

Neural quantum states in condensed matter: advances, best practices, and prospects

Jonas B. Rigo¹, Björn J. Wurst³, Rajah Nutakki⁴, Markus Schmitt^{1,5}, Dante Kennes^{2,3}

¹Institute of Theoretical Physics, University of Regensburg, 93053 Regensburg, Germany

²Institut für Theorie der Statistischen Physik, RWTH Aachen, 52056 Aachen, Germany and JARA - Fundamentals of Future Information Technology

³Max Planck Institute for the Structure and Dynamics of Matter, Center for Free-Electron Laser Science (CFEL), Luruper Chaussee 149, 22761 Hamburg, Germany

⁴École Polytechnique

⁵Forschungszentrum Jülich GmbH, Institute of Quantum Control, Peter Grünberg Institut (PGI-8), 52425 Jülich, Germany

ARTICLE HISTORY

Compiled August 24, 2026

1. Introduction

The investigation of emergent phenomena in condensed matter occupies a distinctive position in modern physics: the microscopic laws governing electrons and nuclei are known, yet describing their collective consequences remains notoriously difficult [1, 2]. In practice, one therefore works with low-energy effective Hamiltonians or reduced descriptions that retain only the degrees of freedom relevant to the phenomena of interest [3–7]. Even after this reduction, most lattice models remain non-integrable, and exact solutions are rapidly ruled out by the exponential growth of the Hilbert-space dimension with system size [8–10]. Analytical progress is possible in special limits, through low-dimensional effective theories, field-theoretic mappings, and integrable constructions [11, 12], but the generic strongly correlated problem must ultimately be addressed numerically.

From a broad perspective, classical simulations of quantum many-body systems can be arranged as in Fig. 1, according to two dominant obstacles: severity of the sign problem and spatial dimension. The dimension-related difficulty is most prominent in wave-function-based approaches. These approaches represent the quantum state explicitly or in compressed form, and include exact diagonalization and Krylov methods [9, 10], configuration interaction [13], coupled-cluster methods [13, 14], and tensor-network states [15, 16]. Although important exceptions exist [17], a common limitation among wave-function-based methods is entanglement: as entanglement grows with system size, dimension, or time evolution, the cost of representing the state accurately rises rapidly, making generic two- and three-dimensional problems difficult.

Stochastic approaches, most notably quantum Monte Carlo [18, 19], avoid representing the full wave function explicitly and instead estimate observables by statistically accumulating contributions from sampled configurations. Their main limitation is the




Figure 1. The figure portrays the schematic landscape of the main computational approaches to strongly correlated quantum many-body systems. The vertical axis indicates the severity of the sign problem, while the horizontal axis shows spatial dimension, used here as a practical proxy for the growing entanglement complexity that limits wave-function-based methods. Tensor-network methods are most effective in low dimensions, quantum Monte Carlo performs best when the sign problem is weak or absent, and DMFT is accurate in the high-dimensional limit. Neural quantum states (NQS) are shown as targeting the broad intermediate regime between these established approaches. The illustrated NQS is a *convolutional neural network* on a 4×4 spin lattice.

sign problem, which can make the variance grow exponentially and thereby render the simulation impractical in precisely the physically interesting regimes [20]. This same stochastic machinery also underlies many of the impurity solvers used in dynamical mean-field theory (DMFT), whose central approximation maps the lattice problem onto a self-consistent quantum impurity model [21–24]. DMFT becomes exact in the infinite-dimensional limit, but in the physically important regime of intermediate dimensions (2D and 3D), quantum Monte Carlo is hindered by the sign problem while DMFT and its extensions cede the advantage of the underlying infinite-dimensional approximation.

Thus, the intermediate regime remains unsettled in the sense that neither tensor-network methods nor conventional quantum Monte Carlo are uniformly satisfactory there. Many central condensed-matter problems are believed to fall into this regime, including frustrated magnets, doped fermionic lattice models, and non-equilibrium dynamics beyond one dimension [25–27]. Neural quantum states—which we discuss in the following—constitute an emerging numerical technique aimed at filling this gap in the computational toolbox by introducing artificial neural network representations of the quantum wave function [28].

The central idea of NQS is to combine a flexible parametrization of the many-body wave function with stochastic evaluation of observables. Artificial neural networks are used for dimensional reduction to achieve a compressed representation of the wave function. Monte Carlo sampling of computational states makes it possible to estimate observables without summing over the exponentially large Hilbert space. This approach does not exhibit the conventional Monte Carlo sign problem [29]; however, when the ground state has a non-trivial sign or phase structure, the NQS must learn it explicitly. This can be difficult, but it is not a fundamental obstruction [30, 31]. At the same time, it has been shown that NQS wave functions can support entanglement structures that go beyond the standard restrictions of low-bond-dimension tensor networks [32–35]. In fact, the universal approximation capabilities of artificial neural networks [36, 37] guarantee that arbitrary wave functions can be represented as an NQS in the limit of large network sizes. This constitutes the foundation for self-consistent accuracy checks, rendering the approach systematically controlled. While tensor networks share this




Figure 2. Schematic illustration of a feed-forward neural network, the prototypical artificial neural-network architecture used in NQS. The network maps an input configuration through a sequence of hidden layers to an output wave-function amplitude. Each layer consists of an affine linear transformation followed by an activation function f , which is nonlinear and typically strictly monotonic.

property, a key advantage of NQS is that different neural network architectures can be used to confirm convergence under varying approximation bias.

In this perspective, we review recent progress in applying NQS to lattice quantum many-body problems in condensed matter. We discuss where NQS are already competitive with alternative state-of-the-art methods, why NQS are suitable for those settings, which obstacles still limit their broader use, and what future directions appear promising at this point. Recent reviews have introduced the general NQS framework, architectures, and applications, and we refer the reader to them for more technical discussions [38, 39]. For current state-of-the-art NQS work in continuum electronic structure, we refer to Refs. [40, 41]. The scope of the present work is to highlight the most successful NQS architectures and best practices for optimization and sampling strategies. To support these choices, we discuss recent results for frustrated spin systems, interacting fermions, and real-time dynamics that are state-of-the-art not only within NQS. Throughout, comparisons to tensor-network, quantum Monte Carlo, and dynamical mean-field approaches serve as a reference for identifying the class of problems for which NQS are most likely to extend the reach of classical simulations of strongly correlated quantum matter.

2. Neural Quantum States in a nutshell

In NQS, the wave-function compression is achieved by representing the wave function in a *fixed local product basis*, often referred to as the *computational basis*. For instance, for a spin-1/2 system of N sites one commonly uses the joint eigenbasis of the local spin- z operator ($\hat{S}_i^z$), that is the product basis $|\mathbf{s}\rangle$ with $\mathbf{s} \in \{-1, 1\}^N$. In this basis the exact wave-function amplitudes $\psi_{s_1 s_2 \dots s_N} = \langle \mathbf{s} | \psi \rangle$ are replaced by a parameterized *ansatz* $\psi_\theta(\mathbf{s})$ defining the variational state

$$|\psi_\theta\rangle = \sum_{\mathbf{s}} \psi_\theta(\mathbf{s}) |\mathbf{s}\rangle . \quad (1)$$

The ansatz can be understood both as a *function* $\psi_\theta : \mathbf{s} \mapsto \psi_\theta(\mathbf{s})$ that maps configurations to amplitudes, and as a *continuous function* $\theta \mapsto |\psi_\theta\rangle$ that maps *variational parameters* $\theta \in \mathbb{R}^{N_P}$ to a quantum state, as shown in Fig. 2. Introducing the parametrized ansatz replaces the exponential storage of $\dim \mathcal{H}$ amplitudes by N_P parameters, which is the key step toward a compressed approximation of quantum many-body states.

For NQS, this map from configurations to amplitudes is represented by an ANN. The simplest example is the feed-forward network sketched in Fig. 2, where the input configuration is transformed through successive compositions of linear and non-linear transformations into an output amplitude. This elementary construction provides the starting point for more problem-specific choices and design principles for NQS architectures, which we discuss in Sec. 4.2.

After choosing an ansatz and defining the NQS $|\psi_\theta\rangle$, one has access to any expectation value of the form

$$\langle \hat{O} \rangle \equiv \frac{\langle \psi_\theta | \hat{O} | \psi_\theta \rangle}{\langle \psi_\theta | \psi_\theta \rangle} = \sum_{\mathbf{s}} p_\theta(\mathbf{s}) O_{\text{loc}}(\mathbf{s}), \quad (2)$$

where $p_\theta(\mathbf{s}) \equiv \frac{|\psi_\theta(\mathbf{s})|^2}{\sum_{\mathbf{s}} |\psi_\theta(\mathbf{s})|^2}$ is the Born distribution defined by the current variational state, and

$$O_{\text{loc}}(\mathbf{s}) \equiv \sum_{\mathbf{s}': O_{\mathbf{s}\mathbf{s}'} \neq 0} O_{\mathbf{s}\mathbf{s}'} \frac{\psi_\theta(\mathbf{s}')}{\psi_\theta(\mathbf{s})}, \quad O_{\mathbf{s}\mathbf{s}'} = \langle \mathbf{s} | \hat{O} | \mathbf{s}' \rangle \quad (3)$$

is the estimator or *local observable* for any physical observable $\hat{O}$. The only requirement on $\hat{O}$ is that it be sufficiently sparse [29], so that the sum over all $\mathbf{s}'$ with $O_{\mathbf{s}\mathbf{s}'} \neq 0$ is tractable. The remaining sum over states $\mathbf{s}$ in Eq. (2) cannot in general be evaluated exactly, since the Hilbert-space dimension grows exponentially with N . Instead, NQS uses *Monte Carlo* (MC), replacing the sum by MC estimates obtained from samples drawn from p_θ

$$\langle \hat{O} \rangle \approx \frac{1}{N_S} \sum_{i=1}^{N_S} O_{\text{loc}}(\mathbf{s}^{(i)}), \quad \mathbf{s}^{(i)} \sim p_\theta(\mathbf{s}). \quad (4)$$

The Monte Carlo estimator in Eq. (4) is unbiased and converges to the expectation value of the *variational* state $|\psi_\theta\rangle$ as the number of samples N_S increases, with the statistical error decreasing as $\mathcal{O}(N_S^{-1/2})$.

Although NQS use artificial neural networks, their training is usually not data driven. In conventional supervised learning, a neural network is trained from reference input-output pairs and learns to reproduce an underlying mapping from these examples [42].

For quantum many-body states, this is typically not feasible, as the target wave function is itself unknown and cannot be provided as training data. Instead, one determines the network parameters through mathematical and physical constraints that the state must satisfy. The paradigmatic example is the ground state $|\psi_0\rangle$, which is singled out by the *variational principle*

$$E_0 = \min_{\psi} \frac{\langle \psi | \hat{H} | \psi \rangle}{\langle \psi | \psi \rangle}. \quad (5)$$

If the ground state is not degenerate, the minimum is unique up to an overall phase. Restricting the search to the variational family $|\psi_\theta\rangle$ gives the variational energy

$$E_0 \leq E(\theta) = \frac{\langle \psi_\theta | \hat{H} | \psi_\theta \rangle}{\langle \psi_\theta | \psi_\theta \rangle} . \quad (6)$$

Thus, optimizing θ amounts to minimizing an upper bound to the exact ground-state energy. The variational energy therefore plays the role of the *loss function* [42], and its minimization is discussed in Sec. 4.3.1.

The other widely used principle concerns the time evolution of a variational state $|\psi_{\theta_t}\rangle$, where θ_t are the NQS parameters at time t , which have to satisfy the time-dependent Schrödinger equation up to the residual $|\phi\rangle$

$$|\phi\rangle = i \sum_j [\dot{\theta}_t]_j |\partial_j \psi_{\theta_t}\rangle - \hat{H} |\psi_{\theta_t}\rangle , \quad \text{with } \langle \partial_j \psi_{\theta_t} | \phi \rangle - \frac{\langle \partial_j \psi_{\theta_t} | \psi_{\theta_t} \rangle \langle \psi_{\theta_t} | \phi \rangle}{\langle \psi_{\theta_t} | \psi_{\theta_t} \rangle} = 0 . \quad (7)$$

Thus, the Schrödinger equation cannot usually be satisfied exactly, because $\hat{H} |\psi\rangle$ does not necessarily lie completely in the tangent space $\text{span}\{ |\partial_j \psi_{\theta_t}\rangle \}$. With the condition that the residual is completely outside of the tangent space one arrives at the *time-dependent variational principle* (TDVP)

$$\sum_j [\dot{\theta}_t]_j \langle \partial_i \psi_{\theta_t} | \hat{Q} | \partial_j \psi_{\theta_t} \rangle = -i \langle \partial_i \psi_{\theta_t} | \hat{Q} \hat{H} | \psi_{\theta_t} \rangle , \quad \text{with } \hat{Q} = 1 - \frac{|\psi_{\theta_t}\rangle \langle \psi_{\theta_t}|}{\langle \psi_{\theta_t} | \psi_{\theta_t} \rangle} , \quad (8)$$

which allows the propagation of NQS in time. The matrix $S_{ij} \equiv \langle \partial_i \psi_{\theta_t} | \hat{Q} | \partial_j \psi_{\theta_t} \rangle$ on the left-hand side of Eq. (8) is the *quantum geometric tensor* (QGT) and the vector $F_i \equiv -i \langle \partial_i \psi_{\theta_t} | \hat{Q} \hat{H} | \psi_{\theta_t} \rangle$ on the right-hand side is the *force vector*.

The key property of Eqs. (5) and (8) is that, for local Hamiltonians, the energy, its gradients, the force vector and QGT can be written as MC expectation values over the Born distribution $|\psi(\mathbf{s})|^2$. This makes it possible to optimize $E(\theta)$ with gradient-based methods and to evaluate the parameter time derivatives required by TDVP (as discussed in more detail in Sec. 4.3.1 and Sec. 4.3.3). In this respect, NQS optimization resembles the minimization of loss functions in machine learning, although it operates without reference data. In practice, this Variational Monte Carlo (VMC)—or *tVMC* in the context of time evolution—formulation is what makes both ground-state search and time evolution feasible in Hilbert spaces too large for explicit summation.

3. The state of the art

In this section we discuss examples of NQS applied to condensed matter problems, obtaining state-of-the-art numerical results when compared to other methods. This illustrates the types of problems where NQS can overcome the limitations inherent in other numerical approaches.




Figure 3. Numerical results for the ground state of the $J_1 - J_2$ model. (Left) Timeline of ground state phase diagrams using various numerical methods since 2010. The type of method / ansatz is indicated along with the lead author of each study. NQS results are indicated by stars. The works shown are compiled from Refs. [43–59]. A consensus has been reached that the ground state phase diagram consists of Néel, gapless spin liquid, valence bond solid and collinear phases. (Right) Comparison of variational energies for $J_2/J_1 = 0.5$ using different NQS ansätze (adapted from Ref. [57] under CC BY 4.0).

3.1. Ground states of frustrated magnets

Our first example is the study of ground states of frustrated quantum magnets [60], where NQS have recently developed into a powerful tool. Frustrated quantum magnets have been well studied over the past few decades as systems which may host exotic states of matter such as quantum spin liquids [61–64]. The success of NQS is best illustrated by work on the $J_1 - J_2$ model on the square lattice,

$$\hat{H} = J_1 \sum_{\langle ij \rangle} \hat{\mathbf{S}}_i \cdot \hat{\mathbf{S}}_j + J_2 \sum_{\langle\langle ij \rangle\rangle} \hat{\mathbf{S}}_i \cdot \hat{\mathbf{S}}_j, \quad (9)$$

where J_1 couples nearest and J_2 next-nearest neighbors. In this and many other condensed-matter Hamiltonians, the matrix elements of the Hamiltonian are real, and the off-diagonal elements have inhomogeneous signs in the sampling basis $\{|\mathbf{s}\rangle\}$, so the ground state has real components with a non-trivial sign structure. Therefore, the components $\psi(\mathbf{s}) = |\psi(\mathbf{s})|e^{i\phi(\mathbf{s})}$ of the exact ground state will have quantum phases $\phi(\mathbf{s}) \in \{0, \pi\}$, and the ansatz should not only approximate the magnitude $|\psi(\mathbf{s})|$, but also the non-trivial sign structure, $e^{i\phi(\mathbf{s})} \in \{-1, +1\}$.

For antiferromagnetic, non-zero J_1, J_2 , quantum Monte Carlo has a sign problem and the Marshall sign rule [65] does not solve the ground-state sign structure. Furthermore, around $J_2/J_1 = 0.5$, competing ground-state orders make variational approaches challenging.

Despite efforts using various methods, including density matrix renormalization group (DMRG) [44], conventional VMC [47] and projected entangled pair states (PEPS) [51, 52], the ground-state phase diagram in the quantum paramagnetic regime (around $J_2/J_1 = 0.5$), remained controversial (see Fig. 3). One possible scenario [48, 49, 53, 54] was that the region consisted of two phases: a gapless quantum spin liquid and a valence bond solid. However, resolving whether the spin liquid was indeed a phase or a quantum critical point required precise finite-size scaling, with indirect evidence coming from the locations of level crossings [53, 54]. Nomura and Imada [55] used a




Figure 4. a) Energies reached on 8×8 for various architectures at half filling. b-c) Stripe pattern at $\frac{1}{8}$ hole doping for only nearest neighbor hopping (b) and including next to nearest neighbor hopping. The top panel shows cuts of the spin correlation function, while the bottom panel shows the staggered spin correlation function (red) and density (blue). All panels are adapted from Ref. [68] and published under CC BY 4.0.

RBM+PP wave function to simulate up to 18×18 lattices with periodic boundaries, obtaining the most accurate ground-state solutions yet. The *RBM+PP* wave function constructs an ansatz

$$\psi_{\theta}^{\text{RBM+PP}}(\mathbf{s}) = \mathcal{J}_{\theta}^{\text{RBM}}(\mathbf{s})\psi_{\theta}^{\text{PP}}(\mathbf{s}) \quad (10)$$

from the combination of a shallow neural network (a restricted Boltzmann machine [28]), $\mathcal{J}_{\theta}^{\text{RBM}}(\mathbf{s})$, and a more conventional pair-product wave function, $\psi_{\theta}^{\text{PP}}(\mathbf{s})$ [66]. The authors also enforced symmetries on the ansatz in order to obtain accurate results. They showed that the level crossings did indeed correspond to the claimed phase transitions and with careful finite-size scaling of the order parameters and correlation ratios confirmed this ground-state phase diagram. Nevertheless, the *RBM+PP* wave function relies on a parton mean-field decoupling of the Hamiltonian to derive the pair-product part, therefore introducing a source of physical bias. Similar ground-state energies were later obtained using a group-convolution neural network (*GCNN*) [56], retaining the role of symmetries while removing the mean-field-derived component of the wave function. Subsequently, the introduction of *minSR* [57] (see Sec. 4.3.1) allowed the use of much larger neural networks of up to $\mathcal{O}(10^6)$ parameters, resulting in even more accurate ground states with translationally-equivariant neural networks [57–59]. This constituted a breakthrough in terms of tackling ground states with non-trivial sign structure in two dimensions, using wave functions free from mean-field bias and applicable to systems with periodic boundary conditions.

As a result, these neural network architectures have been successfully applied to study the ground states of other models, such as the Shastry-Sutherland model [67].

3.2. Ground States of the Hubbard Model

In the study of itinerant, interacting lattice fermions, the square-lattice Hubbard model

$$\hat{H} = -t \sum_{\langle ij \rangle, \sigma} (\hat{c}_{i\sigma}^{\dagger} \hat{c}_{j\sigma} + \hat{c}_{j\sigma}^{\dagger} \hat{c}_{i\sigma}) + U \sum_i \hat{n}_{i\uparrow} \hat{n}_{i\downarrow}, \quad (11)$$

gives rise to a rich variety of physical phenomena from simple microscopic rules, including qualitative features of cuprate superconductors [69]. It consists of a kinetic term, with hopping $t > 0$, and an onsite interaction term of strength U . A particularly numerically challenging problem is finding the ground state of the hole doped square lattice model with filling fraction around $n = 0.875$ and repulsive, moderately large interaction U as *QMC* solutions are hindered by the sign problem [70, 71] and various near degenerate low-energy solutions make the problem also hard to deal with variationally [26, 72–74]. By now a consensus on length $\lambda = 8$ filled stripes as the ground state has been reached in this regime by various numerical techniques [74–77].

More recently, NQS were applied to this part of the ground-state phase diagram [68, 78–80], achieving the lowest variational energies yet, with clear signatures of the $\lambda = 8$ striped phase, on large lattices of up to 16×16 with periodic boundary conditions.

Previous work on the $t - t'$ Hubbard model with next nearest neighbor hopping t' suggests that the ground state at $t'/t = -0.2$, $U/t = 8$ and $n = 0.875$ is a partially filled superconducting state [81] in coexistence with spin density stripes with length λ fluctuating around $\lambda = 4$ depending on the system size. Gu *et al.* [78] focused on the details of the stripes using NQS, showing that this fluctuation, along with the crossover in stripe direction, were boundary effects which could be eliminated using periodic boundary conditions. The ability of NQS to simulate genuinely two-dimensional systems with periodic boundaries offers a major advantage over other methods, allowing one to identify boundary effects at smaller system sizes.

Using a different NQS architecture based on Pfaffians, a superconducting $\lambda = 4$ partially-filled stripe phase was found [79], qualitatively consistent with the results in Ref. [81]. Despite various works showing comparable variational energies for the groundstate in this regime [68, 78–80], they do not agree on the nature of the ground state. Viteritti *et al.* reported convergence between different NQS architectures after both symmetry restoration and energy-variance reduction [82]. Rende *et al.* introduced a symmetry-preserving ansatz intended to avoid broken-symmetry local minima [80]. They achieved the lowest variational energies up to date for systems up to 24×24 , finding superconducting order.

3.3. Dynamics

While the ground state encodes the static zero-temperature properties of a Hamiltonian, out-of-equilibrium dynamics probes a much larger portion of its spectrum. Access to dynamics therefore reveals quasiparticles and collective modes through response functions [84, 85], tests the stability of ergodicity-breaking phases [86, 87], and can even expose intrinsically dynamical phases of matter such as time crystals [88, 89].

This importance is matched by the numerical difficulty, as real-time evolution is among the hardest problems in many-body physics. The introduction of NQS opened a promising new route to this long-standing problem. Early time-evolution studies based on the TDVP formulation of Eq. (8) already demonstrated the viability of this approach for the transverse-field Ising model (TFIM) in one dimension. The substantially more demanding two-dimensional TFIM,

$$\hat{H} = -J \sum_{\langle ij \rangle} \hat{S}_i^z \hat{S}_j^z - h \sum_i \hat{S}_i^x, \quad (12)$$

is of greater physical interest and has therefore become a particularly useful bench-




Figure 5. The schematic a) shows a space-time picture of a quench across a critical point at $t = 0$. During a ramp parameterized by the dimensionless distance to the critical Hamiltonian parameter $\varepsilon(t) = t/\tau_Q$, critical slowing down prevents the correlation length from following the equilibrium divergence of the correlation length, so ξ saturates at a finite freeze-out scale $\hat{\xi}$. A freeze-out length scale together with the freeze-out time $\hat{t}$ form a *sonic horizon*, so that the system breaks symmetry in quasi-independent domains of size $\hat{\xi}$, producing *defects* and thus *excitations*. Figure b) summarizes the results of the finite size examination of the quantum Kibble-Zurek mechanism in Ref. [83]. It shows the finite-size scaling of the excitation energy density $Q = \frac{1}{L^2} (\langle \hat{H}(t) \rangle - E_0(t))$ at the critical point $t = 0$, shown as $Q L^3$ vs $\hat{\xi}/L$ for different $L (= \sqrt{N})$ and boundary conditions. For $\hat{\xi} \ll L$, the data collapse follows the *quantum Kibble-Zurek* prediction $Q \propto \hat{\xi}^{-(d+z)} = \hat{\xi}^{-3}$ ($d = 2, z = 1$). As $\hat{\xi}$ approaches L , the finite-size gap induces a crossover to an adiabatic regime with $Q \propto \tau_Q^{-2}$. The collapse across periodic (NQS) and open-boundary tensor-network simulation demonstrates universality beyond ramp protocol and boundary details, and the system sizes reached in the NQS simulations are essential to identify the scaling behavior in the non-adiabatic regime. Both figures are reproduced from Ref. [83] and published under CC BY 4.0.

mark. In an early benchmark, Schmitt and Heyl showed that NQS can simulate quench dynamics in this model on lattices as large as 10×10 , reaching timescales comparable to or beyond *then*-state-of-the-art tensor-network methods [27]. The techniques introduced in Ref. [27] were subsequently used, with problem-specific modifications, to compute the spectral function of the TFIM close to criticality [90]. This unprecedented access to late-time dynamics also enabled a clear numerical test of the quantum Kibble-Zurek mechanism through a combination of TN and NQS simulations, again in the TFIM [83]. In this picture, a finite-rate sweep across a quantum critical point becomes nonadiabatic because of critical slowing down, generating defects in the symmetry-broken phase, here ferromagnetic domains separated by domain walls. Ref. [83] thereby confirmed the predicted finite-size scaling of the excitation-energy density with the correlation length, as illustrated in Fig. 5. Recent work extended this approach to the case of the three-dimensional quantum Ising model with up to 1000 qubits [91].

In a similar vein, Medvidović et al. computed the Loschmidt echo in the two-dimensional quantum rotor model on systems of up to 8×8 sites, showing that NQS combined with Hamiltonian Monte Carlo yields stable unitary real-time dynamics and quantitatively accurate echoes [92].

Beyond standard TDVP integration, an important alternative is to propagate the state over a short time step and then variationally project the result back onto the ansatz manifold $\mathcal{M} = \{|\psi_\theta\rangle \mid \theta \in \mathbb{R}^{N_P}\}$. In *projected time-dependent variational Monte Carlo* (p-tVMC), this projection is formulated as an optimization problem at each time step, for example by minimizing the infidelity between $|\psi_{\theta_{t+1}}\rangle$ and the propagated state $e^{-i\delta t H} |\psi_{\theta_t}\rangle$ [93]. This is particularly useful when the wave function develops exact or approximate zeros, where standard TDVP-based propagation can struggle [93]. Closely

related projection-based ideas have also been successfully applied to *ab-initio* electron dynamics by Nys et al. [41].

A second alternative abandons stepwise propagation altogether and instead optimizes the full trajectory at once. In the explicitly time-dependent t-NQS approach, the network represents $\psi_\theta(\mathbf{s}, t)$ directly and its parameters are trained over a whole time interval by minimizing a global residual (Eq. (7)) of the Schrödinger equation [94]. Closely related in spirit, the *time-dependent neural Galerkin* method of Ref. [95] represents the trajectory as a time-dependent linear combination of time-independent NQS basis states and optimizes it through a global-in-time variational principle. Both approaches have been benchmarked successfully on the TFIM, including long-time dynamics.

Open quantum systems constitute another established application for dynamic NQS methods. The unitary evolution is replaced by a Lindblad master equation for the density operator. Early work introduced neural density operators based on latent-space purification, restricted Boltzmann machine density matrices and purified steady-state ansätze, and used them to compute non-equilibrium steady states or dissipative dynamics of spin models [96–100]. The derivation of a TDVP directly for Lindblad evolution with autoregressive neural density operators enabled real-time dissipative simulations of one- and two-dimensional Heisenberg systems as well as confinement dynamics with losses [101]. More recent formulations either represent the open-system state through measurement probabilities in a POVM basis or employ deeper density-operator architectures, including Liouville-space, convolutional, and transformer ansätze, broadening the class of accessible Lindbladian steady states [102–105]. Related applications now extend to lattice-gauge and quantum-optical settings, including the Schwinger model, waveguide-QED emitter arrays, and many-body super- and subradiant light-matter dynamics [106–108].

4. Best Practices

Over recent years, several best practices for applying NQS have been identified, mostly empirically:

- (1) **Symmetries.** Enforce as many symmetries as possible (Sec. 4.1).
- (2) **Spin architectures.** For spin models, use a translationally-equivariant or group-equivariant architecture such as a CNN or GCNN (Sec. 4.2).
- (3) **Fermionic signs.** For fermionic models, use a determinant- or Pfaffian-based ansatz (Sec. 4.2).
- (4) **Ground states.** For ground-state calculations, optimize the wave function using minSR, enabling the use of deep neural networks (Sec. 4.3.1).
- (5) **Dynamics.** For dynamics with TDVP, use as many samples as possible to estimate the QGT and regularize its spectrum carefully for the inversion (Sec. 4.3.3).
- (6) **Parallelization.** Use GPUs to parallelize sampling and observable estimation (Sec. 4.4).

In the following, we discuss these points in more detail.

4.1. Physical Bias

The central idea of NQS is to use asymptotically unbiased and expressive wave functions such that an accurate approximation of the true solution can be found via opti-

mization, in a wide range of different phases of matter. However, using very unstructured wave functions makes finding a nearly-optimal solution difficult [109], so some physical bias should be included to guide the optimization towards physically-relevant solutions. Some of these are exact constraints, such as symmetry, or sign structure in the cases when they are exactly known. Others are more general notions of the functional dependence of the wave function on its input variables, such as locality. In this subsection, we discuss aspects of these physical biases which are important factors in NQS design.

First, since condensed matter Hamiltonians in most cases exhibit local interactions, this leads to the notion of locality in their wave functions, for example in spin systems [110–112] and, from an entanglement perspective, the area law for non-critical systems [112, 113]. The need to work with local structures arises not only in *NQS*, but also naturally in a wide range of machine-learning tasks, including image classification, time-series analysis, and natural language processing. Accordingly, a variety of architectures originally developed for such applications, including *CNNs* [114], *RNNs* [42], and *GCNNs* [115], have also found use as NQS. We discuss these architectures in more detail in Sec. 4.2.

Next, condensed matter Hamiltonians possess symmetries, which means that eigenstates must belong to irreducible representations (irreps) of the symmetry group [116]. Enforcing this structure explicitly can improve the accuracy of variational energies and other observables [49, 82, 117, 118]. A recent example is provided by Viteritti *et al.* [82], who compared Pfaffian- and determinant-based states with nearly identical variational energies but different correlation functions, especially in the pairing channel. Including translational and rotational symmetries substantially reduced these discrepancies, illustrating that reliable physical observables may require symmetry enforcement even when the variational energies already agree.

In particular, for discrete symmetries, this can be done via a quantum number projection of the form

$$\psi_{\text{symm}}(\mathbf{s}) = \frac{d_\alpha}{|G|} \sum_{g \in G} \chi_g^* \psi(g^{-1}\mathbf{s}), \quad (13)$$

where G is the symmetry group to be symmetrized over, d_α the irrep dimension and characters χ_g corresponding to the group element g , projecting the wave function to a desired irreducible representation. Typically, Eq. (13) is used to enforce lattice symmetries (the space group), as well as discrete spin symmetries, such as spin-parity symmetry in the computational basis. However, as the number of translation group elements scales with the number of lattice sites, naively applying this quantum number projection is prohibitively expensive for large lattices. This limitation can be overcome by a judicious choice of neural network architecture, which we discuss in Sec. 4.2.

Beyond projection, symmetries such as particle number or magnetization can be easily enforced via sampling [55–57, 67]. Other approaches such as enforcing $\text{SU}(2)$ symmetry by working in a total momentum basis [119, 120], or translational symmetries via use of representative samples [121, 122] or through restricting sampling to a particular set [123], have been proposed, although they have not been widely adopted.

Learning the sign structure (see Sec. 3.1) can be particularly challenging for NQS [124, 125]. For this reason, it has been common to bias trial states towards known proxy sign structures for spin systems [55, 66, 126]. For lattice fermions, successful wave functions are based on determinants and Pfaffians [68, 78, 127–131], biasing towards




Figure 6. A local, translationally-equivariant architecture. Since the output of each layer depends on local information, locally identical environments are mapped to the same features (labelled cells). Since each layer is translationally-equivariant, translating the output is equivalent to translating the input (compare top and bottom rows). Therefore the wave function can be projected to a desired translational irreducible representation in a single evaluation by appropriate choice of output function.

mean field sign structures. From this point of view, approaches that minimize built-in bias while retaining high accuracy are preferable.

4.2. Architectures

A good choice of NQS architecture for a specific problem incorporates the physical biases discussed in the previous section in an efficient way. Fortunately, some of these requirements are encountered in well-established machine learning applications, so neural network components developed in these contexts can be directly applied in NQS.

Both locality and translational symmetry arise in computer vision tasks, resulting in the development of deep convolutional neural networks (CNNs) [114, 132, 133]. These are built out of convolutional layers which, for a two-dimensional system, perform the operation

$$z_{ijm}^{(n)} = b_m^{(n)} + \sum_{\delta_x \in \mathcal{K}_x} \sum_{\delta_y \in \mathcal{K}_y} \sum_{c=1}^{C_{n-1}} K_{\delta_x \delta_y c m}^{(n)} h_{i+\delta_x, j+\delta_y, c}^{(n-1)}, \quad (14)$$

$$h_{ijm}^{(n)} = f(z_{ijm}^{(n)}).$$

Here $h^{(0)} \equiv \mathbf{s}$ is the input configuration, $m = 1, \dots, C_n$ labels output channels, and $\mathcal{K}_x, \mathcal{K}_y$ are the spatial kernel offsets of extents k_x and k_y . The convolutional kernel $K^{(n)} \in \mathbb{R}^{k_x \times k_y \times C_{n-1} \times C_n}$ is the local, weight-shared analogue of the dense matrix $W^{(n)}$ of the linear transformation in Fig. 2, with bias $b_m^{(n)}$; spatial indices are interpreted using the boundary convention of the lattice, for example modulo N_x, N_y for periodic boundaries. Finally, the nonlinearity f acts element-wise on its input.

These layers incorporate locality as the output at a given position (i, j) depends only on the input variables in a given region around it (see Fig. 6). In order to exploit this, the input to convolutional layers in an NQS [27, 109, 134, 135] should be shaped with the locality of the problem in mind, for example ensuring that the positions of variables in x are related to their positions on the lattice, and using *patching* to ensure

up and down spin electrons [68, 136] or variables within a unit cell [67] are associated with the same spatial position.

Convolutional layers are also translationally equivariant [115, 137],

$$y(T_g x) = T_{g'} y(x), \quad (15)$$

where on the left-hand side a translation, T_g , is applied to the input and on the right-hand side, $T_{g'}$, to the output of the convolutional layer (see Fig. 6). Using Eq. (13), an NQS constructed from translationally-equivariant layers can be projected to a momentum sector (translation group irrep) during a single network evaluation. The wave function can be further symmetrized over the point group by quantum number projection, projecting the wavefunction efficiently to a specific irrep of the lattice space group. CNNs can be generalized to be symmetric under the full space-group in GCNNs [115, 138], projecting to a space group irrep in a single evaluation.

Combined with the development of minSR (see Sec. 4.3), allowing optimization of neural networks with large numbers of parameters, state-of-the-art deep translationally-equivariant NQS for ground states of spin problems have been developed [57, 59, 67, 139, 140], inspired by highly optimized architectures used for computer vision tasks [141–143]. These are either explicitly convolutional networks [57] or vision transformers with factored or spatial attention [59, 139], which have broadly similar structures [140]. The structure of these networks consists of an embedding layer, used to set the translational unit cell of the NQS, a deep encoder which extracts a high-dimensional hidden representation [67], and a shallow, complex output block.

For fermions, the advantage of using lattice symmetry-equivariant networks is not as clear-cut as for spins. As discussed in Section 4.1, high-accuracy fermionic NQS use the output of neural networks as the elements of a determinant or Pfaffian state. While it is possible to use the translationally-equivariant architectures described above with determinants [144] and Pfaffians [68, 80], the equivariant encoder appears to be performing worse than symmetrization via quantum number projection (Eq. 13) [136]. Nonetheless, this approach becomes unfeasible for large systems, due to the scaling of the number of the number of elements in the space group symmetry group, tipping the scale again towards equivariant architectures [80]. Instead of enforcing the full lattice symmetries via quantum number projection, Chen *et al.* only restored sublattice symmetries instead in order to overcome the computational cost of quantum number projection achieving state of the art energies [68]. On the other hand, in Ref. [78] the authors were able to obtain state-of-the-art energies using a neural backflow transformer without enforcing any lattice symmetries, highlighting the fact that there is more work to be done in understanding the most effective way to implement lattice symmetries for these systems.

For dynamics, state-of-the-art results have been achieved using CNNs [27, 92]; however, as discussed in Section 4.3, the time-evolution algorithm restricts the size of neural networks which can be used. Therefore, lightweight versions of the translationally-equivariant NQS used for ground states can be applied to dynamics [145, 146].

The notion of locality can alternatively be included using autoregressive architectures [42], which model the output as

$$\psi_\theta(\mathbf{s}) = \sqrt{|\psi_\theta(s_1)|^2 |\psi_\theta(s_2|s_1)|^2 \dots |\psi_\theta(s_N|s_1, \dots, s_{N-1})|^2} e^{i\phi(\mathbf{s})}, \quad (16)$$

where the conditional distributions $|\psi_\theta(s_i|s_1, \dots, s_{i-1})|^2$ are computed in a spatially

sequential fashion starting from $\psi_\theta(s_1)$. Examples of autoregressive NQS include recurrent neural networks (RNNs) [147] and transformers with masked self-attention [148, 149]. They have the advantage that uncorrelated samples can be drawn efficiently and directly from the wave function, rather than being generated through an MCMC process, as discussed in Sec. 4.3. For spin models, the use of autoregressive architectures has largely been restricted to models where the ground state is real and positive (e.g., ground states of stoquastic Hamiltonians) [147–151]. This difficulty in learning non-trivial sign structures could stem from the lack of expressivity of autoregressive wave functions [152].

Other types of networks have been developed to incorporate problem-specific biases, such as gauge symmetries for lattice gauge theories [120, 153–155] and quantum spin liquids [156].

4.3. Optimization and Sampling

We now turn to the algorithmic question of how to *optimize* these states for ground-state searches and how to *propagate* them in time. While the underlying variational principles are conceptually the same as in traditional VMC, the practical setting is remarkably different. Modern NQS ansätze contain far more parameters and exhibit greater expressive power than the strong physically biased wave functions historically used in VMC, which often involve only $\mathcal{O}(N)$ variational parameters. Their greater expressive power broadens the class of states that can be represented, but it also makes optimization more demanding, increases memory requirements, and often requires specialized hardware for efficient training and evaluation. These challenges have motivated the development of algorithms tailored to large neural-network ansätze, often drawing directly on the software ecosystem and optimization strategies of modern machine learning. In this section, we discuss the resulting best practices for ground-state optimization and time evolution, as well as the sampling procedures used to estimate the quantities required for both.

4.3.1. Ground state search

Starting from the variational principle in Eq. (5), ground-state search reduces to the numerical minimization of the variational energy with respect to the NQS parameters θ . In practice, this is carried out with first-order gradient-based optimization. A simple and computationally cheap choice is the *Adam* optimizer, an adaptive stochastic optimization method widely used in machine learning [157]. Adam and related optimizers are robust, easy to implement, and perform well for moderately sized networks. In the NQS context they have been used successfully especially for autoregressive architectures such as RNNs [158–160]. However, autoregressive architectures are by no means limited to Adam-style first-order optimization, and can likewise be optimized with stochastic reconfiguration [161, 162].

In fact, most state-of-the-art NQS ground-state results are obtained with *stochastic reconfiguration* (SR) methods [19, 163], which can be viewed as a stochastic realization of *natural gradient descent* [164]. Physically, SR corresponds to imaginary-time evolution projected onto the variational manifold, in close analogy to the TDVP equation in Eq. (8). This connection is reflected by the close similarity between the TDVP equation




Figure 7. Here we compare gradient descent vs natural-gradient descent on the *Rosenbrock landscape*. (a) Standard gradient descent $-\nabla R(x_1, x_2)$ produces a flow field with opposing directions that meet in a long valley. The opposing gradient directions produce a zig-zag trajectory as it repeatedly overshoots across the narrow curved valley of the Rosenbrock function (white contours), slowing convergence toward the minimum (magenta star). (b) Natural-gradient descent $-S^{-1}\nabla R(x_1, x_2)$, with QGT S as preconditioner, reshapes the flow to better align with the valley geometry, yielding a smoother path that follows the valley more directly. Background colors indicate the gradient magnitude (left color bar), while the grayscale bar encodes the objective value $R(x_1, x_2)$.

for the parameter time derivative $\dot{\theta}$ and the SR equation for the parameter update $\Delta\theta$,

$$\sum_j S_{ij} \Delta\theta_j = -\partial_i E(\theta). \quad (17)$$

The parameters are then updated according to

$$\theta \leftarrow \theta + \eta \Delta\theta, \quad (18)$$

with learning rate η . The key idea is that the update is not chosen with respect to the Euclidean geometry of parameter space, corresponding to the identity metric $S_{ij} = \delta_{ij}$ used in ordinary gradient descent, but with respect to the geometry induced by the quantum state itself. This geometry is encoded in the *quantum geometric tensor* (QGT),

$$S_{ij} = \text{Re} \langle \partial_i \psi_\theta | \hat{Q} | \partial_j \psi_\theta \rangle, \quad \hat{Q} = 1 - \frac{|\psi_\theta\rangle \langle \psi_\theta|}{\langle \psi_\theta | \psi_\theta \rangle}, \quad (19)$$

or, equivalently,

$$S_{ij} = \text{Re} \mathbb{E}_{\mathbf{s} \sim p_\theta} \left[(\Gamma_i(\mathbf{s}) - \mathbb{E}[\Gamma_i(\mathbf{s})])^\dagger (\Gamma_j(\mathbf{s}) - \mathbb{E}[\Gamma_j(\mathbf{s})]) \right], \quad (20)$$

where $\Gamma_i(\mathbf{s}) = \partial_{\theta_i} \log \psi_\theta(\mathbf{s})$ and $p_\theta(\mathbf{s}) \propto |\psi_\theta(\mathbf{s})|^2$. In this way, SR computes the projected imaginary-time update for which the residual of Eq. (19) is orthogonal to the tangent space of the variational manifold $\text{span}(\{|\partial_i \psi_\theta\rangle\})$, in direct analogy to the TDVP construction in Eq. (8). As illustrated in Fig. 7, this does not change the optimization landscape itself, but rescales updates so that broad weak-gradient directions do not slow down the optimization, thereby guiding the trajectory more directly toward the minimum.

The drawback of SR is its higher computational cost. In its standard form, one has to estimate and solve a linear system involving the $N_P \times N_P$ QGT, so the optimization bottleneck scales with the number of variational parameters. Since the Monte Carlo estimate of the QGT has rank at most N_S , this parameter-space formulation becomes prohibitive for modern deep architectures with $N_P \gg N_S$. Recent low-rank reformulations of SR, in particular *minSR* [57, 58], exploit the structure of the stochastic QGT expression in Eq. (20) and replace the $N_P \times N_P$ linear system by an $N_S \times N_S$ problem (or $2N_S \times 2N_S$ for complex wave functions), thereby shifting the bottleneck from parameter count to sample count. Empirically, this reformulation is often already favorable for moderate sample numbers, and in many recent applications stable training was achieved with $N_S \sim 10^3$, enabling deep architectures with up to 10^6 parameters [57, 58, 79, 165].

In any formulation, the SR equation is typically ill-conditioned because the QGT contains very small eigenvalues arising from statistical noise or redundant parameter directions. In practice, this is mitigated by regularization, most simply through a diagonal shift $S \rightarrow S + \lambda \mathbb{I}$ with $\lambda > 0$. Stability and convergence can be further improved by more refined regularization schemes. For example, SPRING (*sub-sampled projected-increment natural gradient descent*) [166] smooths the effective eigenvalue cutoff of the QGT and combines this with momentum, reducing sensitivity to Monte Carlo noise while accelerating convergence. Gu et al. [78] later extended this idea with a dynamically adapted variant that performed particularly well for the Hubbard model, further improving the robustness and scalability of SR-based optimization.

4.3.2. Convergence Metrics

For the ground-state search with NQS, the most common figures of merit are the variational energy and the *energy variance* $\mathbb{V}[H_{\text{loc}}] = \langle (H_{\text{loc}} - \langle H_{\text{loc}} \rangle)^2 \rangle$. By the variational principle, the variational energy provides an upper bound on the exact ground-state energy, while the energy variance vanishes for any exact eigenstate and therefore supplies an additional stringent consistency check. A small variance alone does not certify that the ground state has been found, but Ref. [167] empirically observed that the *dimensionless variance* (V-score) is approximately proportional to the true energy error, namely the relative deviation between the exact and variational ground-state energies. When MCMC is the sampler of choice, final accuracy checks should also be repeated with fresh, independently initialized chains to verify that they thermalize and reproduce the reported observables, otherwise apparently accurate energies can be artifacts of frozen or non-ergodic sampling rather than of an improved variational state [168].

4.3.3. Dynamics

In time-dependent VMC (tVMC), all derivatives are estimated using VMC as described above. Unlike ground-state search, real-time dynamics does not steadily approach a fixed point. Instead, integration errors are propagated forward and can accumulate over time, so only very little noise can be tolerated in the MC estimate of the QGT and in the solution of the TDVP linear system Eq. (8). In practice, increasing the number of samples is generally not sufficient to obtain a well-conditioned full-rank QGT, making careful regularization essential [27, 169]. Since the hard cutoff of a plain Moore-Penrose pseudo-inverse is incompatible with smooth solutions of the TDVP, established regularization schemes rely on generalized pseudo-inverses with smooth truncations of the eigenvalue spectrum [27, 92, 170]. As optimal values for the cutoff may vary with

time, recent work introduced a scheme to adapt the regularization parameter dynamically while fixing the residual [171]. In addition, Ref. [27] introduced an additional regularization to eliminate particularly noisy contributions from the TDVP equation. Due to the non-linearity of the TDVP equation (8), it is generally imperative to employ integrators with adaptive time step. In many cases, unnecessarily small time steps can be avoided by employing the quantum geometric tensor to quantify the integration error in state space instead of parameter space [27].

By contrast, best practices for more recent alternatives such as p-tVMC [93], t-NQS [94], and neural Galerkin [95] approaches are still much less established, so we restrict our discussion to TDVP-based dynamics. Among these newer methods, some practical guidance is beginning to emerge for p-tVMC, for instance in the recent work of Gravina [145].

4.3.4. Sampling

The NQS formalism avoids explicit summation over the full many-body Hilbert space by replacing it with averages over a set of configurations $\{\mathbf{s}_i\}_{i=1}^{N_S} \sim p_\theta(\mathbf{s})$, drawn from the Born distribution of the variational wave function. Observables and derived quantities are then obtained as MC averages of their respective estimators Eq. (4) over these samples. Importantly, any quantity whose estimator can be written in the form of the local observable estimator in Eq. (3) can be estimated efficiently, with the required number of samples increasing only polynomially with the desired precision [29].

By contrast, quantities entering the ground-state search or TDVP, such as the energy gradient and the QGT, do not admit a formulation as local observable estimators. These quantities can nevertheless be estimated by rewriting them as importance-sampled averages with respect to the Born distribution. For a quantity $X = \sum_{\mathbf{s}} x(\mathbf{s})$, one may write

$$X = \sum_{\mathbf{s}: p_\theta(\mathbf{s}) \neq 0} p_\theta(\mathbf{s}) \frac{x(\mathbf{s})}{p_\theta(\mathbf{s})}, \quad (21)$$

but the resulting estimator is only unbiased if $x(\mathbf{s}) = 0$ whenever $p_\theta(\mathbf{s}) = 0$. If $\psi_\theta(\mathbf{s})$ has roots on configurations where $x(\mathbf{s}) \neq 0$, the Born distribution does not sample all relevant contributions and the estimator is therefore biased [38, 93]. This issue is often negligible in ground-state calculations when one works in sectors where the variational state is non-zero everywhere, but it can become important in real-time evolution where nodes may proliferate during the dynamics [93]. Recent work proposes to sidestep biased estimators by employing importance sampling techniques [172].

Even when no formal bias is present, Monte Carlo estimators can have a very large variance. This applies both to local observables and to general importance sampled quantities. While the resulting estimates remain efficient in the complexity-theoretic sense, a large variance increases the practical sampling cost, because the statistical error decreases only as $N_S^{-1/2}$. A particularly unfavorable situation arises when the Born distribution is strongly peaked in the chosen basis. Then the MC samples explore only a small part of configuration space, while the estimator may receive appreciable contributions from a much broader region. In that case, many samples are required before averages converge reliably. For derivative-based quantities such as the QGT and the energy gradient this problem is often more severe, because their significant contributions in configuration space are typically broader than those sampled efficiently by the Born distribution itself [162].

Equally important as the estimation itself is the generation of the underlying samples. For that purpose, Markov-chain Monte Carlo (MCMC) is most commonly employed. In the standard Metropolis–Hastings scheme, illustrated schematically in Fig. 8, one constructs a Markov chain by proposing updates of the current configuration $\mathbf{s} \rightarrow \mathbf{s}'$ and accepting or rejecting them probabilistically. In lattice models with discrete degrees of freedom, the most common proposals are local moves such as single-spin flips, local particle hops, or pair updates chosen to preserve conserved quantum numbers [28, 38]. These updates are simple and architecture agnostic, but their efficiency depends strongly on how rapidly they explore the relevant part of configuration space. MCMC produces correlated samples; in practice, one must also monitor thermalization and autocorrelation times and discard an initial warm-up period before accumulating measurements.

An important alternative is available for autoregressive NQS [147]. There, the joint probability of a configuration is factorized into conditionals, as indicated by Eq. (16)

$$p_\theta(\mathbf{s}) = |\psi_\theta(s_1)|^2 |\psi_\theta(s_2|s_1)|^2 \dots |\psi_\theta(s_N|s_1, \dots, s_{N-1})|^2 = \prod_{i=1}^N p_\theta(s_i|s_{<i}), \quad (22)$$

which allows one to generate configurations sequentially by ancestral sampling

$$s_i \sim p_\theta(s_i|s_{<i}) \quad \text{with } p_\theta(s_1) = p_\theta(s_1|\text{const.}) . \quad (23)$$

The resulting samples are independent by construction and therefore avoid both thermalization and autocorrelation issues. In addition, autoregressive models permit sampling strategies such as *Gumbel top-K sampling* [173], which generate samples without repetition and can improve sample efficiency when the Born distribution is strongly concentrated, thereby alleviating the variance problem [162].

4.4. Complexity and GPU parallelization

The algorithm underlying ground-state optimization with NQS, and likewise TDVP for real-time evolution, consists of three core stages: *sampling*, *estimation of observables and gradients*, and the *optimizer step*. The total wall time of an algorithm iteration is therefore composed of sampling-and-estimation part and an optimization part,

$$T_{\text{ALG}} \sim T_{\text{sample+obs}} + T_{\text{opt}} , \quad (24)$$

where $T_{\text{sample+obs}} = n_{\text{eval}}(N, N_S) T_{\text{NQS}}(N)$ and $T_{\text{NQS}}(N)$ denotes the cost of one *batched* evaluation of the ansatz, that is the simultaneous evaluation of $\psi_\theta(\mathbf{s})$ for a whole set of configurations, and n_{eval} the number of such evaluations required per iteration.

The abstract decomposition above becomes more transparent by considering a typical example like a dense feed-forward ansatz (Fig. 2) with width proportional to system size and depth D , for which one network evaluation scales as

$$T_{\text{NQS}}(N) \sim \mathcal{O}(DN^2). \quad (25)$$

If sampling is performed with local Metropolis updates, then obtaining one effectively new sample requires roughly one sweep of $\mathcal{O}(N)$ proposed moves, so the sampling




Figure 8. The figures show a schematic of the multi-GPU-accelerated workflow for NQS-based VMC or TDVP. The computationally dominant NQS evaluation is parallelized on GPUs during (left) Markov-chain Monte Carlo sampling and (right) local observable estimation. Multiple independent Markov chains are advanced concurrently by batching NQS evaluations on each device to better saturate the GPU. The Metropolis-Hastings accept/reject step is depicted as a probability-weighted binary decision tree, emphasizing the stochastic decision that turns a proposed configuration into the next configuration of the chain. The resulting sample pool $\{\mathbf{s}_i\}_{i=1}^{N_S/N_{\text{rank}}}$ is distributed across N_{rank} independent processes, which then independently evaluate local observables of the form $O_{\mathbf{s}\mathbf{s}'} \frac{\psi_\theta(\mathbf{s}')}{\psi_\theta(\mathbf{s})}$. Finally, per-process contributions are reduced to obtain the global Monte Carlo averages.

cost per sample is $\mathcal{O}(N T_{\text{NQS}})$. For a local Hamiltonian on a d -dimensional lattice, the local-energy estimator couples a configuration to $\mathcal{O}(dN)$ connected configurations, so observable estimation likewise requires $\mathcal{O}(dN T_{\text{NQS}})$ work per sample. The energy gradient can be obtained at essentially the same cost as the energy itself, since the required logarithmic derivatives are evaluated through automatic differentiation [42]. Under local Metropolis sampling and without architecture-specific fast updates, this gives the illustrative estimate

$$T_{\text{sample+obs}} \sim \mathcal{O}(N_S D N^3), \quad (26)$$

up to model-dependent prefactors and possible implementation-dependent speedups from caching, symmetry reduction, or architecture-specific fast updates [174]. At the same time, these stages are embarrassingly parallel, since each chain and each configuration can be processed independently, as sketched in Fig. 8. Distributing the work over N_{rank} parallel processes therefore reduces the wall time approximately to

$$T_{\text{sample+obs}}^{\text{wall}} \sim \mathcal{O}\left(\frac{N_S D N^3}{N_{\text{rank}}}\right), \quad (27)$$

up to minor overhead. For MCMC, each chain incurs a warmup cost before measurements are accumulated. In optimization runs, however, chains are usually initialized from samples of the previous iteration, so once the parameters evolve smoothly only a short rethermalization is often needed, which is why parallelization becomes more effective as the sample count N_S increases. Autoregressive samplers have the further advantage that they generate i.i.d. samples directly and therefore benefit from parallelization immediately.

The optimizer step can become a separate bottleneck when performing TDVP or

ground-state search using (min)SR, since one must solve a linear system involving the quantum geometric tensor (QGT). In a parameter-space formulation with N_P variational parameters, a dense QGT solve scales as $\mathcal{O}(N_P^3)$, which can become prohibitive for modern neural-network ansätze. Matrix-free iterative solvers avoid storing and factorizing the full QGT and instead reduce the problem to repeated QGT-vector products, substantially lowering the per-iteration cost. Their total runtime, however, still depends on the number of Krylov iterations and therefore on the conditioning of the problem and the regularization scheme [28]. The minSR formulation exchanges the parameter dimension N_P for the sample dimension N_S , which can be advantageous when $N_S \ll N_P$, but the same basic tradeoff between dense and matrix-free linear solvers remains. In practice, SR is therefore most effective when the linear solve is kept subdominant compared to sampling and observable estimation.

NQS are particularly well suited to hardware accelerators such as graphics processing units (GPUs), because the dominant workload consists of large batches of dense linear algebra and tensor contractions in the network evaluation and derivative computation. A single Markov chain or a single local-observable evaluation rarely saturates a modern accelerator, so many such tasks can be batched on one or multiple devices with little additional overhead, as illustrated in Fig. 8. This is precisely the execution model exploited by modern NQS packages such as NETKET [175], JVMC [170], and QUANTAX [176], which build on machine-learning frameworks like JAX [177] to provide distributed execution and accelerator support with comparatively little problem-specific engineering.

5. Future Directions

As we have discussed in the previous sections, NQS is a versatile technique, able to address open questions at the forefront of condensed matter research. The major advantages of NQS are:

- (1) *Versatile applicability across lattice geometries, entanglement structures, and irrespective of the QMC sign problem.*
- (2) *Universal approximation capabilities enabling controlled accuracy and non-perturbative treatment of strong correlations.*

These two attributes define the realm where the most impactful future applications of NQS may be expected: Strongly correlated systems exhibiting intricate entanglement structures and a QMC sign problem. Thereby, NQS could break new ground on a variety of current frontiers of condensed matter physics summarized in Fig. 9. However, the most demanding applications pose recurring challenges that will require further methodological advances, namely (A) learning sign and phase structures of the wave function, (B) enforcing symmetries and representation constraints, and (C) accessing dynamical response and non-equilibrium evolution. In the following, we discuss these challenges and we outline possible future applications of NQS that would be enabled by overcoming them.

5.1. (A) Learning sign and phase structures

Sign and phase learning is a central bottleneck whenever the ground state has a non-trivial sign or phase structure, which is the case when the Hamiltonian, expressed in the sampling basis, has both negative and positive off-diagonal entries, a property known as




Figure 9. Signpost of future directions for NQS in condensed matter and related lattice problems. The central panel summarizes three current roadblocks that limit the application of NQS to key physical systems: (A) learning sign and phase structures, (B) enforcing symmetries and entanglement constraints, and (C) accessing real-time propagation and dynamical response. The surrounding panels illustrate problem classes that fall naturally within the scope suggested by versatile applicability and universal approximation capabilities, but whose treatment is still held back by one or more of these bottlenecks.

non-stoquasticity [178]. This offers a useful distinction when discussing sign problems of QMC versus NQS, since the non-stoquastic attractive Hubbard model can present a significant challenge for general NQS, but has a sign-problem-free auxiliary-field QMC formulation [19, 179].

Non-stoquasticity is a natural feature of fermionic Hamiltonians because of fermionic anticommutation relations, and it is also a common property of frustrated magnets. In some important cases, however, the non-stoquasticity can be cured. For example, in the square-lattice Heisenberg antiferromagnet, the Marshall sign rule fixes the ground-state sign structure [65]. For free fermions, the antisymmetric signs are not a fundamental obstacle because the problem reduces to single-particle diagonalization. The difficulty is that there is no generally efficient way to find such a cure: stoquastic local Hamiltonians form a special complexity class [178], and finding even restricted local basis rotations that cure a non-stoquastic Hamiltonian is NP-complete [179]. This leaves the vast majority of Hamiltonians in the seemingly untouchable class of non-stoquastic problems.

The situation is in practice more nuanced than that, as shown by a recent study of the isolated sign-learning problem, which demonstrates that frustrated sign structures can contain exploitable Boolean-Fourier structure and therefore need not behave like featureless worst-case learning problems [31]. This nuance is further underlined by high-quality NQS ground states that have been obtained for non-stoquastic spin Hamiltonians close to simple proxy signs such as the Marshall sign rule [56, 57, 67].

The epitome of non-stoquastic systems currently treated with NQS is the class of fermionic Hamiltonians. Here the standard strategy is to build antisymmetry into determinant, Pfaffian, or backflow-inspired layers [68, 127, 129, 131]. These approaches represent the current state of the art for Hubbard-model benchmarks [68, 78, 79, 82], but they introduce a mean-field bias and a raw determinant or Pfaffian scaling cost of $\mathcal{O}(N_e^3)$, which low-rank-update techniques mitigate but do not remove [174].

This leaves two connected goals: (i) an unbiased treatment of sign and phase structures, and (ii) better scaling for controlled antisymmetric layers. Progress on (ii) would already solve pressing problems that require large simulation sizes, like three-band cuprate models [77, 180], multilayer nickelates [181–183], and Moiré Hubbard models [184, 185]. Progress on (i) would broaden access to elusive phases like quantum spin liquids [25, 60, 63], topological order [186], and anyonic lattice models [119, 120]. Solving both goals is especially important near quantum criticality, where small energy differences between competing phases make careful finite-size scaling decisive [187, 188].

5.2. (B) *Enforcing symmetries and representation constraints*

Historically, variational Monte Carlo wave functions were often built from physically motivated ansätze designed for a specific phase, such as Jastrow-correlated, Gutzwiller-projected, or mean-field states. This makes their symmetry content and representation limits relatively transparent, but it also builds in strong assumptions about the phase being described [19]. That bias becomes limiting when the goal is to identify an unknown phase or resolve a delicate critical point. A concrete example is the square-lattice Néel-VBS deconfined quantum critical point, where spin-rotation symmetry, lattice symmetries, and emergent critical symmetries all matter, while numerical scaling has required nonstandard interpretations such as two divergent length scales [189, 190]. NQS offer a possible route through this problem, since their large parameter spaces make systematic improvement with architecture and parameter count possible, but increasing the number of parameters does not fully address the issue. For that, one must also know whether the architecture can represent the required symmetry sector and how its parameter count controls representational power. This leaves two closely related challenges, namely symmetrizing NQS so that the Hilbert-space sector of interest is represented directly and finding laws that connect architecture and parameter count to representational power.

Recently, progress has been made on the latter in two significant ways. First, the entanglement that NQS can represent is now better understood through bounds such as $S_A = \mathcal{O}(k \log N)$ in the number of scalar nonlinearities k and system size N [35]. Second, transformer-NQS calculations on the square and triangular J_1 - J_2 Heisenberg models found, within the accessible precision window, that the V -score decays approximately as a power law in training compute, with an exponent that decreases with frustration [191]. Taken together, these results suggest that increasing model size and training compute can systematically improve NQS accuracy, while also indicating that frustration can make a given target accuracy more expensive. The latter observation may be a symptom of an increasing number of near-degenerate low-energy states or symmetry-sector mixing [30, 192].

Solving the problem of imposing symmetry constraints on NQS would make calculations more reliable when low-energy orders in different symmetry sectors compete [118, 121]. Currently, two complementary approaches exist for enforcing symmetries: project-

ing the ansatz into the right subspace with Eq. (13), or building the symmetry into an equivariant architecture such as a GCNN [115, 138]. The architecture-based route can avoid explicit sums over all group elements and has been successful for lattice symmetries and, with more specialized equivariant constructions, also for non-Abelian gauge symmetries [138, 144, 193]. Its accuracy and applicability, however, still depend on the expressive power of the chosen ansatz and on which representations or symmetry constraints are built in. Projection-based approaches remain broadly applicable, including to non-Abelian or anyonic symmetry sectors, but they add a symmetry-averaging cost and can complicate optimization [118, 119, 121, 194]. Thus, whether symmetries are enforced projectively or by the underlying ansatz, the interaction between symmetry enforcement, expressivity, and optimization does not automatically improve results, but requires special care, as illustrated by the escalating symmetry projections used in [165].

Better handling of symmetry and representation constraints would make finite-size scaling more reliable near quantum criticality, where competing states can be very close in energy [187, 188]. It would also sharpen calculations of quantum spin-liquid candidates and topologically ordered or anyonic phases, where the relevant sector is selected by symmetry, gauge, or fusion constraints as much as by energy [25, 63, 186, 195]. For lattice gauge theories, the same progress would let NQS work directly in the physical gauge-invariant sector for condensed-matter gauge descriptions and for Hamiltonian gauge dynamics in high-energy physics, including finite-density regimes where conventional Monte Carlo faces a sign problem [25, 196, 197].

5.3. (C) Accessing dynamical response and non-equilibrium evolution

Investigating quantum many-body systems in non-equilibrium situations is a multifaceted research area addressing questions that range from the foundations of statistical physics to technological applications. Theoretical studies are motivated by the development of increasingly more sophisticated experimental platforms for quantum simulation and quantum computing [198–200] as well as ultrafast techniques to probe solid state systems [201, 202]. Such experiments have, for example, revealed the possibility of light-induced superconductivity [203] or discrete time crystals [204], and in many cases, the experimental observations challenge our theoretical understanding.

Time evolution of quantum states combines different challenges discussed in the previous sections. In generic cases, the wave function rapidly develops non-trivial phase structures and the target state is typically surrounded by a continuum of mid-spectrum states. Moreover, while the search for stationary low-energy states is often contractive and thereby forgiving deviations in individual optimization steps, forward-propagation in time requires high accuracy in every time step to avoid the accumulation of errors. Simulations of linear response to weak perturbations as probed in various forms of spectroscopy, however, allow the exploitation of additional structure present within the low energy sector. Promising NQS directions include real-time variational propagation, variational access to excited states, and direct reconstruction of dynamical correlators [90, 161, 205]. Beyond ground-state response functions, scalable multi-state optimization, for example through orthogonalization, penalty terms, or enlarged-space formulations [40, 206], can turn NQS into a tool for computing experimentally relevant response functions in correlated electron materials, even for finite temperatures.

While various works have exploited the capabilities of NQS for novel insights into the non-equilibrium dynamics of quantum many-body systems beyond linear response

[83, 91, 207–209], severe bottlenecks have been reported in other cases [210, 211]—the precise origin of which, however, remains unclear. For further progress, it is therefore pivotal to develop a better methodological understanding and to address the identified limitations. Sinibaldi et al. recently pointed out that conventional tVMC is formally ill-defined [93], triggering various efforts to circumvent the problem by devising alternative ways of solving the time-dependent variational principle [41, 145, 172]. Developing a reliable toolbox with well-understood scope will be crucial for broader future applications. Besides effective optimization, there is a substantial gap between the NQS network sizes used for time evolution (between 10^3 and 10^5 parameters) and the state of the art for ground-state problems (around 10^6 parameters). As optimization cost is the limiting factor for parameter numbers in time evolution, scaling up network sizes is among the outstanding challenges—recall that enabling larger parameter counts by more efficient optimization was a key breakthrough for ground-state search [57]. The common approach to assess accuracy of time-dependent simulations is testing agreement of observable quantities under varying network sizes and architectures and analyzing estimated deviations from the exact solutions (if available) in each time step. These are, however, often difficult to interpret. It would be desirable to develop an expressive figure of merit comparable across NQS architectures and model systems similar to the V-score for ground-state search [167].

Recent works have started exploring alternatives to the plain forward propagation of NQS in time. The t-NQS [94] and neural Galerkin [95] methods incorporate time as an explicit parameter of the wave function ansatz to solve Schrödinger’s equation (or a Lindblad equation [107]) in the spirit of physics-informed neural networks [212]. Besides possible routes to circumvent bottlenecks of conventional forward propagation, these methods may serve as a starting point for even more ambitious NQS modeling, such as foundation models that simultaneously solve the dynamics across a range of varying physical parameters [213].

Despite substantial room for methodological advances, NQS approaches are beginning to reveal new insights into the dynamical properties of quantum many-body systems. Particularly promising directions connect to the strengths of NQS in targeting arbitrary lattice geometries and strong correlations—recall that dynamical situations are generally plagued by the MC sign problem. The ability to obtain spectral functions in two- or three-dimensional systems could open new perspectives on the characteristics of unconventional equilibrium states of matter such as quantum spin liquids [214], unconventional superconductors [215] and quantum critical points [188]. Once pushed out of equilibrium, quantum many-body systems exhibit a rich phenomenology on transient timescales. Light-induced order has been observed in various forms, but investigating the underlying microscopic mechanisms and developing a notion of transient non-equilibrium phases of matter remains an active area of research [202]. The ability to realize artificial quantum many-body systems in two or three dimensions poses numerous physical questions that advanced NQS techniques may help to shed light on, e.g., concerning phase ordering dynamics [216], complex relaxation dynamics of topological defects [217], or time evolution subject to strong dynamical constraints [218, 219]—including lattice gauge theories [220, 221]. The ongoing developments in quantum computing and quantum simulation demand numerical reference simulations. While recent attempts have remained limited [222–224], NQS provide a possible route to further raise the bar of classical simulation capabilities. Besides benchmarking, NQS may serve as a tool to develop strategies for many-body quantum control and the stabilization of non-equilibrium states of matter in quantum simulators.

6. Conclusion

The state-of-the-art applications discussed in this perspective show that NQS have developed from a proof-of-principle variational ansatz into a competitive tool for selected strongly correlated lattice problems. Their successes in frustrated magnets, doped Hubbard models, and non-equilibrium dynamics point to two useful organizing principles. First, NQS have *versatile applicability* across lattice geometries, entanglement structures, and Hamiltonians for which conventional QMC suffers from a sign problem. Second, they offer *universal approximation capabilities*, so that the approximation error can in principle be reduced arbitrarily by enlarging and refining the variational family. In combination with the varying bias of different architectures, this property allows for the certification of obtained results via self-consistent convergence checks.

The same examples also make clear why NQS are not yet a black-box solver. The main roadblocks are learning non-trivial sign and phase structures, enforcing symmetries and physical constraints without sacrificing expressivity, and obtaining stable real-time evolution and dynamical response. These are serious algorithmic challenges, but—to the best of today’s knowledge—they differ in character from the conventional QMC sign problem or the entanglement bottleneck of tensor networks. In QMC, the sign problem can make statistical errors grow exponentially for a fixed formulation, and in tensor networks, high-dimensional or highly entangled states can exceed the practical capacity of the representation. For NQS, the corresponding obstacles are more often questions of architecture, optimization, sampling, and diagnostics. They can still be prohibitive in practice, but current results suggest that they are not universal no-goes in the same sense.

This distinction is the main reason to continue working on the roadblocks summarized in Fig. 9. Progress would let NQS shed light on long-standing problems such as spin-liquid diagnostics in frustrated magnets, competing orders in multi-orbital and Moiré correlated materials, heavy-fermion quantum criticality, lattice gauge dynamics, and light-induced or transient phases of matter. It would also make NQS valuable as classical verification tools for programmable quantum simulators and quantum computing platforms [198–200], and as a route to theoretical predictions for ultrafast probes of quantum materials [201, 202].

Acknowledgments— The authors acknowledge valuable input and contributions by Filippo Vicentini during early stages of writing. This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – Cluster of Excellence Matter and Light for Quantum Computing (ML4Q) EXC2004/2 – 390534769. JBR and MS were supported via the Helmholtz Initiative and Networking Fund, grant no. VH-NG-1711.

References

- [1] A. Sommerfeld. “Zur Quantentheorie der Spektrallinien”. In: *Annalen der Physik* 356.17 (Jan. 1916), pp. 1–94. ISSN: 0003-3804, 1521-3889. DOI: 10.1002/andp.19163561702. (Visited on 12/11/2025).
- [2] P. W. Anderson. “More Is Different”. In: *Science* 177.4047 (1972), pp. 393–396. DOI: 10.1126/science.177.4047.393. eprint: <https://www.science.org/doi/pdf/10.1126/science.177.4047.393>. URL: <https://www.science.org/doi/abs/10.1126/science.177.4047.393>.
- [3] Walter Metzner et al. “Functional Renormalization Group Approach to Correlated Fermion Systems”. In: *Reviews of Modern Physics* 84.1 (Mar. 2012), pp. 299–352. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.84.299. (Visited on 11/03/2025).
- [4] M. Born and R. Oppenheimer. “Zur Quantentheorie der Molekeln”. In: *Annalen der Physik* 389.20 (Jan. 1927), pp. 457–484. ISSN: 0003-3804, 1521-3889. DOI: 10.1002/andp.19273892002. (Visited on 12/11/2025).
- [5] Nicola Marzari et al. “Maximally Localized Wannier Functions: Theory and Applications”. In: *Reviews of Modern Physics* 84.4 (Oct. 2012), pp. 1419–1475. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.84.1419. (Visited on 12/11/2025).
- [6] Yueqing Chang et al. “Downfolding from Ab Initio to Interacting Model Hamiltonians: Comprehensive Analysis and Benchmarking of the DFT+cRPA Approach”. In: *npj Computational Materials* 10.1 (June 2024), p. 129. ISSN: 2057-3960. DOI: 10.1038/s41524-024-01314-6. (Visited on 12/11/2025).
- [7] F. Aryasetiawan et al. “Frequency-Dependent Local Interactions and Low-Energy Effective Models from Electronic Structure Calculations”. In: *Physical Review B* 70.19 (Nov. 2004), p. 195104. ISSN: 1098-0121, 1550-235X. DOI: 10.1103/PhysRevB.70.195104. (Visited on 12/11/2025).
- [8] Andreas M. Läuchli, Julien Sudan, and Erik S. Sørensen. “Ground-State Energy and Spin Gap of Spin-1/2 Kagomé-Heisenberg Antiferromagnetic Clusters: Large-scale Exact Diagonalization Results”. In: *Physical Review B* 83.21 (June 2011), p. 212401. ISSN: 1098-0121, 1550-235X. DOI: 10.1103/PhysRevB.83.212401. (Visited on 12/11/2025).
- [9] Alexander Weiße and Holger Fehske. “Exact Diagonalization Techniques”. In: 739 (2008). Ed. by H. Fehske, R. Schneider, and A. Weiße, pp. 529–544. DOI: 10.1007/978-3-540-74686-7_18. (Visited on 12/11/2025).
- [10] Pratik Nandy et al. “Quantum Dynamics in Krylov Space: Methods and Applications”. In: *Physics Reports* 1125–1128 (June 2025), pp. 1–82. ISSN: 03701573. DOI: 10.1016/j.physrep.2025.05.001. (Visited on 12/11/2025).
- [11] Andrew J. A. James et al. “Non-Perturbative Methodologies for Low-Dimensional Strongly-Correlated Systems: From Non-Abelian Bosonization to Truncated Spectrum Methods”. In: *Reports on Progress in Physics* 81.4 (2018), p. 046002. DOI: 10.1088/1361-6633/aa91ea. URL: <https://doi.org/10.1088/1361-6633/aa91ea>.
- [12] H. B. Thacker. “Exact Integrability in Quantum Field Theory and Statistical Systems”. In: *Reviews of Modern Physics* 53.2 (1981), pp. 253–285. DOI: 10.1103/RevModPhys.53.253. URL: <https://link.aps.org/doi/10.1103/RevModPhys.53.253>.
- [13] Attila Szabo and Neil S Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. Courier Corporation, 2012.

- [14] Hendrik J. Monkhorst. “Calculation of properties with the coupled-cluster method”. In: *International Journal of Quantum Chemistry* 12.S11 (1977), pp. 421–432. DOI: <https://doi.org/10.1002/qua.560120850>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/qua.560120850>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/qua.560120850>.
- [15] Ulrich Schollwöck. “The density-matrix renormalization group in the age of matrix product states”. In: *Annals of Physics* 326.1 (Jan. 2011), pp. 96–192. ISSN: 0003-4916. DOI: 10.1016/j.aop.2010.09.012. URL: <http://dx.doi.org/10.1016/j.aop.2010.09.012>.
- [16] Román Orús. “Tensor networks for complex quantum systems”. In: *Nature Reviews Physics* 1.9 (Sept. 2019), pp. 538–550. ISSN: 2522-5820. DOI: 10.1038/s42254-019-0086-7. URL: <https://doi.org/10.1038/s42254-019-0086-7>.
- [17] Paul M. Schindler et al. “Variational Ansatz for the Ground State of the Quantum Sherrington-Kirkpatrick Model”. In: *Physical Review Letters* 129.22 (2022), p. 220401. DOI: 10.1103/PhysRevLett.129.220401. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.220401>.
- [18] F. F. Assaad and H. G. Evertz. “World-Line and Determinantal Quantum Monte Carlo Methods for Spins, Phonons and Electrons”. In: *Computational Many-Particle Physics*. Vol. 739. Lecture Notes in Physics. Berlin, Heidelberg: Springer, 2008, pp. 277–356. DOI: 10.1007/978-3-540-74686-7_10. URL: https://doi.org/10.1007/978-3-540-74686-7_10.
- [19] Federico Becca and Sandro Sorella. *Quantum Monte Carlo Approaches for Correlated Systems*. 1st ed. Cambridge: Cambridge University Press, Nov. 2017. ISBN: 978-1-107-12993-1 978-1-316-41704-1. DOI: 10.1017/9781316417041. (Visited on 04/18/2024).
- [20] Matthias Troyer and Uwe-Jens Wiese. “Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations”. In: *Physical Review Letters* 94.17 (May 2005), p. 170201. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.94.170201. (Visited on 11/03/2025).
- [21] Antoine Georges et al. “Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions”. In: *Rev. Mod. Phys.* 68 (1 Jan. 1996), pp. 13–125. DOI: 10.1103/RevModPhys.68.13. URL: <https://link.aps.org/doi/10.1103/RevModPhys.68.13>.
- [22] Thomas Maier et al. “Quantum cluster theories”. In: *Rev. Mod. Phys.* 77 (3 Oct. 2005), pp. 1027–1080. DOI: 10.1103/RevModPhys.77.1027. URL: <https://link.aps.org/doi/10.1103/RevModPhys.77.1027>.
- [23] A. Toschi, A. A. Katanin, and K. Held. “Dynamical vertex approximation: A step beyond dynamical mean-field theory”. In: *Phys. Rev. B* 75 (4 Jan. 2007), p. 045118. DOI: 10.1103/PhysRevB.75.045118. URL: <https://link.aps.org/doi/10.1103/PhysRevB.75.045118>.
- [24] G. Rohringer et al. “Diagrammatic routes to nonlocal correlations beyond dynamical mean field theory”. In: *Rev. Mod. Phys.* 90 (2 May 2018), p. 025003. DOI: 10.1103/RevModPhys.90.025003. URL: <https://link.aps.org/doi/10.1103/RevModPhys.90.025003>.
- [25] Leon Balents. “Spin Liquids in Frustrated Magnets”. In: *Nature* 464.7286 (Mar. 2010), pp. 199–208. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature08917. (Visited on 10/12/2020).
- [26] J. P. F. LeBlanc et al. “Solutions of the Two-Dimensional Hubbard Model: Benchmarks and Results from a Wide Range of Numerical Algorithms”. In:

- Phys. Rev. X* 5 (4 Dec. 2015), p. 041041. DOI: 10.1103/PhysRevX.5.041041. URL: <https://link.aps.org/doi/10.1103/PhysRevX.5.041041>.
- [27] Markus Schmitt and Markus Heyl. “Quantum Many-Body Dynamics in Two Dimensions with Artificial Neural Networks”. In: *Physical Review Letters* 125.10 (Sept. 2020), p. 100503. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.125.100503. (Visited on 11/14/2025).
 - [28] Giuseppe Carleo and Matthias Troyer. “Solving the Quantum Many-Body Problem with Artificial Neural Networks”. In: *Science* 355.6325 (Feb. 2017), pp. 602–606. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.aag2302. (Visited on 11/03/2025).
 - [29] Maarten Van den Nest. “Simulating quantum computers with probabilistic methods”. In: *Quantum Inf. Comput.* 11 (2009), pp. 784–812. URL: <https://api.semanticscholar.org/CorpusID:16943113>.
 - [30] Tom Westerhout, Mikhail I. Katsnelson, and Andrey A. Bagrov. “Many-Body Quantum Sign Structures as Non-Glassy Ising Models”. In: *Communications Physics* 6.1 (Sept. 2023), p. 275. ISSN: 2399-3650. DOI: 10.1038/s42005-023-01388-6. (Visited on 02/10/2026).
 - [31] Ilya Schurov et al. *Learning Complexity of Many-Body Quantum Sign Structures through the Lens of Boolean Fourier Analysis*. Aug. 2025. DOI: 10.48550/arXiv.2508.09870. arXiv: 2508.09870 [cond-mat]. (Visited on 10/28/2025).
 - [32] Dong-Ling Deng, Xiaopeng Li, and S. Das Sarma. “Quantum Entanglement in Neural Network States”. In: *Physical Review X* 7.2 (May 2017), p. 021021. ISSN: 2160-3308. DOI: 10.1103/PhysRevX.7.021021. (Visited on 12/10/2025).
 - [33] Yoav Levine et al. “Quantum Entanglement in Deep Learning Architectures”. In: *Physical Review Letters* 122.6 (Feb. 2019), p. 065301. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.122.065301. (Visited on 11/13/2025).
 - [34] Or Sharir, Amnon Shashua, and Giuseppe Carleo. “Neural tensor contractions and the expressive power of deep neural quantum states”. In: *Physical Review B* 106.20 (2022), p. 205136.
 - [35] Nisarga Paul. “Bound on Entanglement in Neural Quantum States”. In: *Physical Review Letters* 136.12 (Mar. 2026), p. 120403. DOI: 10.1103/rpj5-cns6. URL: <https://doi.org/10.1103/rpj5-cns6>.
 - [36] G. Cybenko. “Approximation by superpositions of a sigmoidal function”. In: *Math. Control. Signals, Syst.* 2 (4 1989), pp. 303–314. DOI: 10.1007/BF02551274. URL: <https://doi.org/10.1007/BF02551274>.
 - [37] Kurt Hornik. “Approximation capabilities of multilayer feedforward networks”. In: *Neural Netw* 4.2 (1991), pp. 251–257. ISSN: 0893-6080. DOI: [https://doi.org/10.1016/0893-6080\(91\)90009-T](https://doi.org/10.1016/0893-6080(91)90009-T). URL: <https://www.sciencedirect.com/science/article/pii/089360809190009T>.
 - [38] Matija Medvidović and Javier Robledo Moreno. “Neural-Network Quantum States for Many-Body Physics”. In: *The European Physical Journal Plus* 139.7 (July 2024), p. 631. ISSN: 2190-5444. DOI: 10.1140/epjp/s13360-024-05311-y. arXiv: 2402.11014 [cond-mat]. (Visited on 11/03/2025).
 - [39] Hannah Lange et al. “From Architectures to Applications: A Review of Neural Quantum States”. In: *Quantum Science and Technology* 9.4 (Sept. 2024), p. 040501. DOI: 10.1088/2058-9565/ad7168. arXiv: 2402.09402 [cond-mat.dis-nn]. URL: <https://doi.org/10.1088/2058-9565/ad7168>.
 - [40] David Pfau et al. “Accurate computation of quantum excited states with neural networks”. In: *Science* 385.6711 (Aug. 2024). ISSN: 1095-9203. DOI: 10.1126/science.adn0137. URL: <http://dx.doi.org/10.1126/science.adn0137>.

- [41] Jannes Nys et al. “Ab-initio variational wave functions for the time-dependent many-electron Schrödinger equation”. In: *Nature communications* 15.1 (2024), p. 9404.
- [42] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Cambridge, MA: MIT Press, 2016. URL: <https://www.deeplearningbook.org/> (visited on 05/14/2024).
- [43] J. Richter and J. Schulenburg. “The Spin-1/2 J_1 - J_2 Heisenberg Antiferromagnet on the Square Lattice: Exact Diagonalization for $N=40$ Spins”. In: *The European Physical Journal B* 73.1 (Jan. 2010), pp. 117–124. ISSN: 1434-6028, 1434-6036. DOI: 10.1140/epjb/e2009-00400-4. (Visited on 09/29/2025).
- [44] Hong-Chen Jiang, Hong Yao, and Leon Balents. “Spin Liquid Ground State of the Spin- 1 2 Square $J_1 - J_2$ Heisenberg Model”. In: *Physical Review B* 86.2 (July 2012), p. 024424. ISSN: 1098-0121, 1550-235X. DOI: 10.1103/PhysRevB.86.024424. (Visited on 09/29/2025).
- [45] Tao Li et al. “Gapped Spin Liquid Phase in the J_1 - J_2 Heisenberg Model by a Bosonic Resonating Valence-Bond Ansatz”. In: *Physical Review B* 86.7 (Aug. 2012), p. 075111. DOI: 10.1103/PhysRevB.86.075111. arXiv: 1205.3838 [cond-mat.str-el]. URL: <https://arxiv.org/abs/1205.3838>.
- [46] Ling Wang et al. “Constructing Gapless Spin Liquid State for the Spin-1/2 J_1 - J_2 Heisenberg Model on a Square Lattice”. In: *Physical Review Letters* 111.3 (July 2013), p. 037202. DOI: 10.1103/PhysRevLett.111.037202. arXiv: 1301.4492 [cond-mat.str-el]. URL: <https://arxiv.org/abs/1301.4492>.
- [47] Wen-Jun Hu et al. “Direct Evidence for a Gapless Z_2 Spin Liquid by Frustrating Néel Antiferromagnetism”. In: *Physical Review B* 88.6 (Aug. 2013), p. 060402. ISSN: 1098-0121, 1550-235X. DOI: 10.1103/PhysRevB.88.060402. (Visited on 09/29/2025).
- [48] Shou-Shu Gong et al. “Plaquette Ordered Phase and Quantum Phase Diagram in the Spin- 1 2 $J_1 - J_2$ Square Heisenberg Model”. In: *Physical Review Letters* 113.2 (July 2014), p. 027201. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.113.027201. (Visited on 09/29/2025).
- [49] Satoshi Morita, Ryui Kaneko, and Masatoshi Imada. “Quantum Spin Liquid in Spin 1/2 $J_1 - J_2$ Heisenberg Model on Square Lattice: Many-Variable Variational Monte Carlo Study Combined with Quantum-Number Projections”. In: *Journal of the Physical Society of Japan* 84.2 (Feb. 2015), p. 024720. ISSN: 0031-9015, 1347-4073. DOI: 10.7566/JPSJ.84.024720. (Visited on 09/29/2025).
- [50] Ling Wang et al. “Tensor-Product State Approach to Spin-1/2 Square J_1 - J_2 Antiferromagnetic Heisenberg Model: Evidence for Deconfined Quantum Criticality”. In: *Physical Review B* 94.7 (Aug. 2016), p. 075143. DOI: 10.1103/PhysRevB.94.075143. arXiv: 1112.3331 [cond-mat.str-el]. URL: <https://arxiv.org/abs/1112.3331>.
- [51] R. Haghshenas and D. N. Sheng. “ $U(1)$ -Symmetric Infinite Projected Entangled-Pair States Study of the Spin-1/2 Square $J_1 - J_2$ Heisenberg Model”. In: *Physical Review B* 97.17 (May 2018), p. 174408. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.97.174408. (Visited on 09/29/2025).
- [52] Wen-Yuan Liu et al. “Gapless Spin Liquid Ground State of the Spin- 1 2 $J_1 - J_2$ Heisenberg Model on Square Lattices”. In: *Physical Review B* 98.24 (Dec. 2018), p. 241109. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.98.241109. (Visited on 09/29/2025).
- [53] Ling Wang and Anders W. Sandvik. “Critical Level Crossings and Gapless Spin Liquid in the Square-Lattice Spin- 1 / 2 $J_1 - J_2$ Heisenberg Antiferromagnet”.

- In: *Physical Review Letters* 121.10 (Sept. 2018), p. 107202. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.121.107202. (Visited on 09/29/2025).
- [54] Francesco Ferrari and Federico Becca. “Gapless Spin Liquid and Valence-Bond Solid in the $J_1 - J_2$ Heisenberg Model on the Square Lattice: Insights from Singlet and Triplet Excitations”. In: *Physical Review B* 102.1 (July 2020), p. 014417. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.102.014417. (Visited on 09/29/2025).
 - [55] Yusuke Nomura and Masatoshi Imada. “Dirac-Type Nodal Spin Liquid Revealed by Refined Quantum Many-Body Solver Using Neural-Network Wave Function, Correlation Ratio, and Level Spectroscopy”. In: *Physical Review X* 11.3 (Aug. 2021), p. 031034. ISSN: 2160-3308. DOI: 10.1103/PhysRevX.11.031034. (Visited on 04/04/2024).
 - [56] Christopher Roth, Attila Szabó, and Allan H. MacDonald. “High-Accuracy Variational Monte Carlo for Frustrated Magnets with Deep Neural Networks”. In: *Physical Review B* 108.5 (Aug. 2023), p. 054410. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.108.054410. (Visited on 06/13/2024).
 - [57] Ao Chen and Markus Heyl. “Empowering deep neural quantum states through efficient optimization”. In: *Nature Physics* 20.9 (Sept. 1, 2024), pp. 1476–1481. ISSN: 1745-2481. DOI: 10.1038/s41567-024-02566-1. URL: <https://doi.org/10.1038/s41567-024-02566-1>.
 - [58] Riccardo Rende et al. “A simple linear algebra identity to optimize large-scale neural network quantum states”. In: *Communications Physics* 7.1 (Aug. 2024). ISSN: 2399-3650. DOI: 10.1038/s42005-024-01732-4. URL: <http://dx.doi.org/10.1038/s42005-024-01732-4>.
 - [59] Luciano Loris Viteritti et al. *Approaching the Thermodynamic Limit with Neural-Network Quantum States*. Feb. 2026. DOI: 10.48550/arXiv.2602.02665. arXiv: 2602.02665 [cond-mat.str-el]. URL: <https://arxiv.org/abs/2602.02665>.
 - [60] C. Lhuillier and G. Misguich. “Introduction to Quantum Spin Liquids”. In: *Introduction to Frustrated Magnetism*. Ed. by Claudine Lacroix, Philippe Mendels, and Frédéric Mila. Vol. 164. Springer Series in Solid-State Sciences. Heidelberg: Springer Berlin Heidelberg, 2011, pp. 23–44. ISBN: 978-3-642-10588-3 978-3-642-10589-0. DOI: 10.1007/978-3-642-10589-0. (Visited on 06/18/2020).
 - [61] Yi Zhou, Kazushi Kanoda, and Tai-Kai Ng. “Quantum Spin Liquid States”. In: *Reviews of Modern Physics* 89.2 (Apr. 2017), p. 025003. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.89.025003. (Visited on 03/01/2022).
 - [62] Lucile Savary and Leon Balents. “Quantum Spin Liquids”. In: *Reports on Progress in Physics* 80.1 (Jan. 2017), p. 016502. ISSN: 0034-4885, 1361-6633. DOI: 10.1088/0034-4885/80/1/016502. arXiv: 1601.03742. (Visited on 06/08/2020).
 - [63] J. Knolle and R. Moessner. “A Field Guide to Spin Liquids”. In: *Annual Review of Condensed Matter Physics* 10.1 (Mar. 2019), pp. 451–472. ISSN: 1947-5454, 1947-5462. DOI: 10.1146/annurev-conmatphys-031218-013401. (Visited on 06/08/2020).
 - [64] C. Broholm et al. “Quantum Spin Liquids”. In: *Science* 367.6475 (Jan. 2020), eaay0668. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.aay0668. (Visited on 08/29/2024).
 - [65] W. Marshall. “Antiferromagnetism”. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 232 (1955), pp. 48–68.

- [66] Yusuke Nomura et al. “Restricted Boltzmann machine learning for solving strongly correlated quantum systems”. In: *Phys. Rev. B* 96 (20 Nov. 2017), p. 205152. DOI: 10.1103/PhysRevB.96.205152. URL: <https://link.aps.org/doi/10.1103/PhysRevB.96.205152>.
- [67] Luciano Loris Viteritti et al. “Transformer Wave Function for Two Dimensional Frustrated Magnets: Emergence of a Spin-Liquid Phase in the Shastry-Sutherland Model”. In: *Physical Review B* 111.13 (Apr. 2025), p. 134411. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.111.134411. (Visited on 06/12/2025).
- [68] Ao Chen et al. *Neural Network-Augmented Pfaffian Wave-functions for Scalable Simulations of Interacting Fermions*. 2025. arXiv: 2507.10705 [cond-mat.str-el]. URL: <https://arxiv.org/abs/2507.10705>.
- [69] B. Keimer et al. “From Quantum Matter to High-Temperature Superconductivity in Copper Oxides”. In: *Nature* 518.7538 (Feb. 2015), pp. 179–186. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature14165. (Visited on 01/29/2026).
- [70] T. A. Maier et al. “Systematic Study of d -Wave Superconductivity in the 2D Repulsive Hubbard Model”. In: *Phys. Rev. Lett.* 95 (23 Nov. 2005), p. 237001. DOI: 10.1103/PhysRevLett.95.237001. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.95.237001>.
- [71] D. J. Scalapino. “Numerical Studies of the 2D Hubbard Model”. In: *Handbook of High-Temperature Superconductivity: Theory and Experiment*. Ed. by J. Robert Schrieffer and James S. Brooks. New York, NY: Springer New York, 2007, pp. 495–526. ISBN: 978-0-387-68734-6. DOI: 10.1007/978-0-387-68734-6_13. URL: https://doi.org/10.1007/978-0-387-68734-6_13.
- [72] Mingpu Qin et al. “Absence of Superconductivity in the Pure Two-Dimensional Hubbard Model”. In: *Phys. Rev. X* 10 (3 July 2020), p. 031016. DOI: 10.1103/PhysRevX.10.031016. URL: <https://link.aps.org/doi/10.1103/PhysRevX.10.031016>.
- [73] Andrew S. Darmawan et al. “Stripe and superconducting order competing in the Hubbard model on a square lattice studied by a combined variational Monte Carlo and tensor network method”. In: *Phys. Rev. B* 98 (20 Nov. 2018), p. 205132. DOI: 10.1103/PhysRevB.98.205132. URL: <https://link.aps.org/doi/10.1103/PhysRevB.98.205132>.
- [74] Bo-Xiao Zheng et al. “Stripe order in the underdoped region of the two-dimensional Hubbard model”. In: *Science* 358.6367 (2017), pp. 1155–1160. DOI: 10.1126/science.aam7127. eprint: <https://www.science.org/doi/pdf/10.1126/science.aam7127>. URL: <https://www.science.org/doi/abs/10.1126/science.aam7127>.
- [75] Kota Ido, Takahiro Ohgoe, and Masatoshi Imada. “Competition among Various Charge-Inhomogeneous States and d -Wave Superconducting State in Hubbard Models on Square Lattices”. In: *Physical Review B* 97.4 (Jan. 2018), p. 045138. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.97.045138. (Visited on 10/13/2025).
- [76] Luca Fausto Tocchio, Arianna Montorsi, and Federico Becca. “Metallic and Insulating Stripes and Their Relation with Superconductivity in the Doped Hubbard Model”. In: *SciPost Physics* 7.2 (Aug. 2019), p. 021. ISSN: 2542-4653. DOI: 10.21468/SciPostPhys.7.2.021. (Visited on 10/13/2025).
- [77] Mingpu Qin et al. “The Hubbard Model: A Computational Perspective”. In: *Annual Review of Condensed Matter Physics* 13. Volume 13, 2022 (2022), pp. 275–302. ISSN: 1947-5462. DOI: <https://doi.org/10.1146/annurev-conmatphys->

- 090921-033948. URL: <https://www.annualreviews.org/content/journals/10.1146/annurev-conmatphys-090921-033948>.
- [78] Yuntian Gu et al. *Solving the Hubbard model with Neural Quantum States*. 2025. arXiv: 2507.02644 [cond-mat.str-el]. URL: <https://arxiv.org/abs/2507.02644>.
 - [79] Christopher Roth et al. *Superconductivity in the Two-Dimensional Hubbard Model Revealed by Neural Quantum States*. 2025. DOI: 10.48550/ARXIV.2511.07566. (Visited on 12/08/2025).
 - [80] Riccardo Rende, Luciano Loris Viteritti, and Antoine Georges. *Superconductivity in the t - t' Hubbard Model from Symmetry-Preserving Neural-Network Quantum States*. Aug. 12, 2026. DOI: 10.48550/arXiv.2608.12465. arXiv: 2608.12465 [cond-mat.str-el]. URL: <http://arxiv.org/abs/2608.12465> (visited on 08/20/2026).
 - [81] Hao Xu et al. “Coexistence of Superconductivity with Partially Filled Stripes in the Hubbard Model”. In: *Science* 384.6696 (May 2024), eadh7691. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.adh7691. (Visited on 02/26/2026).
 - [82] Luciano Loris Viteritti et al. *Beyond Variational Bias: Resolving Intertwined Orders in the Hubbard Model*. 2026. arXiv: 2604.21978 [cond-mat.str-el]. URL: <https://arxiv.org/abs/2604.21978>.
 - [83] Markus Schmitt et al. “Quantum Phase Transition Dynamics in the Two-Dimensional Transverse-Field Ising Model”. In: *Science Advances* 8.37 (Sept. 2022), eabl6850. ISSN: 2375-2548. DOI: 10.1126/sciadv.abl6850. (Visited on 11/03/2025).
 - [84] Andrea Damascelli, Zahid Hussain, and Zhi-Xun Shen. “Angle-Resolved Photoemission Studies of the Cuprate Superconductors”. In: *Reviews of Modern Physics* 75.2 (Apr. 2003), pp. 473–541. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.75.473. (Visited on 02/19/2026).
 - [85] Renate Landig et al. “Measuring the Dynamic Structure Factor of a Quantum Gas Undergoing a Structural Phase Transition”. In: *Nature Communications* 6.1 (May 2015), p. 7046. ISSN: 2041-1723. DOI: 10.1038/ncomms8046. (Visited on 04/07/2026).
 - [86] Dmitry A. Abanin et al. “Many-Body Localization, Thermalization, and Entanglement”. In: *Reviews of Modern Physics* 91.2 (May 2019), p. 021001. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.91.021001. arXiv: 1804.11065 [cond-mat]. (Visited on 11/03/2025).
 - [87] Wojciech De Roeck and François Huveneers. “Stability and Instability towards Delocalization in Many-Body Localization Systems”. In: *Physical Review B* 95.15 (Apr. 2017), p. 155129. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.95.155129. (Visited on 04/07/2026).
 - [88] Frank Wilczek. “Quantum Time Crystals”. In: *Physical Review Letters* 109.16 (Oct. 2012), p. 160401. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.109.160401. (Visited on 04/07/2026).
 - [89] J. Zhang et al. “Observation of a Discrete Time Crystal”. In: *Nature* 543.7644 (Mar. 2017), pp. 217–220. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature21413. (Visited on 04/07/2026).
 - [90] Tiago Mendes-Santos, Markus Schmitt, and Markus Heyl. “Highly Resolved Spectral Functions of Two-Dimensional Systems with Neural Quantum States”. In: *Phys. Rev. Lett.* 131 (4 July 2023), p. 046501. DOI: 10.1103/PhysRevLett.131.046501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.046501>.

- [91] Vighnesh Dattatraya Naik, Zheng-Hang Sun, and Markus Heyl. “Real-time Dynamics in 3D for up to 1000 Qubits with Neural Quantum States: Quenches and the Quantum Kibble–Zurek Mechanism”. In: *arXiv:2604.05032* (2026). URL: <https://arxiv.org/abs/2604.05032>.
- [92] Matija Medvidović and Dries Sels. “Variational quantum dynamics of two-dimensional rotor models”. In: *PRX Quantum* 4.4 (2023), p. 040302.
- [93] Alessandro Sinibaldi et al. “Unbiasing Time-Dependent Variational Monte Carlo by Projected Quantum Evolution”. In: *Quantum* 7 (Oct. 2023), p. 1131. ISSN: 2521-327X. DOI: 10.22331/q-2023-10-10-1131. (Visited on 11/13/2025).
- [94] Anka Van De Walle, Markus Schmitt, and Annabelle Bohrdt. “Many-Body Dynamics with Explicitly Time-Dependent Neural Quantum States”. In: *Machine Learning: Science and Technology* 6.4 (Dec. 2025), p. 045011. ISSN: 2632-2153. DOI: 10.1088/2632-2153/ae0f39. (Visited on 04/01/2026).
- [95] Alessandro Sinibaldi et al. “Time-Dependent Neural Galerkin Method for Quantum Dynamics”. In: *Physical Review Letters* 136.12 (Mar. 2026), p. 120402. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/kqvx-d154. (Visited on 04/01/2026).
- [96] Giacomo Torlai and Roger G. Melko. “Latent Space Purification via Neural Density Operators”. In: *Physical Review Letters* 120.24 (June 2018), p. 240503. DOI: 10.1103/PhysRevLett.120.240503.
- [97] Michael J Hartmann and Giuseppe Carleo. “Neural-network approach to dissipative quantum many-body dynamics”. In: *Physical review letters* 122.25 (2019), p. 250502.
- [98] Alexandra Nagy and Vincenzo Savona. “Variational Quantum Monte Carlo Method with a Neural-Network Ansatz for Open Quantum Systems”. In: *Physical Review Letters* 122.25 (June 2019), p. 250501. DOI: 10.1103/PhysRevLett.122.250501.
- [99] Filippo Vicentini et al. “Variational Neural-Network Ansatz for Steady States in Open Quantum Systems”. In: *Physical Review Letters* 122.25 (June 2019). ISSN: 1079-7114. DOI: 10.1103/physrevlett.122.250503. URL: <http://dx.doi.org/10.1103/PhysRevLett.122.250503>.
- [100] Nobuyuki Yoshioka and Ryusuke Hamazaki. “Constructing neural stationary states for open quantum many-body systems”. In: *Physical Review B* 99.21 (June 2019). ISSN: 2469-9969. DOI: 10.1103/physrevb.99.214306. URL: <http://dx.doi.org/10.1103/PhysRevB.99.214306>.
- [101] Moritz Reh, Markus Schmitt, and Martin Gärttner. “Time-dependent variational principle for open quantum systems with artificial neural networks”. In: *Physical Review Letters* 127.23 (2021), p. 230501.
- [102] Di Luo et al. “Autoregressive Neural Network for Simulating Open Quantum Systems via a Probabilistic Formulation”. In: *Phys. Rev. Lett.* 128 (9 Feb. 2022), p. 090501. DOI: 10.1103/PhysRevLett.128.090501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.128.090501>.
- [103] Simon Kothe and Peter Kirton. “Liouville-Space Neural Network Representation of Density Matrices”. In: *Physical Review A* 109.6 (June 2024), p. 062215. DOI: 10.1103/PhysRevA.109.062215.
- [104] Johannes Mellak, Enrico Arrigoni, and Wolfgang von der Linden. “Deep Neural Networks as Variational Solutions for Correlated Open Quantum Systems”. In: *Communications Physics* 7.1 (Aug. 2024), p. 268. DOI: 10.1038/s42005-024-01757-9.

- [105] Lu Wei et al. “Variational Transformer Ansatz for the Density Operator of Steady States in Dissipative Quantum Many-Body Systems”. In: *Physical Review B* 112.6 (Aug. 2025), p. 064303. DOI: 10.1103/bpf7-d1yp.
- [106] Joshua Lin et al. “Real-time dynamics of the Schwinger model as an open quantum system with Neural Density Operators”. In: *Journal of High Energy Physics* 2024.6 (June 2024). ISSN: 1029-8479. DOI: 10.1007/jhep06(2024)211. URL: [http://dx.doi.org/10.1007/JHEP06\(2024\)211](http://dx.doi.org/10.1007/JHEP06(2024)211).
- [107] Tatiana Vovk et al. “Neural quantum states for emitter dynamics in waveguide QED”. In: *Phys. Rev. A* 113 (5 May 2026), p. 053708. DOI: 10.1103/qccj-6vyt. URL: <https://link.aps.org/doi/10.1103/qccj-6vyt>.
- [108] Gianluca Lagnese et al. *Neural Network Modeling of Many-Body Super- and Sub-Radiant Dynamics*. 2026. arXiv: 2605.04640 [quant-ph].
- [109] Moritz Reh, Markus Schmitt, and Martin Gärttner. “Optimizing Design Choices for Neural Quantum States”. In: *Physical Review B* 107.19 (May 2023), p. 195115. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.107.195115. (Visited on 04/19/2024).
- [110] Chi-Fang (Anthony) Chen, Andrew Lucas, and Chao Yin. “Speed limits and locality in many-body quantum dynamics”. In: *Reports on Progress in Physics* 86.11 (Sept. 2023), p. 116001. DOI: 10.1088/1361-6633/acfaae. URL: <https://doi.org/10.1088/1361-6633/acfaae>.
- [111] Elliott H. Lieb and Derek W. Robinson. “The finite group velocity of quantum spin systems”. In: *Communications in Mathematical Physics* 28.3 (1972), pp. 251–257. ISSN: 1432-0916. DOI: 10.1007/BF01645779. URL: <https://doi.org/10.1007/BF01645779>.
- [112] M. B. Hastings. *Locality in Quantum Systems*. 2010. arXiv: 1008.5137 [math-ph]. URL: <https://arxiv.org/abs/1008.5137>.
- [113] J. Eisert, M. Cramer, and M. B. Plenio. “Colloquium: Area laws for the entanglement entropy”. In: *Rev. Mod. Phys.* 82 (1 Feb. 2010), pp. 277–306. DOI: 10.1103/RevModPhys.82.277. URL: <https://link.aps.org/doi/10.1103/RevModPhys.82.277>.
- [114] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. “Convolutional Networks”. In: *Deep Learning*. Cambridge, MA: MIT Press, 2016. Chap. 9. URL: <https://www.deeplearningbook.org/contents/convnets.html> (visited on 05/14/2024).
- [115] Taco Cohen and Max Welling. “Group Equivariant Convolutional Networks”. In: *Proceedings of the 33rd International Conference on Machine Learning*. New York, NY, 2016.
- [116] Michael El-Batanouny and Frederick Wooten. *Symmetry and Condensed Matter Physics: A Computational Approach*. Cambridge, UK New York: Cambridge University Press, 2008. ISBN: 978-0-511-39482-9 978-0-511-39417-1.
- [117] Daisuke Tahara and Masatoshi Imada. “Variational Monte Carlo Method Combined with Quantum-Number Projection and Multi-Variable Optimization”. In: *Journal of the Physical Society of Japan* 77.11 (Nov. 2008), p. 114701. ISSN: 0031-9015, 1347-4073. DOI: 10.1143/JPSJ.77.114701. (Visited on 07/16/2024).
- [118] Yusuke Nomura. “Helping Restricted Boltzmann Machines with Quantum-State Representation by Restoring Symmetry”. In: *Journal of Physics: Condensed Matter* 33.17 (Apr. 2021), p. 174003. ISSN: 0953-8984, 1361-648X. DOI: 10.1088/1361-648X/abe268. (Visited on 04/24/2024).
- [119] Tom Vieijra et al. “Restricted Boltzmann Machines for Quantum States with Non-Abelian or Anyonic Symmetries”. In: *Phys. Rev. Lett.* 124 (9 Mar. 2020),

- p. 097201. DOI: 10.1103/PhysRevLett.124.097201. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.097201>.
- [120] Di Luo et al. “Gauge-Invariant and Anyonic-Symmetric Autoregressive Neural Network for Quantum Lattice Models”. In: *Physical Review Research* 5.1 (Mar. 2023), p. 013216. ISSN: 2643-1564. DOI: 10.1103/PhysRevResearch.5.013216. (Visited on 05/07/2024).
 - [121] Kenny Choo et al. “Symmetries and Many-Body Excitations with Neural-Network Quantum States”. In: *Physical Review Letters* 121.16 (Oct. 2018), p. 167204. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.121.167204. (Visited on 10/28/2025).
 - [122] Marin Bukov, Markus Schmitt, and Maxime Dupont. “Learning the Ground State of a Non-Stoquastic Quantum Hamiltonian in a Rugged Neural Network Landscape”. In: *SciPost Physics* 10.6 (June 2021), p. 147. ISSN: 2542-4653. DOI: 10.21468/SciPostPhys.10.6.147. (Visited on 10/30/2025).
 - [123] Lixing Zhang and Di Luo. *Neural Transformer Backflow for Solving Momentum-Resolved Ground States of Strongly Correlated Materials*. Sept. 2025. DOI: 10.48550/arXiv.2509.09275. arXiv: 2509.09275 [cond-mat]. (Visited on 09/24/2025).
 - [124] Tom Westerhout et al. “Generalization properties of neural network approximations to frustrated magnet ground states”. In: *Nature Communications* 11.1 (Mar. 27, 2020), p. 1593. ISSN: 2041-1723. DOI: 10.1038/s41467-020-15402-w. URL: <https://doi.org/10.1038/s41467-020-15402-w>.
 - [125] Attila Szabó and Claudio Castelnovo. “Neural network wave functions and the sign problem”. In: *Physical Review Research* 2.3 (2020), p. 033075.
 - [126] Ao Chen et al. “Neural network evolution strategy for solving quantum sign structures”. In: *Phys. Rev. Res.* 4 (2 May 2022), p. L022026. DOI: 10.1103/PhysRevResearch.4.L022026. URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.4.L022026>.
 - [127] Di Luo and Bryan K. Clark. “Backflow Transformations via Neural Networks for Quantum Many-Body Wave Functions”. In: *Phys. Rev. Lett.* 122 (22 June 2019), p. 226401. DOI: 10.1103/PhysRevLett.122.226401. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.226401>.
 - [128] James Stokes et al. “Phases of two-dimensional spinless lattice fermions with first-quantized deep neural-network quantum states”. In: *Phys. Rev. B* 102 (20 Nov. 2020), p. 205122. DOI: 10.1103/PhysRevB.102.205122. URL: <https://link.aps.org/doi/10.1103/PhysRevB.102.205122>.
 - [129] Javier Robledo Moreno et al. “Fermionic wave functions from neural-network constrained hidden states”. In: *Proceedings of the National Academy of Sciences* 119.32 (2022), e2122059119. DOI: 10.1073/pnas.2122059119. eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.2122059119>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.2122059119>.
 - [130] Chloé Gauvin-Ndiaye et al. “Mott Transition and Volume Law Entanglement with Neural Quantum States”. In: *Phys. Rev. Lett.* 134 (7 Feb. 2025), p. 076502. DOI: 10.1103/PhysRevLett.134.076502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.134.076502>.
 - [131] Hannah Lange et al. “Simulating the Two-Dimensional t - J Model at Finite Doping with Neural Quantum States”. In: *Phys. Rev. Lett.* 135 (13 Sept. 2025), p. 136504. DOI: 10.1103/rc31-5h19. URL: <https://link.aps.org/doi/10.1103/rc31-5h19>.

- [132] Fei-Fei Li et al. *CS231n Convolutional Neural Networks for Visual Recognition*. <https://cs231n.github.io/convolutional-networks/#overview>. (Visited on 04/19/2024).
- [133] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. “ImageNet Classification with Deep Convolutional Neural Networks”. In: *Advances in Neural Information Processing Systems*. Vol. 25. Curran Associates, Inc., 2012. (Visited on 01/10/2025).
- [134] Xiao Liang et al. “Solving Frustrated Quantum Many-Particle Models with Convolutional Neural Networks”. In: *Physical Review B* 98.10 (Sept. 2018), p. 104426. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.98.104426. (Visited on 11/13/2025).
- [135] Kenny Choo, Titus Neupert, and Giuseppe Carleo. “Two-Dimensional Frustrated J 1 - J 2 Model Studied with Neural Network Quantum States”. In: *Physical Review B* 100.12 (Sept. 2019), p. 125124. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.100.125124. (Visited on 11/13/2025).
- [136] Louis Sharma et al. *Comparing Symmetrized Determinant Neural Quantum States for the Hubbard Model*. Oct. 2025. DOI: 10.48550/arXiv.2510.11710. arXiv: 2510.11710 [cond-mat]. (Visited on 11/07/2025).
- [137] Michael M. Bronstein et al. *Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges*. May 2021. DOI: 10.48550/arXiv.2104.13478. arXiv: 2104.13478 [cs]. (Visited on 03/24/2026).
- [138] Christopher Roth and Allan H. MacDonald. *Group Convolutional Neural Networks Improve Quantum State Accuracy*. May 2021. arXiv: 2104.05085 [cond-mat, physics:quant-ph]. (Visited on 04/19/2024).
- [139] Luciano Loris Viteritti, Riccardo Rende, and Federico Becca. “Transformer Variational Wave Functions for Frustrated Quantum Spin Systems”. In: *Physical Review Letters* 130.23 (June 2023), p. 236401. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.130.236401. (Visited on 04/19/2024).
- [140] Rajah P. Nutakki, Ahmedeo Shokry, and Filippo Vicentini. *Design Principles of Deep Translationally-Symmetric Neural Quantum States for Frustrated Magnets*. May 2025. DOI: 10.48550/arXiv.2505.03466. arXiv: 2505.03466 [cond-mat]. (Visited on 10/23/2025).
- [141] Kaiming He et al. *Deep Residual Learning for Image Recognition*. Dec. 2015. DOI: 10.48550/arXiv.1512.03385. arXiv: 1512.03385 [cs]. (Visited on 11/14/2025).
- [142] Alexey Dosovitskiy et al. *An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale*. June 2021. arXiv: 2010.11929 [cs]. (Visited on 05/10/2024).
- [143] Zhuang Liu et al. “A ConvNet for the 2020s”. In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. New Orleans, LA, USA: IEEE, June 2022, pp. 11966–11976. ISBN: 978-1-6654-6946-3. DOI: 10.1109/CVPR52688.2022.01167. (Visited on 10/06/2025).
- [144] Imelda Romero, Jannes Nys, and Giuseppe Carleo. “Spectroscopy of two-dimensional interacting lattice electrons using symmetry-aware neural backflow transformations”. In: *Communications Physics* 8.1 (2025), p. 46. ISSN: 2399-3650. DOI: 10.1038/s42005-025-01955-z. URL: <https://doi.org/10.1038/s42005-025-01955-z>.
- [145] Luca Gravina, Vincenzo Savona, and Filippo Vicentini. “Neural Projected Quantum Dynamics: A Systematic Study”. In: *Quantum* 9 (July 2025), p. 1803.

- ISSN: 2521-327X. DOI: 10.22331/q-2025-07-22-1803. arXiv: 2410.10720 [quant-ph]. (Visited on 11/14/2025).
- [146] Ao Chen, Vighnesh Dattatraya Naik, and Markus Heyl. *Convolutional Transformer Wave Functions*. Mar. 2025. DOI: 10.48550/arXiv.2503.10462. arXiv: 2503.10462 [cond-mat]. (Visited on 03/17/2025).
 - [147] Mohamed Hibat-Allah et al. “Recurrent Neural Network Wave Functions”. In: *Physical Review Research* 2.2 (June 2020), p. 023358. ISSN: 2643-1564. DOI: 10.1103/PhysRevResearch.2.023358. (Visited on 10/29/2025).
 - [148] Yuan-Hang Zhang and Massimiliano Di Ventra. “Transformer Quantum State: A Multipurpose Model for Quantum Many-Body Problems”. In: *Physical Review B* 107.7 (Feb. 2023), p. 075147. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.107.075147. (Visited on 10/30/2025).
 - [149] Kyle Sprague and Stefanie Czischek. “Variational Monte Carlo with Large Patched Transformers”. In: *Communications Physics* 7.1 (Mar. 2024), p. 90. ISSN: 2399-3650. DOI: 10.1038/s42005-024-01584-y. (Visited on 10/29/2025).
 - [150] Mohamed Hibat-Allah, Roger G. Melko, and Juan Carrasquilla. “Investigating Topological Order Using Recurrent Neural Networks”. In: *Physical Review B* 108.7 (Aug. 2023), p. 075152. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.108.075152. (Visited on 10/29/2025).
 - [151] M. Schuyler Moss et al. *Leveraging Recurrence in Neural Network Wavefunctions for Large-Scale Simulations of Heisenberg Antiferromagnets on the Square Lattice*. Oct. 2025. DOI: 10.48550/arXiv.2502.17144. arXiv: 2502.17144 [cond-mat]. (Visited on 10/28/2025).
 - [152] Massimo Bortone, Yannic Rath, and George H. Booth. “Impact of Conditional Modelling for a Universal Autoregressive Quantum State”. In: *Quantum* 8 (Feb. 2024), p. 1245. ISSN: 2521-327X. DOI: 10.22331/q-2024-02-08-1245. arXiv: 2306.05917 [quant-ph]. (Visited on 10/30/2025).
 - [153] Di Luo et al. “Gauge Equivariant Neural Networks for Quantum Lattice Gauge Theories”. In: *Physical Review Letters* 127.27 (Dec. 2021), p. 276402. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.127.276402. (Visited on 08/23/2024).
 - [154] Di Luo et al. *Gauge Equivariant Neural Networks for 2+1D U(1) Gauge Theory Simulations in Hamiltonian Formulation*. Nov. 2022. arXiv: 2211.03198 [cond-mat, physics:hep-lat, physics:quant-ph]. (Visited on 08/23/2024).
 - [155] Matteo Favoni et al. “Lattice Gauge Equivariant Convolutional Neural Networks”. In: *Physical Review Letters* 128.3 (Jan. 2022), p. 032003. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.128.032003. (Visited on 08/23/2024).
 - [156] Dominik S Kufel et al. “Approximately Symmetric Neural Networks for Quantum Spin Liquids”. In: *Physical Review Letters* 135.5 (2025), p. 056702.
 - [157] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. Jan. 2017. DOI: 10.48550/arXiv.1412.6980. arXiv: 1412.6980 [cs]. (Visited on 01/28/2026).
 - [158] Mohamed Hibat-Allah et al. “Recurrent Neural Network Wave Functions for Rydberg Atom Arrays on Kagome Lattice”. In: *Communications Physics* 8.1 (July 2025), p. 308. ISSN: 2399-3650. DOI: 10.1038/s42005-025-02226-7. (Visited on 11/13/2025).
 - [159] Hannah Lange et al. “Neural Network Approach to Quasiparticle Dispersions in Doped Antiferromagnets”. In: *Communications Physics* 7.1 (June 2024), p. 187.

- ISSN: 2399-3650. DOI: 10.1038/s42005-024-01678-7. arXiv: 2310.08578 [cond-mat]. (Visited on 10/30/2025).
- [160] Andrew Jreissaty et al. “Entanglement and Optimization within Autoregressive Neural Quantum States”. In: *Physical Review Research* 8.1 (Feb. 2026), p. 013147. ISSN: 2643-1564. DOI: 10.1103/t2cg-kr7y. (Visited on 03/24/2026).
 - [161] Jonas B. Rigo and Markus Schmitt. *Operator Lanczos Approach Enabling Neural Quantum States as Real-Frequency Impurity Solvers*. Dec. 2025. DOI: 10.48550/arXiv.2512.08624. arXiv: 2512.08624 [cond-mat]. (Visited on 02/11/2026).
 - [162] Aleksei Malyshev, Markus Schmitt, and A. I. Lvovsky. *Neural Quantum States and Peaked Molecular Wave Functions: Curse or Blessing?* Aug. 2024. DOI: 10.48550/arXiv.2408.07625. arXiv: 2408.07625 [quant-ph]. (Visited on 10/30/2025).
 - [163] Sandro Sorella. “Green Function Monte Carlo with Stochastic Reconfiguration”. In: *Physical Review Letters* 80.20 (May 1998), pp. 4558–4561. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.80.4558. (Visited on 04/04/2024).
 - [164] S. Amari and S. C. Douglas. “Why Natural Gradient?” In: *Proceedings 1998 IEEE International Conference on Acoustics, Speech and Signal Processing*. Vol. 2. 1998, pp. 1213–1216.
 - [165] Luciano Loris Viteritti et al. *Transformer Wave Function for the Shastry-Sutherland Model: Emergence of a Spin-Liquid Phase*. Feb. 2024. arXiv: 2311.16889 [cond-mat]. (Visited on 04/19/2024).
 - [166] Gil Goldshlager, Nilin Abrahamsen, and Lin Lin. “A Kaczmarz-inspired approach to accelerate the optimization of neural network wavefunctions”. In: *Journal of Computational Physics* 516 (Nov. 2024), p. 113351. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2024.113351. URL: <http://dx.doi.org/10.1016/j.jcp.2024.113351>.
 - [167] Dian Wu et al. “Variational benchmarks for quantum many-body problems”. In: *Science* 386.6719 (2024), pp. 296–301. DOI: 10.1126/science.adg9774. eprint: <https://www.science.org/doi/pdf/10.1126/science.adg9774>. URL: <https://www.science.org/doi/abs/10.1126/science.adg9774>.
 - [168] Helia Kamal et al. *Comment on “Spin-1/2 Kagome Heisenberg Antiferromagnet: Machine Learning Discovery of the Spinon Pair-Density-Wave Ground State”*. May 2026. arXiv: 2605.28861 [cond-mat.str-el]. URL: <https://arxiv.org/abs/2605.28861>.
 - [169] Damian Hofmann et al. “Role of stochastic noise and generalization error in the time propagation of neural-network quantum states”. en. In: *SciPost Physics* 12.5 (May 2022), p. 165. ISSN: 2542-4653. DOI: 10.21468/SciPostPhys.12.5.165. URL: <https://scipost.org/10.21468/SciPostPhys.12.5.165> (visited on 07/22/2024).
 - [170] Markus Schmitt and Moritz Reh. “jVMC: Versatile and Performant Variational Monte Carlo Leveraging Automated Differentiation and GPU Acceleration”. In: *SciPost Physics Codebases* (Aug. 2022), p. 2. DOI: 10.21468/SciPostPhysCodeb.2. (Visited on 11/17/2025).
 - [171] Gianluca Lagnese et al. “Neural network modeling of many-body super- and sub-radiant dynamics”. In: *arXiv:2605.04640* (2026). DOI: 10.48550/arXiv.2605.04640. URL: <http://arxiv.org/abs/2605.04640> (visited on 05/07/2026).
 - [172] Wladislaw Krinitsin and Markus Schmitt. “Time-dependent variational Monte Carlo without bias”. In: *arXiv:2605.03930* (2026). URL: <https://arxiv.org/abs/2605.03930>.

- [173] Wouter Kool, Herke van Hoof, and Max Welling. *Stochastic Beams and Where to Find Them: The Gumbel-Top-k Trick for Sampling Sequences Without Replacement*. May 2019. DOI: 10.48550/arXiv.1903.06059. arXiv: 1903.06059 [cs]. (Visited on 01/28/2026).
- [174] Ao Chen and Christopher Roth. *Lrux: Fast Low-Rank Updates of Determinants and Pfaffians in JAX*. Feb. 2026. DOI: 10.48550/arXiv.2602.05255. arXiv: 2602.05255 [cond-mat]. (Visited on 04/05/2026).
- [175] Filippo Vicentini et al. “NetKet 3: Machine Learning Toolbox for Many-Body Quantum Systems”. In: *SciPost Physics Codebases* (Aug. 2022), p. 7. DOI: 10.21468/SciPostPhysCodeb.7. (Visited on 11/17/2025).
- [176] Chen Ao, chrisrothUT, and Zhanghao Zhouyin. *ChenAo-Phys/quantax*. <https://github.com/ChenAo-Phys/quantax>. Apr. 2026. URL: <https://github.com/ChenAo-Phys/quantax>.
- [177] James Bradbury et al. *JAX: composable transformations of Python+NumPy programs*. Version 0.3.13. 2018. URL: <http://github.com/jax-ml/jax>.
- [178] Sergey Bravyi et al. *The Complexity of Stoquastic Local Hamiltonian Problems*. Oct. 2007. DOI: 10.48550/arXiv.quant-ph/0606140. arXiv: quant-ph/0606140. (Visited on 11/03/2025).
- [179] Milad Marvian, Daniel A. Lidar, and Itay Hen. “On the Computational Complexity of Curing Non-Stoquastic Hamiltonians”. In: *Nature Communications* 10 (2019), p. 1571. DOI: 10.1038/s41467-019-09501-6. URL: <https://doi.org/10.1038/s41467-019-09501-6>.
- [180] V. J. Emery. “Theory of High- T_c Superconductivity in Oxides”. In: *Physical Review Letters* 58.26 (June 1987), pp. 2794–2797. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.58.2794. (Visited on 02/03/2026).
- [181] Danfeng Li et al. “Superconductivity in an Infinite-Layer Nickelate”. In: *Nature* 572.7771 (Aug. 2019), pp. 624–627. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-019-1496-5. (Visited on 02/03/2026).
- [182] Yusuke Nomura and Ryotaro Arita. “Superconductivity in Infinite-Layer Nickelates”. In: *Reports on Progress in Physics* 85.5 (May 2022), p. 052501. ISSN: 0034-4885, 1361-6633. DOI: 10.1088/1361-6633/ac5a60. (Visited on 02/03/2026).
- [183] Yusuke Nomura et al. “Strong-Coupling High- T_c Superconductivity in Doped Correlated Band Insulators”. In: *Physical Review B* 112.2 (July 2025), p. L020504. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/dygc-94fq. (Visited on 01/29/2026).
- [184] Shaofeng Wang et al. “Emerging Characteristics and Properties of Moiré Materials”. In: *Nanomaterials* 13.21 (Oct. 2023), p. 2881. ISSN: 2079-4991. DOI: 10.3390/nano13212881. (Visited on 02/03/2026).
- [185] Di Luo, David D. Dai, and Liang Fu. *Simulating moiré quantum matter with neural network*. 2024. arXiv: 2406.17645 [cond-mat.str-el]. URL: <https://arxiv.org/abs/2406.17645>.
- [186] Chetan Nayak et al. “Non-Abelian Anyons and Topological Quantum Computation”. In: *Reviews of Modern Physics* 80.3 (2008), pp. 1083–1159. DOI: 10.1103/RevModPhys.80.1083. URL: <https://link.aps.org/doi/10.1103/RevModPhys.80.1083>.
- [187] Philipp Gegenwart, Qimiao Si, and Frank Steglich. “Quantum Criticality in Heavy-Fermion Metals”. In: *Nature Physics* 4.3 (Mar. 2008), pp. 186–197. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/nphys892. (Visited on 01/29/2026).
- [188] Stefan Kirchner et al. “Colloquium : Heavy-electron Quantum Criticality and Single-Particle Spectroscopy”. In: *Reviews of Modern Physics* 92.1 (Mar. 2020),

- p. 011002. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.92.011002. (Visited on 01/29/2026).
- [189] Chong Wang et al. “Deconfined Quantum Critical Points: Symmetries and Dualities”. In: *Physical Review X* 7.3 (Sept. 2017), p. 031051. DOI: 10.1103/PhysRevX.7.031051. arXiv: 1703.02426 [cond-mat.str-el]. URL: <https://link.aps.org/doi/10.1103/PhysRevX.7.031051>.
 - [190] Hui Shao, Wenan Guo, and Anders W. Sandvik. “Quantum Criticality with Two Length Scales”. In: *Science* 352.6282 (Apr. 2016), pp. 213–216. DOI: 10.1126/science.aad5007. URL: <https://www.science.org/doi/10.1126/science.aad5007>.
 - [191] Riccardo Rende et al. *Scaling Laws for Neural-Network Quantum States*. June 2026. DOI: 10.48550/arXiv.2606.02794. arXiv: 2606.02794 [cond-mat.dis-nn]. URL: <https://arxiv.org/abs/2606.02794> (visited on 06/03/2026).
 - [192] Chae-Yeun Park and Michael J. Kastoryano. “Expressive Power of Complex-Valued Restricted Boltzmann Machines for Solving Nonstoquastic Hamiltonians”. In: *Physical Review B* 106.13 (Oct. 2022), p. 134437. ISSN: 2469-9950, 2469-9969. DOI: 10.1103/PhysRevB.106.134437. (Visited on 03/25/2026).
 - [193] Thomas Spriggs et al. *Accurate Ground States of $SU(2)$ Lattice Gauge Theory in $2 + 1D$ and $3 + 1D$* . Mar. 2026. DOI: 10.1103/wsk2-qcvy. URL: <https://link.aps.org/doi/10.1103/wsk2-qcvy>.
 - [194] Tom Vieijra and Jannes Nys. “Many-body quantum states with exact conservation of non-Abelian and lattice symmetries through variational Monte Carlo”. In: *Phys. Rev. B* 104 (4 July 2021), p. 045123. DOI: 10.1103/PhysRevB.104.045123. URL: <https://link.aps.org/doi/10.1103/PhysRevB.104.045123>.
 - [195] Xiao-Gang Wen. “Colloquium: Zoo of Quantum-Topological Phases of Matter”. In: *Reviews of Modern Physics* 89.4 (2017), p. 041004. DOI: 10.1103/RevModPhys.89.041004. URL: <https://link.aps.org/doi/10.1103/RevModPhys.89.041004>.
 - [196] N. Brambilla et al. “QCD and Strongly Coupled Gauge Theories: Challenges and Perspectives”. In: *The European Physical Journal C* 74.10 (Oct. 2014), p. 2981. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-014-2981-5. (Visited on 12/10/2025).
 - [197] Keitaro Nagata. “Finite-Density Lattice QCD and Sign Problem: Current Status and Open Problems”. In: *Progress in Particle and Nuclear Physics* 127 (Nov. 2022), p. 103991. ISSN: 01466410. DOI: 10.1016/j.pnpnp.2022.103991. (Visited on 12/10/2025).
 - [198] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger. “Many-Body Physics with Ultracold Gases”. In: *Reviews of Modern Physics* 80.3 (2008), pp. 885–964. DOI: 10.1103/RevModPhys.80.885. URL: <https://doi.org/10.1103/RevModPhys.80.885>.
 - [199] I. M. Georgescu, S. Ashhab, and Franco Nori. “Quantum Simulation”. In: *Reviews of Modern Physics* 86.1 (2014), pp. 153–185. DOI: 10.1103/RevModPhys.86.153. URL: <https://doi.org/10.1103/RevModPhys.86.153>.
 - [200] Ehud Altman et al. “Quantum Simulators: Architectures and Opportunities”. In: *PRX Quantum* 2.1 (2021), p. 017003. DOI: 10.1103/PRXQuantum.2.017003. URL: <https://doi.org/10.1103/PRXQuantum.2.017003>.
 - [201] D. N. Basov, R. D. Averitt, and D. Hsieh. “Towards Properties on Demand in Quantum Materials”. In: *Nature Materials* 16 (2017), pp. 1077–1088. DOI: 10.1038/nmat5017. URL: <https://doi.org/10.1038/nmat5017>.

- [202] Alberto de la Torre et al. “Colloquium: Nonthermal pathways to ultrafast control in quantum materials”. In: *Rev. Mod. Phys.* 93 (4 Oct. 2021), p. 041002. DOI: 10.1103/RevModPhys.93.041002. URL: <https://link.aps.org/doi/10.1103/RevModPhys.93.041002>.
- [203] D. Fausti et al. “Light-Induced Superconductivity in a Stripe-Ordered Cuprate”. In: *Science* 331.6014 (Jan. 2011), pp. 189–191. DOI: 10.1126/science.1197294.
- [204] Soonwon Choi et al. “Observation of discrete time-crystalline order in a disordered dipolar many-body system”. In: *Nature* 543.7644 (Mar. 2017), pp. 221–225. DOI: 10.1038/nature21426.
- [205] Matija Medvidović et al. *Adiabatic transport of neural network quantum states*. 2025. arXiv: 2510.15030 [quant-ph]. URL: <https://arxiv.org/abs/2510.15030>.
- [206] Agnes Valenti et al. “Correlation-enhanced neural networks as interpretable variational quantum states”. In: *Phys. Rev. Res.* 4 (1 Jan. 2022), p. L012010. DOI: 10.1103/PhysRevResearch.4.L012010. URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.4.L012010>.
- [207] T. Mendes-Santos et al. “Wave-Function Network Description and Kolmogorov Complexity of Quantum Many-Body Systems”. In: *Phys. Rev. X* 14 (2 May 2024), p. 021029. DOI: 10.1103/PhysRevX.14.021029. URL: <https://link.aps.org/doi/10.1103/PhysRevX.14.021029>.
- [208] Tobias Wiener, Laurin Brunner, and Markus Heyl. “Universal Symmetry-Breaking Dynamics at Continuous Phase Transitions: Evidence for a New Dynamical Critical Exponent”. In: *arXiv:2605.07753* (2026). URL: <https://arxiv.org/abs/2605.07753>.
- [209] Vighnesh Dattatraya Naik and Markus Heyl. “Persistent coherent quantum dynamics in 2D long-range magnets via magnon binding”. In: *arXiv:2512.09037* (2026). URL: <https://arxiv.org/abs/2512.09037>.
- [210] Andrew D. King et al. “Beyond-classical computation in quantum simulation”. In: *Science* 388.6743 (2025), pp. 199–204. DOI: 10.1126/science.ado6285. eprint: <https://www.science.org/doi/pdf/10.1126/science.ado6285>. URL: <https://www.science.org/doi/abs/10.1126/science.ado6285>.
- [211] Joseph Vovrosh et al. “Simulating dynamics of the two-dimensional transverse-field Ising model: a comparative study of large-scale classical numerics”. In: *arXiv:2511.19340* (2025). URL: <https://arxiv.org/abs/2511.19340>.
- [212] M. Raissi, P. Perdikaris, and G.E. Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational Physics* 378 (2019), pp. 686–707. ISSN: 0021-9991. DOI: <https://doi.org/10.1016/j.jcp.2018.10.045>. URL: <https://www.sciencedirect.com/science/article/pii/S0021999118307125>.
- [213] Zihao Qi, Christopher Earls, and Yang Peng. “Neural Operator Quantum State: A Foundation Model for Quantum Dynamics”. In: *arXiv:2603.25066* (2026). URL: <https://arxiv.org/abs/2603.25066>.
- [214] Arnab Banerjee et al. “Neutron Scattering in the Proximate Quantum Spin Liquid α -RuCl₃”. In: *Science* 356.6342 (June 2017), pp. 1055–1059. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.aah6015. (Visited on 12/09/2025).
- [215] G. R. Stewart. “Unconventional Superconductivity”. In: *Advances in Physics* 66.2 (Apr. 2017), pp. 75–196. ISSN: 0001-8732. DOI: 10.1080/00018732.2017.1331615. (Visited on 02/03/2026).

- [216] Tom Manovitz et al. “Quantum coarsening and collective dynamics on a programmable simulator”. en. In: *Nature* 638.8049 (Feb. 2025), pp. 86–92. ISSN: 1476-4687. DOI: 10.1038/s41586-024-08353-5. URL: <https://www.nature.com/articles/s41586-024-08353-5> (visited on 07/01/2025).
- [217] Wladislaw Krinitsin et al. “Roughening Dynamics of Interfaces in the Two-Dimensional Quantum Ising Model”. In: *Physical Review Letters* 134.24 (June 2025), p. 240402. DOI: 10.1103/9bsk-x9rw. URL: <https://link.aps.org/doi/10.1103/9bsk-x9rw> (visited on 06/25/2025).
- [218] Sanjay Moudgalya, B. Andrei Bernevig, and Nicolas Regnault. “Quantum many-body scars and Hilbert space fragmentation: a review of exact results”. In: *Reports on Progress in Physics* 85.8 (2022), p. 086501. DOI: 10.1088/1361-6633/ac73a0. URL: <https://iopscience.iop.org/article/10.1088/1361-6633/ac73a0>.
- [219] Daniel Adler et al. “Observation of Hilbert space fragmentation and fractonic excitations in 2D”. In: *Nature* 636.8041 (2024), pp. 80–85. DOI: 10.1038/s41586-024-08188-0. URL: <https://www.nature.com/articles/s41586-024-08188-0>.
- [220] Monika Aidelsburger et al. “Cold atoms meet lattice gauge theory”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 380.2216 (2022), p. 20210064. DOI: 10.1098/rsta.2021.0064. URL: <https://royalsocietypublishing.org/doi/10.1098/rsta.2021.0064>.
- [221] Michael Meth et al. “Simulating two-dimensional lattice gauge theories on a qudit quantum computer”. In: *Nature Physics* 21.4 (2025), pp. 570–576. DOI: 10.1038/s41567-025-02797-w. URL: <https://www.nature.com/articles/s41567-025-02797-w>.
- [222] Andrew D. King et al. “Beyond-classical computation in quantum simulation”. In: *Science* 388.6743 (2025), pp. 199–204. DOI: 10.1126/science.ado6285. eprint: <https://www.science.org/doi/pdf/10.1126/science.ado6285>. URL: <https://www.science.org/doi/abs/10.1126/science.ado6285>.
- [223] Joseph Vovrosh et al. “Simulating dynamics of the two-dimensional transverse-field Ising model: A comparative study of large-scale classical numerics”. In: *Phys. Rev. Res.* 8 (2 June 2026), p. 023311. DOI: 10.1103/x4gs-g63p. URL: <https://link.aps.org/doi/10.1103/x4gs-g63p>.
- [224] R. Haghshenas et al. “Digital quantum magnetism on a trapped-ion quantum computer”. In: *Nature* 653.8113 (2026), pp. 56–62. DOI: 10.1038/s41586-026-10445-3. URL: <https://www.nature.com/articles/s41586-026-10445-3>.