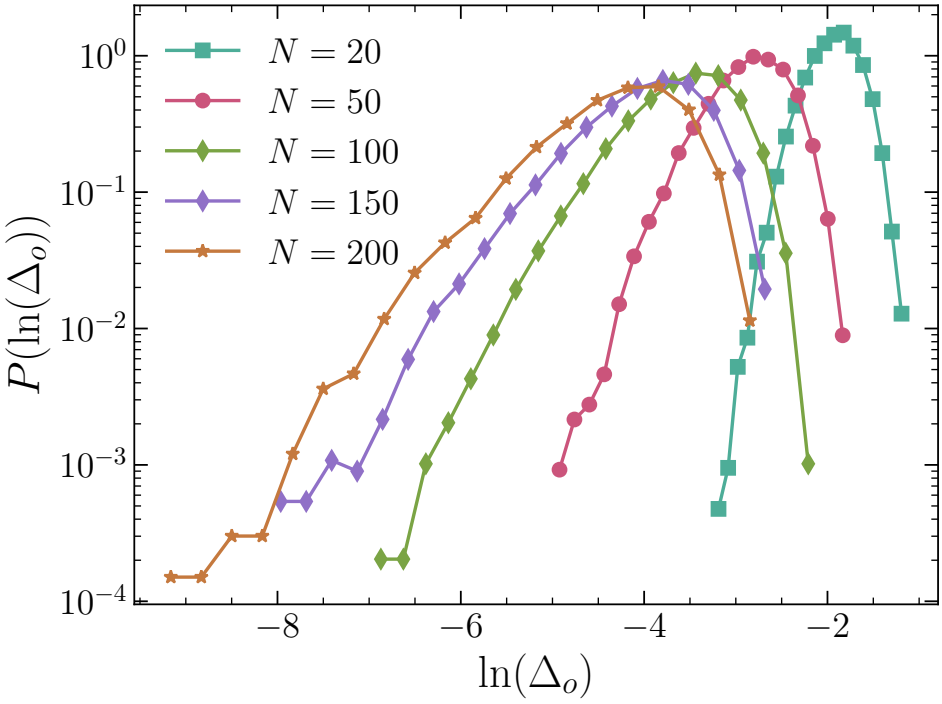

FIGURE 7.1: Distribution of the logarithm of the odd energy gaps, $\ln(\Delta_o)$, for the one-dimensional random transverse-field Ising model with open boundary conditions and ferromagnetic couplings $J_i \in [0.75, 1.25]$. The data correspond to different system sizes $N = 20, 50, 100, 150, 200$ and are obtained from 5×10^4 disorder realizations at $\Gamma = 1.0$, close to the critical point $\Gamma_c \simeq 0.989$. As the system size increases, the distributions broaden and develop a heavier tail toward small gap values, indicating the growing weight of rare low-energy excitations as the thermodynamic limit is approached.

critical point of the transition between the paramagnetic and spin-glass phases, which for this model is well established and corresponds to $\Gamma_c = 1.5$.

We then extend the analysis to different points of the phase diagram, considering both the spin-glass phase and the paramagnetic phase. Finally, we present preliminary results of an analysis focusing on the phenomenon of Many-Body Localization (MBL). For this part, we do not rely on the ED code described in the previous section, since the full spectrum of the Hamiltonian is required. Instead, we analyze systems of size up to $N = 16$ by computing the level-spacing statistics of the energy spectrum. This diagnostic is particularly suited to distinguish between the ergodic regime, where level statistics follow the Wigner-Dyson distribution, and the localized regime, characterized by Poissonian statistics [172]. The comparison between these two regimes provides insight into the possible emergence of localization phenomena in the disordered SK model.

Gap scaling with system size

Our analysis of the gap scaling as a function of system size focuses on the region of the phase diagram around the quantum phase transition. In particular, we concentrate on the critical point $\Gamma_c = 1.5$, since it represents the most relevant regime where correlations are strongest and signatures of both the spin-glass and the paramagnetic phases coexist. We simulate the SK Hamiltonian for system sizes ranging from $N = 15$ to $N = 26$, averaging over a large number of disorder realizations N_r, specifically from $N_r \simeq 2000$ up to $N_r \simeq 10000$, depending on system size.

In Fig. (7.2), we show the scaling of the average even-parity gap Δ_e as a function of N in log-log scale. The data points align along a straight line in this representation, indicating a polynomial scaling of the gap. Indeed, by fitting the results with a power-law function of the form $f(N) = aN^{-\theta}$, we obtain an exponent $\theta = 0.335(2)$, corresponding to a behavior $\propto N^{-1/3}$, which is consistent with the scaling suggested in Ref. [166] based on a theory for the quantum Hopfield network.

Similarly, in Fig. (7.3), we display the fundamental (odd-parity) gap Δ_o, averaged over the same number of disorder realizations. Once again, the log-log plot shows the same linear trend as in the even sector. A power-law fit using the same functional form yields $\theta = 0.332(4)$, fully compatible with the even-parity case within the error bars.

In addition to the scaling of the mean gap, we examine the full probability distribution of the low-energy gaps $P(\Delta)$ at the critical point. In Fig. 7.4, we show the distributions of the even-parity gap Δ_e (panel (a)) and the odd gap Δ_o (panel (b)) for system sizes $N = 18, 22, 26$. Both distributions display a broad shape, with significant weight at small values of the gap, indicating that rare disorder realizations can produce particularly small gaps. This feature is consistent with the physical picture of criticality in disordered systems, where sample-to-sample fluctuations are large and dominate the low-energy behavior. For odd gaps, as N increases, the distributions become sharper around their typical scale but retain a long tail toward small gaps, which is of particular relevance for adiabatic dynamics: such rare small-gap instances represent the hardest cases for quantum annealing algorithms. Overall, the decay of the tails towards very small gap values does not appear to be particularly severe for the system sizes considered here. This is likely a consequence of the limited range of accessible sizes in the present analysis. A more comprehensive characterization of the

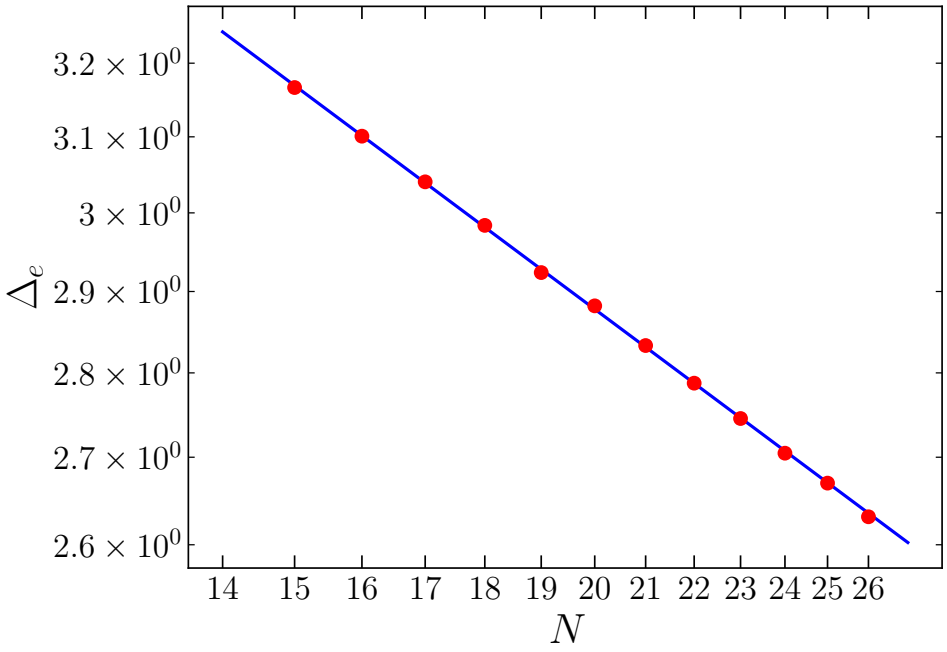

FIGURE 7.2: Even-parity gap Δ_e averaged over 2000-10000 disorder realizations, plotted as a function of system size N in log-log scale at the critical point. The data points align along a straight line, indicating polynomial scaling of the gap. The solid blue line shows a power-law fit of the form $f(N) = aN^{-\theta}$, yielding an exponent $\theta = 0.335(2)$, consistent with $\Delta_e \propto N^{-1/3}$.

small gap behavior will be presented in the next chapter, where we employ PQMC simulations to access significantly larger system sizes.

In order to further investigate the behavior of the low energy gaps away from criticality, we performed the same finite-size analysis for several values of the transverse field in the vicinity of the critical point, namely $\Gamma = 1.3, 1.4, 1.6, 1.7$. In all cases the scaling remains polynomial with system size, $\Delta \propto N^{-\theta(\Gamma)}$, but the extracted exponents $\theta(\Gamma)$ show a characteristic dependence on the transverse field. When plotting $\theta(\Gamma)$, corresponding to the even and odd parity sectors, we observe that the two curves intersect at $\Gamma_c \simeq 1.5$, where both exponents converge to $\theta \simeq 1/3$. This crossing point provides an independent numerical signature of the quantum critical point of the SK model.

Physically, this behavior reflects the fact that, away from criticality, the finite-size gaps are dominated by different mechanisms in the two parity sectors (with the odd sector typically controlling the fundamental excitation), whereas at the transition the critical collective fluctuations impose the same universal scaling on both sectors.

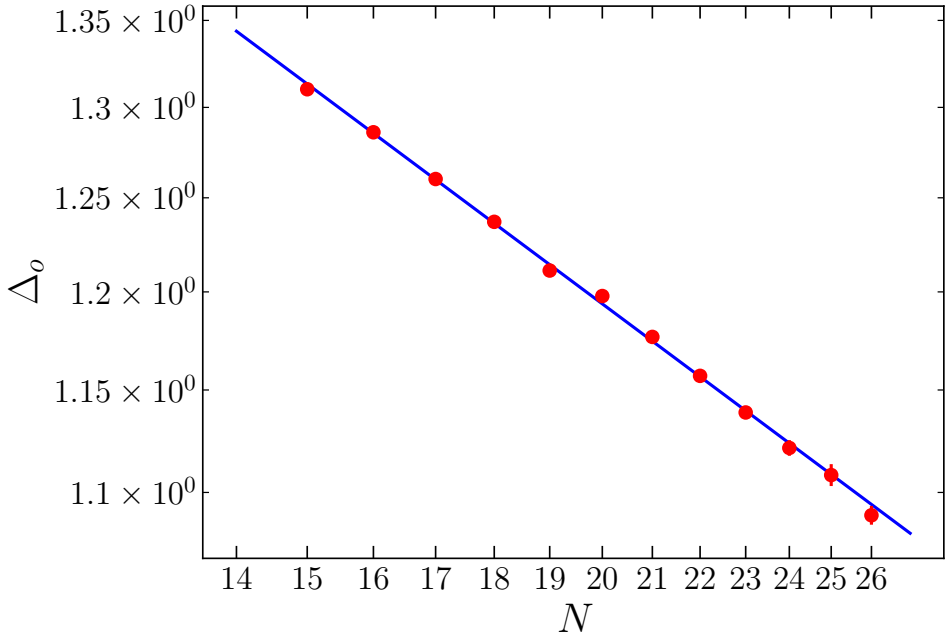

FIGURE 7.3: Fundamental (odd-parity) gap Δ_o averaged over 2000-10000 disorder realizations, plotted as a function of system size N in log–log scale at the critical point. The same linear trend is observed as in the even sector. A power-law fit of the form $f(N) = aN^z$ gives $\theta 0.332(4)$, fully consistent with the even-parity case within error bars, confirming the universal $N^{-1/3}$ scaling of the critical gap.

To further verify the scaling behavior of the low energy gaps, we examined how the rescaled quantity $\Delta_e N^{1/3}$ varies with the transverse field Γ. If the exponent $\theta = 1/3$ correctly captures the critical scaling, the curves corresponding to different system sizes should intersect at the critical point. As shown in Fig. 7.6, this is indeed the case: all curves cross at a common value of $\Gamma_c \simeq 1.5$. This crossing provides a clear and independent confirmation of the critical exponent extracted from the finite-size scaling analysis, as well as a direct numerical signature of the quantum phase transition in the SK model.

Many-body localization in SK model

Beyond the properties of the ground state and low-lying excitations, the statistical properties of the entire energy spectrum provide important insights into the quantum dynamics of disordered systems. A key phenomenon in this context is Many-Body Localization (MBL), which prevents the system from thermalizing. To investigate the emergence of MBL in the quantum Sherrington-Kirkpatrick model, we analyze the

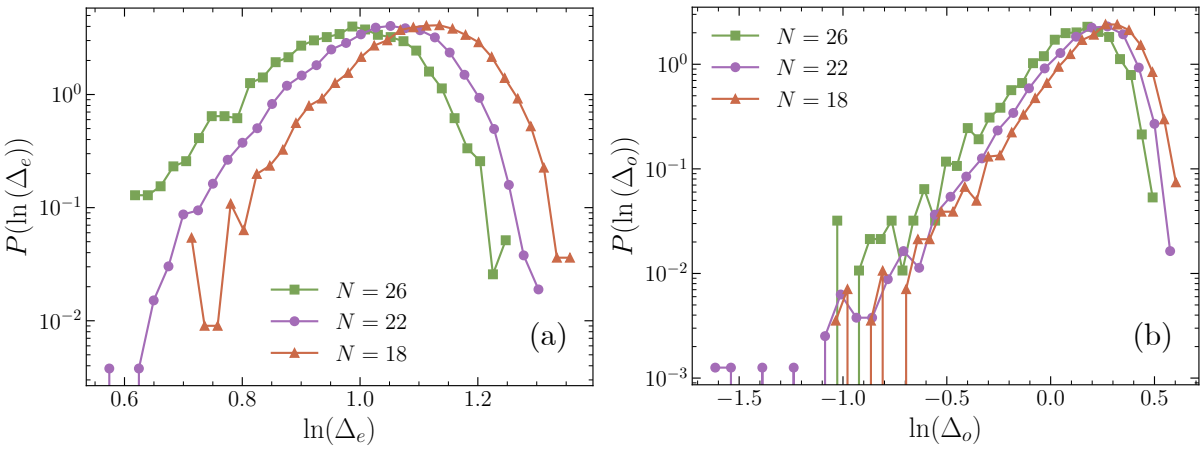

FIGURE 7.4: Probability distributions of the log of low-energy gaps at the critical point $\Gamma = 1.5$. (a) Log of even gap $\ln(\Delta_e)$, (b) log of fundamental (odd-parity) gap $\ln(\Delta_o)$, in log-scale for system sizes $N = 18, 22, 26$, for ~ 1000 to ~ 5000 realization of disorder. Both distributions are broad and exhibit significant sample-to-sample fluctuations, with tails extending to small gaps.

statistics of energy level spacings, a powerful diagnostic for distinguishing between thermal (ergodic) and localized (non-ergodic) phases.

In a generic, isolated quantum many-body system, interactions are expected to drive the system towards thermal equilibrium. This behavior is encapsulated by the Eigenstate Thermalization Hypothesis (ETH) [173]. According to the ETH, individual high-energy eigenstates of a chaotic (ergodic) Hamiltonian already encode the properties of a statistical ensemble. Consequently, the expectation value of any local observable in such an eigenstate is equivalent to its microcanonical average, meaning the system acts as its own heat bath and thermalizes. A key signature of this quantum chaos is found in the energy spectrum: the energy levels exhibit strong correlations and repulsion, leading to level spacing statistics described by the Wigner-Dyson distribution. For Hamiltonians with time-reversal symmetry, such as the model under study, this corresponds to the Gaussian Orthogonal Ensemble (GOE). An interested reader can find a detailed explanation in Ref. [174]. However, in the presence of strong disorder, interactions may fail to thermalize the system. This leads to the Many-Body Localized (MBL) phase, a quantum phenomenon where the system remains localized and retains memory of its initial conditions [175]. The MBL phase represents a breakdown of the ETH; it is a non-ergodic state of matter where transport is absent, and entanglement growth is suppressed. Spectrally, the MBL phase is characterized by

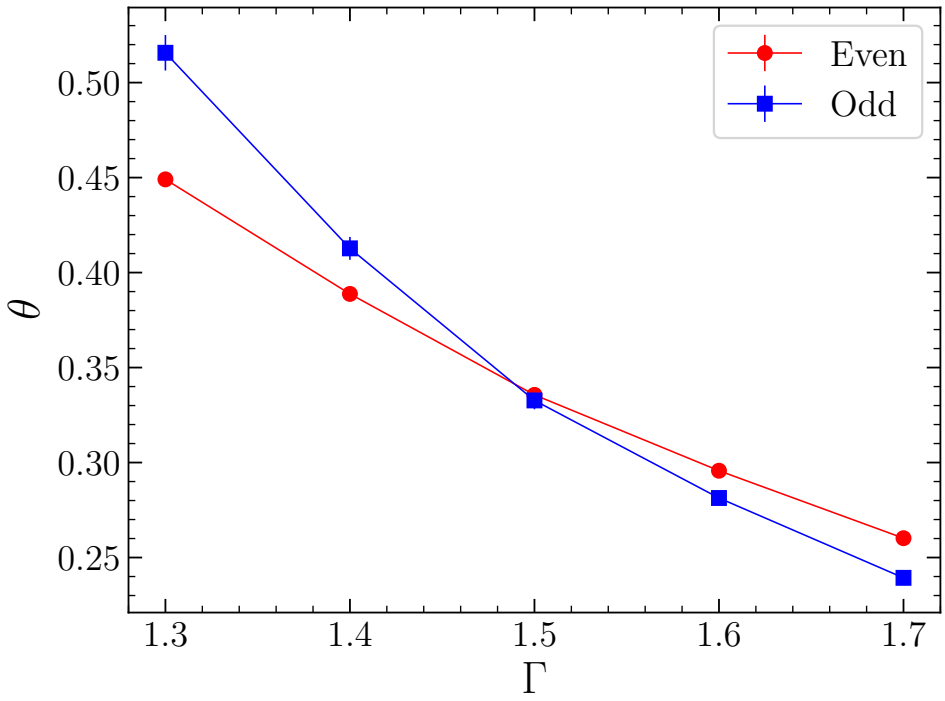

FIGURE 7.5: Scaling exponents θ extracted from the size analysis of the even-parity gap Δ_e (red) and the odd gap Δ_o (blue) as a function of the transverse field Γ. The two curves exhibit a clear crossing at $\Gamma_c \simeq 1.5$, where both exponents converge to $\theta \approx 1/3$. This crossing provides an independent numerical signature of the quantum critical point of the SK model and indicates that, at criticality, the scaling of the low-energy excitations becomes universal across parity sectors.

uncorrelated energy levels, akin to an integrable system. The level spacing statistics therefore follow a Poissonian distribution, signal of the absence of chaos and level repulsion [174].

A standard metric for distinguishing these two phases is the dimensionless ratio of consecutive energy level spacings, or gap ratio, defined as

$$r_n = \frac{\min(\delta_n, \delta_{n+1})}{\max(\delta_n, \delta_{n+1})}, \tag{7.7}$$

where $\delta_n = E_{n+1} - E_n$ is the gap between adjacent energy eigenvalues. For an ergodic system described by the GOE, the disorder-averaged gap ratio is $\langle r \rangle \approx 0.5307$, while for a localized system with Poissonian statistics, it is $\langle r \rangle = 2\ln 2 - 1 \approx 0.3853$ [173]. In Fig. 7.7, we present the average gap ratio $\langle r(\varepsilon) \rangle$ as a function of the normalized energy density $\varepsilon = (E - E_{\text{min}})/(E_{\text{max}} - E_{\text{min}})$ for various system sizes from $N = 8$

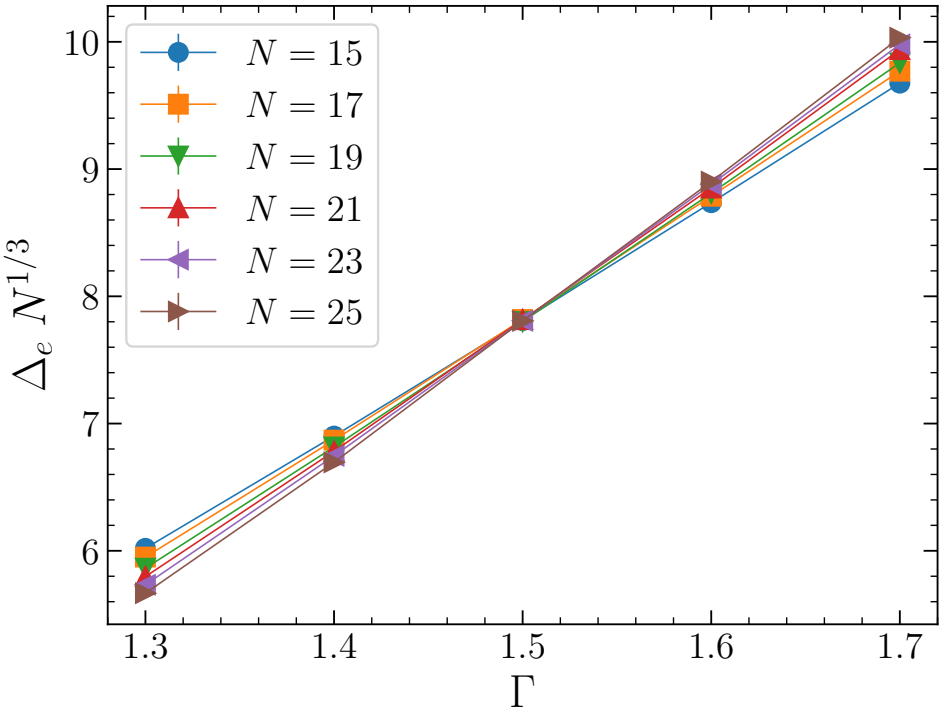

FIGURE 7.6: Crossing of the rescaled even-parity gaps $\Delta_e N^{1/3}$ as a function of the transverse field Γ for different system sizes. All curves intersect at $\Gamma_c \simeq 1.5$, confirming the critical exponent $\theta 1/3$ and providing an independent numerical signature of the quantum phase transition in the SK model.

to $N = 16$. The panels show the evolution of the spectral statistics as the transverse field Γ is tuned from a moderately high value down to the strongly interacting limit. Note that, in this case, the sizes under consideration are smaller compared with those in the previous analysis. The reason behind this restriction relies in Eq. (7.7). In fact, to compute the full energy statistics, we need to diagonalize the full Hamiltonian. In this case, the Lanczos algorithm is not efficient.

In panel (a), at $\Gamma = 0.5$, the system exhibits clear signs of ergodicity. Across the bulk of the spectrum ($\varepsilon \approx 0.5$), the average gap ratio approaches the GOE value, particularly for larger system sizes. This indicates that the quantum fluctuations induced by the transverse field are strong enough to drive the system into a thermal phase consistent with the ETH. The slight decrease of $\langle r \rangle$ at the spectral edges is a common feature in finite-sized systems with few-body interactions. As the transverse field is reduced to $\Gamma = 0.2$ in panel (b), the system begins to show signatures of localization. While the center of the spectrum still tends towards the GOE limit, the value of $\langle r \rangle$ at the spectral

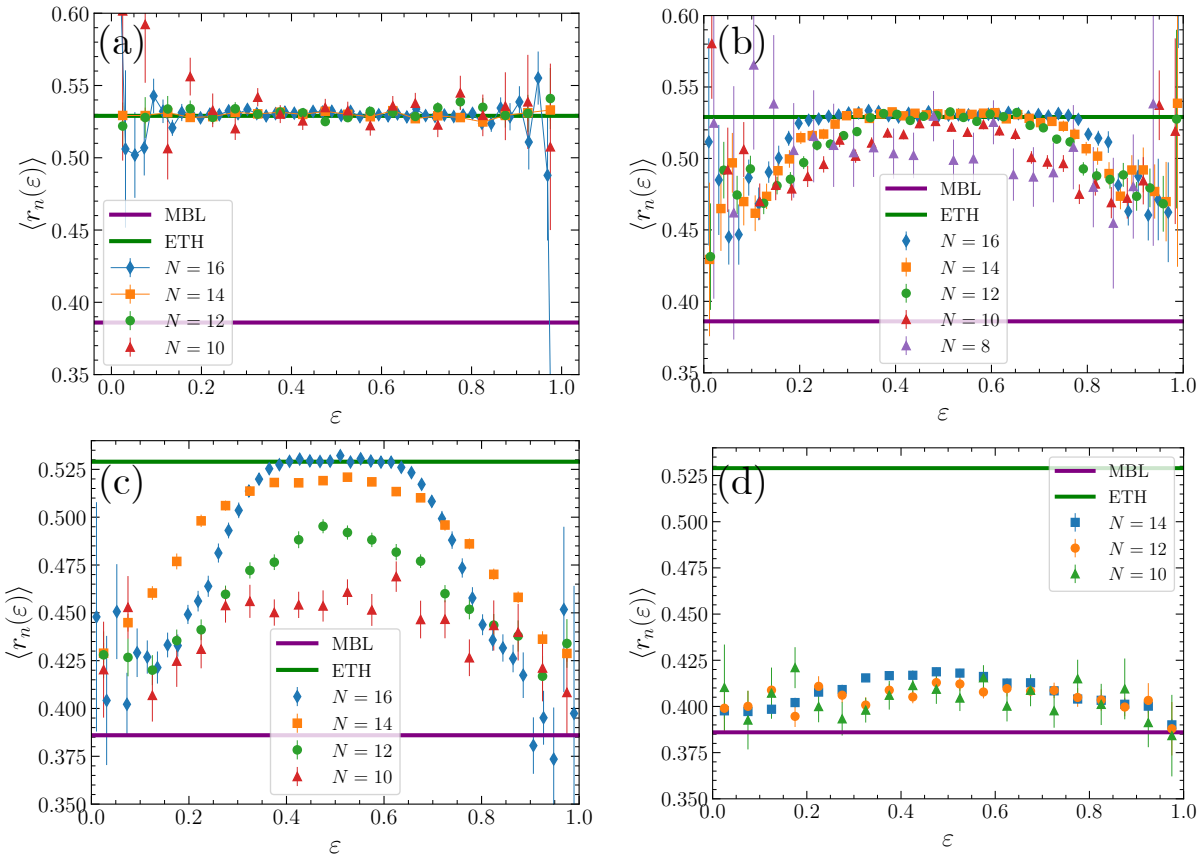

FIGURE 7.7: Average gap ratio $\langle r(\varepsilon)\rangle$ as a function of normalized energy density ε for system sizes $N = 8, 10, 12, 14, 16$. The horizontal dashed lines indicate the theoretical values for the Gaussian Orthogonal Ensemble (GOE, $\langle r\rangle \approx 0.53$) and the Poisson distribution ($\langle r\rangle \approx 0.39$). The panels show results for different transverse field strengths: (a) $\Gamma = 0.5$, (b) $\Gamma = 0.2$, (c) $\Gamma = 0.1$, and (d) $\Gamma = 0.005$. The data show a transition from an ergodic phase at larger Γ to a localized phase at small transverse field.

edges drops significantly towards the Poissonian value. This suggests the formation of many-body mobility edges, which separate localized states at the edges of the spectrum from delocalized thermal states in the middle. Furthermore, a distinct flow with system size becomes apparent: for a fixed energy density in the bulk, $\langle r\rangle$ decreases as N increases, signaling a tendency towards localization in the thermodynamic limit. Finally, in Fig. 7.7 (d), at a very small transverse field $\Gamma = 0.005$, the system is deep in the localized regime. The average gap ratio is very close to the Poisson value of ≈ 0.39 across the entire spectrum and for all system sizes studied. The weak dependence on system size suggests that the localization length is small and the system is firmly in the MBL phase. In this regime, the strong, disordered Ising interactions dominate, leading to the breakdown of ergodicity and the emergence of a non-thermal, localized state.

In summary, our analysis of the gap ratio statistics provides some indications for a many-body localization-delocalization transition in the quantum SK model. The transverse field Γ acts as a control parameter that drives the system from a fully localized MBL phase at small transverse field to a fully ergodic, thermal phase at large transverse field. For intermediate values of Γ, our results are consistent with the existence of many-body mobility edges, a hallmark of the MBL transition in finite-energy-density systems. This investigation complements our study of the ground-state quantum phase transition by revealing a rich dynamical phase structure across the entire many-body spectrum.

7.3.3 Edwards-Anderson model

Having established the critical gap scaling behavior for the fully connected SK model, we now turn to the short-range disordered system, the two-dimensional Edwards-Anderson model. Here, spatial locality and geometric frustration play a central role, introducing significant differences in the low-energy spectrum compared to the infinite-range case.

We study system sizes from $N = 4$ to $N = 25$ spins arranged on a square lattice with periodic boundary conditions. For each system size, we compute both the even and odd parity gaps over $\sim 10^3$ disorder realizations at the critical transverse field $\Gamma_c = 1.98$, referring to the results obtained in [32].

Gap scaling with system size

In Fig. 7.8 we show the average even-parity gap Δ_e as a function of linear size of the system L in log-log scale at the quantum critical point $\Gamma_c = 1.98$. The results exhibit a polynomial decay of the gap with increasing system size, consistent with the power-law behavior $\Delta_e(N) \propto N^{-z}$. A best-fit analysis yields an exponent $z = 1.073(3)$. A similar scaling is observed for the fundamental (odd-parity) gap Δ_o, shown in Fig. 7.9, with an exponent $z = 1.17(1)$, close to the even-sector exponent.

These results demonstrate that the gap closing in the EA model follows a polynomial law analogous to the SK case, confirming the robustness of the N^{-z} scaling behavior across disordered spin systems with very different connectivity structures.

Of course, for a deeper understanding of the size-scaling we need to address larger sizes, as discussed in the next chapter.

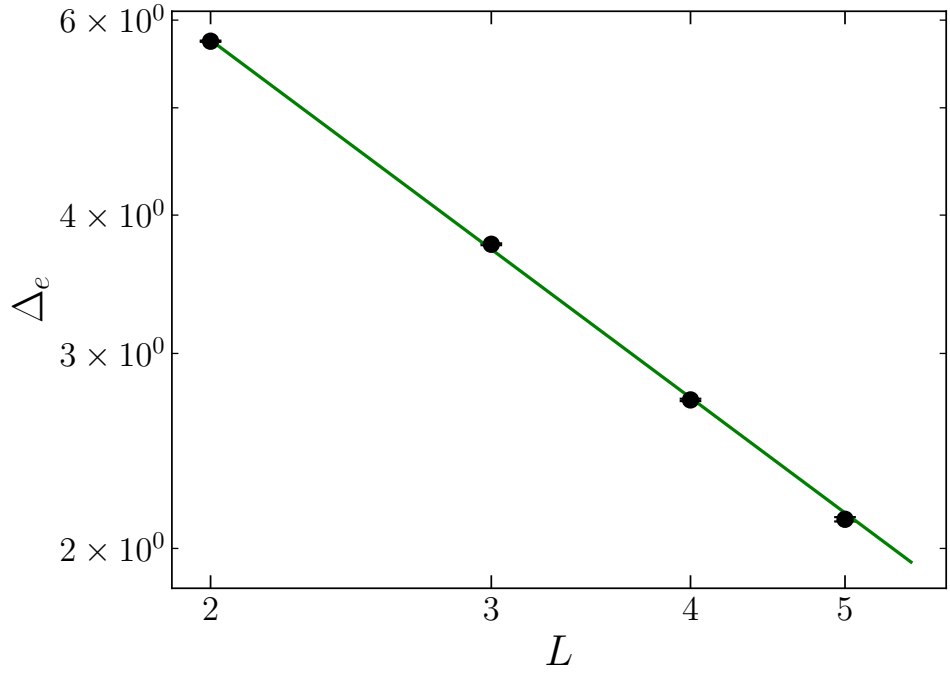

FIGURE 7.8: Even-parity gap Δ_e for the 2D Edwards–Anderson model as a function of the linear size L at the critical field $\Gamma_c \approx 1.98$. The power-law fit yields an exponent $z = 1.073(3)$, confirming polynomial scaling similar to the SK case.

Gap distribution

The full distributions of Δ_e and Δ_o at criticality are shown in Fig. 7.10. Both distributions are broad and asymmetric, with long tails extending toward small gap values. However, in contrast to the SK case, the EA model exhibits a more pronounced skewness and a stronger dependence on system size, reflecting the role of geometric frustration and spatial disorder. This might be due to the difference in the system sizes under investigation. As in the previous case, we show the logarithm of the gaps to enhance the behavior of the distribution for small values of Δ.

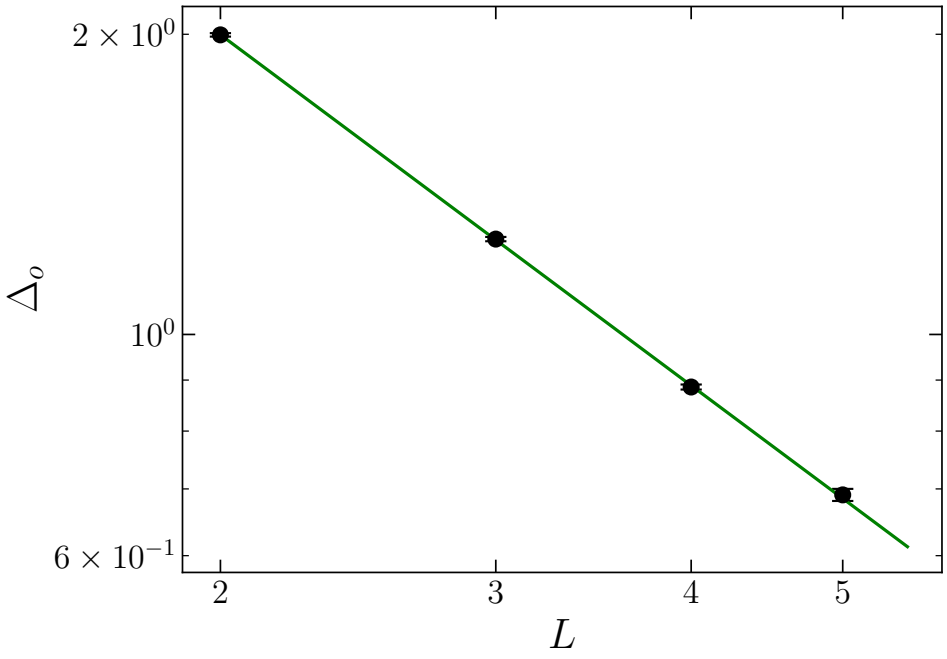

FIGURE 7.9: Odd-parity (fundamental) gap Δ_o for the 2D Edwards–Anderson model at $\Gamma \approx 1.98$. The fitted exponent $z = 1.17(1)$ agrees with the even sector, suggesting universal polynomial gap closing across parity sectors.

7.4 Conclusions

In this chapter, we have presented a comprehensive numerical study of the low-energy spectral properties of disordered quantum Ising models in a transverse field, focusing on the fully-connected Sherrington-Kirkpatrick model and the short-range two-dimensional Edwards-Anderson model. By combining Exact Diagonalization with GPU-accelerated Lanczos iterations, we were able to access system sizes up to $N = 26$ spins for the SK model and $N = 25$ spins for the EA model, pushing the computational boundaries of exact methods for this class of disordered many-body Hamiltonians.

For the SK model, our ED results shows a consistent polynomial scaling of both the even and odd parity gaps at the quantum critical point $\Gamma_c \simeq 1.5$, with an exponent $z \approx 1/3$, providing a sound confirmation to the arguments presented in Ref. [166]. The gap distributions show tails toward small gap values, indicating significant sample-to-sample fluctuations, though the decay of the tails remains moderate within the accessible size range. In addition, the analysis of level-spacing statistics provided signatures

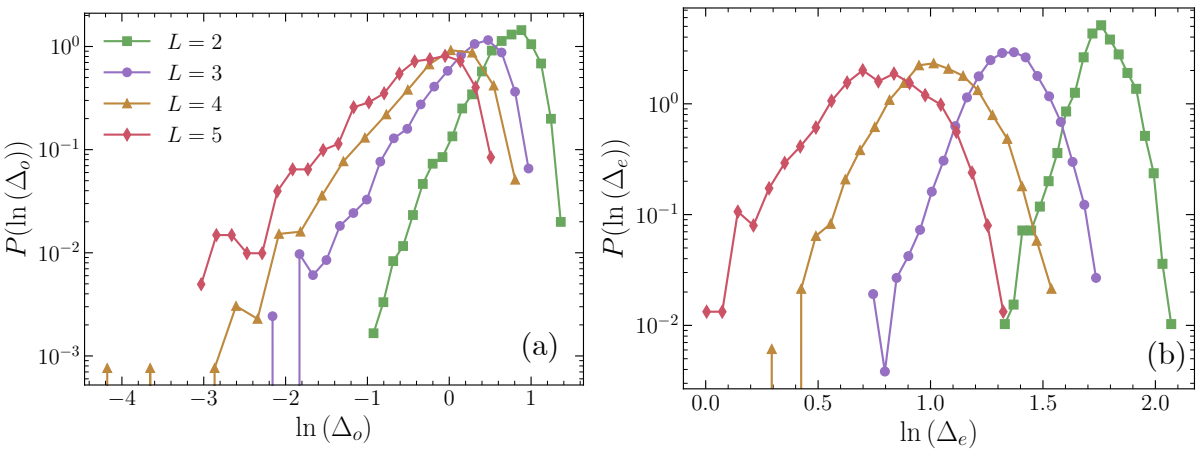

FIGURE 7.10: Probability distributions of the log of odd (a) and even (b) parity gaps $\ln(\Delta_e)$ and $\ln(\Delta_o)$ in log-scale for ~ 5000 realization of the 2D EA model at $\Gamma \approx 1.98$. Both display broad, non-Gaussian shapes with tails toward low values of Δ, signing that a few instances presents small gaps, especially for larger system sizes

of a many-body localization-delocalization transition in the SK model, with the transverse field acting as a tuning parameter between a localized, nonergodic regime and an ergodic, thermal phase.

Turning to the short-range Edwards-Anderson model, we observed qualitatively similar behavior: both parity gaps display polynomial scaling in the considered size range, characterized by exponents $z_e \approx 1.073(3)$ and $z_o \approx 1.17(1)$. The corresponding gap distributions are broad and asymmetric, with more pronounced tails than in the SK case, reflecting the stronger impact of geometric frustration and finite connectivity. While the data accessible via ED suggest algebraic gap closing at criticality, the limited system sizes prevent a definitive conclusion. Indeed, recent works [27] report indications that, in the EA model, the scaling of the odd (fundamental) gap may actually be super-algebraic, decaying faster than any power law in system size.

The results presented in this chapter should therefore be regarded as preliminary, constrained by the finite sizes accessible to ED. Nevertheless, they establish a robust and fully controlled benchmark for the spectral properties of disordered quantum spin systems, and provide crucial reference data for the larger-scale simulations that follow.

In the next chapter, we will extend this study using PQMC simulations. This approach will enable us to explore significantly larger system sizes, refine the finite-size scaling analysis, and assess in detail whether the apparent super-algebraic behavior of the odd-parity gap in the EA model arises in the thermodynamic limit [27].

Chapter 8

Energy gap estimation with Projection Quantum Monte Carlo

In the previous chapter, we analyzed the low-energy properties of disordered quantum Ising models through Exact Diagonalization (ED), obtaining detailed information on the distribution and scaling of excitation gaps for both fully connected and short-range systems. While ED provides numerically exact results, its applicability is restricted to small system sizes due to the exponential growth of the Hilbert space with the number of spins. As a consequence, several important questions remained open, in particular, whether the algebraic scaling of the energy gaps observed in small systems persists in the thermodynamic limit, and how the distribution tails evolve as the system size increases.

To address these questions, we present a novel PQMC technique that allow us to compute energy gaps for larger system sizes and directly extract low-lying excitations from imaginary-time correlation functions. By combining PQMC with a high-quality variational guiding wave function, we can obtain statistically accurate estimates of both even and odd parity gaps for disordered spin models.

8.1 Introduction

The goal of this chapter is to present a scalable methodology for estimating excitation gaps in disordered quantum spin systems using the PQMC technique. The method allows us to overcome the size limitations of Exact Diagonalization and to extend the analysis of even and odd parity gaps to larger systems.

We start by outlining the theoretical framework that underpins the PQMC algorithm, emphasizing how the energy gap can be extracted from the asymptotic decay of imaginary-time correlation functions. We then describe the numerical procedure used to fit these correlators and obtain accurate estimates of the fundamental gaps.

To validate the method, we perform several benchmark tests. First, we study the one-dimensional transverse-field Ising model with open boundary conditions, for which exact results can be obtained via the Jordan-Wigner transformation. This benchmark confirms that PQMC correctly reproduces the expected scaling of the odd-parity gap across different points near the critical region of phase diagram. As a final validation, we compare PQMC results for some instances of disorder of Edwards-Anderson model against those obtained from Exact Diagonalization, showing excellent agreement across overlapping system sizes and confirming that the method is reliable in the presence of strong frustration.

In the second part of the chapter, we extend our study to paradigmatic spin-glass models, namely the fully connected Sherrington-Kirkpatrick (SK) model and the short-range Edwards-Anderson (EA) model in order to investigate the scaling behavior of the low-energy gaps. For the SK model, we find that the distribution of odd-parity gaps broadens with increasing system size but does not develop extremely heavy tails; the average gap follows a polynomial scaling with exponent $\theta \simeq 1/3$, consistent with the ED results extended to larger sizes ($N = 50, 100$). In contrast, the EA model displays a more intricate scenario: while the mean gap still exhibits an algebraic scaling, the tails of its distribution become progressively heavier, suggesting a possible super-algebraic decay of the fundamental gap, in line with recent findings in the literature [27].

Overall, this chapter establishes PQMC as a robust and scalable tool for probing the low-energy physics of disordered quantum systems and for clarifying the asymptotic behavior of energy gaps beyond the limitations of Exact Diagonalization. In the next section, we present the mathematical derivation that allows us to extract the energy gap from PQMC simulations. We show that the use of importance sampling, through a guiding wave function that modifies the PQMC algorithm as introduced in Chap. 1, does not affect the gap estimation, since this observable is not a mixed measurement.

8.2 PQMC correlation function

Our aim is to calculate the following correlation function in imaginary time

$$F(\tau) = \frac{\langle \psi_0 | O(\tau) O(0) | \psi_0 \rangle}{\langle \psi_0 | \psi_0 \rangle}, \tag{8.1}$$

where $|\psi_0\rangle$ is the ground state and $O(\tau)$ is a generic operator acting on a spin configuration $\mathbf{S} = (s_1, s_2, \ldots, s_N)$ where $s_i = \pm 1$ and $i = 1, \ldots, N$. Here, $\psi(\mathbf{S}) = \langle \mathbf{S} | \psi \rangle$ with $|\mathbf{S}\rangle$ eigenstates of σ_i^z.

In a general Projection Monte Carlo calculation one has access to the quantity

$$F_{\mathrm{DMC}}(\tau) = \frac{1}{N_w(\tau)} \sum_{\mathbf{S}} f(\mathbf{S}, \tau) O(\tau) O(0), \tag{8.2}$$

obtained by saving the spin configurations $\mathbf{S}_i(0)$ at time $\tau = 0$ and averaging the operator product over the probability distribution function $f(\mathbf{S}, \tau)$ of the N_w walkers at time τ.

We now suppose that the distribution $f(\mathbf{S}, \tau)$ is already well equilibrated and is given by the product $f(S, \tau) \simeq \psi_T(\mathbf{S}) \psi_0(\mathbf{S})$. The correlation function in Eq. 8.2 takes then the form of a mixed estimator

$$F_{\mathrm{DMC}}(\tau) = \frac{\langle \psi_T | O(\tau) O(0) | \psi_0 \rangle}{\langle \psi_T | \psi_0 \rangle}. \tag{8.3}$$

Due to symmetry, if the operator O is odd under the symmetry transformation (for instance, spin inversion or parity) has zero average over the ground state, i.e, $\langle \psi_0 | O | \psi_0 \rangle = 0$. Hence, the state $O|\psi_0\rangle$ must be orthogonal to the ground state and admits the expansion over the subspace of excited states with odd symmetry

$$O|\psi_0\rangle = \sum_{n \neq 0} c_n |\psi_n\rangle, \tag{8.4}$$

for each $n \neq 0$ such that $|\psi_n\rangle$ belongs to the odd-parity subspace. In the following, we will also discuss the case of an even operator O for which $\langle \psi_0 | O | \psi_0 \rangle \neq 0$ and additional considerations are required.

For the operator in real time t one obtains

$$O(t)|\psi_0\rangle = e^{iHt}Oe^{-iHt}|\psi_0\rangle = \sum_{n\neq 0} c_n e^{i(E_n-E_0)t}|\psi_n\rangle. \tag{8.5}$$

Switching to the imaginary time $\tau = it$, the matrix element involving the ground state becomes

$$\langle\psi_0|O(\tau)O(0)|\psi_0\rangle = \sum_{n\neq 0} |c_n|^2 e^{-(E_n-E_0)\tau} \simeq |c_n|^2 e^{-(E_n-E_0)\tau}, \tag{8.6}$$

where the last approximation holds for a sufficiently large imaginary time τ, such that only the contribution of the first excited state survives in the sum. The same reasoning applies for the trial state $|\psi_T\rangle$ such that

$$O|\psi_T\rangle = \sum_{n\neq 0} d_n|\psi_n\rangle, \tag{8.7}$$

in terms of different expansion coefficients d_n compared to the linear combination in Eq. 8.4. As a result we can conclude that the mixed estimator in Eq. 8.3, for large enough time windows $[0,\tau]$, reduces to the contribution from the first excited state with energy gap $E_1 - E_0$:

$$F_{\mathrm{DMC}}(\tau) \simeq \frac{d_1^* c_1}{\langle\psi_T|\psi_0\rangle} e^{-(E_1-E_0)\tau}. \tag{8.8}$$

Notice that the time dependence in the above equation is the same as the one expected in the correct correlation function in Eq. 8.6.

Before proceeding, it is important to stress that the goal of the following derivation is to demonstrate that the extraction of the excitation gap from the imaginary-time correlation function is independent on the presence of importance sampling. In particular, we show that introducing a guiding (trial) wave function $\psi_T(\mathbf{S})$ does not modify the imaginary-time dependence of the correlator from which the energy gap is extracted. To this aim, we examine whether the modified time evolution induced by importance sampling needs to be taken into account in the evaluation of the correlation function. For this reason we analyze the matrix element

$$\langle\psi_T|O(\tau)O(0)|\psi_0\rangle = \langle\psi_T|e^{\tau H}Oe^{-\tau H}O|\psi_0\rangle, \tag{8.9}$$

between the ground state and the trial state $|\psi_T\rangle$. Due to time invariance, the above matrix element must be equal to the corresponding correlation function with a shifted time interval,

$$\langle\psi_T|O(0)O(-\tau)|\psi_0\rangle = \langle\psi_T|Oe^{-\tau H}e^{\tau H}O|\psi_0\rangle. \tag{8.10}$$

As we said in Chap. 1, the calculation of the ground state $|\psi_0\rangle$ is the result of a sufficiently long projection in imaginary time starting from the trial state

$$|\psi_0\rangle = \lim_{n\to\infty} e^{-n\delta_\tau H}|\psi_T\rangle. \tag{8.11}$$

If the projection time is long compared to the time interval τ then $e^{\tau H}|\psi_0\rangle = |\psi_0\rangle$ and the relevant matrix element can be rewritten as

$$\langle\psi_T|Oe^{-\tau H}O|\psi_0\rangle = \sum_{\mathbf{S},\mathbf{S}_0} \psi_T(\mathbf{S})\langle\mathbf{S}|Oe^{-\tau H}O|\mathbf{S}_0\rangle\psi_0(\mathbf{S}_0). \tag{8.12}$$

Since the operator O is diagonal in the coordinate basis, the latter equation takes the form

$$\langle\psi_T|Oe^{-\tau H}O|\psi_0\rangle = \sum_{\mathbf{S},\mathbf{S}_0} O(\mathbf{S})O(\mathbf{S}_0)\psi_T(\mathbf{S})\langle\mathbf{S}|e^{-\tau H}|\mathbf{S}_0\rangle\psi_0(\mathbf{S}_0). \tag{8.13}$$

At this point, one can notice that the Green's function associated with the distribution function $f(\mathbf{S},\tau)$ is

$$G_T(\mathbf{S},\mathbf{S}_0,\tau) = \psi_T(\mathbf{S})\langle\mathbf{S}|e^{-\tau H}|\mathbf{S}_0\rangle\psi_T^{-1}(\mathbf{S}). \tag{8.14}$$

We can use it to rewrite the matrix element as

$$\begin{aligned}\langle\psi_T|Oe^{-\tau H}O|\psi_0\rangle &= \sum_{\mathbf{S},\mathbf{S}_0} O(\mathbf{S})O(\mathbf{S_0})G_T(\mathbf{S},\mathbf{S}_0,\tau)f(\mathbf{S}_0) \\ &= \sum_{\mathbf{S},\mathbf{S}_0} O(\mathbf{S})O(\mathbf{S_0})f(\mathbf{S},\tau),\end{aligned} \tag{8.15}$$

where the average is carried out over the walkers $\mathbf{S}$ at imaginary time τ which were located in $\mathbf{S}_0$ at $\tau = 0$. Notice that the latter equation coincides with Eq. 8.2.

8.2.1 Extraction of the energy gap from PQMC correlators

Once the PQMC simulation has converged to the correct ground state through sufficient imaginary-time projection, we have access to the imaginary time correlation function

$$C(\tau) = \langle O(0)O(\tau)\rangle. \tag{8.16}$$

Let us focus on the case where O is an odd parity flipping operator that couples the even-parity ground state $|E_o^{\text{even}}\rangle$ to the lowest odd-parity excited state $|E_1^{\text{odd}}\rangle$. In the long-time limit, this correlator decays exponentially as

$$C(\tau) \propto e^{-\Delta_o \tau}, \tag{8.17}$$

where $\Delta_o = E_1^{\text{odd}} - E_0^{\text{even}}$ is the fundamental gap. However, in practical PQMC simulations, the correlator signal spans only a finite imaginary-time interval before being suppressed by statistical noise. For this reason a selection of the fitting window and an adaptive fitting strategy are essential to obtain a reliable estimate of Δ_o. In an analogous way, one can consider an even-parity operator O for which the expectation value over the ground state does not vanish, i.e. $\langle\psi_0|O|\psi_0\rangle \neq 0$. In this case, the imaginary-time correlation function develops a constant contribution that dominates at large τ, corresponding to the product of the asymptotic averages of the operator at different imaginary times. To isolate the exponential decay associated with the lowest even-parity excitation, we define

$$\tilde{C}_{\text{even}} = C_{\text{even}} - \langle O(\tau)\rangle\langle O(0)\rangle. \tag{8.18}$$

Under this assumption, for sufficiently long imaginary times the correlation factorizes as

$$\langle O(\tau)\rangle\langle O(0)\rangle \simeq \langle O\rangle^2, \tag{8.19}$$

and the subtracted correlator $\tilde{C}_{\text{even}}$ decays exponentially as

$$\tilde{C}_{\text{even}}(\tau) \propto e^{-\Delta_e \tau}, \tag{8.20}$$

with $\Delta_e = E_1^{\text{even}} - E_0^{\text{even}}$. During the PQMC simulation, the quantities $\langle O(\tau)\rangle$ and $\langle O(\tau)\rangle$ are evaluated independently, and their product is subtracted from the raw correlator in the analysis phase. This procedure allows one to extract the even-parity gap using the same exponential fitting strategy employed for the odd-parity case. An example of even and odd correlators in imaginary time τ is provided in Fig. 8.1 (a), where the odd-parity correlator exhibits a natural exponential decay, while the even-parity correlator approaches a constant asymptotic value. In the analysis, the constant contribution in the even-parity correlator is subtracted to isolate the exponential decay, allowing the extraction of the corresponding gap in the same way as for the odd-parity case.

Choice of parity operators

We specify here the concrete choices of the operators used to probe the two parity sectors, and the procedure adopted to select, during the simulation, the operator that yields the cleanest imaginary-time signal.

For the odd-parity sector we employ the local spin-flip operator

$$O_i^{\text{odd}} = \sigma_i^z, \tag{8.21}$$

which couples the even-parity ground state to odd-parity excited states. At the beginning of each simulation we evaluate, over a number of initial blocks, the correlators associated with a large set of sites i. For each site we compute the corresponding signal-to-noise ratio (SNR) in the relevant imaginary-time window. The site that exhibits the best SNR is then selected on-the-fly and used for the remainder of the simulation. The final estimate of the excitation gap is extracted from this optimal correlator.

For the even-parity sector we analogously consider two-site operators of the form

$$O_{ij}^{\text{even}} = \sigma_i^z \sigma_j^z, \tag{8.22}$$

which couple the ground state to even-parity excitations. Again, during an initial calibration stage we compute correlators for a large ensemble of pairs (i, j) and evaluate the SNR for each of them. The pair providing the highest SNR is automatically selected

and used for all subsequent imaginary-time measurements, and the final even-parity gap is obtained from this selected correlator.

We emphasize that this on-the-fly operator selection is crucial for maximizing the quality of the exponential decay regime used in the fits. Both σ_i^z and $\sigma_i^z\sigma_j^z$ are diagonal in the computational basis and therefore straightforward to measure within PQMC; the adaptive SNR-based choice ensures that the numerical extraction of the excitation gaps is as stable and noise-resilient as possible.

Selection of the fitting window

The exponential decay regime typically appears after an initial transient, during which higher excitations still contribute. At very long imaginary times, however, the signal-to-noise ratio (SNR) deteriorates rapidly due to the exponential damping of the signal, eventually making the correlator indistinguishable from noise $\sigma(\tau)$. To isolate the region dominated by a single exponential decay, we progressively restrict the fitting window $[\tau_{\text{min}}, \tau_{\text{max}}]$, discarding data points where the SNR falls below a predefined threshold (typically $C(\tau)/\sigma(\tau) \sim 8$). In practice, in a semi-logarithmic scale, this regime appears as a linear portion of the curve, corresponding to an exponential decay.

Iterative fitting procedure

Within the selected window, we perform a nonlinear least-squares fit of the form $C(\tau) = A\, e^{-\Delta\tau}$, where A and Δ are free parameters. The fit is carried out iteratively by progressively reducing the fitting range from short times: at each iteration, the first data point on the left of the fitting window is removed, and a new fit is performed on the remaining data. For each fit, we record the best-fit gap $\Delta(\tau_{\text{min}})$, its uncertainty from the covariance matrix, and the corresponding reduced chi-squared value. In Fig. 8.1 (b) we provide a representation of this fit procedure for the even-parity gap of $N = 100$ spins in a ferromagnetic Ising chain at transverse field $\Gamma = 1.1$.

Stopping criterion and extraction of the optimal gap

As the left bound τ_{min} of the fitting window increases, short-time transients are progressively excluded, and the fitted gap $\Delta(\tau)$ converges toward a plateau value. The

plateau region signals that the single-exponential assumption is valid and that the extracted Δ corresponds to the true physical gap. Once τ_{min} is increased beyond this region, the SNR deteriorates and the fit begins to capture statistical noise. The procedure stops when the fitted gap remains constant within statistical uncertainty for consecutive windows or begins to increase.

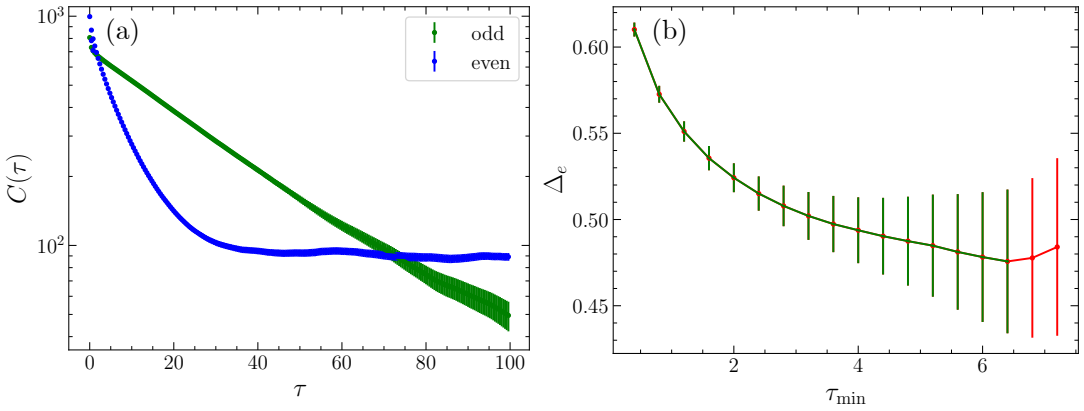

FIGURE 8.1: Example of gap extraction from the imaginary-time correlation function for the one-dimensional transverse field Ising model with $N = 100$ spins. (a) Even (blue) and odd (green) Correlation functions in semi-logarithmic scale in the selected time window. The odd-parity correlation function is approximately linear in this scale, while the even-parity correlation function approaches a constant value, which is subtracted in order to extract the even-parity gap using the same procedure as for the odd gap. (b) Estimated even-parity gap Δ_e as a function of the minimum fitting time τ_{min}. The plateau region corresponds to the optimal fitting window from which the odd-parity energy gap is extracted.

8.3 Benchmark results on testbed models

Before addressing disordered and frustrated systems, we first validate the PQMC methodology on paradigmatic testbed models for which reliable reference results are available. This step is crucial to assess the accuracy and robustness of the approach in estimating excitation gaps from imaginary-time correlations. In particular, we benchmark the method on the one-dimensional transverse-field Ising chain, for which the energy spectrum can be exactly computed via the Jordan-Wigner transformation.

8.3.1 One-dimensional Ising chain

Let us consider the one-dimensional ferromagnetic transverse-field Ising model of Eq. (2.22) with open boundary conditions and system size $N = 100$. Figure 8.2 shows the comparison between the odd-parity energy gaps obtained from PQMC simulations and the exact results computed using the Jordan-Wigner transformation as a function of the transverse field Γ across the quantum critical point $\Gamma_c = 1$ [171].

The PQMC results are in excellent agreement with the exact values throughout the entire range of Γ, capturing both the closing of the gap near the critical point and its subsequent reopening in the paramagnetic phase. This validates the reliability of the imaginary-time fitting procedure discussed in the previous section. Notably, even in the vicinity of the critical region, where the correlation time diverges, the extracted gaps remain quantitatively consistent with the exact solution, confirming that the method can accurately resolve small excitation energies.

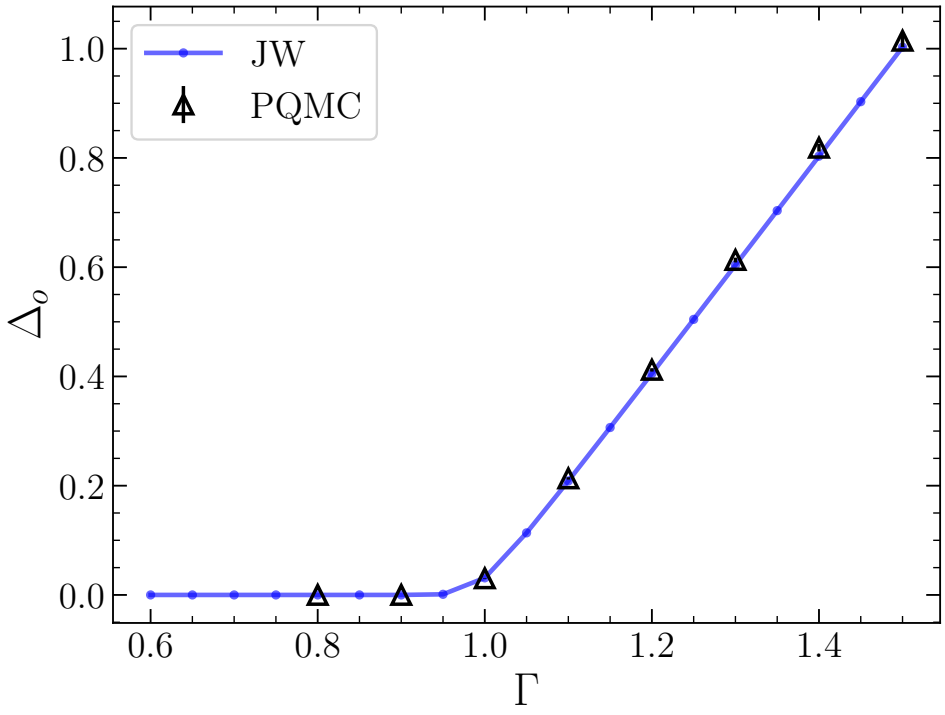

FIGURE 8.2: Odd-parity energy gap Δ_o for the one-dimensional ferromagnetic transverse-field Ising model with $N = 100$ spins as a function of the transverse field Γ. The black triangles show the PQMC results, while the blue dotted-line corresponds to the exact solution obtained from the Jordan-Wigner transformation. The agreement across the entire range of Γ demonstrates the accuracy of the PQMC gap estimation method, even near the quantum critical point at $\Gamma_c = 1$.

In addition to the field dependence at fixed system size, we analyze the scaling of the excitation gaps with system size, which provides a stringent validation of the PQMC approach. To this end, we computed both the even- and odd-parity energy gaps for the one-dimensional transverse field Ising chain at the critical point $\Gamma_c = 1.0$, considering system sizes $N = 20, 40, 60, 80$, and 100. Figure 8.3 reports the corresponding values of Δ_e and Δ_o as obtained from PQMC simulations, compared to the exact results from the Jordan-Wigner solution. The excellent agreement between PQMC and the exact spectrum across all system sizes confirms that the imaginary-time correlation analysis yields quantitatively accurate gap estimates not only for fixed-size systems, but also in the scaling regime. Both parity sectors display the expected inverse-size dependence, demonstrating that PQMC faithfully captures the finite-size scaling behavior of low-energy excitations. These results provide further evidence of the robustness and precision of the method, establishing a reliable baseline for its application to more complex disordered and frustrated models discussed in the following sections.

8.4 Gap scaling in spin glass models

Now, let us investigate the behavior of the low-energy excitation gaps in paradigmatic disordered spin-glass models. In the previous chapter, we have seen that Exact Diagonalization provides accurate insights into small systems. However, here we extend the analysis with the PQMC technique introduced in this chapter which allows us to explore significantly larger system sizes, thus enabling a more reliable scaling analysis. In this context, the accurate estimation of the lowest excitation gaps could requires particularly long imaginary-time evolutions, since the contribution of excited states decays slowly and must be carefully filtered out. To make these long runs computationally efficient, we replaced the self-learning paradigm employed in the previous chapters with an optimized guiding wave function obtained through a separate variational optimization. Specifically, we used the NetKet framework [61], which allows us to train a Restricted Boltzmann Machine directly on the target Hamiltonian prior to the PQMC projection. The resulting optimized wave function provides an excellent approximation of the true ground state and therefore serves as an efficient guiding function. This strategy drastically reduces the variance of the walker population and suppresses population-control bias, allowing unbiased PQMC simulations even with

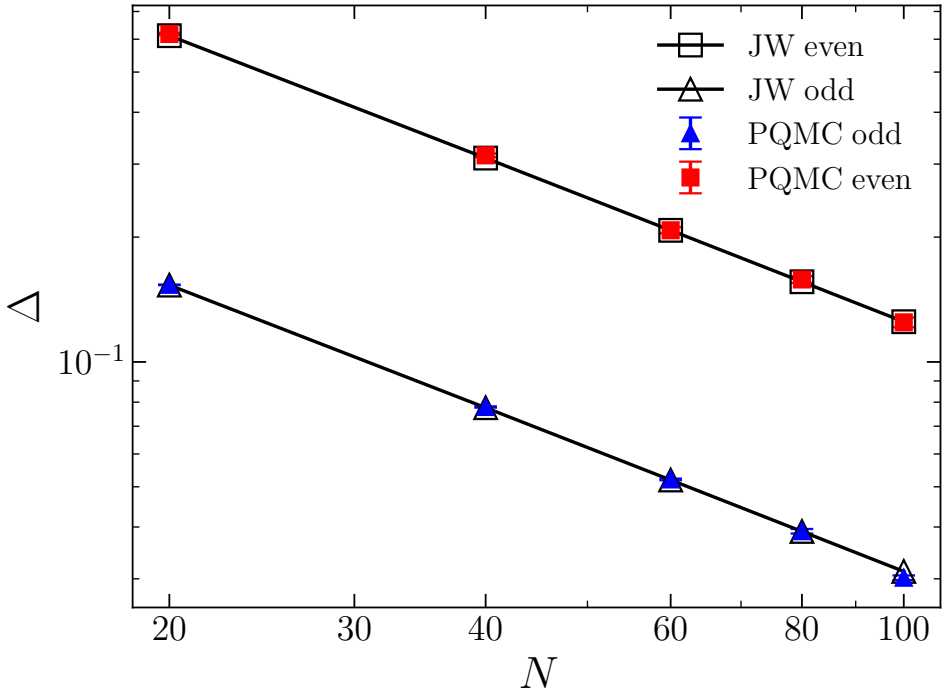

FIGURE 8.3: Size scaling of the even-parity (Δ_e) and odd-parity (Δ_o) excitation gaps for the one-dimensional transverse-field Ising chain at the critical point $\Gamma_c = 1$. PQMC results (red and blue markers) are compared with exact values obtained from the Jordan-Wigner transformation (solid lines). The quantitative agreement across all system sizes $N = 20, 40, 60, 80$, and 100 demonstrates the accuracy of the PQMC gap estimation procedure and its ability to reproduce the expected finite-size scaling of the critical spectrum.

a relatively small number of walkers. In practice, we used a walker population of $N_w = 400$ and a hidden-layer size $N_h = N$, where N denotes the total number of spins in the system. This setup ensures an optimal balance between accuracy and computational cost, while maintaining the reliability of the gap estimation procedure.

8.4.1 Results: the Edwards-Anderson model

As we have seen in Chap. 3, the Edwards-Anderson model is the standard realization of a short-range spin-glass system. The Hamiltonian is the one of Eq. (3.1), where the couplings J_{ij} are randomly sampled from a Gaussian distribution with zero mean and unit variance, namely $\mathcal{N}(0, 1)$. A central open question concerns the scaling behavior of the excitation gap Δ with the system size L at the critical point. In a system with polynomial scaling, one expects $\Delta \sim L^{-z}$, where z is the dynamical critical exponent.

However, recent findings report that for the quantum EA model, the situation is more subtle: while the even-parity gap follows an algebraic scaling with $z = 2.46(17)$, the odd-parity gap exhibits indications of a super-algebraic decay, associated with the emergence of fat-tails in gaps distribution $P(\Delta_o)$ as the system size increases [27]. This behavior is associated to rare low-energy excitations that become increasingly dominant in the thermodynamic limit. For this reason, our study focuses on the odd-parity gap Δ_o.

Before addressing the scaling behavior, once again, we validate the PQMC technique in this highly frustrated models by directly comparing its results with those obtained from ED for small system sizes. Fig. (8.4) shows the odd-parity gaps computed for different disorder realizations of the EA model with lattice size $L = 5$ at fixed critical transverse field $\Gamma_c = 1.98$. Each point corresponds to a single realization of the couplings $\{J_{ij}\}$. The agreement between PQMC and ED confirms that the method reproduces individual excitation energies even in the presence of strong disorder, without systematic bias across samples. To further assess the statistical consis-

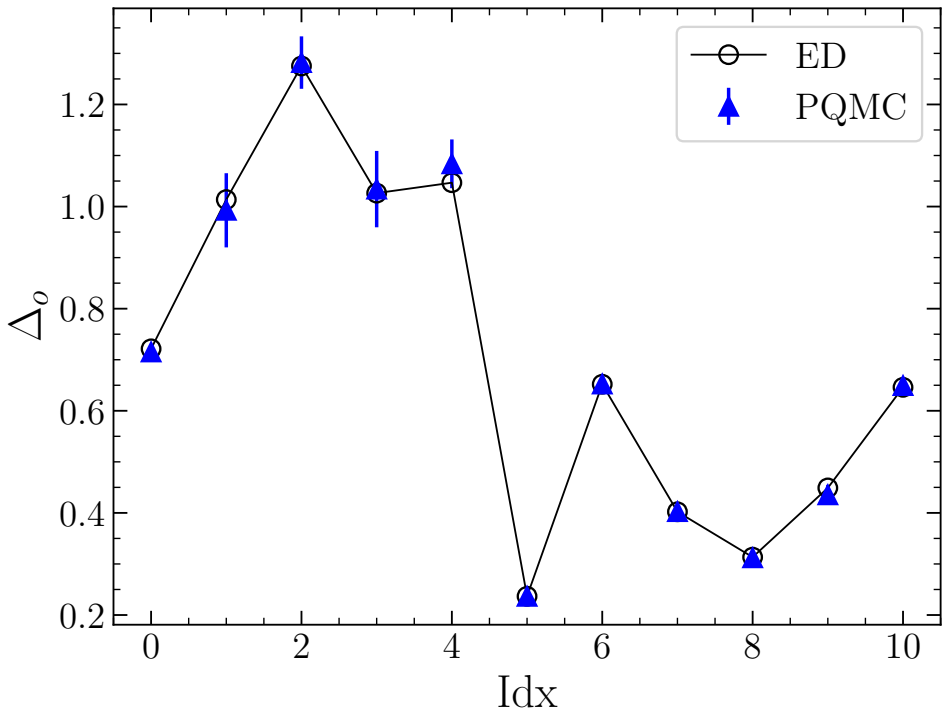

FIGURE 8.4: Odd-parity energy gaps Δ_o for different disorder realizations of the Edwards–Anderson model with system size $L = 5$ at $\Gamma_c = 1.98$. Blue triangles denote PQMC results, while black empty circles represent ED values.

tency between the two approaches, we compare the corresponding probability distributions of the logarithmic gaps. Fig. (8.5) displays the histograms of $\ln(\Delta_o)$ obtained from 5000 disorder realizations using both ED and PQMC. The two distributions are indistinguishable within statistical uncertainty, confirming that the PQMC correctly samples the disorder ensemble and reproduces not only the average properties but also the full shape of the gap distribution, including its tails.

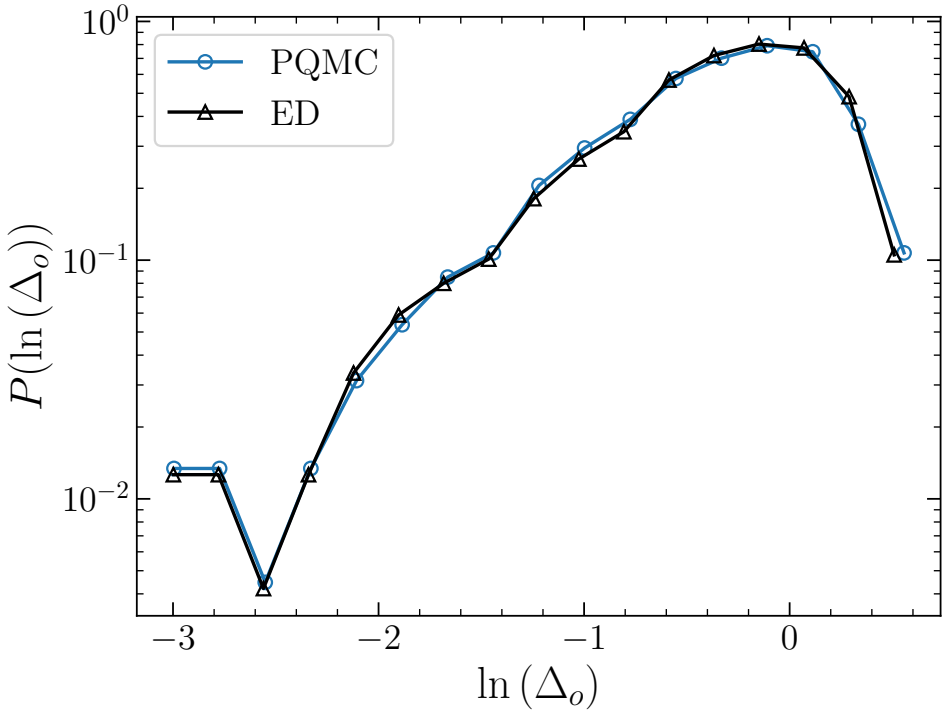

FIGURE 8.5: Probability distribution of the logarithm of the odd-parity gap, $\ln(\Delta_o)$, for the Edwards-Anderson model with $L = 5$ and $\Gamma_c = 1.98$. The histograms obtained from PQMC (light blue circles) and ED (black triangles) coincide within statistical uncertainty.

Scaling behavior of the Edwards-Anderson gap at the critical point

Having validated the method against Exact Diagonalization for small system sizes, we now extend the analysis to larger lattices in order to investigate the scaling behavior of the excitation gap in the quantum EA model. In Fig. (8.6) we augment the ED data in Fig. (7.9) by adding PQMC results for $L = 7, 10$. Despite the inclusion of the PQMC data still support the algebraic behavior of the scaling, giving a dynamical exponent of $z = 1.18(1)$, a closer inspection reveals that the data for $L = 10$ lie slightly below the

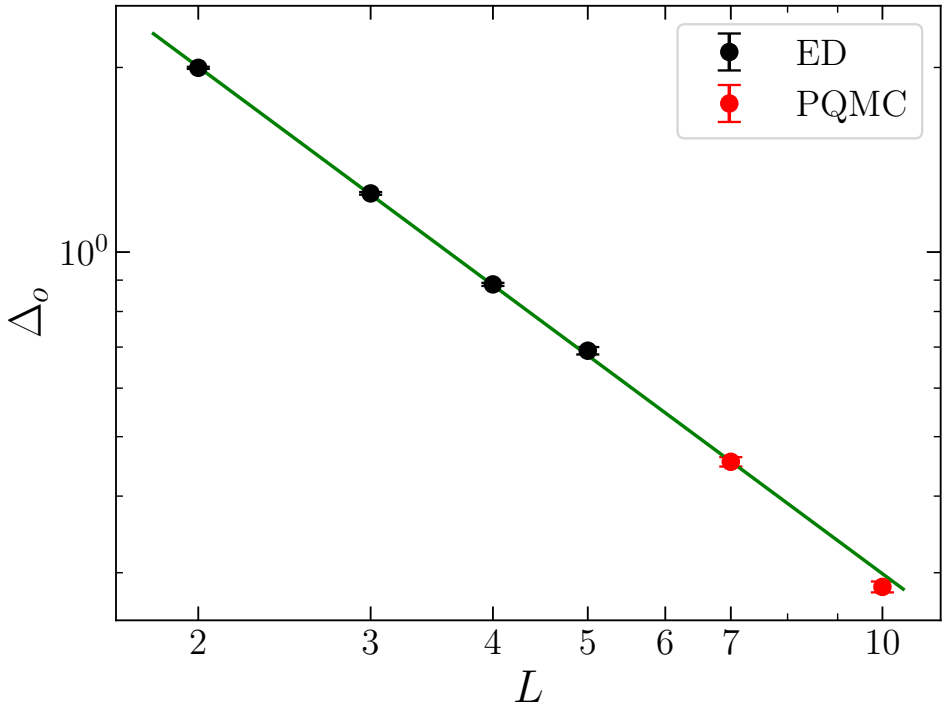

FIGURE 8.6: Scaling of the mean odd-parity gap Δ_o for the quantum Edwards-Anderson model at the critical point $\Gamma_c = 1.98$. Black symbols denote Exact Diagonalization data from Chap. 7, while red points correspond to new PQMC results for $L = 7$ and $L = 10$. The green line represents a power-law fit $\Delta_o \sim L^{-z}$ with exponent $z = 1.18(1)$, again coherent with previous results. The slight deviation observed for $L = 10$ may indicate the onset of a crossover to a different, possibly super-algebraic, scaling regime at larger sizes.

extrapolated power law, suggesting a possible crossover towards a different scaling regime for larger systems.

To better investigate this hypothesis, it is useful to examine the distribution of the logarithmic gap $P(\ln(\Delta_0))$. Once again, we recall the results obtained in the previous chapter (See Fig. 7.10) and add the data for the new sizes. The results in Fig. (8.7) show that the tails toward small gaps become progressively broader and heavier as the system size increases.

The emergence of such tails implies that the arithmetic mean of the gap is increasingly dominated by a few atypical realizations, making it an inaccurate estimator of the observable. In this case, a more representative quantity is the typical gap, defined as the geometric average

$$\Delta_o^{\text{typ}} = \exp(\langle \ln \Delta_o \rangle)\,, \tag{8.23}$$

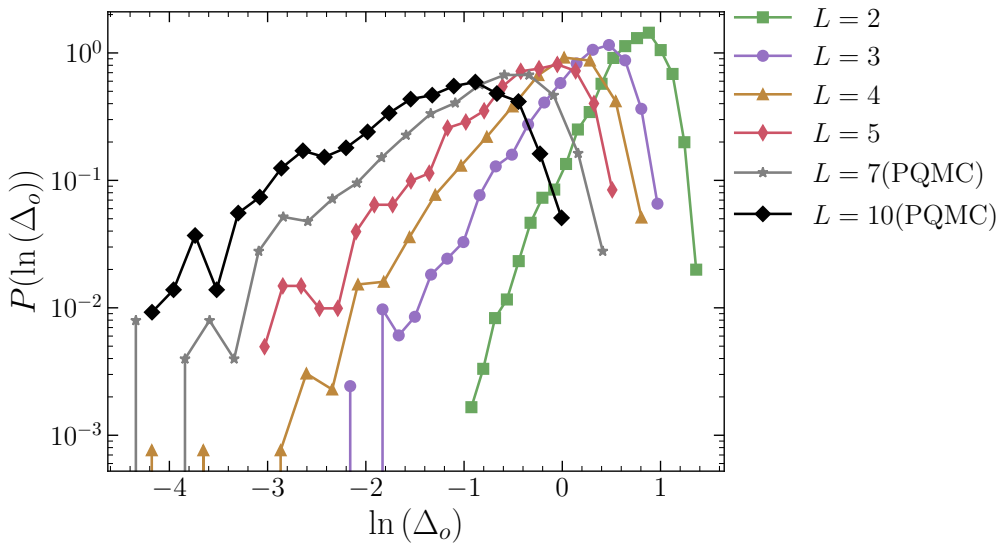

FIGURE 8.7: Probability distributions of $\ln(\Delta_o)$ for the Edwards-Anderson model at critical point for various system sizes, combining ED data (L = 3-6) and PQMC results ($L = 7, 10$). The progressive broadening and enhancement of the tails toward small gap values indicate the increasing relevance of rare low-energy excitations as the system size grows.

which suppresses the influence of rare extreme events and captures the behavior of the bulk of the distribution. In Fig. (8.8) we show that for $L = 7$ and $L = 10$, the typical gap decreases much faster than the power-law prediction, deviating from the fitted trend. This behavior suggests that, while the arithmetic mean remains roughly algebraic, the typical gap enters a distinct regime where super-algebraic scaling could becomes dominant, consistent with the emergence of fat-tailed gap distributions discussed earlier.

Heavy-tail distributions and tail-exponent analysis

A further way to investigate the emergence of rare, small-gap events is to analyze the tail behavior of the inverse-gap distribution. Following Ref. [27], we introduce the variable

$$\eta = \frac{1}{\Delta_o} \tag{8.24}$$

and observe the complementary cumulative distribution function $1 - F[\eta]$ for different system sizes. As the sizes increases, the curves systematically shift upward, indicating an enhancement of the weight in the tail of the distribution. In the asymptotic regime

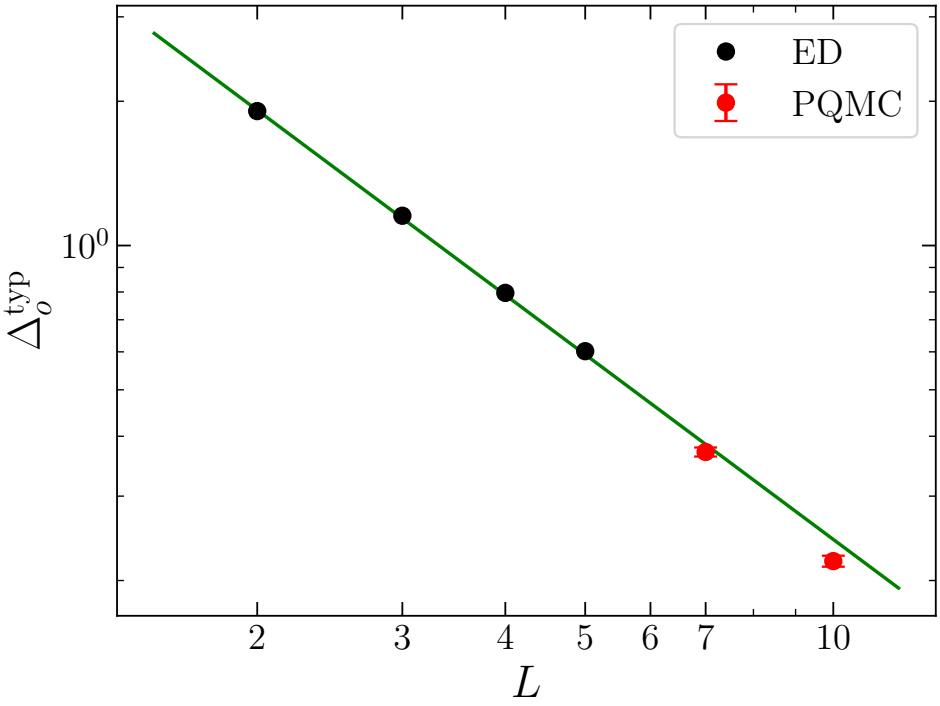

FIGURE 8.8: Scaling of the typical odd-parity gap Δ_o^{typ}. Black points correspond to typical averages obtained from Exact Diagonalization, while red points indicate PQMC results for larger system sizes. The green line shows a power-law fit, which captures well the small-size behavior but fails to reproduce the rapid suppression observed at larger sizes. This deviation signals the possible onset of a super-algebraic scaling regime for the typical gap.

of a power-law tail, one expects $1 - \mathrm{F}[\eta] \sim \eta^{-\alpha}$, so that a straight line in a log-log plot signals the presence of a possibly fat-tailed regime characterized by a finite tail exponent α. The results in Fig. (8.9) resemble those reported in Ref. [27].

Quantitative characterization of tail behavior can be achieved through the Hill estimator, a standard statistical tool to estimate the tail exponent of possibly heavy-tailed distributions [176]. In our case, given a sample of inverse gaps, sorted in descending order $\eta_{(1)} \geq \eta_{(2)} \geq \cdots \geq \eta_{(M)}$, the Hill estimator α is defined as

$$\alpha^{-1}(k) = \frac{1}{k-1} \sum_{i=1}^{k-1} \ln \left(\frac{\eta_{(i)}}{\eta_{(k)}} \right), \tag{8.25}$$

where k denotes the number of top-order statistics included in the estimate, i.e., the fraction of largest η values used to probe the asymptotic tail distribution. In the case

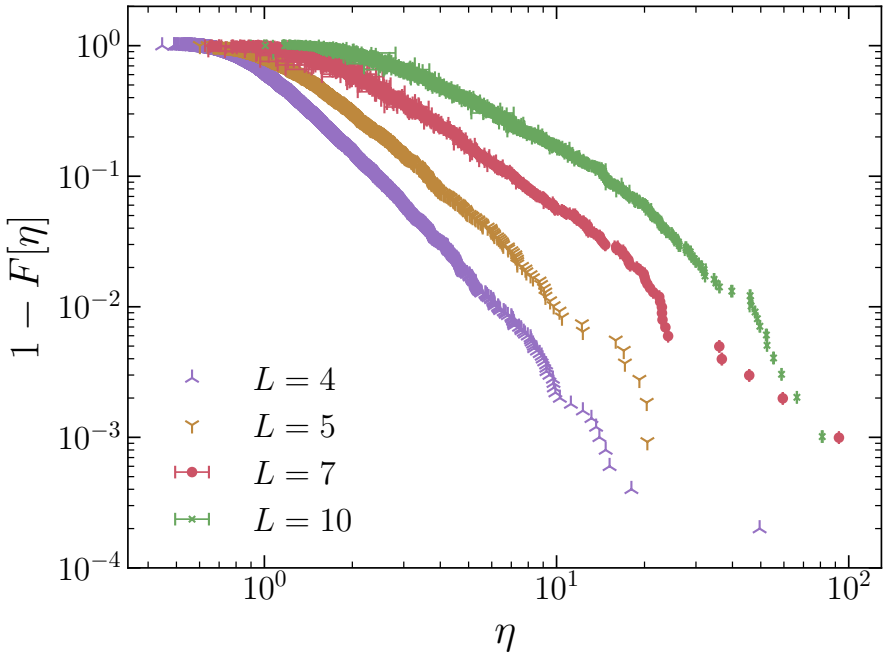

FIGURE 8.9: Complementary cumulative distribution function $1 - \mathrm{F}[\eta]$ of the inverse odd-parity gaps $\eta = 1/\Delta_o$ for the Edwards-Anderson model at different system sizes $L = 4, 5, 7, 10$. The upward shift of the curves with increasing L reveals the progressive accumulation of rare realizations with very small gaps, indicative of a fat-tailed regime consistent with Ref. [27].

of a true power-law tail, the Hill estimator converges to the true exponent α for sufficiently large M and properly chosen k.

The values of α extracted from our simulations, shown in Fig. 8.10, exhibit a systematic downward trend as the system size increases. This behavior indicates a progressive enhancement of the probability weight in the far tail of the distribution, consistent with an increasingly broad, heavy-tailed regime. In particular, the decrease of α below certain threshold values has clear statistical implications. For a power-law tail $1 - F[\eta] \sim \eta^{-\alpha}$, the n-th moment $\langle \eta^n \rangle$ diverges for $\alpha \leq n$. Recalling that $n = 1$ and $n = 2$ are the moment associated to the mean and the variance respectively, we get that:

- for $\alpha > 2$, both the mean and the variance of the distribution are finite;
- for $1 < \alpha \leq 2$, the mean remains finite, but the variance diverges;
- for $\alpha \leq 1$, even the mean diverges, implying that the distribution is dominated by rare, extreme events.

In our data, $\alpha(L)$ belongs to the the second case for the largest simulated sizes. This trend is confirmed by the asymptotic values extrapolated via the power-law fit of the form $aL^b + c$ which yield negative b values and $c = 1.7(3)$, $c = 1.3(1)$ and $c = 1.06(7)$ for $k = 0.1, 0.2$ and 0.3, respectively. These results suggest that the variance of the inverse-gap distribution, and therefore the fluctuations of the gap itself, diverge in the thermodynamic limit. This observation provides strong quantitative evidence that the Edwards-Anderson model at criticality enters a fat-tailed regime, where the dynamics are increasingly governed by rare disorder realizations associated with exponentially small excitation gaps. Such behavior supports the scenario of a super-algebraic closure of the odd-parity gap, as previously inferred from the scaling analysis of Δ_o^{typ} and from the broadening of the $\ln(\Delta_o)$ distributions. These preliminary results find fair agreement with those in Ref. [27].

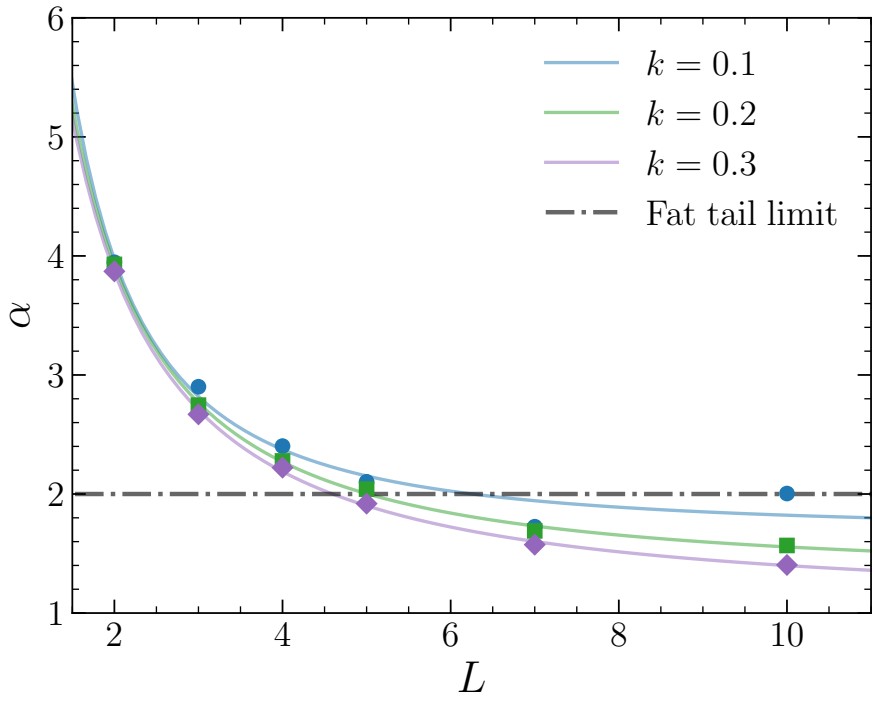

FIGURE 8.10: Hill estimator α of the tail exponent of the inverse-gap distribution as a function of system size L, computed for different fractions $k = 0.1$ (blue), 0.2 (green), and 0.3 (purple) of the largest η values. The systematic decrease of α with L indicates progressively heavier tails, supporting the interpretation of a super-algebraic suppression of the typical gap in the thermodynamic limit. The asymptotic values for $k = 0.1, 0.2, 0.3$ are $c = 1.7(3), 1.3(1), 1.06(7)$ respectively.

Scaling behavior of the Sherringhton-Kirkpatrick gap at the critical point

We now turn our attention to the fully connected Sherrington-Kirkpatrick (SK) spin-glass model. As discussed in Chap. 7, the SK model is characterized by infinite-range couplings $J_{ij} \sim \mathcal{N}(0, 1/N)$. In the previous chapter, we analyzed the low-energy spectrum of small systems using Exact Diagonalization, finding that the odd-parity energy gaps follow a polynomial scaling with system size, characterized by the exponent $\theta \simeq 1/3$. Here we extend this analysis to larger systems by employing the PQMC method, which allows us to reach sizes up to $N = 100$ spins.

We begin by examining the probability distributions of the logarithmic odd-parity gaps, $P(\ln(\Delta_o))$, at the critical transverse field $\Gamma_c = 1.5$. Fig. 8.11 shows the evolution of these distributions as the system size increases, combining the ED results in Fig. (7.4) from Chap. 7 with new PQMC data for $N = 50$ and $N = 100$.

While the distributions gradually broaden as N increases, indicating the presence of a wider range of low-energy excitations, the tails towards small Δ_o remain relatively moderate compared to those observed in the Edwards-Anderson model. This suggests that the SK model does not develop huge heavy tails in the thermodynamic limit, and that the statistics of the gap remain well-behaved even at large sizes.

Scaling of the mean gap

Given the absence of signs of heavy tails in the SK model, the arithmetic mean of the odd-parity gap, provides a meaningful and representative measure of the excitation energy. In Fig. 8.12, we report the mean value of Δ_o as a function of the system size, combining ED data with new PQMC results for $N = 50$ and $N = 100$. The combined dataset is well described by a power-law scaling yielding an exponent $\theta = 0.33(2)$, in good agreement with the pure ED estimation from smaller systems.

This result confirms that the polynomial closure of the excitation gap persists up to much larger system sizes, reinforcing the interpretation that the critical dynamics of the SK model are governed by a well-defined algebraic scaling rather than by rare-event dominated processes.

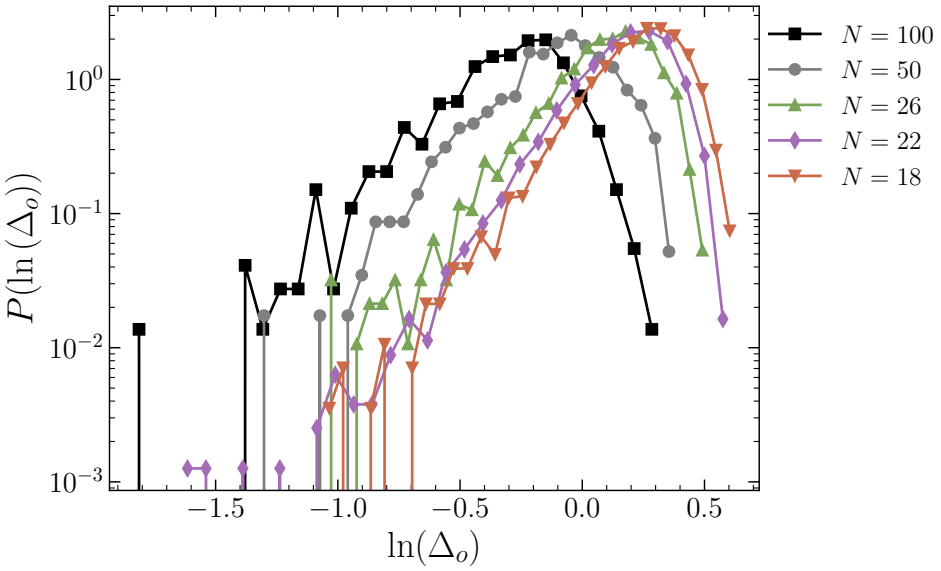

FIGURE 8.11: Probability distributions of $\ln(\Delta_o)$ for the Sherrington–Kirkpatrick model at the critical transverse field $\Gamma_c = 1.5$ for various system sizes. Exact Diagonalization data (black) are complemented by PQMC results for $N = 50$ (gray) and $N = 100$ (black). Although the distributions broaden with increasing N, the tails towards small gaps remain significantly lighter than in the Edwards-Anderson model, suggesting the absence of strong fat-tail effects.

Tail-exponent analysis

To further quantify the nature of the distribution tails, we again evaluate the Hill estimator [176] for the inverse-gap variable $\eta = 1/\Delta_o$, using the same definition as in Eq. (8.25). The results, reported in Fig.8.13, display a slow, nearly power-law decrease of the estimated tail exponent $\alpha(N)$ with system size. However, in all cases-and for all tested fractions of top-order statistics $k = 0.1$, 0.2, and 0.3 the asymptotic values of α remain safely above 2. An extrapolation of this trend via a power-law fit $aN^b + c$ confirms this, yielding asymptotic values $c = 3.5(8)$ (for $k = 0.2$), $c = 4.3(4)$ (for $k = 0.3$), and $c = 4.1(5)$ (for $k = 0.4$). This implies that both the mean and the variance of the gap distribution remain finite, confirming the absence of a true fat-tailed regime.

8.5 Summary of the chapter

In this chapter, we have described and validated a Projection Quantum Monte Carlo (PQMC) methodology for the quantitative estimation of excitation gaps in disordered

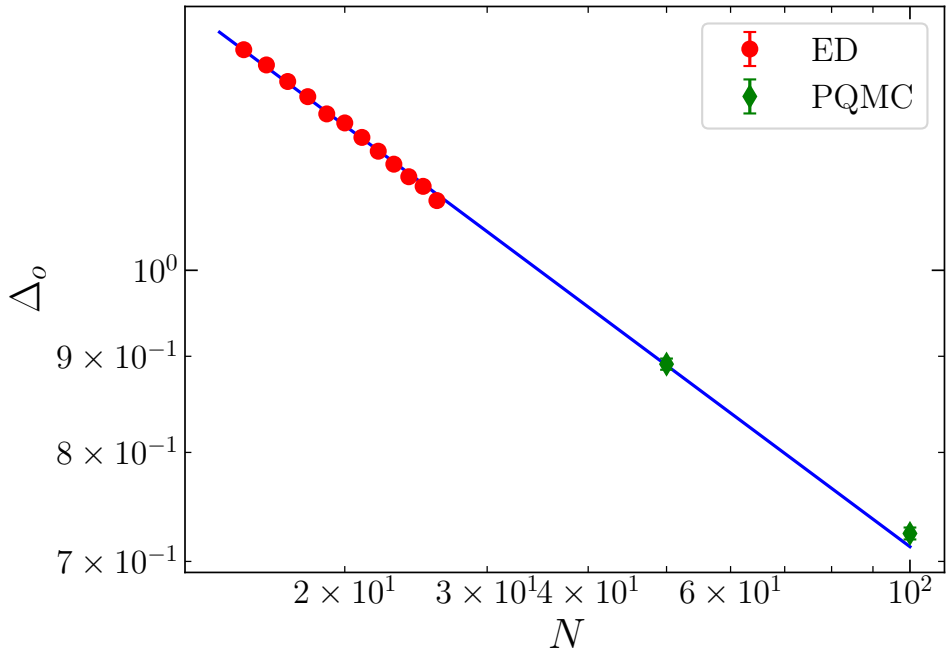

FIGURE 8.12: Scaling of the mean odd-parity gap Δ_o in the Sherrington–Kirkpatrick model at $\Gamma_c = 1.5$. Red points correspond to Exact Diagonalization data, while green markers denote PQMC results for $N = 50$ and $N = 100$. The blue line represents a power-law fit $\Delta_o \sim N^{-\theta}$ with $\theta = 0.321(4)$, confirming the persistence of the algebraic scaling observed in ED up to larger sizes.

quantum spin systems. The approach enables the extraction of both even- and odd-parity low excitations from imaginary-time correlation functions, thus overcoming the severe size limitations of Exact Diagonalization while maintaining a high level of statistical accuracy.

We first benchmarked the method on models with known exact solutions, such as the one-dimensional transverse-field Ising chain, showing that PQMC reproduces the full field dependence of the energy gap, including its closure at the quantum critical point, with excellent agreement with exact results. This validation confirmed the robustness of the iterative fitting strategy used to extract the asymptotic exponential decay from noisy imaginary-time correlators.

We then applied the method to paradigmatic disordered systems, focusing on the short-range Edwards-Anderson model and the fully connected Sherrington-Kirkpatrick model. By directly comparing PQMC results with ED data for small systems, we demonstrated that the method is capable of accurately reproducing individual excitation energies and full gap distributions even in the presence of strong quenched

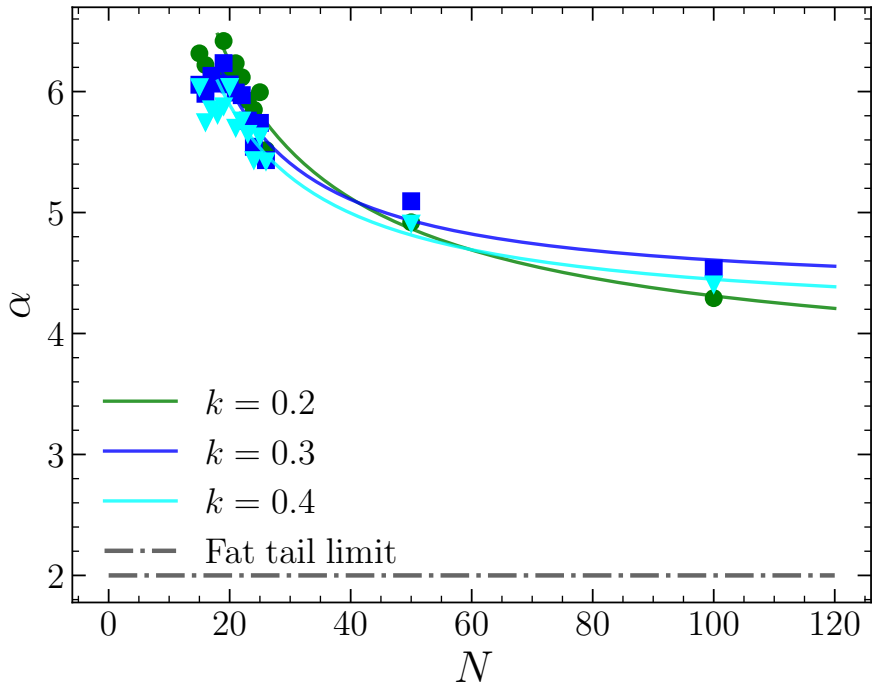

FIGURE 8.13: Hill estimator α of the tail exponent of the inverse-gap distribution for the SK model at $\Gamma_c = 1.5$, as a function of system size N. Different colors correspond to fractions $k = 0.2$, 0.3, and 0.4 of the largest η values used in the estimate. Although $\alpha(N)$ decreases mildly with increasing N, all values remain well above the fat-tail threshold $\alpha = 2$. This is confirmed by the extrapolated asymptotic values $c = 3.5(8)$ for $k = 0.2$, $c = 4.3(4)$ for $k = 0.3$, and $c = 4.1(5)$ for $k = 0.4$, indicating that the distributions retain finite mean and variance.

disorder and frustration.

Extending the analysis to larger system sizes revealed qualitatively distinct scaling behaviors between the two models.

In the SK model, the distribution of odd-parity gaps broadens only moderately with increasing system size, and the scaling of the mean gap follows a robust algebraic law $\Delta_o \sim N^{-\theta}$ with exponent $\theta \simeq 1/3$, consistent with previous ED findings. The tail-exponent analysis based on the Hill estimator confirms that the inverse-gap distribution retains finite moments ($\alpha > 2$), indicating the absence of true fat-tailed behavior. These results suggest that the critical dynamics of the SK model are governed by a well-defined algebraic scaling, unaffected by rare-event statistics.

In the EA model, by contrast, the situation is different. As the lattice size increases, the distributions of $\ln \Delta_o$ broaden significantly and develop heavy tails towards small

gap values, signaling the emergence of rare low-energy excitations. While the arithmetic mean of the gap still displays an approximately algebraic scaling, the typical gap exhibits a faster, super-algebraic suppression, suggesting that the gap distribution becomes increasingly dominated by rare disorder realizations. The analysis of the tail exponent α using the Hill estimator confirms this picture: α systematically decreases with system size and enters the regime $1 \lesssim \alpha \lesssim 2$, where the variance of the inverse-gap distribution diverges. This provides quantitative evidence that the quantum EA model at criticality develops a fat-tailed distribution of inverse gaps.

Overall, the PQMC framework introduced in this chapter establishes a scalable and reliable zero-temperature approach to probe the low-energy spectrum of disordered quantum systems far beyond the reach of Exact Diagonalization. These findings strengthen the view that rare-region effects and sample-to-sample fluctuations play a fundamental role in shaping the quantum critical dynamics of disordered systems.

Conclusions and outlook

The work presented in this thesis explored the integration of neural network representations of quantum states within the Projection Quantum Monte Carlo (PQMC) framework, with the overarching goal of extending the applicability of this method to complex, disordered, and frustrated quantum systems. Through the application of a self-learning PQMC algorithm to a variety of physical models, the thesis demonstrates how neural-guided Monte Carlo methods can overcome the traditional limitations of stochastic quantum simulations, enabling unbiased and high-precision investigations of strongly correlated systems. In the first part of the thesis, the theoretical foundations of PQMC and Neural Quantum States (NQS) were established. The PQMC algorithm was described in detail, emphasizing the role of importance sampling, continuous-time formulation, and forward-walking techniques in obtaining unbiased estimators of ground-state observables. The Restricted Boltzmann Machine (RBM) architecture was then introduced as a compact and expressive neural ansatz capable of approximating complex quantum correlations. Building upon these elements, a self-learning PQMC scheme was described, where the RBM guiding wave function is trained iteratively on the walker population. This approach eliminates the need for a separate variational optimization and leads to a fully adaptive, data-driven simulation framework. Benchmark tests on the one-dimensional Transverse Field Ising Model demonstrated that the method achieves exact convergence to the known analytical results, confirming its stability and accuracy. The algorithm was then applied to a series of increasingly complex systems. In two-dimensional quantum Edwards-Anderson spin glasses, a modified version of neural-guided PQMC, in which two replica of the same system is simulated in a single quantum Monte Carlo simulation, enabled the study of zero-temperature quantum phase transitions, allowing the computation of replica overlap distributions and the scaling of the spin-glass susceptibility. The results unveiled critical behavior and provided further insight into the interplay between disorder and quantum fluctuations in this model. The framework was then extended to

a realistic Hamiltonian describing amorphous arrays of Rydberg atoms where, as we show, positional disorder induces glassy magnetic phases. The simulations revealed the emergence of a quantum spin-glass transition in these amorphous geometries, an effect absent in periodic Rydberg lattices, thus offering a concrete theoretical benchmark for experiments on modern programmable quantum simulators. In frustrated triangular lattices, the method captured the subtle interplay between geometry, interaction range, and frustration, revealing the clock magnetic order through the order-by-disorder mechanism. These results underline the flexibility of the neural PQMC approach in handling both short-range and long-range interacting systems, which are of central interest in current quantum simulation platforms. A significant methodological advancement was introduced in the development of a hybrid quantum-classical pipeline, where experimental data from a Rydberg-atom quantum computer were incorporated into the PQMC simulations. By using the projective measures from a real quantum simulator based on neutral atoms technology as input for the guiding function, the PQMC algorithm achieved improved convergence, effectively combining classical Monte Carlo power with real quantum hardware information. This result illustrates a concrete pathway toward symbiotic quantum-classical simulation strategies and demonstrates how noisy quantum devices can already enhance classical algorithms in simulating complex quantum systems. Furthermore, a GPU-accelerated exact diagonalization (ED) code was designed to investigate the spectral properties of disordered quantum systems. Finite-size scaling of energy gaps in the Sherrington-Kirkpatrick (SK) and Edwards-Anderson (EA) models provided benchmarks for small systems and highlighted the presence of many-body localization effects at criticality. Finally, a PQMC-based energy gap estimator was developed and validated against exact diagonalization. This new tool allowed the extension of gap scaling analyses to larger system sizes, revealing heavy-tailed gap distributions in paradigmatic glassy systems and enabling precise extraction of critical exponents. The ability to characterize the low-energy spectrum within PQMC represents a major step forward, since it opens the possibility of investigating adiabatic dynamics and quantum annealing processes from purely classical simulations.

The results obtained in this thesis demonstrate that neural-guided PQMC methods can substantially enhance the efficiency and accuracy of quantum Monte Carlo simulations for systems affected by disorder, frustration, or experimental noise. The

combination of machine learning, high-performance computing, and stochastic algorithms provides a powerful framework capable of bridging the gap between theoretical models and quantum simulation experiments. From a broader perspective, the self-learning PQMC method represents a general strategy for embedding adaptive, data-driven representations of quantum states into traditional Monte Carlo frameworks. The hybrid quantum-classical protocol developed in this thesis could be generalized to other quantum hardware platforms, serving as a benchmark and calibration tool for near-term quantum devices. Moreover, the PQMC-based gap estimator provides a foundation for studying adiabatic quantum computation and the scaling of minimal gaps in realistic disordered Hamiltonians, directly relevant for the performance of quantum annealers. In summary, this thesis establishes a versatile and extensible framework for combining neural-network representations, Monte Carlo sampling, and quantum experimental data into a unified approach for the study of complex quantum systems.

The methodological framework and the results presented in this thesis open several promising avenues for future research. A direct extension would be to investigate the spin-glass phase transition in the presence of a longitudinal magnetic field, which explicitly breaks the Z_2 simmetry of the Edwards-Anderson model. This remains a central open question in the field of disordered quantum systems. This study is not only of theoretical interest but is also relevant for benchmarking experimental quantum simulators, in particular those based on Rydberg atoms arrays, where the longitudinal field term naturally appears in the system's Hamiltonian. Beyond the specific models studied here, the algorithmic frameworks presented in this work can be directly applied for the study and benchmark other quantum simulators. A natural next step is the study of systems with even longer-range interactions, such as those realized in trapped-ions platforms, where the interactions can exhibit a tunable long-range power-law decay.

Publications and related works

The research presented in this thesis is based on a series of works carried out during the doctoral program, some of which have already led to peer-reviewed publications and open-access data releases.

- Chapter 3 builds upon the results reported in Ref. [147], which investigates the zero-temperature properties of disordered quantum spin systems using neural-network-based Projection Quantum Monte Carlo methods. In connection with this work, we have released an open-access dataset on Zenodo [69] containing ground-state energies for multiple disorder realizations of the Edwards–Anderson model within the spin-glass phase. The same computational framework was also employed as a benchmark in Ref. [177], where I am listed as second author; however, the corresponding results are not presented in this thesis.

  - **Brodoloni, L.**, and S. Pilati. "Zero-temperature Monte Carlo simulations of two-dimensional quantum spin glasses guided by neural network states." Physical Review E 110.6 (2024): 065305.
  - **Brodoloni, L.**, Pilati, S. (2024). Ground-state dataset "Zero-temperature Monte Carlo simulations of two-dimensional quantum spin glasses guided by neural network states" [Data set]. Zenodo.
  - Cantori, S., **Brodoloni, L.**, Recchi, E., Costa, E., Juliá-Díaz, B., Pilati, S. (2025). Adaptive-basis sample-based neural diagonalization for quantum many-body systems. Methods, 4, 4.

- Chapter 4 is based on the study presented in Ref. [101], accepted for publication as a Letter in *Physical Review A*, which explores the emergence of glassy phases in amorphous Rydberg-atom arrays. The simulated amorphous lattice configurations discussed therein are openly available as a Zenodo dataset [143].

  - **Brodoloni, L.**, Vovrosh, J., Julià-Farré, S., Dauphin, A., Pilati, S. (2025). Spin-glass quantum phase transition in amorphous arrays of Rydberg atoms. arXiv preprint arXiv:2505.05117. (Accepted as a Letter in PRA)
  - **Brodoloni, L.**, Vovrosh, J., Julià-Farré, S., Dauphin, A., Pilati, S. (2025). Amorphous arrays used in "Spin-glass quantum phase transition in amorphous arrays of Rydberg atoms" [Data set]. Zenodo.

- The results of Chapter 6 were obtained during my research visit at the *Perimeter Institute for Theoretical Physics* (Ontario, Canada) and form the basis of a manuscript currently in preparation.

- The results concerning the gap analysis reported in Chapters 7 and 8 are part of an ongoing study, for which a dedicated manuscript is currently in preparation.

Funding and Acknowledgments

The works presented in this thesis were supported by the PNRR MUR Project No. PE0000023-NQSTI and by the Italian MUR under the PRIN2022 project "Hybrid algorithms for quantum simula- tors" No. 2022H77XB7. Partial contributions from the PRIN-PNRR 2022 MUR project "UEFA" (P2022NMBAJ) are also acknowledged. We acknowledge support from the CINECA ISCRA award IsCb2 NEMCASRA and the CINECA-INFN agreement, for the availability of high-performance computing resources and support and National Centre for HPC, Big Data and Quantum Computing (ICSC), CN00000013 Spoke 7 – Materials Molecular Sciences; CINECA awards ISCRA IsCc2_REASON, INF24_lincoln, and INF25_lincoln. We also acknowledge the EuroHPC Joint Undertaking for awarding this project access to the EuroHPC supercomputer LUMI, hosted by CSC (Finland) and the LUMI consortium through a EuroHPC Regular Access call.

I would like to express my special thanks to the Perimeter Institute for Theoretical Physics for hosting me as a visiting PhD researcher. The results presented in Chapter 6 were obtained during this research visit. I also wish to extend my sincere gratitude to PASQAL for hosting me and inviting me to give a seminar at their headquarters in Massy. The work presented in Chapter 4 was carried out in close collaboration with the PASQAL simulation group.

Bibliography

[1] W. M. C. Foulkes et al. "Quantum Monte Carlo simulations of solids". In: *Rev. Mod. Phys.* 73 (1 Jan. 2001), pp. 33–83. DOI: 10.1103/RevModPhys.73.33. URL: https://link.aps.org/doi/10.1103/RevModPhys.73.33.

[2] M. Troyer and U. Wiese. "Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations". In: *Physical review letters* 94.17 (2005), p. 170201.

[3] David M Ceperley. "An overview of quantum Monte Carlo methods". In: *English. In: Reviews in Mineralogy and Geochemistry* 71 (2018), pp. 129–135.

[4] I. Kosztin, B. Faber, and K. Schulten. "Introduction to the diffusion Monte Carlo method". In: *arXiv preprint physics/9702023* (1997).

[5] S. Giorgini, J. Boronat, and J. Casulleras. "Ground state of a homogeneous Bose gas: A diffusion Monte Carlo calculation". In: *Physical Review A* 60.6 (1999), p. 5129.

[6] J. Boronat. "Monte Carlo simulations at zero temperature: helium in one, two, and three dimensions". In: *Microscopic Approaches to Quantum Liquids in Confined Geometries*. Ed. by Eckhard Krotscheck and Jesus Navarro. World Scientific, Singapore, 2002, 2002. Chap. 2, pp. 21–90. DOI: https://doi.org/10.1142/9789812778475_0002.

[7] T. Parolini et al. "Tunneling in projective quantum Monte Carlo simulations with guiding wave functions". In: *Phys. Rev. B* 100 (21 Dec. 2019), p. 214303. DOI: 10.1103/PhysRevB.100.214303. URL: https://link.aps.org/doi/10.1103/PhysRevB.100.214303.

[8] M. Boninsegni and S. Moroni. "Population size bias in diffusion Monte Carlo". In: *Phys. Rev. E* 86 (5 Nov. 2012), p. 056712. DOI: 10.1103/PhysRevE.86.056712. URL: https://link.aps.org/doi/10.1103/PhysRevE.86.056712.

[9] K. Ghanem, N. Liebermann, and A. Alavi. "Population control bias and importance sampling in full configuration interaction quantum Monte Carlo". In: *Phys. Rev. B* 103 (15 Apr. 2021), p. 155135. DOI: 10.1103/PhysRevB.103.155135. URL: https://link.aps.org/doi/10.1103/PhysRevB.103.155135.

[10] J. Brand, M. Yang, and E. Pahl. "Stochastic differential equation approach to understanding the population control bias in full configuration interaction quantum Monte Carlo". In: *Phys. Rev. B* 105 (23 June 2022), p. 235144. DOI: 10.1103/PhysRevB.105.235144. URL: https://link.aps.org/doi/10.1103/PhysRevB.105.235144.

[11] G. Carleo and M. Troyer. "Solving the quantum many-body problem with artificial neural networks". In: *Science* 355.6325 (2017), pp. 602–606. DOI: 10.1126/science.aag2302.

[12] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski. "A learning algorithm for Boltzmann machines". In: *Cogn. Sci.* 9.1 (1985), pp. 147–169. ISSN: 0364-0213. DOI: https://doi.org/10.1016/S0364-0213(85)80012-4. URL: https://www.sciencedirect.com/science/article/pii/S0364021385800124.

[13] W. Zhang et al. "Ground state search by local and sequential updates of neural network quantum states". In: *Phys. Rev. B* 107 (16 Apr. 2023), p. 165149. DOI: 10.1103/PhysRevB.107.165149. URL: https://link.aps.org/doi/10.1103/PhysRevB.107.165149.

[14] W. Ren et al. "Towards the ground state of molecules via diffusion Monte Carlo on neural networks". In: *Nat. Commun.* 14.1 (2023), p. 1860. DOI: https://doi.org/10.1038/s41467-023-37609-3.

[15] L. L. Viteritti, R. Rende, and F. Becca. "Transformer variational wave functions for frustrated quantum spin systems". In: *Physical Review Letters* 130.23 (2023), p. 236401.

[16] S. Pilati, E. M. Inack, and P. Pieri. "Self-learning projective quantum Monte Carlo simulations guided by restricted Boltzmann machines". In: *Phys. Rev. E* 100 (4 Oct. 2019), p. 043301. DOI: 10.1103/PhysRevE.100.043301. URL: https://link.aps.org/doi/10.1103/PhysRevE.100.043301.

[17] E. M. Inack et al. "Projective quantum Monte Carlo simulations guided by unrestricted neural network states". In: *Physical Review B* 98.23 (2018), p. 235145.

[18] R. Blatt and C. F. Roos. "Quantum simulations with trapped ions". In: *Nat. Phys.* 8.4 (2012), pp. 277–284. DOI: https://doi.org/10.1038/nphys2252.

[19] A. A. Houck, H. E. Türeci, and J. Koch. "On-chip quantum simulation with superconducting circuits". In: *Nat. Phys.* 8.4 (2012), pp. 292–299. DOI: https://doi.org/10.1038/nphys2251.

[20] A. Browaeys and T. Lahaye. "Many-body physics with individually controlled Rydberg atoms". In: *Nature Physics* 16.2 (2020), pp. 132–142. DOI: https://doi.org/10.1038/s41567-019-0733-z.

[21] M. Lewenstein, A. Sanpera, and V. Ahufinger. *Ultracold Atoms in Optical Lattices: Simulating quantum many-body systems*. Oxford University Press, Oxford (UK), 2012.

[22] I. Bloch, J. Dalibard, and S. Nascimbene. "Quantum simulations with ultracold quantum gases". In: *Nat. Phys.* 8.4 (2012), pp. 267–276. DOI: https://doi.org/10.1038/nphys2259.

[23] L. Henriet et al. "Quantum computing with neutral atoms". In: *Quantum* 4 (Sept. 2020), p. 327. ISSN: 2521-327X. DOI: 10.22331/q-2020-09-21-327. URL: http://arxiv.org/abs/2006.12326 (visited on 08/17/2023).

[24] E. Farhi, J. Goldstone, and S. Gutmann. "A quantum approximate optimization algorithm". In: *arXiv preprint arXiv:1411.4028* (2014).

[25] J. Cohen, A. Khan, and C. Alexander. "Portfolio optimization of 40 stocks using the dwave quantum annealer". In: *arXiv preprint arXiv:2007.01430* (2020).

[26] P. Hauke et al. "Perspectives of quantum annealing: methods and implementations". In: *Rep. Prog. Phys.* 83.5 (May 2020), p. 054401. DOI: 10.1088/1361-6633/ab85b8. URL: https://dx.doi.org/10.1088/1361-6633/ab85b8.

[27] M. Bernaschi et al. "The quantum transition of the two-dimensional Ising spin glass". In: *Nature* (July 2024). ISSN: 1476-4687. DOI: 10.1038/s41586-024-07647-y. URL: https://doi.org/10.1038/s41586-024-07647-y.

[28] G. Mazzola, V. N. Smelyanskiy, and M. Troyer. "Quantum Monte Carlo tunneling from quantum chemistry to quantum annealing". In: *Phys. Rev. B* 96 (13 Oct. 2017), p. 134305. DOI: 10.1103/PhysRevB.96.134305. URL: https://link.aps.org/doi/10.1103/PhysRevB.96.134305.

[29] L. T. Brady and W. van Dam. "Quantum Monte Carlo simulations of tunneling in quantum adiabatic optimization". In: *Phys. Rev. A* 93 (3 Mar. 2016), p. 032304. DOI: 10.1103/PhysRevA.93.032304. URL: https://link.aps.org/doi/10.1103/PhysRevA.93.032304.

[30] K. Binder and A. P. Young. "Spin glasses: Experimental facts, theoretical concepts, and open questions". In: *Rev. Mod. Phys.* 58 (4 Oct. 1986), pp. 801–976. DOI: 10.1103/RevModPhys.58.801. URL: https://link.aps.org/doi/10.1103/RevModPhys.58.801.

[31] F. Becca and S. Sorella. *Quantum Monte Carlo approaches for correlated systems*. Cambridge University Press, 2017.

[32] L. Brodoloni and S. Pilati. "Zero-temperature Monte Carlo simulations of two-dimensional quantum spin glasses guided by neural network states". In: *Physical Review E* 110.6 (2024), p. 065305.

[33] A. D. King et al. "Quantum critical dynamics in a 5,000-qubit programmable spin glass". In: *Nature* 617.7959 (2023), pp. 61–66. DOI: https://doi.org/10.1038/s41586-023-05867-2.

[34] S. Mukherjee, A. Rajak, and B. K. Chakrabarti. "Classical-to-quantum crossover in the critical behavior of the transverse-field Sherrington-Kirkpatrick spin glass model". In: *Phys. Rev. E* 92 (4 Oct. 2015), p. 042107. DOI: 10.1103/PhysRevE.92.042107. URL: https://link.aps.org/doi/10.1103/PhysRevE.92.042107.

[35] T. Parolini et al. "Tunneling in projective quantum Monte Carlo simulations with guiding wave functions". In: *Physical Review B* 100.21 (2019), p. 214303.

[36] E. M. Inack et al. "Projective quantum Monte Carlo simulations guided by unrestricted neural network states". In: *Phys. Rev. B* 98 (23 Dec. 2018), p. 235145. DOI: 10.1103/PhysRevB.98.235145. URL: https://link.aps.org/doi/10.1103/PhysRevB.98.235145.

[37] J. Thijssen. *Computational physics*. Cambridge university press, 2007.

[38] N. Nemec. "Diffusion Monte Carlo: Exponential scaling of computational cost for large systems". In: *Physical Review B—Condensed Matter and Materials Physics* 81.3 (2010), p. 035119.

[39] L. Pollet, N. V. Prokof'ev, and B. V. Svistunov. "Stochastic lists: Sampling multivariable functions with population methods". In: *Phys. Rev. B* 98 (8 Aug. 2018), p. 085102. DOI: 10.1103/PhysRevB.98.085102. URL: https://link.aps.org/doi/10.1103/PhysRevB.98.085102.

[40] S. Sorella and L. Capriotti. "Green function Monte Carlo with stochastic reconfiguration: An effective remedy for the sign problem". In: *Phys. Rev. B* 61 (4 Jan. 2000), pp. 2599–2612. DOI: 10.1103/PhysRevB.61.2599. URL: https://link.aps.org/doi/10.1103/PhysRevB.61.2599.

[41] N. Trivedi and D. M. Ceperley. "Ground-state correlations of quantum antiferromagnets: A Green-function Monte Carlo study". In: *Physical Review B* 41.7 (1990), p. 4552.

[42] J. Casulleras and J. Boronat. "Unbiased estimators in quantum Monte Carlo methods: Application to liquid ^{4}He". In: *Phys. Rev. B* 52 (5 Aug. 1995), pp. 3654–3661. DOI: 10.1103/PhysRevB.52.3654. URL: https://link.aps.org/doi/10.1103/PhysRevB.52.3654.

[43] G. Cybenko. "Approximation by superpositions of a sigmoidal function". In: *Mathematics of control, signals and systems* 2.4 (1989), pp. 303–314.

[44] M. S. Moss et al. "Enhancing variational Monte Carlo simulations using a programmable quantum simulator". In: *Physical Review A* 109.3 (2024), p. 032410.

[45] S. Czischek et al. "Data-enhanced variational Monte Carlo simulations for Rydberg atom arrays". In: *Physical Review B* 105.20 (2022), p. 205108.

[46] K. Sprague and S. Czischek. "Variational Monte Carlo with large patched transformers". In: *Communications Physics* 7.1 (2024), p. 90.

[47] H. Lange et al. "From architectures to applications: A review of neural quantum states". In: *Quantum Science and Technology* (2024).

[48] E. M. Inack et al. "Understanding quantum tunneling using diffusion Monte Carlo simulations". In: *Phys. Rev. A* 97 (3 Mar. 2018), p. 032307. DOI: 10.1103/PhysRevA.97.032307. URL: https://link.aps.org/doi/10.1103/PhysRevA.97.032307.

[49] P. Charbonneau et al. *Spin Glass Theory and Far Beyond: Replica Symmetry Breaking after 40 Years*. World Scientific, 2023.

[50] S. Tanaka, R. Tamura, and B. K Chakrabarti. *Quantum spin glasses, annealing and computation*. Cambridge University Press, 2017.

[51] W. Guo G. Schumm H. Shao, F. Mila, and A. W. Sandvik. "Primary and secondary order parameters in the fully frustrated transverse-field Ising model on the square lattice". In: *Phys. Rev. B* 109 (14 Apr. 2024), p. L140408. DOI: 10.1103/PhysRevB.109.L140408. URL: https://link.aps.org/doi/10.1103/PhysRevB.109.L140408.

[52] P. Lai and Y. Y. Goldschmidt. "Monte Carlo Studies of the Ising Spin-Glass in a Transverse Field". In: *Europhys. Lett.* 13.4 (Oct. 1990), p. 289. DOI: 10.1209/0295-5075/13/4/001. URL: https://dx.doi.org/10.1209/0295-5075/13/4/001.

[53] H. Rieger and A. P. Young. "Zero-temperature quantum phase transition of a two-dimensional Ising spin glass". In: *Phys. Rev. Lett.* 72 (26 June 1994), pp. 4141–4144. DOI: 10.1103/PhysRevLett.72.4141. URL: https://link.aps.org/doi/10.1103/PhysRevLett.72.4141.

[54] M. Guo, R. N. Bhatt, and D. A. Huse. "Quantum critical behavior of a three-dimensional Ising spin glass in a transverse magnetic field". In: *Phys. Rev. Lett.* 72 (26 June 1994), pp. 4137–4140. DOI: 10.1103/PhysRevLett.72.4137. URL: https://link.aps.org/doi/10.1103/PhysRevLett.72.4137.

[55] J. V. Alvarez and F. Ritort. "Quantum Monte Carlo study of the infinite-range Ising spin glass in a transverse field". In: *J. Phys. A Math. Gen.* 29.23 (Dec. 1996), p. 7355. DOI: 10.1088/0305-4470/29/23/006. URL: https://dx.doi.org/10.1088/0305-4470/29/23/006.

[56] G. Mossi et al. "On the quantum spin glass transition on the Bethe lattice". In: *J. Stat. Mech.: Theory Exp.* 2017.1 (Jan. 2017), p. 013102. DOI: 10.1088/1742-5468/aa5286. URL: https://dx.doi.org/10.1088/1742-5468/aa5286.

[57] A. P. Young. "Stability of the quantum Sherrington-Kirkpatrick spin glass model". In: *Phys. Rev. E* 96 (3 Sept. 2017), p. 032112. DOI: 10.1103/PhysRevE.96.032112. URL: https://link.aps.org/doi/10.1103/PhysRevE.96.032112.

[58] A. Kiss, G. Zaránd, and I. Lovas. "Complete replica solution for the transverse field Sherrington-Kirkpatrick spin glass model with continuous-time quantum Monte Carlo method". In: *Phys. Rev. B* 109 (2 Jan. 2024), p. 024431. DOI: 10.1103/PhysRevB.109.024431. URL: https://link.aps.org/doi/10.1103/PhysRevB.109.024431.

[59] D. Lancaster and F. Ritort. "Solving the Schrödinger equation for the Sherrington - Kirkpatrick model in a transverse field". In: *J. Phys. A Math. Gen.* 30.4 (Feb. 1997), p. L41. DOI: 10.1088/0305-4470/30/4/001. URL: https://dx.doi.org/10.1088/0305-4470/30/4/001.

[60] P. M. Schindler et al. "Variational Ansatz for the Ground State of the Quantum Sherrington-Kirkpatrick Model". In: *Phys. Rev. Lett.* 129 (22 Nov. 2022), p. 220401. DOI: 10.1103/PhysRevLett.129.220401. URL: https://link.aps.org/doi/10.1103/PhysRevLett.129.220401.

[61] G. Carleo et al. "NetKet: A machine learning toolkit for many-body quantum systems". In: *SoftwareX* 10 (2019), p. 100311. ISSN: 2352-7110. DOI: https://doi.org/10.1016/j.softx.2019.100311. URL: https://www.sciencedirect.com/science/article/pii/S2352711019300974.

[62] S. Pilati and P. Pieri. "Simulating disordered quantum Ising chains via dense and sparse restricted Boltzmann machines". In: *Phys. Rev. E* 101 (6 June 2020), p. 063308. DOI: 10.1103/PhysRevE.101.063308. URL: https://link.aps.org/doi/10.1103/PhysRevE.101.063308.

[63] G. Parisi. "Order Parameter for Spin-Glasses". In: *Phys. Rev. Lett.* 50 (24 June 1983), pp. 1946–1948. DOI: 10.1103/PhysRevLett.50.1946. URL: https://link.aps.org/doi/10.1103/PhysRevLett.50.1946.

[64] J. D. Reger, R. N. Bhatt, and A. P. Young. "Monte Carlo study of the order-parameter distribution in the four-dimensional Ising spin glass". In: *Phys. Rev. Lett.* 64 (16 Apr. 1990), pp. 1859–1862. DOI: 10.1103/PhysRevLett.64.1859. URL: https://link.aps.org/doi/10.1103/PhysRevLett.64.1859.

[65] E. Marinari and F. Zuliani. "Numerical simulations of the four-dimensional Edwards-Anderson spin glass with binary couplings". In: *J. Phys. A Math. Gen.* 32.43 (Oct. 1999), p. 7447. DOI: 10.1088/0305-4470/32/43/301. URL: https://dx.doi.org/10.1088/0305-4470/32/43/301.

[66] E. Marinari et al. "Replica symmetry breaking in short-range spin glasses: Theoretical foundations and numerical evidences". In: *J. Stat. Phys.* 98 (2000), pp. 973–1074. DOI: https://doi.org/10.1023/A:1018607809852.

[67] H. G. Katzgraber and A. P. Young. "Monte Carlo studies of the one-dimensional Ising spin glass with power-law interactions". In: *Phys. Rev. B* 67 (13 Apr. 2003), p. 134410. DOI: 10.1103/PhysRevB.67.134410. URL: https://link.aps.org/doi/10.1103/PhysRevB.67.134410.

[68] M. A. Moore. "Droplet-scaling versus replica symmetry breaking debate in spin glasses revisited". In: *Phys. Rev. E* 103 (6 June 2021), p. 062111. DOI: 10.1103/PhysRevE.103.062111. URL: https://link.aps.org/doi/10.1103/PhysRevE.103.062111.

[69] L. Brodoloni and S. Pilati. *Ground-state dataset "Zero-temperature Monte Carlo simulations of two-dimensional quantum spin glasses guided by neural network states"*. Zenodo, June 2024. DOI: 10.5281/zenodo.11534963. URL: https://doi.org/10.5281/zenodo.11534963.

[70] L. Stella and G. E. Santoro. "Quantum annealing of an Ising spin-glass by Green's function Monte Carlo". In: *Phys. Rev. E* 75 (3 Mar. 2007), p. 036703. DOI: 10.1103/PhysRevE.75.036703. URL: https://link.aps.org/doi/10.1103/PhysRevE.75.036703.

[71] C. Krämer et al. "Quantum-critical properties of the one- and two-dimensional random transverse-field Ising model from large-scale quantum Monte Carlo simulations". In: *SciPost Phys.* 17 (2024), p. 061. DOI: 10.21468/SciPostPhys.17.2.061.

[72] H. W. J. Blöte and Y. Deng. "Cluster Monte Carlo simulation of the transverse Ising model". In: *Phys. Rev. E* 66 (6 Dec. 2002), p. 066110. DOI: 10.1103/PhysRevE.66.066110. URL: https://link.aps.org/doi/10.1103/PhysRevE.66.066110.

[73] R. Oliveira S. Bravyi D.P. Divincenzo and B. M. Terhal. "The complexity of stoquastic local Hamiltonian problems". In: *Quantum Info. Comput.* 8.5 (May 2008), pp. 361–385. ISSN: 1533-7146. DOI: https://doi.org/10.26421/QIC8.5-1.

[74] E. Marinari, G. Parisi, and J. J. Ruiz-Lorenzo. "Numerical simulations of spin glass systems". In: *Spin Glasses and Random Fields*. World Scientific, 1997, pp. 59–98. DOI: 10.1142/9789812819437_0003.

[75] A. Lage-Castellanos et al. "Inference algorithm for finite-dimensional spin glasses: Belief propagation on the dual lattice". In: *Phys. Rev. E* 84 (4 Oct. 2011), p. 046706. DOI: 10.1103/PhysRevE.84.046706. URL: https://link.aps.org/doi/10.1103/PhysRevE.84.046706.

[76] H. Leschke et al. "Existence of Replica-Symmetry Breaking in Quantum Glasses". In: *Phys. Rev. Lett.* 127 (20 Nov. 2021), p. 207204. DOI: 10.1103/PhysRevLett.127.207204. URL: https://link.aps.org/doi/10.1103/PhysRevLett.127.207204.

[77] O. Melchert. "autoScale.py - A program for automatic finite-size scaling analyses: A user's guide". In: *arXiv:0910.5403* (2009). DOI: https://doi.org/10.48550/arXiv.0910.5403.

[78] J. T. Chayes et al. "Finite-Size Scaling and Correlation Lengths for Disordered Systems". In: *Phys. Rev. Lett.* 57 (24 Dec. 1986), pp. 2999–3002. DOI: 10.1103/PhysRevLett.57.2999. URL: https://link.aps.org/doi/10.1103/PhysRevLett.57.2999.

[79] M. A. Moore, H. Bokil, and B. Drossel. "Evidence for the Droplet Picture of Spin Glasses". In: *Phys. Rev. Lett.* 81 (19 Nov. 1998), pp. 4252–4255. DOI: 10.1103/PhysRevLett.81.4252. URL: https://link.aps.org/doi/10.1103/PhysRevLett.81.4252.

[80] H. Bokil, B. Drossel, and M. A. Moore. "Influence of critical behavior on the spin-glass phase". In: *Phys. Rev. B* 62 (2 July 2000), pp. 946–951. DOI: 10.1103/

PhysRevB.62.946. URL: https://link.aps.org/doi/10.1103/PhysRevB.62.946.

[81] H. G. Katzgraber, M. Palassini, and A. P. Young. "Monte Carlo simulations of spin glasses at low temperatures". In: *Phys. Rev. B* 63 (18 Apr. 2001), p. 184422. DOI: 10.1103/PhysRevB.63.184422. URL: https://link.aps.org/doi/10.1103/PhysRevB.63.184422.

[82] M. Bernaschi et al. "The QISG suite: High-performance codes for studying quantum Ising spin glasses". In: *Comput. Phys. Commun.* 298 (2024), p. 109101. ISSN: 0010-4655. DOI: https://doi.org/10.1016/j.cpc.2024.109101. URL: https://www.sciencedirect.com/science/article/pii/S0010465524000249.

[83] G. E. Santoro et al. "Theory of Quantum Annealing of an Ising Spin Glass". In: *Science* 295.5564 (2002), pp. 2427–2430. DOI: 10.1126/science.1068774.

[84] S. V. Isakov et al. "Understanding Quantum Tunneling through Quantum Monte Carlo Simulations". In: *Phys. Rev. Lett.* 117 (18 Oct. 2016), p. 180402. DOI: 10.1103/PhysRevLett.117.180402. URL: https://link.aps.org/doi/10.1103/PhysRevLett.117.180402.

[85] E. Andriyash and M. H. Amin. "Can quantum Monte Carlo simulate quantum annealing?" In: *arXiv preprint arXiv:1703.09277* (2017). DOI: https://doi.org/10.48550/arXiv.1703.09277.

[86] J. Smith et al. "Many-body localization in a quantum simulator with programmable random disorder". In: *Nat. Phys.* 12.10 (2016), pp. 907–911. DOI: https://doi.org/10.1038/nphys3783.

[87] H. Labuhn et al. "Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum Ising models". In: *Nature* 534.7609 (2016), pp. 667–670. DOI: https://doi.org/10.1038/nature18274.

[88] S. Ebadi et al. "Quantum phases of matter on a 256-atom programmable quantum simulator". In: *Nature* 595.7866 (2021), pp. 227–232. DOI: https://doi.org/10.1038/s41586-021-03582-4.

[89] A. Angelone, F. Mezzacapo, and G. Pupillo. "Superglass Phase of Interaction-Blockaded Gases on a Triangular Lattice". In: *Phys. Rev. Lett.* 116 (13 Apr. 2016), p. 135303. DOI: 10.1103/PhysRevLett.116.135303. URL: https://link.aps.org/doi/10.1103/PhysRevLett.116.135303.

[90] Z. Yan et al. "Emergent Glassy Behavior in a Kagome Rydberg Atom Array". In: *Phys. Rev. Lett.* 130 (20 May 2023), p. 206501. DOI: 10.1103/PhysRevLett.130.206501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.130.206501.

[91] T. Zhang and Z. Cai. "Quantum Slush State in Rydberg Atom Arrays". In: *Phys. Rev. Lett.* 132 (20 May 2024), p. 206503. DOI: 10.1103/PhysRevLett.132.206503. URL: https://link.aps.org/doi/10.1103/PhysRevLett.132.206503.

[92] P. He et al. "Floquet Amorphous Topological Orders in a Rydberg Glass". In: *arXiv:2404.18512* (2024). DOI: https://doi.org/10.48550/arXiv.2404.18512.

[93] M. Hibat-Allah et al. "Recurrent neural network wave functions for Rydberg atom arrays on kagome lattice". In: *arXiv:2405.20384* (2024). DOI: https://doi.org/10.48550/arXiv.2405.20384.

[94] S. Julià-Farré, J. Vovrosh, and A. Dauphin. "Amorphous quantum magnets in a two-dimensional Rydberg atom array". In: *Phys. Rev. A* 110 (1 July 2024), p. 012602. DOI: 10.1103/PhysRevA.110.012602.

[95] L. Sanchez-Palencia and L. Santos. "Bose-Einstein condensates in optical quasicrystal lattices". In: *Phys. Rev. A* 72 (5 Nov. 2005), p. 053607. DOI: 10.1103/PhysRevA.72.053607. URL: https://link.aps.org/doi/10.1103/PhysRevA.72.053607.

[96] S. Iyer et al. "Many-body localization in a quasiperiodic system". In: *Phys. Rev. B* 87 (13 Apr. 2013), p. 134202. DOI: 10.1103/PhysRevB.87.134202. URL: https://link.aps.org/doi/10.1103/PhysRevB.87.134202.

[97] S. Pilati and V. K. Varma. "Localization of interacting Fermi gases in quasiperiodic potentials". In: *Phys. Rev. A* 95 (1 Jan. 2017), p. 013613. DOI: 10.1103/PhysRevA.95.013613. URL: https://link.aps.org/doi/10.1103/PhysRevA.95.013613.

[98] R. Menu and T. Roscilde. "Anomalous Diffusion and Localization in a Positionally Disordered Quantum Spin Array". In: *Phys. Rev. Lett.* 124 (13 Apr. 2020), p. 130604. DOI: 10.1103/PhysRevLett.124.130604. URL: https://link.aps.org/doi/10.1103/PhysRevLett.124.130604.

[99] S. Ghosh et al. "Many-body localization due to correlated disorder in Fock space". In: *Phys. Rev. B* 99 (16 Apr. 2019), p. 165131. DOI: 10.1103/PhysRevB.99.165131. URL: https://link.aps.org/doi/10.1103/PhysRevB.99.165131.

[100] M. Ciardi et al. "Quasicrystalline Bose Glass in the Absence of Disorder and Quasidisorder". In: *Phys. Rev. Lett.* 131 (17 Oct. 2023), p. 173402. DOI: 10.1103/PhysRevLett.131.173402. URL: https://link.aps.org/doi/10.1103/PhysRevLett.131.173402.

[101] L. Brodoloni et al. "Spin-glass quantum phase transition in amorphous arrays of Rydberg atoms". In: *Phys. Rev. A* 112 (5 Nov. 2025), p. L051303. DOI: 10.1103/r6dq-2ztr. URL: https://link.aps.org/doi/10.1103/r6dq-2ztr.

[102] F. Barahona. "On the computational complexity of Ising spin glass models". In: *Journal of Physics A: Mathematical and General* 15.10 (Oct. 1982), p. 3241. DOI: 10.1088/0305-4470/15/10/028. URL: https://dx.doi.org/10.1088/0305-4470/15/10/028.

[103] Y. Y. Goldschmidt. "Solvable model of the quantum spin glass in a transverse field". In: *Phys. Rev. B* 41 (7 Mar. 1990), pp. 4858–4861. DOI: 10.1103/PhysRevB.41.4858. URL: https://link.aps.org/doi/10.1103/PhysRevB.41.4858.

[104] G. Büttner and K. D. Usadel. "Stability analysis of an Ising spin glass with transverse field". In: *Phys. Rev. B* 41 (1 Jan. 1990), pp. 428–431. DOI: 10.1103/PhysRevB.41.428. URL: https://link.aps.org/doi/10.1103/PhysRevB.41.428.

[105] D. Thirumalai, Q. Li, and T. R. Kirkpatrick. "Infinite-range Ising spin glass in a transverse field". In: *J. Phys. A* 22.16 (Aug. 1989), p. 3339. DOI: 10.1088/0305-4470/22/16/023. URL: https://dx.doi.org/10.1088/0305-4470/22/16/023.

[106] N. Read, S. Sachdev, and J. Ye. "Landau theory of quantum spin glasses of rotors and Ising spins". In: *Phys. Rev. B* 52 (1 July 1995), pp. 384–410. DOI: 10.

1103/PhysRevB.52.384. URL: https://link.aps.org/doi/10.1103/PhysRevB.52.384.

[107] P. Ray, B. K. Chakrabarti, and A. Chakrabarti. "Sherrington-Kirkpatrick model in a transverse field: Absence of replica symmetry breaking due to quantum fluctuations". In: *Phys. Rev. B* 39 (16 June 1989), pp. 11828–11832. DOI: 10.1103/PhysRevB.39.11828. URL: https://link.aps.org/doi/10.1103/PhysRevB.39.11828.

[108] C. Manai and S. Warzel. "The de Almeida–Thouless line in hierarchical quantum spin glasses". In: *Journal of Statistical Physics* 186.1 (2022), p. 14.

[109] C. R. Laumann, A. Pal, and A. Scardicchio. "Many-Body Mobility Edge in a Mean-Field Quantum Spin Glass". In: *Phys. Rev. Lett.* 113 (20 Nov. 2014), p. 200405. DOI: 10.1103/PhysRevLett.113.200405. URL: https://link.aps.org/doi/10.1103/PhysRevLett.113.200405.

[110] J. A. Kjäll, J. H. Bardarson, and F. Pollmann. "Many-Body Localization in a Disordered Quantum Ising Chain". In: *Phys. Rev. Lett.* 113 (10 Sept. 2014), p. 107204. DOI: 10.1103/PhysRevLett.113.107204. URL: https://link.aps.org/doi/10.1103/PhysRevLett.113.107204.

[111] L. Rademaker and D. A. Abanin. "Slow Nonthermalizing Dynamics in a Quantum Spin Glass". In: *Phys. Rev. Lett.* 125 (26 Dec. 2020), p. 260405. DOI: 10.1103/PhysRevLett.125.260405. URL: https://link.aps.org/doi/10.1103/PhysRevLett.125.260405.

[112] P. Schauss. "Quantum simulation of transverse Ising models with Rydberg atoms". In: *Quantum Sci. Technol.* 3.2 (Jan. 2018), p. 023001. DOI: 10.1088/2058-9565/aa9c59. URL: https://dx.doi.org/10.1088/2058-9565/aa9c59.

[113] S. Humeniuk. "Quantum Monte Carlo study of long-range transverse-field Ising models on the triangular lattice". In: *Phys. Rev. B* 93 (10 Mar. 2016), p. 104412. DOI: 10.1103/PhysRevB.93.104412. URL: https://link.aps.org/doi/10.1103/PhysRevB.93.104412.

[114] R. Samajdar et al. "Complex Density Wave Orders and Quantum Phase Transitions in a Model of Square-Lattice Rydberg Atom Arrays". In: *Phys. Rev. Lett.*

124 (10 Mar. 2020), p. 103601. DOI: 10.1103/PhysRevLett.124.103601. URL: https://link.aps.org/doi/10.1103/PhysRevLett.124.103601.

[115] M. Kalinowski et al. "Bulk and boundary quantum phase transitions in a square Rydberg atom array". In: *Phys. Rev. B* 105 (17 May 2022), p. 174417. DOI: 10.1103/PhysRevB.105.174417. URL: https://link.aps.org/doi/10.1103/PhysRevB.105.174417.

[116] M. J. O'Rourke and G. K. L. Chan. "Entanglement in the quantum phases of an unfrustrated Rydberg atom array". In: *Nat. Commun.* 14.1 (2023), p. 5397. DOI: https://doi.org/10.1038/s41467-023-41166-0.

[117] S. Guo et al. "Order by disorder and an emergent Kosterlitz-Thouless phase in a triangular Rydberg array". In: *Phys. Rev. A* 108 (5 Nov. 2023), p. 053314. DOI: 10.1103/PhysRevA.108.053314. URL: https://link.aps.org/doi/10.1103/PhysRevA.108.053314.

[118] S. Yang and J. B. Xu. "Density-wave-ordered phases of Rydberg atoms on a honeycomb lattice". In: *Phys. Rev. E* 106 (3 Sept. 2022), p. 034121. DOI: 10.1103/PhysRevE.106.034121. URL: https://link.aps.org/doi/10.1103/PhysRevE.106.034121.

[119] M. Sarkar et al. "Quantum order-by-disorder induced phase transition in Rydberg ladders with staggered detuning". In: *SciPost Phys.* 14 (2023), p. 004. DOI: 10.21468/SciPostPhys.14.1.004. URL: https://scipost.org/10.21468/SciPostPhys.14.1.004.

[120] R. Samajdar et al. "Quantum phases of Rydberg atoms on a kagome lattice". In: *Proc. Natl. Acad. Sci. U.S.A.* 118.4 (2021), e2015785118. DOI: 10.1073/pnas.2015785118. URL: https://www.pnas.org/doi/abs/10.1073/pnas.2015785118.

[121] P. Scholl et al. "Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms". In: *Nature* 595.7866 (2021), pp. 233–238. DOI: https://doi.org/10.1038/s41586-021-03585-1.

[122] C. Chen et al. "Continuous symmetry breaking in a two-dimensional Rydberg array". In: *Nature* 616.7958 (2023), pp. 691–695. DOI: https://doi.org/10.1038/s41586-023-05859-2.

[123] H. Weimer and H. P. Büchler. "Two-Stage Melting in Systems of Strongly Interacting Rydberg Atoms". In: *Phys. Rev. Lett.* 105 (23 Nov. 2010), p. 230403. DOI: 10.1103/PhysRevLett.105.230403.

[124] M. Rader and A. M. Läuchli. "Floating phases in one-dimensional Rydberg Ising chains". In: *arXiv:1908.02068* (2019). DOI: https://doi.org/10.48550/arXiv.1908.02068.

[125] X. J. Yu et al. "Fidelity susceptibility as a diagnostic of the commensurate-incommensurate transition: A revisit of the programmable Rydberg chain". In: *Phys. Rev. B* 106 (16 Oct. 2022), p. 165124. DOI: 10.1103/PhysRevB.106.165124. URL: https://link.aps.org/doi/10.1103/PhysRevB.106.165124.

[126] J. Soto-Garcia and N. Chepiga. "Numerical investigation of quantum phases and phase transitions in a two-leg ladder of Rydberg atoms". In: *Phys. Rev. Res.* 7 (1 Feb. 2025), p. 013215. DOI: 10.1103/PhysRevResearch.7.013215. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.7.013215.

[127] J. Zhang et al. "Probing quantum floating phases in Rydberg atom arrays". In: *Nat. Commun.* 16.1 (2025), p. 712. DOI: https://doi.org/10.1038/s41467-025-55947-2.

[128] M. Hibat-Allah et al. "Recurrent neural network wave functions for Rydberg atom arrays on kagome lattice". In: *arXiv:2405.20384* (2024). DOI: https://doi.org/10.48550/arXiv.2405.20384.

[129] S. Ebadi et al. "Quantum optimization of maximum independent set using Rydberg atom arrays". In: *Science* 376.6598 (2022), pp. 1209–1215. DOI: 10.1126/science.abo6587. eprint: https://www.science.org/doi/pdf/10.1126/science.abo6587.

[130] M. T. Nguyen et al. "Quantum Optimization with Arbitrary Connectivity Using Rydberg Atom Arrays". In: *PRX Quantum* 4 (1 Feb. 2023), p. 010316. DOI: 10.1103/PRXQuantum.4.010316. URL: https://link.aps.org/doi/10.1103/PRXQuantum.4.010316.

[131] M. Kim et al. "Rydberg quantum wires for maximum independent set problems". In: *Nature Physics* 18.7 (2022), pp. 755–759. DOI: https://doi.org/10.1038/s41567-022-01629-5.

[132] L. Leclerc et al. "Implementing transferable annealing protocols for combinatorial optimization on neutral-atom quantum processors: A case study on smart charging of electric vehicles". In: *Phys. Rev. A* 111 (3 Mar. 2025), p. 032611. DOI: 10.1103/PhysRevA.111.032611. URL: https://link.aps.org/doi/10.1103/PhysRevA.111.032611.

[133] P. Cazals et al. "Identifying hard native instances for the maximum independent set problem on neutral atoms quantum processors". In: *arXiv:2502.04291* (2025). DOI: https://doi.org/10.48550/arXiv.2502.04291.

[134] B. Albrecht et al. "Quantum feature maps for graph machine learning on a neutral atom quantum processor". In: *Phys. Rev. A* 107 (4 Apr. 2023), p. 042615. DOI: 10.1103/PhysRevA.107.042615. URL: https://link.aps.org/doi/10.1103/PhysRevA.107.042615.

[135] M. Kornjača et al. "Large-scale quantum reservoir learning with an analog quantum computer". In: *arXiv preprint arXiv:2407.02553* (2024).

[136] L. Leclerc et al. "Financial risk management on a neutral atom quantum processor". In: *Phys. Rev. Res.* 5 (4 Nov. 2023), p. 043117. DOI: 10.1103/PhysRevResearch.5.043117. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.5.043117.

[137] G. Bornet et al. "Enhancing a Many-Body Dipolar Rydberg Tweezer Array with Arbitrary Local Controls". In: *Phys. Rev. Lett.* 132 (26 June 2024), p. 263601. DOI: 10.1103/PhysRevLett.132.263601. URL: https://link.aps.org/doi/10.1103/PhysRevLett.132.263601.

[138] H. Yeo, H. E. Kim, and K. Jeong. "Approximating Maximum Independent Set on Rydberg Atom Arrays Using Local Detunings". In: *Adv. Quantum Technol.* 8.1 (2025), p. 2400291. DOI: https://doi.org/10.1002/qute.202400291.

[139] K. Goswami et al. "Solving optimization problems with local light-shift encoding on Rydberg quantum annealers". In: *Phys. Rev. Res.* 6 (2 Apr. 2024), p. 023031. DOI: 10.1103/PhysRevResearch.6.023031. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.6.023031.

[140] M. Kornjača et al. *Large-scale quantum reservoir learning with an analog quantum computer*. 2024. arXiv: 2407.02553 [quant-ph]. URL: https://arxiv.org/abs/2407.02553.

[141] G. M. Crosswhite, A. C. Doherty, and G. Vidal. "Applying matrix product operators to model systems with long-range interactions". In: *Phys. Rev. B* 78 (3 July 2008), p. 035116. DOI: 10.1103/PhysRevB.78.035116. URL: https://link.aps.org/doi/10.1103/PhysRevB.78.035116.

[142] S. Li, H. Shao, and A. W. Sandvik. "Ground State of the $S = 1/2$ Heisenberg Spin Chain with Random Ferromagnetic and Antiferromagnetic Couplings". In: *Phys. Rev. Lett.* 134 (8 Feb. 2025), p. 086501. DOI: 10.1103/PhysRevLett.134.086501. URL: https://link.aps.org/doi/10.1103/PhysRevLett.134.086501.

[143] L. Brodoloni et al. *Amorphous arrays used in "Spin-glass quantum phase transition in amorphous arrays of Rydberg atoms"*. Zenodo, Apr. 2025. DOI: 10.5281/zenodo.15276169. URL: https://doi.org/10.5281/zenodo.15276169.

[144] R. Moessner, S. L. Sondhi, and P. Chandra. "Two-Dimensional Periodic Frustrated Ising Models in a Transverse Field". In: *Phys. Rev. Lett.* 84 (19 May 2000), pp. 4457–4460. DOI: 10.1103/PhysRevLett.84.4457. URL: https://link.aps.org/doi/10.1103/PhysRevLett.84.4457.

[145] P. Narasimhan et al. "Simulating the transverse-field Ising model on the kagome lattice using a programmable quantum annealer". In: *Phys. Rev. B* 110 (5 Aug. 2024), p. 054432. DOI: 10.1103/PhysRevB.110.054432. URL: https://link.aps.org/doi/10.1103/PhysRevB.110.054432.

[146] A.P. Young. "Numerical Simulations of Spin Glasses: Methods and Some Recent Results". In: *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology Volume 2*. Ed. by Mauro Ferrario, Giovanni Ciccotti, and Kurt Binder. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 31–44. ISBN: 978-3-540-35284-6. DOI: 10.1007/3-540-35284-8_2.

[147] L. Brodoloni and S. Pilati. "Zero-temperature Monte Carlo simulations of two-dimensional quantum spin glasses guided by neural network states". In: *Phys. Rev. E* 110 (6 Dec. 2024), p. 065305. DOI: 10.1103/PhysRevE.110.065305. URL: https://link.aps.org/doi/10.1103/PhysRevE.110.065305.

[148] A. Sorge. *pyfssa 0.7.6*. Zenodo. 2025. DOI: https://zenodo.org/badge/latestdoi/6089/andsor/pyfssa.

[149] J. Houdayer and A. K. Hartmann. "Low-temperature behavior of two-dimensional Gaussian Ising spin glasses". In: *Phys. Rev. B* 70 (1 July 2004), p. 014418. DOI: 10.1103/PhysRevB.70.014418. URL: https://link.aps.org/doi/10.1103/PhysRevB.70.014418.

[150] L. Leuzzi et al. "Dilute One-Dimensional Spin Glasses with Power Law Decaying Interactions". In: *Phys. Rev. Lett.* 101 (10 Sept. 2008), p. 107203. DOI: 10.1103/PhysRevLett.101.107203.

[151] H. D. Berry and P. Singh. "Monte Carlo simulation of 2D Ising spin glass with power law decaying interactions". In: *Journal of superconductivity and novel magnetism* 26.4 (2013), pp. 991–994. DOI: https://doi.org/10.1007/s10948-012-1908-3.

[152] M. Qiao et al. "Tunable quantum simulation of spin models with a two-dimensional ion crystal". In: *Nature Physics* 20.4 (2024), pp. 623–630. DOI: https://doi.org/10.1038/s41567-023-02378-9.

[153] H. J. Manetsch et al. "A tweezer array with 6100 highly coherent atomic qubits". In: *arXiv:2403.12021* (2024). DOI: https://doi.org/10.48550/arXiv.2403.12021.

[154] G. Pichard et al. "Rearrangement of individual atoms in a 2000-site optical-tweezer array at cryogenic temperatures". In: *Phys. Rev. Appl.* 22 (2 Aug. 2024), p. 024073. DOI: 10.1103/PhysRevApplied.22.024073. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.22.024073.

[155] F. Liu and L. M. Duan. "Computational characteristics of the random-field Ising model with long-range interaction". In: *Phys. Rev. A* 108 (1 July 2023), p. 012415. DOI: 10.1103/PhysRevA.108.012415. URL: https://link.aps.org/doi/10.1103/PhysRevA.108.012415.

[156] C. Lacroix, P. Mendels, and F. Mila. *Introduction to frustrated magnetism: materials, experiments, theory*. Vol. 164. Springer Science & Business Media, 2011.

[157] H. T. Diep et al. *Frustrated spin systems*. World scientific, 2013.

[158] S. V. Isakov and R. Moessner. "Interplay of quantum and thermal fluctuations in a frustrated magnet". In: *Physical Review B* 68.10 (2003), p. 104409.

[159] Y. Da Liao et al. "Phase diagram of the quantum Ising model on a triangular lattice under external field". In: *Physical Review B* 103.10 (2021), p. 104416.

[160] J. W. Britton et al. "Engineered two-dimensional Ising interactions in a trapped-ion quantum simulator with hundreds of spins". In: *Nature* 484.7395 (2012), pp. 489–492.

[161] SN Saadatmand, SD Bartlett, and IP McCulloch. "Phase diagram of the quantum Ising model with long-range interactions on an infinite-cylinder triangular lattice". In: *Physical Review B* 97.15 (2018), p. 155116.

[162] G. Scriva et al. "Accelerating equilibrium spin-glass simulations using quantum annealers via generative deep learning". In: *SciPost Physics* 15.1 (2023), p. 018.

[163] C. X. Li, S. Yang, and J. B. Xu. "Quantum phases of Rydberg atoms on a frustrated triangular-lattice array". In: *Optics Letters* 47.5 (2022), pp. 1093–1096.

[164] P. Narasimhan et al. "Simulating the transverse-field Ising model on the kagome lattice using a programmable quantum annealer". In: *Physical Review B* 110.5 (2024), p. 054432.

[165] E. Merali, I. J.S. De Vlugt, and R. G. Melko. "Stochastic series expansion quantum Monte Carlo for Rydberg arrays". In: *SciPost Physics Core* 7.2 (2024), p. 016.

[166] S. Knysh. "Computational bottlenecks of quantum annealing". In: *arXiv preprint arXiv:1506.08608* (2015).

[167] P. Weinberg and M. Bukov. "QuSpin: a Python package for dynamics and exact diagonalisation of quantum many body systems. Part II: bosons, fermions and higher spins". In: *SciPost Phys.* 7 (2019), p. 020. DOI: 10.21468/SciPostPhys.7.2.020. URL: https://scipost.org/10.21468/SciPostPhys.7.2.020.

[168] R. Okuta et al. "CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations". In: *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*. 2017. URL: http://learningsys.org/nips17/assets/papers/paper_16.pdf.

[169] Y. W. Koh. "Effects of low-lying excitations on ground-state energy and energy gap of the Sherrington-Kirkpatrick model in a transverse field". In: *Physical Review B* 93.13 (2016), p. 134202.

[170] K. Takahashi and Y. Matsuda. "Effect of random fluctuations on quantum spin-glass transitions at zero temperature". In: *Journal of the Physical Society of Japan* 79.4 (2010), p. 043712.

[171] A. P. Young and H. Rieger. "Numerical study of the random transverse-field Ising spin chain". In: *Physical Review B* 53.13 (1996), p. 8486.

[172] V. Oganesyan and D. A. Huse. "Localization of interacting fermions at high temperature". In: *Physical Review B—Condensed Matter and Materials Physics* 75.15 (2007), p. 155111.

[173] H. Li et al. "Many-body localization in Ising models with random long-range interactions". In: *Physical Review A* 94.6 (2016), p. 063625.

[174] A. Gubin and L. F. Santos. "Quantum chaos: An introduction via chains of interacting spins 1/2". In: *American Journal of Physics* 80.3 (2012), pp. 246–251.

[175] S. Mukherjee, S. Nag, and A. Garg. "Many-body localization-delocalization transition in the quantum Sherrington-Kirkpatrick model". In: *Physical Review B* 97.14 (2018), p. 144202.

[176] B. M. Hill. "A simple general approach to inference about the tail of a distribution". In: *The annals of statistics* (1975), pp. 1163–1174.

[177] S. Cantori et al. "Adaptive-basis sample-based neural diagonalization for quantum many-body systems". In: *Methods* 4 (2025), p. 4.