

Role of stochastic noise and generalization error in the time propagation of neural-network quantum states

D. Hofmann^{1 α} , G. Fabiani^{2 β} , J. H. Mentink^{2 γ} , G. Carleo^{3 δ} , M. A. Sentef^{1 ϵ}

1 Max Planck Institute for the Structure and Dynamics of Matter,
Luruper Chaussee 149, 22761 Hamburg, Germany

2 Radboud University, Institute for Molecules and Materials,
Heyendaalseweg 135, 6525 AJ Nijmegen, The Netherlands

3 Institute of Physics, École Polytechnique Fédérale de Lausanne (EPFL),
1015 Lausanne, Switzerland

^{α} damian.hofmann@mpsd.mpg.de, ^{β} g.fabiani@science.ru.nl, ^{γ} j.mentink@science.ru.nl,
 ^{δ} giuseppe.carleo@epfl.ch, ^{ϵ} michael.sentef@mpsd.mpg.de

Abstract

Neural-network quantum states (NQS) have been shown to be a suitable variational ansatz to simulate out-of-equilibrium dynamics in two-dimensional systems using time-dependent variational Monte Carlo (t-VMC). In particular, stable and accurate time propagation over long time scales has been observed in the square-lattice Heisenberg model using the Restricted Boltzmann machine architecture. However, achieving similar performance in other systems has proven to be more challenging. In this article, we focus on the two-leg Heisenberg ladder driven out of equilibrium by a pulsed excitation as a benchmark system. We demonstrate that unmitigated noise is strongly amplified by the nonlinear equations of motion for the network parameters, which by itself is sufficient to cause numerical instabilities in the time-evolution. As a consequence, the achievable accuracy of the simulated dynamics is a result of the interplay between network expressiveness and regularization required to remedy these instabilities. Inspired by machine learning practice, we propose a validation-set based diagnostic tool to help determining the optimal regularization hyperparameters for t-VMC based propagation schemes. For our benchmark, we show that stable and accurate time propagation can be achieved in regimes of sufficiently regularized variational dynamics.

Contents

1	Introduction	2
2	Model and methods	3
3	Stability and regularization	5
3.1	Numerical instabilities from unmitigated noise	6
3.2	Influence of regularization	9
4	Overfitting to noise and validation error	13

5	Conclusions and outlook	18
A	Variational ansatz	19
B	Time propagation	20
C	Representability of the trajectory	21
D	Alternative regularization schemes	22
	References	23

1 Introduction

In recent times, the application of machine learning methods to problems in quantum physics has received considerable interest [1]. Examples include the use of neural networks for quantum state reconstruction [2], quantum control [3] and feedback [4], as well as classifying phases of matter [5–7]. Due to their success in approximating high-dimensional nonlinear functions in machine learning applications, neural networks were proposed in 2017 as a variational ansatz for the quantum wave function [8]. These neural-network quantum states (NQS) have been applied to a wide variety of problems in quantum many-body physics, including spin [8–21, 46, 47], bosonic [11, 22] and fermionic [23, 24] systems, as well as quantum computation [25, 26] and dissipative systems [27–30]. One particular research area where NQS could prove important in the near future are non-equilibrium quantum many-body problems, which are of interest across research fields, reaching from quantum simulators with cold atoms and trapped ions [31, 32] via arrays of Rydberg atoms [33], and photonic platforms [34] to laser-driven quantum materials [35]. The theoretical investigation of such scenarios is restricted by a lack of computational methods that allow researchers to reliably simulate driven correlated systems, in particular in two dimensions. NQS provide a particularly promising candidate wave function for the purpose of investigating out-of-equilibrium dynamics, in part due to their ability to capture high-entanglement [36] and topological states of matter [9, 10], which may serve to complement other approaches, in particular those based on tensor network states [37].

Typically, NQS are time propagated using time-dependent variational Monte Carlo (t-VMC) [8, 38, 39]. So far, this has been studied in the literature primarily in the context of quenches in the spin-1/2 Ising and Heisenberg models in both one and two dimensions [40–45]. In many cases achieving numerical stability has been identified as the key challenge for the reliable simulation of quantum dynamics [40–42] and also for ground state optimization using imaginary time propagation [19]. In contrast, the capability of the NQS ansatz to represent the relevant quantum states was not the main limiting factor. This combination of the capability of NQS to in principle represent complex quantum dynamics on the one hand and the observation of numerical instabilities on the other hand motivates an in-depth investigation of the various sources of error in t-VMC propagation and their relation to the complexity of the underlying dynamics.

In this work, we therefore take a closer look at dynamics in the antiferromagnetic 2D Heisenberg model, which has previously been studied with t-VMC on a 2D square lattice geometry. There, stable time propagation has been demonstrated using the well-established restricted Boltzmann machine (RBM) architecture [43–45]. We compare this setting to the same Hamiltonian on a two-leg ($L \times 2$) ladder geometry, which features significantly more complicated quantum dynamics and which we find to be much more sensitive to numerical instabilities. This is true already for very small systems which are still accessible by exact numerical time evolution (exact diagonalization, ED) which provides us with reliable data to benchmark the RBM dynamics.

We demonstrate that the observed instabilities arise primarily as a result of stochastic error, which is amplified through the generally ill-conditioned variational equations of motion, and highlight the strong effect of regularization parameters on the quality of the resulting trajectory. Taking inspiration from machine learning terminology, this effect can be described as overfitting to stochastic noise, which leads to poor reliability of the time stepping procedure. As in machine learning, this type of overfitting can be detected and quantified by validating the optimized time step on independently sampled data. To this end, we introduce a validation-set variational error and show that it can help identify unstable regimes and thus optimize regularization hyperparameters.

This article is structured as follows: In Section 2 we define the driven Heisenberg model used as reference in the rest of this paper, in Section 3 we show the influence of regularization on stability and accuracy of the NQS dynamics, and in Section 4 we discuss how this can be quantified using a validation-set approach. Finally, in Section 5, we conclude and provide an outlook on future work.

2 Model and methods

We study excitations in the two-dimensional antiferromagnetic (AFM) Heisenberg model (working in units with $\hbar = 1$),

$$\hat{H}_0 = J_0 \sum_{\{i,j\} \in \mathcal{N}} \sum_{\mu=1}^3 \hat{\sigma}_i^\mu \hat{\sigma}_j^\mu, \quad (1)$$

with the Pauli matrices $\hat{\sigma}_i^\mu$, $\mu \in \{1, 2, 3\}$, and exchange coupling strength $J_0 > 0$. Here, the outer sum runs over the nearest-neighbor bonds \mathcal{N} of a finite-dimensional rectangular lattice \mathcal{L} of size $N = L_x \times L_y$ [Fig. 1(a)]. We consider two different lattice geometries: the square lattice with side length $L := L_x = L_y$ and the ladder geometry with $L := L_x$, $L_y = 2$ sites. In both cases, periodic boundary conditions in x direction are assumed. For the square lattice, we also impose periodic boundary conditions in the y direction.

Starting from the ground state, we study the time evolution of the system under an excitation created by a pulsed modulation of the exchange coupling along the x direction of the lattice (which is the long direction in the ladder system). The y direction coupling is kept constant, $J_y(t) = J_0$. For J_x , we consider a pulsed modulation [Fig. 1(c)] of the form $J_x(t) = J_0(1 + \Delta_x(t))$ for $t \geq 0$, with

$$\Delta_x(t) = A_p \sin(\omega_p t) \exp\left(-\frac{(t - t_p)^2}{2\sigma_p}\right), \quad (2)$$

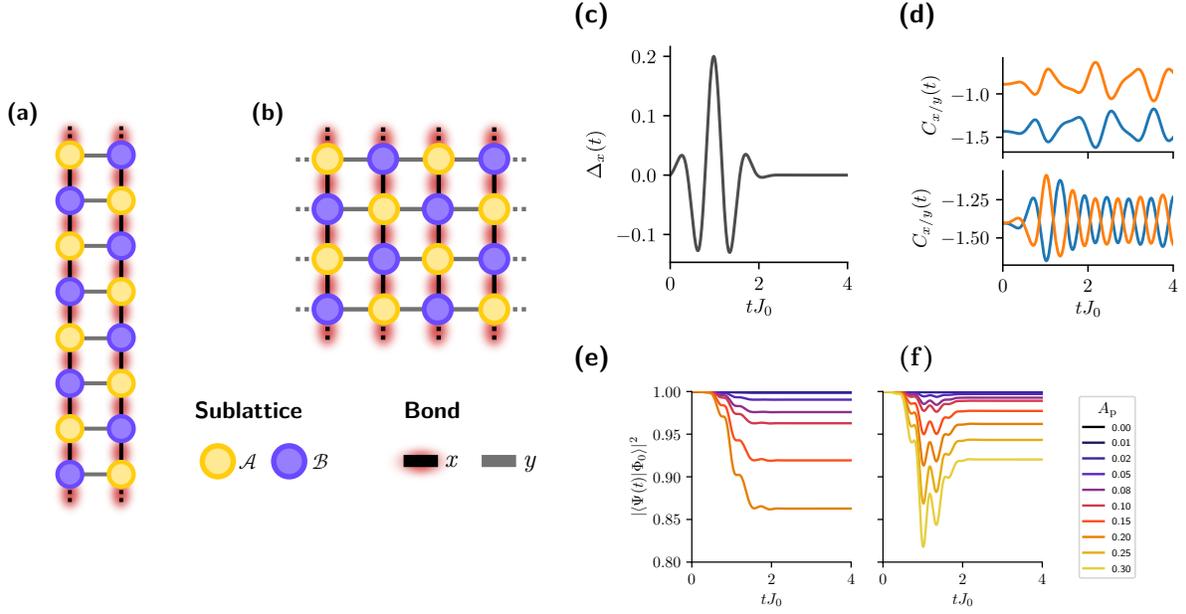


Figure 1: (a),(b) Ladder and square lattice structure for the 2D AFM Heisenberg model. The colors indicate the \mathcal{A} and \mathcal{B} sublattices of the bipartite model. In the ladder geometry, periodic boundary conditions (PBC) in the x direction are imposed. In the square geometry, PBC are imposed in both directions. (c) The Heisenberg system is driven out of its ground state by modulating the x -bond coupling in a single pulse with shape $\Delta_x(t)$ [Eq. (2)], here displayed for $A_p = 0.20$, $t_p = 0.987J_0^{-1}$, $\omega_p = 8.0J_0$, and $\sigma_p = 0.4J_0^{-2}$. (d) Oscillations of the x and y bond spin-spin correlations $C_{x/y}(t)$ [Eq. (4)] caused by the pulse in panel (a) for the 8×2 (top) and 4×4 (bottom) geometry as computed from the exact time-evolved state. (e),(f) Overlap of the initial state $|\Phi_0\rangle$ obtained from ED with the exact time-evolved state $|\Psi(t)\rangle = \hat{U}(t)|\Phi_0\rangle$ for varying pulse amplitude A_p on the 8×2 and 4×4 geometries. The other pulse parameters are the same as in panel (d).

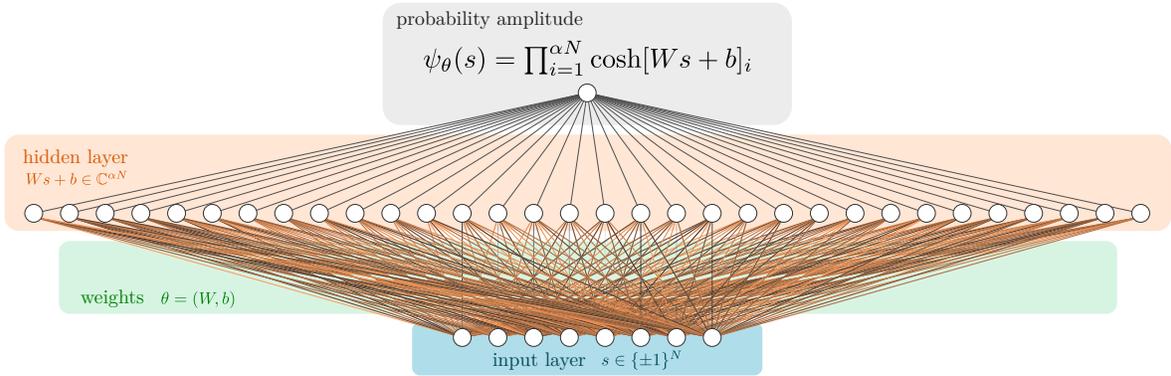


Figure 2: Restricted Boltzmann machine architecture used as a variational quantum state with N visible units corresponding to the lattice size and a hidden unit density α . A detailed description of the ansatz is given in Appendix A.

which is physically motivated and can be viewed as a single-cycle THz pulse polarized along x which drives the exchange coupling through a Raman process [43, 49]. The time-dependent Hamiltonian is then given by

$$\hat{H}(t) = \hat{H}_0 + J_0 \Delta_x(t) \sum_{\{i,j\} \in \mathcal{N}_x} \sum_{\mu=1}^3 \hat{\sigma}_i^\mu \hat{\sigma}_j^\mu \quad (3)$$

where $\mathcal{N}_{x/y} \subseteq \mathcal{N}$ denotes set of x/y bonds. In addition to the energy, we compute the average nearest-neighbor bond correlation

$$C_\nu(t) = \frac{1}{N} \sum_{\{i,j\} \in \mathcal{N}_\nu} \sum_{\mu=1}^3 \langle \hat{\sigma}_i^\mu \hat{\sigma}_j^\mu \rangle \quad (4)$$

as an observable. As reference data, we compute the time evolution under this pulse through exact (ED) propagation. The resulting dynamics are shown in [Fig. 1(d)]. In both systems, the pulse causes oscillations that persist after the pulse. However, while our driving protocol causes only singlet excitations on the square lattice, the ladder model exhibits both singlet and triplet excitations [50–52] and we indeed observe more complex and irregular dynamics in the time-dependent bond correlations (Fig. 1). For equal amplitude A_p of the x -bond modulation, the ladder system is more strongly affected [as evidenced by the higher distance to the initial state [Fig. 1(f),(g)].

As a variational ansatz, we employ the restricted Boltzmann machine (RBM) with complex-valued weights (Fig. 2). The translation symmetries of the lattice are enforced in the manner described in Refs. [8, 43], which reduce the number of variational parameters to $M = \alpha(N+1)$, where α is the hidden unit density. The translation group of an $N = L \times L$ lattice with periodic boundary conditions contains N distinct operations. In the ladder geometry, the notion of periodic boundary conditions only applies to the long direction; however, we include the reflection symmetry along the short direction, so the total order of the translation symmetries is still N . The network architecture is fully described in Appendix A.

The time propagation is done using time-dependent variational Monte Carlo (t-VMC) [8, 38], which corresponds to numerically solving the equation of motion of the time-dependent variational principle (TDVP)

$$\frac{d\theta(t)}{dt} = \dot{\theta} \quad \text{where} \quad \mathfrak{i} S(\theta(t)) \dot{\theta} = F(\theta(t), t) \quad (5)$$

using a stochastic estimate of the quantum Fisher matrix $S(\theta)$ and energy gradient $F(\theta, t)$ for the vector of variational parameters θ . Further details are provided in Appendix B. The initial ground state is prepared by minimizing the energy of a randomly initialized RBM using stochastic reconfiguration [39].

3 Stability and regularization

In this section, we will highlight jump-like numerical instabilities that we find to be caused by stochastic noise from VMC sampling that enters into the nonlinear equation of motion (5). We will then discuss the effect of regularizing this equation on the dynamics.

3.1 Numerical instabilities from unmitigated noise

We first highlight the practical challenge posed by the highly nonlinear and stochastic t-VMC equation of motion, by demonstrating how a change in lattice geometry of an otherwise unaltered physical model can affect the stability of the NQS propagation.

In previous works [43–45], it has been shown that the dynamics of the Heisenberg model on a square lattice can indeed be successfully simulated using t-VMC with RBM quantum states. We obtain equivalent results for our pulsed driving [Fig. 3(d)–(f)]. The main manifestation of the error as compared to the exact dynamics is a continuous decay of amplitude of the resulting oscillations, which is visible from the averaged nearest-neighbor correlation $C_x(t)$. Increasing the width of the network, i.e., the hidden unit density α , both improves the accuracy of the initial (ground) state and reduces the loss of accuracy over the course of the time evolution (data not shown). This shows that the decay is the result of accumulated TDVP error and can be reduced by an increase in network size, which is in agreement with results for the square lattice model presented in Ref. [43]. However, this behavior is markedly different for the ladder geometry (with otherwise unchanged system parameters), where instabilities quickly occur during t-VMC evolution already for weak pulse strengths $A_p \geq 0.02$ [Fig. 3(a),(b)]. Notably, the observed instabilities violate energy conservation, a property that is inherent to the TDVP equation of motion for a static Hamiltonian. Therefore, their origin has to be numerical. In order to better understand which sources of error in the t-VMC method contribute to these instabilities, we compare three different propagation schemes:

1. t-VMC propagation where the components of the equation of motion are estimated stochastically using Markov chain Monte Carlo (MCMC) with the Metropolis algorithm, which is the standard propagation method for NQS [8, 39];
2. autocorrelation-free “exact” Monte Carlo (EMC), where samples are directly drawn from the Born distribution $|\psi(\cdot)|^2$ without using the Metropolis algorithm; and
3. time propagation based on full summation of the t-VMC equation of motion over all spin configurations, which provides a reference free of stochastic noise in the estimate of energy gradient and QFM.

In the Metropolis MCMC scheme, updates to the spin configuration are proposed based on exchanges of spin pairs which preserve the total magnetization and thus the restriction of the ansatz to the zero-magnetization sector. The EMC scheme can only be applied to small systems accessible to ED, as this relies on knowledge of the full Born distribution. Here, it is used strictly as a benchmark to uncover the influence of noise on the dynamics while ruling out errors due to non-convergence of the Metropolis sampling¹. The full summation scheme is similarly limited to small systems with tractable Hilbert space size.

Figure 3 shows results for these different propagation schemes. In both cases, we have used a second-order Runge-Kutta method for time propagation (Heun’s method), which uses two evaluations of the equation of motion (5) per time step, a fixed step size of $\delta t = 0.002$ and $N_s = 7000$ Monte Carlo samples for EMC and t-VMC. The stability does not increase significantly with further reduced integrator step size. While there is a visibly increased level of noise with the MCMC sampling, divergences occur at similar time points of the evolution

¹Note that autocorrelation-free Monte Carlo sampling is practically possible beyond the ED regime for NQS architectures based on autoregressive networks [53], though this is quite different from the benchmark implementation considered here.

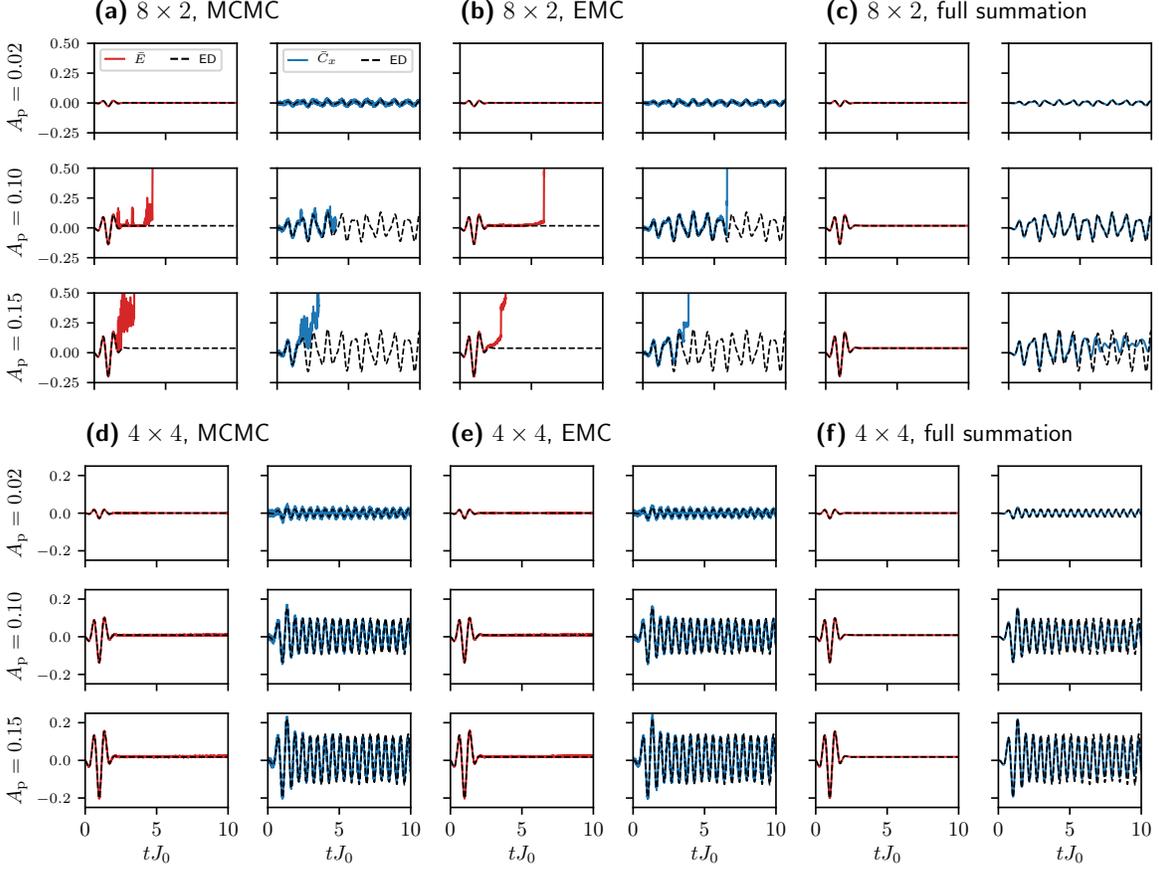


Figure 3: Time-dependent per-site energy change $\bar{E}(t) = [\langle \hat{H}(t) \rangle - E(0)]/N$ and average x -bond correlation $\bar{C}_x(t) = C_x(t) - C_x(0)$ for the 8×2 ladder [panels (a)–(c)] and 4×4 square geometry [panels (d)–(f)] and varying pulse strengths A_p . The trajectories have been obtained from t-VMC evolution using MCMC [panels (a),(d)] and EMC [panels (b),(e)] sampling as well as results based on full summation of the equations of motion [panels (c),(f)]; see Sect. 3.1 for details. The dashed lines show ED results for reference. In all cases using a symmetric RBM with hidden unit density of $\alpha = 10$ has been used. The initial state is the approximate ground state of the respective system obtained by stochastic reconfiguration and is the same for panels (a)–(c) and (d)–(f), respectively. In all cases, the equation of motion is evaluated by singular-value decomposition of S and applying a diagonal shift of $\epsilon = 10^{-3}$ (see Sect. 3.2 and Appendix D) for regularization.

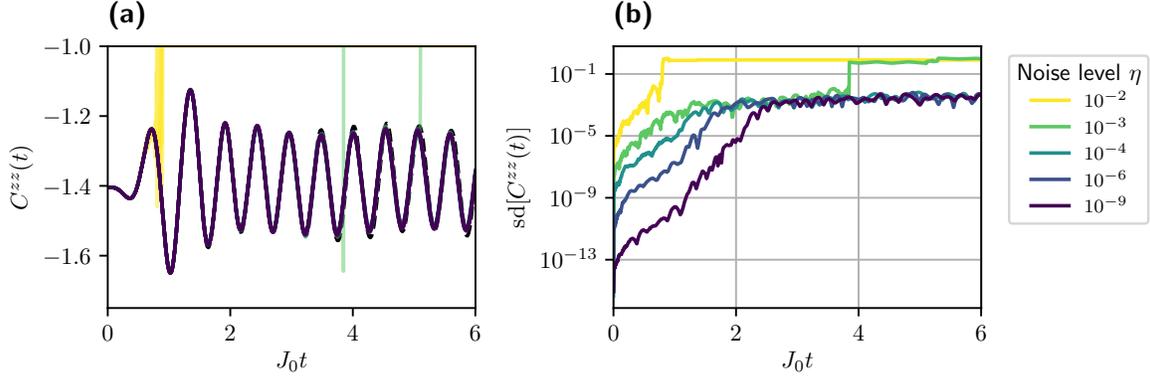


Figure 4: **(a)** Time evolution of the nearest-neighbor coupling term $C^{zz}(t) = \langle \sigma_0^z \sigma_1^z \rangle$ using the TDVP equation of motion evaluated by full summation with added artificial noise [Eq. 6]. We show results for the dynamics on the 4×4 lattice with pulse strength $A_p = 0.20$ for an RBM with hidden unit density $\alpha = 10$. For each value of the noise strength η , five independent trajectories are shown as faint lines. The opaque lines indicate the median of the respective curves. **(b)** Empirical standard deviation between the set of curves in panel (a) grouped by noise strength η .

for both approaches in the ladder system. In the full summation runs, the energy jumps are absent and the dynamics are more accurately reproduced on the ladder with pulses $A_p \leq 0.10$ [Fig. 3(c),(f)]. We note that even in the absence of stochastic noise, the time evolution shown here fails to accurately capture the ladder dynamics for stronger excitations, as can be seen from the $A_p = 0.15$ trajectory. It is clear, however, the occurrence of jump-type instabilities is significantly more likely in the presence of Monte Carlo noise (which affects both EMC and MCMC sampling).

The results shown in this section are for networks with a hidden unit density of $\alpha = 10$. We have observed RBMs with fewer parameters to be generally more stable. This, however, comes at the cost of decreased accuracy. The numerical instability present in larger network thus counteracts the benefits of increased expressiveness, making it particularly important to find ways of alleviating this effect.

Indeed, Monte Carlo noise in the energy gradient amplified through the equation of motion alone is sufficient to cause this type of instabilities. We demonstrate this in an idealized picture as follows: Consider Eq. (5) without any stochastic sampling of QFM and gradient but with an artificial proportional Gaussian white noise term that is added to the energy gradient². This provides a direct way to control the noise without affecting other steps in the propagation. Specifically, we solve

$$S\dot{\theta} = -i(F + \xi) \quad (6)$$

where ξ is a random vector with components drawn from a complex normal distribution

²We note that the proportional Gaussian noise model of Eq. (6) is indeed a simplification. An analysis in Ref. [42] shows that the actual noise level varies between components of the energy gradient and depends on the quantum geometry of the ansatz (through the QFM spectrum) as well as the energy fluctuations. However, already the simple proportional model used here does exhibit the jump-like instabilities when subjected to noise amplified through the equation of motion.

with mean $\mathbb{E}[\xi_i] = 0$ and variance $\text{Var}[\xi_i] = |\eta F_i|^2$. The parameter η determines the relative noise strength and the standard error in each component is proportional to $|F_i|$. In Fig. 4 we show resulting trajectories for varying noise levels η . For each value of η , five independent trajectories are shown. In the absence of instabilities, the standard deviation of the trajectories at first grows with increasing noise. After a time of the order of the pulse length, the spread of the trajectories stabilizes at a value that is independent of the noise strength. For $\eta \geq 10^{-3}$, jump instabilities occur within the simulation time. Whereas for $\eta = 10^{-2}$ all trajectories show this instability already around $J_0 t = 1.0$, the jumps happen more sporadically and at later times for $\eta = 10^{-3}$, with only two of the trajectories exhibiting a jump before $J_0 t = 6.0$. This idealized experiment shows that random noise in the gradient can be amplified through Eq. 5, and this is indeed sufficient to cause jump-type instabilities like we have observed in the t-VMC propagation, instead of a gradually increasing spread of the trajectories that would be expected for a more regular equation of motion.

In agreement with observations made for other systems [42], we find that the expressive capabilities of the network are not a limiting factor. In fact, for individual time points, the exact quantum state shown in Fig. 3 can already be represented to good accuracy using very small network ($\alpha = 2$). Increasing the hidden unit density to the level used in the propagation further improves this accuracy. See Appendix C for detailed results.

While the ladder system is particularly sensitive to the types of instabilities discussed here, we found that by increasing the pulse strength beyond $A_p = 0.30$ also leads to divergences in the square lattice system (data not shown). Therefore, while the Heisenberg ladder is a very suitable benchmark system for these types of numerical issues, the observations made here can be expected to apply to a broader class of systems, especially when they are driven far away from their ground state on short time scales.

In the next section, we will discuss how the choice of regularization scheme affects the stability of the dynamics.

3.2 Influence of regularization

The formal solution of the TDVP equation (5) is given by

$$\dot{\theta} = -iS^+ F(t) \quad (7)$$

where S^+ denotes the Moore-Penrose pseudoinverse of the QFM [8, 40, 42, 44]. Computing the pseudoinverse can be done by singular-value decomposition (SVD). The SVD is equivalent to the eigendecomposition in this case because S is a covariance matrix and therefore positive semi-definite, i.e., all eigenvalues are nonnegative. Thus, in the eigenbasis of $S = V \text{diag}(\{\sigma_j\}_{j=1}^M) V^\dagger$, the TDVP equation reduces to

$$\sigma_j [V^\dagger \dot{\theta}]_j = -i [V^\dagger F]_j. \quad (8)$$

We order the eigenvalues of S by magnitude $\sigma_1 \geq \sigma_2 \dots \geq \sigma_M$ in the following and denote the smallest nonzero eigenvalue by σ_r . Then, for all nonzero eigenvalues corresponding to $j \leq r$, we have $[V^\dagger \dot{\theta}]_j = -i \sigma_j^{-1} [V^\dagger F]_j$. The directions in the null-space of S do not contribute to the physical dynamics; changes of θ in those directions only affect gauge degrees of freedom of the quantum state. In order to obtain the minimum-norm solution, these are set to zero, i.e., $[V^\dagger \dot{\theta}]_j = 0$ for $j > r$.

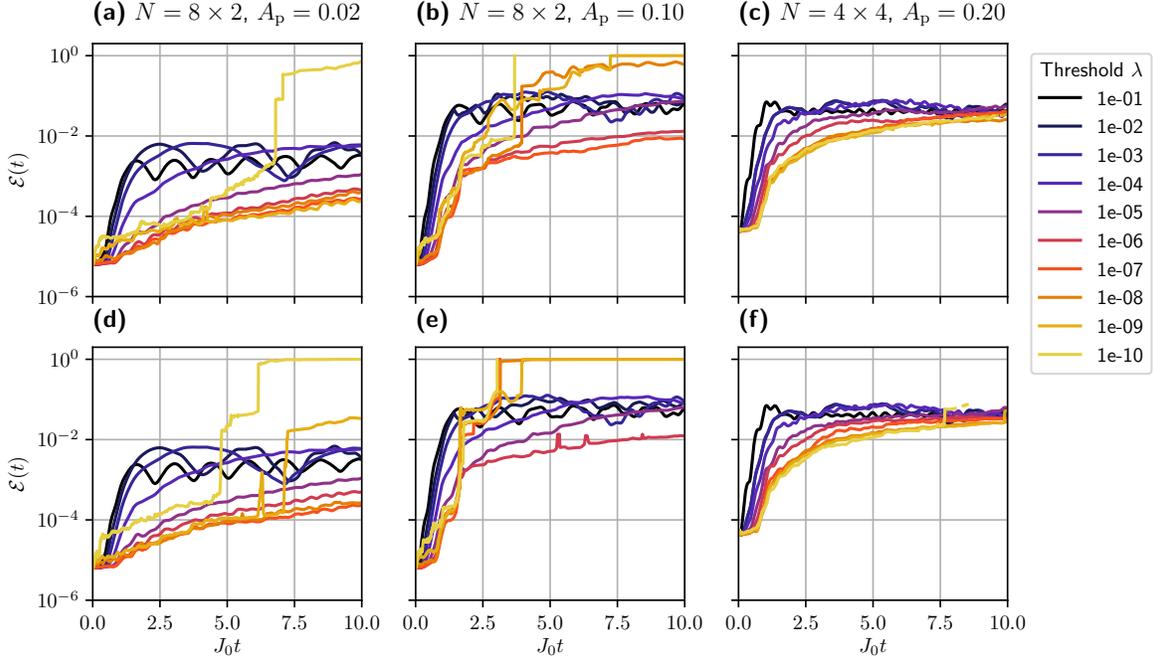


Figure 5: Inverse fidelity of the time-evolved variational state compared to the exact trajectory for varying regularization (in the form of the SVD threshold λ). The trajectories have been computed using t-VMC with EMC [panels (a)–(c)] and Metropolis [panels (d)–(f)] sampling. The columns correspond to a weak excitation $A_p = 0.02$ [panels (a),(d)] and a moderate excitation $A_p = 0.10$ [panels (b),(e)] in the 8×2 ladder, as well as a stronger excitation $A_p = 0.20$ in the 4×4 square lattice [panels (c),(f)]. The initial $\mathcal{E}(t=0)$ of order 10^{-5} is the approximation error of the variational ground state. The time propagation has been computed using t-VMC with EMC sampling with $N_s = 24000$ samples for the ladder [panels (a),(b)] and $N_s = 11200$ for the square lattice [panel (c)]. We have used the symmetrized RBM ansatz with a hidden unit density of $\alpha = 10$.

In practice, the numerical solution of this equation is complicated by the fact that NQS typically possess a non-exponential but still large number of variational parameters compared to more traditional variational wave functions and further allow for redundancy in the parametrization of a specific quantum state. As a consequence, the QFM is singular, and the nonzero part of its spectrum typically spans many orders of magnitude [41, 42, 54]. Therefore, the linear system (5) has a high condition number $\kappa(S) = \sigma_1/\sigma_r$ which, in particular, means that small perturbations in the right-hand side $F(t)$ can be strongly amplified in the solution $\dot{\theta}$ of the EOM, causing the jump instabilities we have empirically observed above. For this reason, it is necessary to regularize the EOM in order to stabilize the dynamics while preserving the physical accuracy of the resulting trajectory. Typically, this is done by truncating eigenvalues below a threshold λ in Eq. (7). Specifically, σ_i is treated as zero when $\sigma_i \leq \lambda\sigma_1$. The effective condition number of S is then bounded by $\kappa \leq \sigma_1/\lambda$. While we focus on this regularization scheme in the remainder, our analysis also applies to a broader class of regularization schemes. In particular, the application of a diagonal shift to S and a signal-to-noise-based regularization scheme proposed in Ref. [42] are briefly discussed in Appendix D.

Naturally, there is a trade-off between stability and accuracy: Too much regularization will suppress crucial parts of the physical dynamics, while the system is susceptible to instabilities without or with only weak regularization. This can be seen in Fig. 5, where we show the inverse fidelity of the time-dependent quantum state relative to the ED time evolution of the system,

$$\mathcal{E}(t) = 1 - \frac{|\langle \Psi_{\text{ED}}(t) | \psi_{\theta}(t) \rangle|^2}{\langle \Psi_{\text{ED}}(t) | \Psi_{\text{ED}} \rangle \langle \psi_{\theta}(t) | \psi_{\theta}(t) \rangle}, \quad (9)$$

for varying regularization strength λ . Specifically for the weak pulse [Fig. 5(a)], we can clearly see a separation of three regimes: an over-regularized regime (for $\lambda \geq 10^{-4}$), where the dynamics are stable but inaccurate; an intermediate stable regime where the accuracy of the observables is good and the regularization still sufficient to stabilize the dynamics; and an unstable regime ($\lambda \leq 10^{-8}$) where jump instabilities occur within the simulation time frame. While for the weak pulse the stable regime spans several orders of magnitude, it becomes smaller with increasing strength [Fig. 5(b)]. By contrast, on the square lattice the dynamics remain stable and largely unaffected by the regularization strength over a wide range of λ even at a pulse strength of $A_p = 0.20$ [Fig. 5(c)]. These have been obtained with EMC sampling, but the same behavior can be observed for the practically relevant case of Metropolis sampling, which is shown in Fig. 5[(d)–(f)]. In this case, the stable regime of regularization becomes smaller for the ladder system. Figure 6 shows the expectation values of energy and bond correlations for several trajectories. These observables show how in the over-regularized regime, numerical stability comes at the cost of physical errors which manifest here in an incorrect reproduction of the oscillation frequencies. At the same time, the square lattice system is almost unaffected by Metropolis sampling, except at very low thresholds. Notably, for both EMC and MCMC sampling strategies, the dynamics converge to a stable trajectory at a much lower number of Monte Carlo samples than in the the ladder system. This hints at an increased sampling complexity of the low-lying ladder excitations in accordance with the higher physical complexity of the ladder excitations, an observation which is corroborated by our results in the next section.

Altogether, our results highlight the delicate balance between stability and accuracy of the dynamics in the presence of stochastic noise and the resulting necessity to fine-tune regularization hyperparameters to reach the optimal regime. From the comparison of ladder and

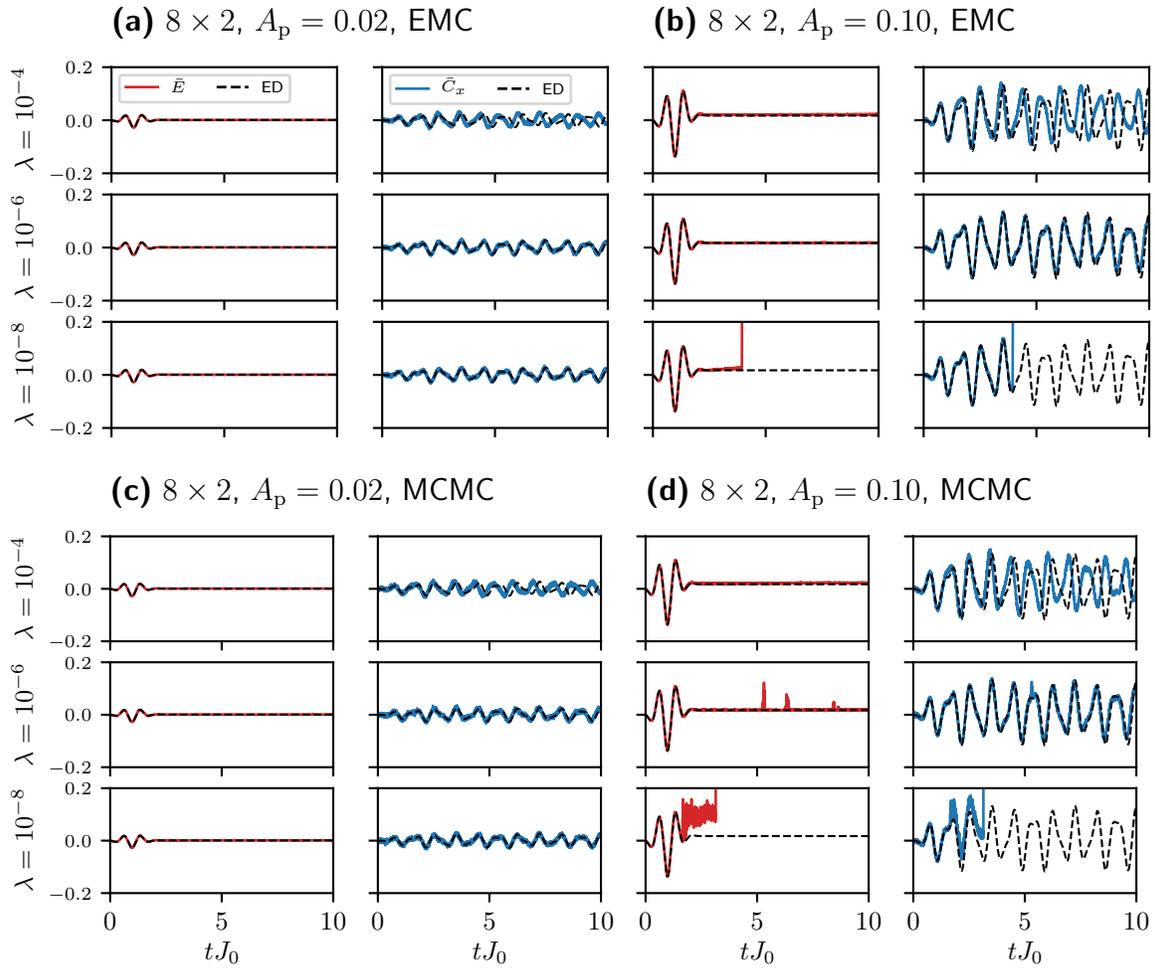


Figure 6: Relative change in energy and x -bond correlation as defined in Fig. 3 for a selection of trajectories shown in Fig. 5.

square geometry, we have further seen that the extent of this behavior depends strongly on the details of the system.

4 Overfitting to noise and validation error

In order to choose an optimal regularization for a given system and excitation scheme, it is important to have access to appropriate diagnostics. While the error relative to ED as shown in the previous section provides a straightforward way to assess the quality of the solution, this option is restricted to small benchmark systems. Here, we therefore propose an alternative diagnostic which is more generally applicable.

The local truncation error resulting from a single time step in the variational approximation is quantified by the TDVP error [8, 42]

$$r^2(t) = \left[\frac{D(\psi[\theta(t) + \dot{\theta} \delta t], \hat{U}_{t+\delta t, t} \psi[\theta(t)])}{D(\psi[\theta(t)], \hat{U}_{t+\delta t, t} \psi[\theta(t)])} \right]^2. \quad (10)$$

Here, $D(\cdot, \cdot)$ denotes the Fubini-Study distance and $\hat{U}_{t', t}$ is the unitary time evolution operator from t to t' . The equation of motion (5) can be derived by locally minimizing the numerator of Eq. (10). The denominator provides a rescaling of the error to account for the varying exact distance between points along the trajectory. This quantity can be estimated to second order in δt as [42]

$$r^2(\dot{\theta}; S, F, \delta E) = 1 + \frac{\dot{\theta}^\dagger (S \dot{\theta} + iF) - iF^\dagger \dot{\theta}}{(\delta E)^2} \quad (11)$$

where $(\delta E)^2 = \text{Var}[\hat{H}(t)]$ and the other quantities are defined as in Eq. (5).

While capturing loss of accuracy due to the variational approximation, the TDVP error does not account for effects caused by the stochastic noise affecting the equation of motion and thus its solution. In order to account for this additional source of errors, we take inspiration from a standard practice of machine learning: the use of a so called validation error to detect failure to generalize beyond the training data caused by overfitting to a specific sample [55]. Adapted to our present purpose, we consider the specific realization of spin configurations used to estimate the EOM as the training set. Solely optimizing the parameter update for this realization bears the risk of overfitting, in which case the solution may be optimal only on the training set but performs badly to independent estimates of the same EOM. In order to detect this, we can compute two updates $\dot{\theta}^{(i)}$, $i = 1, 2$, from independently drawn samples (separately estimating $S^{(i)}$, $F^{(i)}$ for both). While the resulting $\dot{\theta}^{(i)}$ and corresponding error estimates

$$r_{\text{tr}, i}^2 = r^2(\dot{\theta}^{(i)}; S^{(i)}, F^{(i)}, \delta E^{(i)}) \quad (12)$$

are identically distributed, the error of the update $\dot{\theta}^{(2)}$ with respect to the independently estimated equation of motion $S^{(1)} \dot{\theta} = -iF^{(1)}$ can be used to quantify the generalization properties of the parameter derivative. This procedure is illustrated in Fig. 7. Specifically, we define the validation TDVP error as

$$r_{\text{val}}^2 = r^2(\dot{\theta}^{(2)}; S^{(1)}, F^{(1)}, \delta E^{(1)}). \quad (13)$$

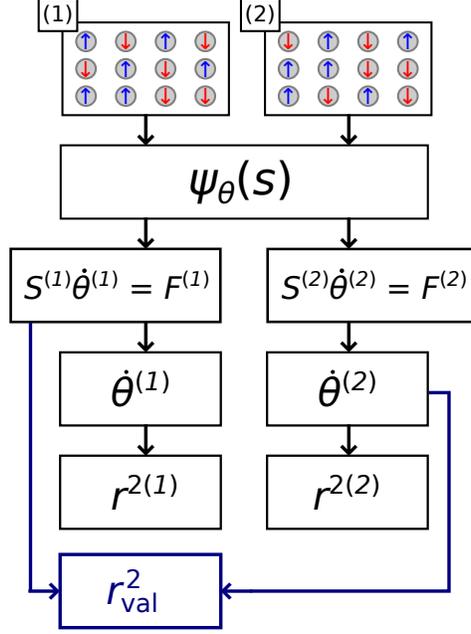


Figure 7: Illustration of the process for computing the validation TDVP error as described in Sect. 4. Two sets of spin configurations are independently generated via Monte Carlo sampling and used to generate two independent derivatives through solving the equation of motion. The validation error (13) is computed as the error of the second update with respect to the first equation of motion.

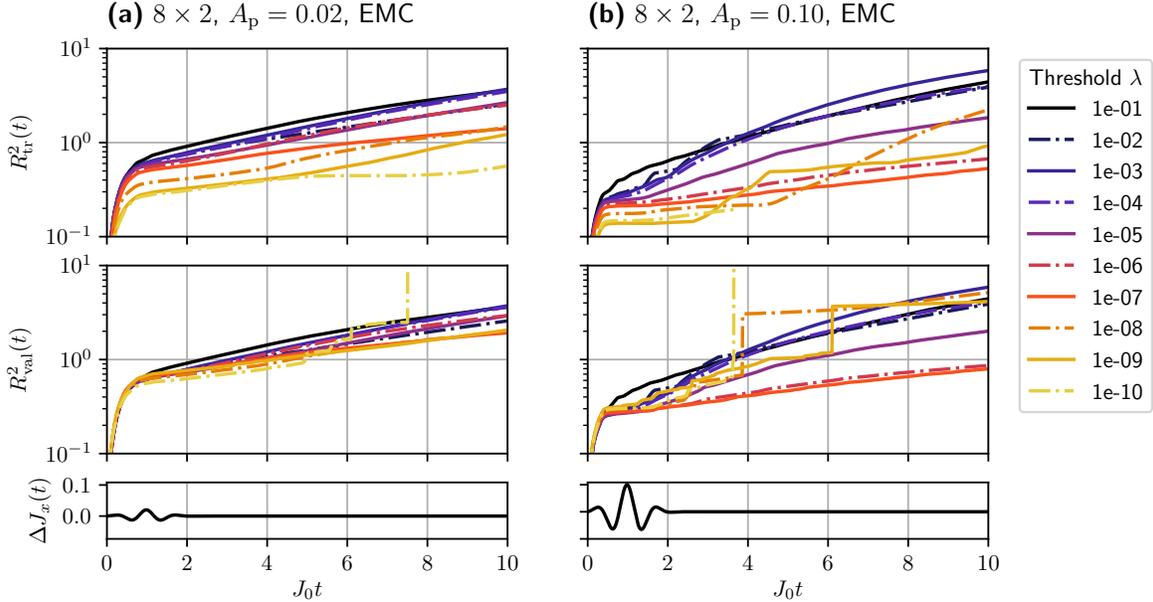


Figure 8: Integrated bare and validation TDVP error (14) over time with varying regularization (SVD threshold λ) for the same trajectories as shown in Fig. 5(a),(b), i.e., using EMC sampling and with pulse strengths of (a) $A_p = 0.02$ and (b) $A_p = 0.10$.

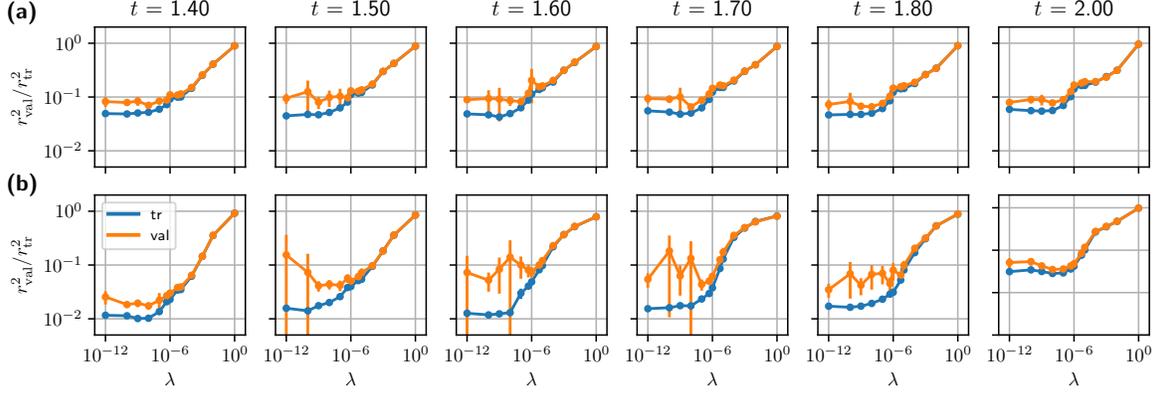


Figure 9: Comparison of the TDVP error r_{tr}^2 and the validation error r_{val}^2 on a logarithmic scale for varying regularization strength in the form of the SVD threshold λ at different points in time for **(a)** a weak $A_p = 0.02$ and **(b)** stronger driving amplitude $A_p = 0.10$ on the 8×2 ladder system. The TDVP error has been computed here by taking variational states $|\psi_{\theta(t)}\rangle$ from the stable $\lambda = 10^{-6}$ trajectory [Fig. 5(d),(e)] and then performing a single step $\delta t = 0.002$ at each displayed time and for each λ using Metropolis sampling with $N_s = 28 \cdot 10^3$ samples. We show here the average error over five independent realizations of the validation error, with error bars indicating the standard deviation, in order to account for variance in the error estimate itself.

Crucially, r_{val}^2 can be estimated using only quantities that are accessible as part of the t-VMC computation. This makes it feasible to use the validation error as a diagnostic for the degree of overfitting and thus reliability of the TDVP solution in systems where a comparison to ED data is no longer possible. If the solution of the EOM is deterministic, we have $\dot{\theta}^{(1)} = \dot{\theta}^{(2)}$ and consequently $r_{\text{val}}^2 = r_{\text{tr},i}^2$. Otherwise, the validation error will be larger, indicating the amount of “overfitting” of the update $\dot{\theta}^{(1)}$ to noise present in the sample. We note that the error estimates are themselves affected by noise in both EOM and energy variance. This is alleviated by considering the integral

$$R_{\text{tr/val}}^2(t) = \int_0^t r_{\text{tr/val}}^2(t') dt' \quad (14)$$

or, if the local quantity is needed, by averaging over additional realizations of r_{val}^2 . Despite its ad-hoc nature, this definition of a validation error does provide a useful way of quantifying how the regularization scheme affects the solution of the EOM in the presence of noise. Figure 8 shows the integrated TDVP and validation error for weak and moderate driving. While the bare TDVP error is insensitive to the Monte Carlo error and corresponding instabilities³, a clearly discernible effect is present in the validation error which therefore shows a much better qualitative agreement with the reference ED error (compare Fig. 5). Note that we show here the error for EMC sampling because, while the utility of the (local) validation error is not limited to this case, the integrated curves are strongly affected by local perturbances present in the MCMC data.

³While R_{tr}^2 does increase for small λ eventually in Fig. 8, this only happens after the jumps in the corresponding trajectories have already occurred, in contrast to R_{val}^2 which detects hints of the instability already before that point.

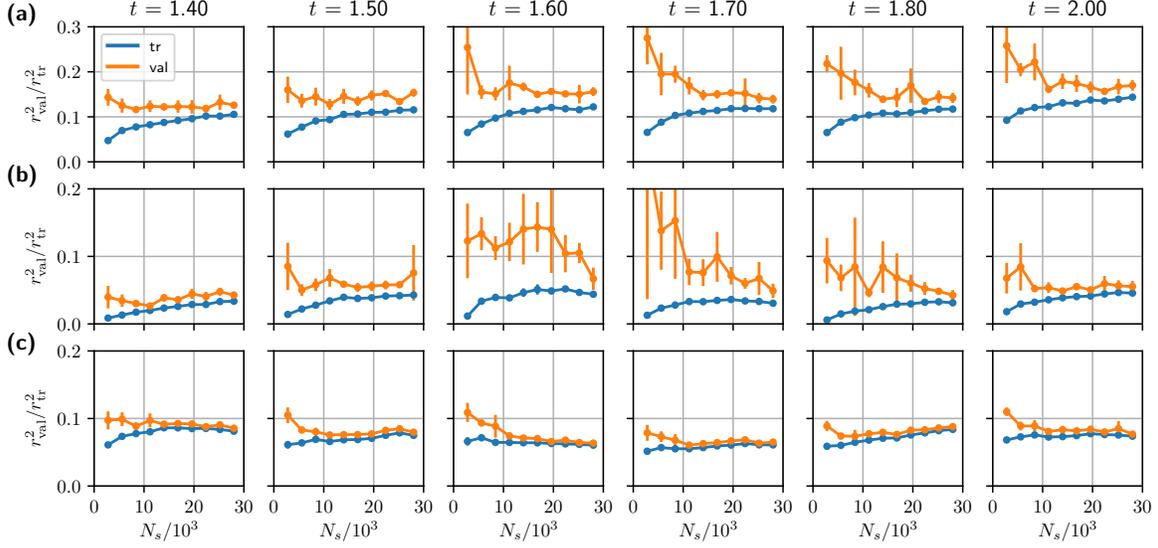


Figure 10: **(a),(b)** Comparison of TDVP and validation error for varying number of Monte Carlo samples on a linear scale for the same trajectory and parameters as in Fig. 9, with $\lambda = 10^{-6}$ and at driving amplitudes $A_p = 0.02$ [panel (a)] and $A_p = 0.10$ [panel (b)] in the 8×2 ladder system. **(c)** The same data for an $A_p = 0.20$ pulse in the 4×4 square geometry. In all panels, the error bars indicate the standard deviation over five independent realizations of the validation error as in Fig. 9.

Figure 9 shows the local TDVP and validation error over a range of thresholds at various times during the duration of the pulse. Here we can see that the unstable regimes of regularization indeed correspond to an increased validation error r_{val}^2 compared to r_{tr}^2 and can thus be interpreted as a consequence of overfitting while in the stable regions almost no overfitting error is observed, indicating a high degree of consistency between updates. Furthermore, this behavior is not uniform over time: For $A_p = 0.10$, overfitting occurs particularly strongly at the waning edge of the pulse around $t = 1.60$. The degree of overfitting and its sensitivity to the regularization strength is significantly lower for the weaker excitation. Note that the absolute magnitude of the TDVP error is not directly comparable between different pulse strengths. This is because the denominator of Eq. (10) depends on the distance between $|\psi(t)\rangle$ and $|\psi(t + \delta t)\rangle$ which decreases with decreasing driving strength and goes to zero for vanishing dynamics. Therefore, r^2 measures the error relative to the magnitude of the physical dynamics which explains the observed higher values of both $r_{\text{tr/val}}^2$ for the weaker excitation strengths. Data for the validation error relative to the baseline TDVP error can be found in Fig. 11.

While sensitive to the regularization, the Monte Carlo error in the update $\dot{\theta}$ is of course also dependent on the number of samples used in the estimate of the equation of motion. The fewer samples are used, the higher the generalization error with respect to the full Hilbert space will be. As with the regularization, this behavior is strongly system and excitation dependent. For the square lattice geometry, convergence with respect to the sample size occurs quickly compared to the ladder system, where a much higher number of samples is needed to obtain reliable estimates. This indicates a higher sampling complexity of the ladder

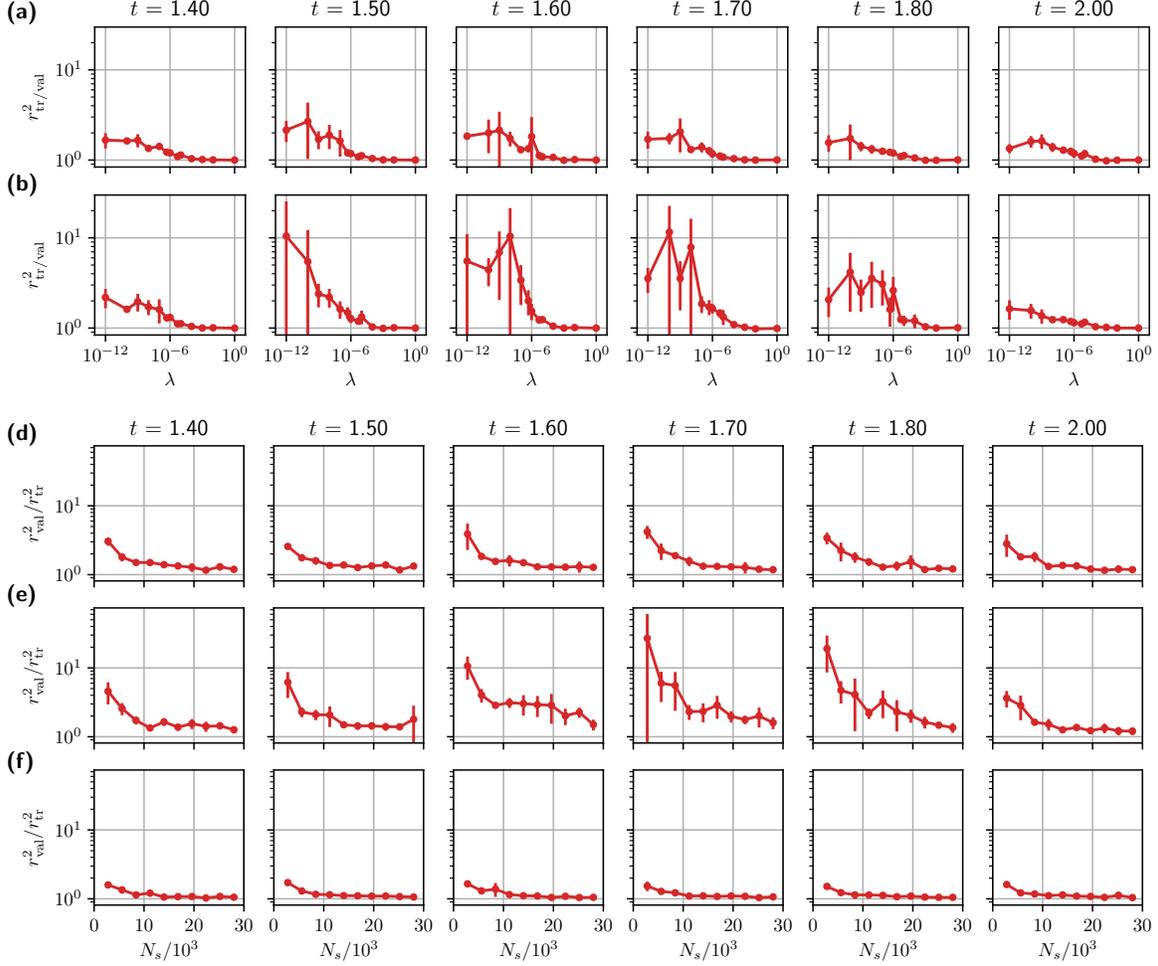


Figure 11: Validation error relative to the TDVP error for the data shown in Fig. 9 [panels (a),(b)] and Fig. 10 [panels (d)–(f)].

states compared to the well-behaved singlet magnon excitation in the square lattice, which is indeed captured by the validation error. This is demonstrated in Figs. 10 and 11[(d)–(f)], which shows the TDVP and validation error as a function of sample size for both geometries and different driving strengths. We see a clear overfitting behavior which is especially strong around the waning edge of the pulse, as has been similarly observed in the regularization dependence, and is only suppressed by increasing the number of samples up to $\sim 28 \cdot 10^3$ for the ladder. In contrast, convergence of the validation error in the square lattice system occurs much faster, even for the stronger excitation shown here.

In summary, we have demonstrated that the sensitivity of the ladder system to the regularization scheme as well as the need for a high number of Monte Carlo samples to accurately estimate dynamics especially around the end of the THz pulse is indeed captured by the proposed TDVP validation error. These effects can thus be seen as a consequence of a lack of generalization of the derivative estimate $\dot{\theta}$ or overfitting to an insufficiently representative sample of spin configurations.

5 Conclusions and outlook

We have presented the time propagation of the Heisenberg model on the two-leg ladder as a key benchmark for neural-network-based methods to simulate quantum many-body dynamics. In line with other studies, we have found that RBM quantum states are in principle capable of representing the relevant quantum states during the simulated time evolutions. However, the combination of (i) numerical instabilities already in small systems and (ii) tunability between relatively well-behaved dynamics on the square lattice and the much more challenging dynamics on the ladder make this model system a suitable case study for t-NQS. Moreover, larger-scale ladders can also be simulated with tensor network states, which makes Heisenberg ladders an ideal *drosophila* for more detailed comparisons between different systematically improvable variational ansätze and propagation schemes beyond system sizes accessible to exact diagonalization.

We have shed light on the delicate balance between stabilizing regularization and physical accuracy of the variational time evolution in the presence of stochastic noise inherent in the t-VMC approach. In particular, motivated by the interpretation of these instabilities as a consequence of overfitting to Monte Carlo noise in the equation of motion, we have introduced a validation-set approach as a quantitative diagnostic of the noise-based error. We have demonstrated that this validation error can be used to aid in the optimization of relevant hyperparameters. While this is particularly relevant for NQS dynamics, the validation-set approach can be applied to t-VMC simulations using other variational states as well as ground state optimization based on imaginary-time propagation.

The ability of quantifying the generalization error in t-VMC propagation also opens up the possibility of devising an adaptive scheme to control regularization hyperparameters and Monte Carlo sampling in order to achieve stable dynamics without the need for manual fine tuning. This is particularly relevant for general NQS software frameworks such as NETKET which strive to be usable in a wide range of physical settings.

Acknowledgments

Simulations have been performed using a t-VMC implementation based on NETKET 2.1b1 [56] as well as code based on ULTRAFast [43]. The exact Schrödinger dynamics have been computed using QUTIP 4.5.0 [57]. We acknowledge support from Flatiron Institute, a division of the Simons Foundation.

Author contributions

D.H. implemented the t-VMC simulation code based on NETKET, performed simulations and data analysis, wrote the initial manuscript, with contributions from M.A.S., and created the figures. G.F. contributed t-VMC and full summation results using an independent implementation (ULTRAFast). J.H.M., G.C., and M.A.S. conceived the project. All authors contributed to discussions throughout the project and the preparation of the final manuscript.

Funding information

G.F. and J.H.M. acknowledge funding from the Shell-NWO/FOM- initiative ‘‘Computational sciences for energy research’’ of Shell and Chemical Sciences, Earth and Life Sciences, Physical Sciences, FOM and STW. M.A.S. acknowledges financial support by Deutsche Forschungsgemeinschaft (German Research Foundation, DFG) under the Emmy Noether program (SE 2558/2).

A Variational ansatz

In our simulations, we employ the translation-invariant RBM ansatz as introduced in Ref. [8]. Explicitly,

$$\ln \psi_\theta(s) = \sum_{j=1}^{N_h} \ln \cosh [\tilde{W}s + \tilde{b}]_j. \quad (15)$$

Here, the full weight matrix $\tilde{W} \in \mathbb{C}^{N_h \times N}$ and hidden bias $\tilde{b} \in \mathbb{C}^{N_h}$ are defined in terms of a smaller number of independent parameters $W \in \mathbb{C}^{\alpha \times N}$ and $b \in \mathbb{C}^\alpha$, ensuring that $\psi(\tau(s)) = \psi(s)$ for all N lattice translations τ . We refer to the independent parameters collectively as $\theta = (W, b) \in \mathbb{C}^M$. This ansatz reduces the number of variational parameters to $M = \alpha(N+1)$ where $\alpha = N_h/N$ is the hidden unit density. Thus, the dimension of the parameter space grows only linearly in the system size for the symmetric RBM ansatz. Note that RBM wave functions can include another term, the visible bias $\tilde{a} \in \mathbb{C}^N$, as $\psi_{a,\theta}(s) = e^{\tilde{a}^\top s} \psi_\theta(s)$. When enforcing translation invariance in the manner described above, only one component of the visible bias $\tilde{a}_i = a \in \mathbb{C}$ remains independent which, in the zero magnetization sector, only affects the global gauge of the quantum state. Therefore, it is not included in our variational ansatz.

In addition to enforcing translation symmetry, we restrict the state space of the model to the zero magnetization subspace of the full Hilbert space. Thus, the ansatz wave function is only evaluated for spin configurations satisfying $\sum_j s_j = 0$ and $\psi_\theta(s) = 0$ is assumed otherwise. For N sites, the dimension of the full Hilbert space is 2^N , the zero-magnetization subspace has dimension $\binom{N}{N/2}$. Note that the driving (3) is compatible with both the translation symmetry and zero magnetization constraints. Furthermore, all of our calculations were performed in a computational basis taking into account the sign structure of the AFM ground state, which is a standard approach for the Heisenberg model [58–60] and helps circumvent the difficulty of learning states with a nontrivial sign structure, which is a more challenging task for NQS [17, 18]. Specifically, this is done as follows: Let $\{|s\rangle \mid s \in \{\pm 1\}^N\}$ denote the $\hat{\sigma}^z$ eigenbasis, so that $\hat{\sigma}_i^z |s\rangle = s_i |s\rangle$. In the AFM phase, the Heisenberg model has a nondegenerate ground state

$$|\Phi_0\rangle = \sum_{s \in \{\pm 1\}^N} \Phi_0(s) |s\rangle \quad (16)$$

which is part of the zero eigenspace of the magnetization $\hat{M}^z = \sum_{i \in \mathcal{Q}} \hat{\sigma}_i^z$. On a bipartite lattice where the sites are partitioned into disjoint subsets \mathcal{A} and \mathcal{B} [compare Fig. 1(a),(b)],

the ground state coefficients have the form

$$\Phi_0(s) = (-1)^{\varpi(s)} A_0(s) \quad (17)$$

where $A_0(s) \in \mathbb{R}_{>0}$ is real and nonnegative. The parity $\varpi(s) = \sum_{i \in \mathcal{A}} s_i$ is determined by the magnetization on the \mathcal{A} sublattice. This property is known as Marshall's sign rule [61] and makes it possible to represent the ground state by a real nonnegative wave function in the computational basis $|s^c\rangle = \prod_{i \in \mathcal{A}} \hat{\sigma}_i^z |s\rangle$, which significantly improves convergence of the NQS ground state optimization as the network only needs to learn a trivial sign structure.

B Time propagation

In t-VMC [38, 39], the time propagation of the variational ansatz is based on the time-dependent variational principle (TDVP). The equation of motion for the vector of variational parameters $\theta \in \mathbb{C}^M$ is given by $d\theta(t)/dt = \hat{\theta}$, where $\hat{\theta}$ is the solution of the linear system Eq. (5). The quantities involved in this equation are the quantum Fisher matrix (QFM)

$$S_{ij}(\theta) = \frac{\langle \partial_i \psi_\theta | \partial_j \psi_\theta \rangle}{\langle \psi_\theta | \psi_\theta \rangle} - \frac{\langle \partial_i \psi_\theta | \psi_\theta \rangle \langle \psi_\theta | \partial_j \psi_\theta \rangle}{\langle \psi_\theta | \psi_\theta \rangle^2} \quad (18)$$

and the energy gradient

$$F_i(\theta, t) = \frac{\partial \langle \hat{H}(t) \rangle}{\partial \theta_i^*} = \frac{\langle \partial_i \psi_\theta | \hat{H}(t) - \langle \hat{H}(t) \rangle | \psi_\theta \rangle}{\langle \psi_\theta | \psi_\theta \rangle}. \quad (19)$$

Here, $\partial_i = \partial/\partial \theta_i$ denotes the complex partial derivative with respect to θ_i . Geometrically, the QFM accounts for the local curvature around $|\psi_\theta\rangle$ on the manifold of variational states. This is analogous to the role of the Fisher information matrix for classical probability distributions [26], which is used in natural gradient descent [62]. The QFM only depends on the form of the variational ansatz and the location in parameter space but not on the Hamiltonian. Still, analysis of its structure and, in particular, its spectrum can give insight into the quantum properties and phase diagram of the system for the RBM ansatz [54].

In t-VMC, both QFM and gradient are estimated as stochastic expectations values $\mathbb{E}[\cdot]$ with respect to the Born probability distribution $\sim |\psi_\theta(\cdot)|^2$. Specifically,

$$S_{ij}(\theta) = \mathbb{E}[\Theta_i^* \Theta_j] - \mathbb{E}[\Theta_i^*] \mathbb{E}[\Theta_j], \quad (20)$$

$$F_i(\theta, t) = \mathbb{E}[\Theta_i^* \mathfrak{H}(t)] - \mathbb{E}[\Theta_i^*] \mathbb{E}[\mathfrak{H}(t)], \quad (21)$$

with log-derivatives $\Theta_i(s) = \partial_i \ln \psi_\theta(s)$ and local energy $\mathfrak{H}(t)(s) = \frac{\langle s | \hat{H}(t) | \psi_\theta \rangle}{\langle s | \psi_\theta \rangle}$. These equations of motion are valid for a complex differentiable (holomorphic) mapping $\theta \mapsto \psi_\theta$ between the parameters and the quantum wave function. This is indeed satisfied by the symmetric RBM ansatz.

In order to obtain the ground states used as initial states for the time propagation, we have used stochastic reconfiguration [8, 39], which is based on an approximation of the imaginary-time Schrödinger equation in the manner of Eq. (5).

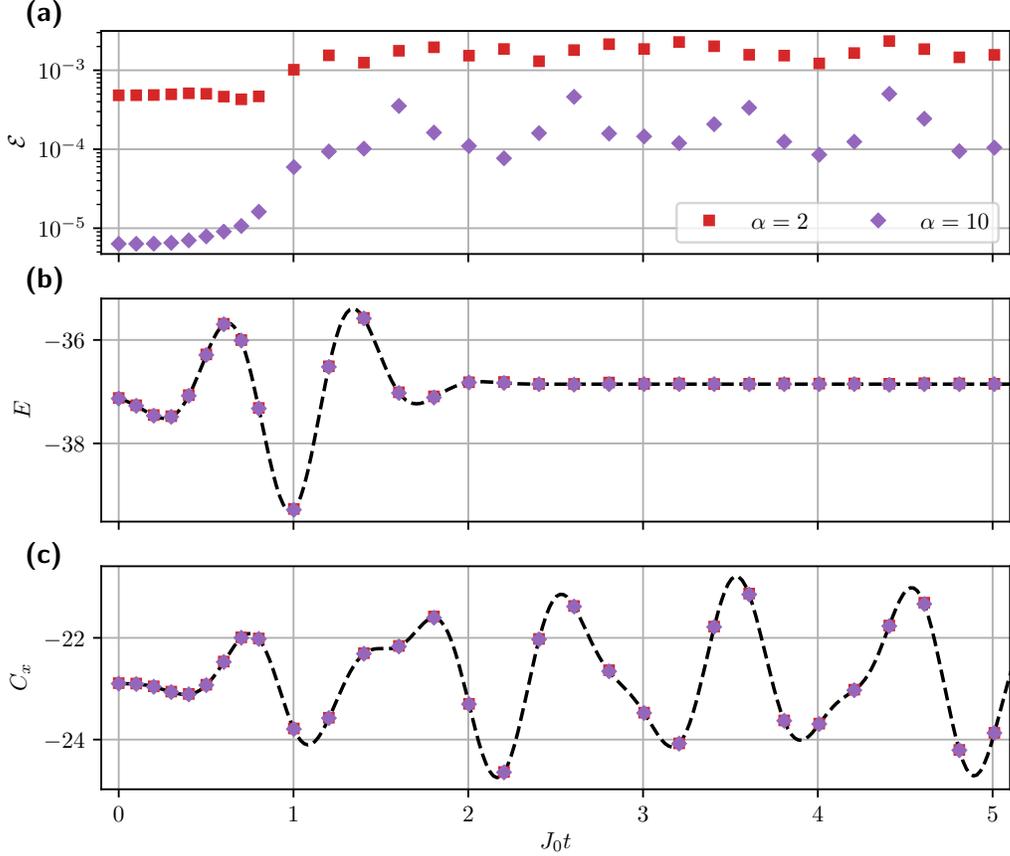


Figure 12: **(a)** Inverse fidelity [Eq. (9)], **(b)** energy, and **(c)** spin correlation for $\alpha = 2$ and $\alpha = 10$ RBM states obtained from supervised learning of the amplitudes along the exact trajectory at an excitation strength of $A_p = 0.10$ (see Appendix C).

C Representability of the trajectory

Even though wider RBMs are necessary in order to better follow the true dynamics via TDVP-based propagation, the states along the trajectories considered here are in general not significantly harder to learn by an RBM than the ground state. In order to test this statement, we fit RBMs with $\alpha = 2$ and $\alpha = 10$ to the exact states along the ED trajectory $|\Psi(t)\rangle$ in the ladder system using a supervised learning approach. To this end, for each given time t , we independently minimize the negative log-overlap loss

$$L_t(\theta) = -\ln \frac{\langle \psi_\theta | \Psi(t) \rangle \langle \Psi(t) | \psi_\theta \rangle}{\langle \psi_\theta | \psi_\theta \rangle \langle \Psi(t) | \Psi(t) \rangle} \quad (22)$$

using the known probability amplitudes of the exact states that we have obtained from ED. The results presented here have been computed using the supervised learning implementation

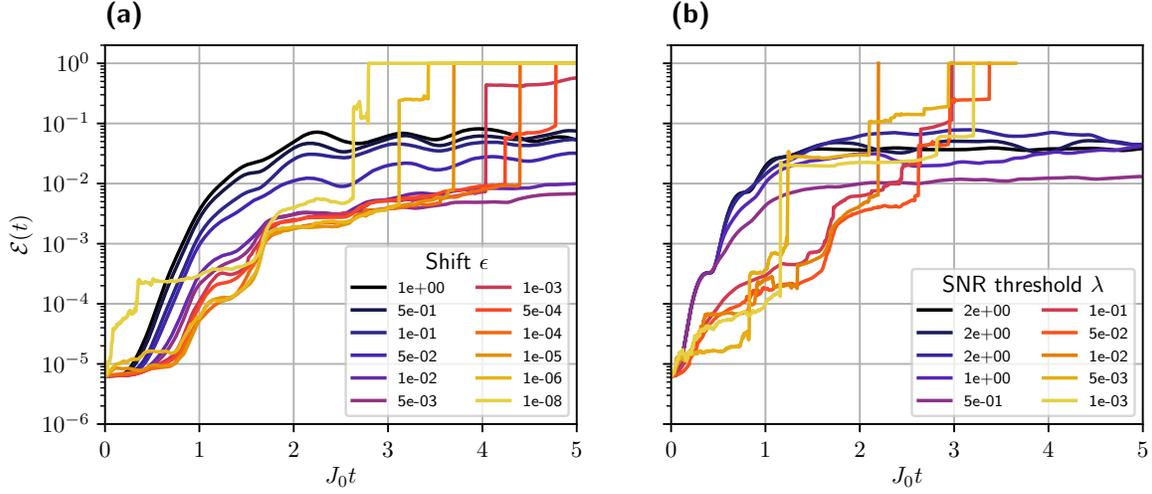


Figure 13: Overlap error of the time-evolved variational state compared to the exact trajectory for **(a)** the diagonal shift regularization with varying strength ϵ and **(b)** the signal-to-noise (SNR) ratio based regularization of Ref. [42] with varying threshold λ_{SNR} . The trajectories have been computed using t-VMC with EMC at a moderate excitation strength of $A_p = 0.10$ in the 8×2 ladder geometry using the symmetrized RBM ansatz with a hidden unit density of $\alpha = 10$.

available in NETKET [56]. We have used stochastic reconfiguration starting from an approximate ground state for 1000 steps at a constant learning rate of $\delta = 0.01$. Results are shown in Fig. 12. Already at a small width of $\alpha = 2$, the trajectory can be represented with an overlap error of the order of 10^{-3} with good accuracy in energy and spin-spin correlations. This error is improved by the wider $\alpha = 10$ RBM. While we can observe an increase in error for the excited states after the pulse, the accuracy is still sufficient to capture the observables under consideration. We note that the instabilities and increased sampling complexity of the states occurring in the ladder dynamics can also affect the convergence of the supervised learning scheme. Thus, the increased error for the excited states is not necessarily an indication of an absolute inability of the RBM ansatz to capture them more accurately but may also be attributed to an increased difficulty of the optimization procedure. In any case, these results indicate that the representability of the time-evolved states is not the key limitation here and an improved time propagation scheme should be expected to be able to reach the accuracy of the supervised learned states.

D Alternative regularization schemes

There are many ways to regularize the linear equation of motion (5) to reduce its susceptibility to noise. We have focused here on the conceptually simple method of truncating the QFM spectrum at a relative threshold as described in Sect. 3.2. Alternatively, the EOM can be regularized by adding a diagonal shift $\tilde{S} = S + \epsilon I$ with $\epsilon > 0$. This is typically done for ground state optimization (see, e.g., Refs. [8, 40, 44]), but can in principle also be applied to the time-

dependent case. Since S is Hermitian and positive semi-definite, the spectrum of the shifted \tilde{S} is bounded from below by ϵ , making the matrix invertible and bounding the condition number by $\kappa(\tilde{S}) \leq (\sigma_1 + \epsilon)/\epsilon$. Therefore, a shift significantly larger than machine precision also serves to improve the condition number and stabilize the propagation in a fashion similar to the SVD cutoff. This can be seen in Fig. 13(a) which shows a similar behavior of the shift regularization when compared to the threshold in Fig. 5.

A more sophisticated regularization strategy has recently been proposed in Ref. [42]. This approach truncates parts of the equations of motion akin to the singular value threshold above but taking into account the strength of VMC noise in different components of the energy gradient. Specifically, directions in the S eigenbasis are discarded based on a softened cutoff λ_{SNR} of the signal-to-noise of the corresponding component of the energy gradient which can be estimated from the t-VMC data. When applying this approach to the ladder geometry, we have found a behavior akin to the other methods discussed above as a function of varying λ_{SNR} [Fig. 13(b)]. This highlights that the trade-off between stability and physical accuracy we have discussed and the need for reliable diagnostics is relevant beyond the simple regularization scheme used in the main text.

References

- [1] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto and L. Zdeborová, *Machine learning and the physical sciences*, Rev. Mod. Phys. **91**(4), 045002 (2019), doi:[10.1103/RevModPhys.91.045002](https://doi.org/10.1103/RevModPhys.91.045002).
- [2] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko and G. Carleo, *Neural-network quantum state tomography*, Nature Physics **14**(5), 447 (2018), doi:[10.1038/s41567-018-0048-5](https://doi.org/10.1038/s41567-018-0048-5).
- [3] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov and P. Mehta, *Reinforcement Learning in Different Phases of Quantum Control*, Phys. Rev. X **8**(3), 031086 (2018), doi:[10.1103/PhysRevX.8.031086](https://doi.org/10.1103/PhysRevX.8.031086).
- [4] T. Fösel, P. Tighineanu, T. Weiss and F. Marquardt, *Reinforcement Learning with Neural Networks for Quantum Feedback*, Phys. Rev. X **8**(3), 031084 (2018), doi:[10.1103/PhysRevX.8.031084](https://doi.org/10.1103/PhysRevX.8.031084).
- [5] J. Carrasquilla and R. G. Melko, *Machine learning phases of matter*, Nature Physics **13**(5), 431 (2017), doi:[10.1038/nphys4035](https://doi.org/10.1038/nphys4035).
- [6] B. S. Rem, N. Käming, M. Tarnowski, L. Asteria, N. Fläschner, C. Becker, K. Sengstock and C. Weitenberg, *Identifying quantum phase transitions using artificial neural networks on experimental data*, Nature Physics (2019), doi:[10.1038/s41567-019-0554-0](https://doi.org/10.1038/s41567-019-0554-0).
- [7] A. Bohrdt, C. S. Chiu, G. Ji, M. Xu, D. Greif, M. Greiner, E. Demler, F. Grusdt and M. Knap, *Classifying snapshots of the doped Hubbard model with machine learning*, Nat. Phys. **15**(9), 921 (2019), doi:[10.1038/s41567-019-0565-x](https://doi.org/10.1038/s41567-019-0565-x).
- [8] G. Carleo and M. Troyer, *Solving the quantum many-body problem with artificial neural networks*, Science **355**(6325), 602 (2017), doi:[10.1126/science.aag2302](https://doi.org/10.1126/science.aag2302).

- [9] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez and J. I. Cirac, *Neural-Network Quantum States, String-Bond States, and Chiral Topological States*, Phys. Rev. X **8**(1), 011006 (2018), doi:[10.1103/PhysRevX.8.011006](https://doi.org/10.1103/PhysRevX.8.011006).
- [10] R. Kaubruegger, L. Pastori and J. C. Budich, *Chiral topological phases from artificial neural networks*, Phys. Rev. B **97**(19), 195136 (2018), doi:[10.1103/PhysRevB.97.195136](https://doi.org/10.1103/PhysRevB.97.195136).
- [11] K. Choo, G. Carleo, N. Regnault and T. Neupert, *Symmetries and Many-Body Excitations with Neural-Network Quantum States*, Phys. Rev. Lett. **121**(16), 167204 (2018), doi:[10.1103/PhysRevLett.121.167204](https://doi.org/10.1103/PhysRevLett.121.167204).
- [12] X. Liang, W.-Y. Liu, P.-Z. Lin, G.-C. Guo, Y.-S. Zhang and L. He, *Solving frustrated quantum many-particle models with convolutional neural networks*, Phys. Rev. B **98**(10), 104426 (2018), doi:[10.1103/PhysRevB.98.104426](https://doi.org/10.1103/PhysRevB.98.104426).
- [13] L. Pastori, R. Kaubruegger and J. C. Budich, *Generalized transfer matrix states from artificial neural networks*, Phys. Rev. B **99**(16), 165123 (2019), doi:[10.1103/PhysRevB.99.165123](https://doi.org/10.1103/PhysRevB.99.165123).
- [14] K. Choo, T. Neupert and G. Carleo, *Two-Dimensional Frustrated J_1 - J_2 Model Studied with Neural Network Quantum States*, Phys. Rev. B **100**(12), 125124 (2019), doi:[10.1103/PhysRevB.100.125124](https://doi.org/10.1103/PhysRevB.100.125124).
- [15] D. Hendry and A. E. Feiguin, *Machine learning approach to dynamical properties of quantum many-body systems*, Phys. Rev. B **100**(24), 245123 (2019), doi:[10.1103/PhysRevB.100.245123](https://doi.org/10.1103/PhysRevB.100.245123).
- [16] T. Vieijra, C. Casert, J. Nys, W. De Neve, J. Haegeman, J. Ryckebusch and F. Verstraete, *Restricted Boltzmann Machines for Quantum States with Non-Abelian or Anyonic Symmetries*, Phys. Rev. Lett. **124**(9), 097201 (2020), doi:[10.1103/PhysRevLett.124.097201](https://doi.org/10.1103/PhysRevLett.124.097201).
- [17] T. Westerhout, N. Astrakhantsev, K. S. Tikhonov, M. I. Katsnelson and A. A. Bagrov, *Generalization properties of neural network approximations to frustrated magnet ground states*, Nat Commun **11**(1), 1 (2020), doi:[10.1038/s41467-020-15402-w](https://doi.org/10.1038/s41467-020-15402-w).
- [18] A. Szabó and C. Castelnovo, *Neural network wave functions and the sign problem*, Phys. Rev. Research **2**(3), 033075 (2020), doi:[10.1103/PhysRevResearch.2.033075](https://doi.org/10.1103/PhysRevResearch.2.033075).
- [19] M. Bukov, M. Schmitt and M. Dupont, *Learning the ground state of a non-stoquastic quantum Hamiltonian in a rugged neural network landscape* (2020), [2011.11214](https://arxiv.org/abs/2011.11214).
- [20] N. Astrakhantsev, T. Westerhout, A. Tiwari, K. Choo, A. Chen, M. H. Fischer, G. Carleo and T. Neupert, *Broken-symmetry ground states of the heisenberg model on the pyrochlore lattice* (2021), [arXiv:2101.08787](https://arxiv.org/abs/2101.08787).
- [21] A. Valenti, E. Greplova, N. H. Lindner and S. D. Huber, *Correlation-enhanced neural networks as interpretable variational quantum states* (2021), [arXiv:2103.05017](https://arxiv.org/abs/2103.05017).
- [22] K. McBrien, G. Carleo and E. Khatami, *Ground state phase diagram of the one-dimensional Bose-Hubbard model from restricted Boltzmann machines*, J. Phys.: Conf. Ser. **1290**, 012005 (2019), doi:[10.1088/1742-6596/1290/1/012005](https://doi.org/10.1088/1742-6596/1290/1/012005).

- [23] Y. Nomura, A. S. Darmawan, Y. Yamaji and M. Imada, *Restricted Boltzmann machine learning for solving strongly correlated quantum systems*, Phys. Rev. B **96**(20), 205152 (2017), doi:[10.1103/PhysRevB.96.205152](https://doi.org/10.1103/PhysRevB.96.205152).
- [24] K. Choo, A. Mezzacapo and G. Carleo, *Fermionic neural-network states for ab-initio electronic structure*, Nature Communications **11**(1), 2368 (2020), doi:[10.1038/s41467-020-15724-9](https://doi.org/10.1038/s41467-020-15724-9).
- [25] B. Jónsson, B. Bauer and G. Carleo, *Neural-network states for the classical simulation of quantum computing* (2018), [arXiv:1808.05232](https://arxiv.org/abs/1808.05232).
- [26] J. Stokes, J. Izaac, N. Killoran and G. Carleo, *Quantum Natural Gradient*, Quantum **4**, 269 (2020), doi:[10.22331/q-2020-05-25-269](https://doi.org/10.22331/q-2020-05-25-269).
- [27] G. Torlai and R. G. Melko, *Latent Space Purification via Neural Density Operators*, Phys. Rev. Lett. **120**(24), 240503 (2018), doi:[10.1103/PhysRevLett.120.240503](https://doi.org/10.1103/PhysRevLett.120.240503).
- [28] M. J. Hartmann and G. Carleo, *Neural-Network Approach to Dissipative Quantum Many-Body Dynamics*, Phys. Rev. Lett. **122**(25), 250502 (2019), doi:[10.1103/PhysRevLett.122.250502](https://doi.org/10.1103/PhysRevLett.122.250502).
- [29] N. Yoshioka and R. Hamazaki, *Constructing neural stationary states for open quantum many-body systems*, Phys. Rev. B **99**(21), 214306 (2019), doi:[10.1103/PhysRevB.99.214306](https://doi.org/10.1103/PhysRevB.99.214306).
- [30] F. Vicentini, A. Biella, N. Regnault and C. Ciuti, *Variational neural-network ansatz for steady states in open quantum systems*, Phys. Rev. Lett. **122**, 250503 (2019), doi:[10.1103/PhysRevLett.122.250503](https://doi.org/10.1103/PhysRevLett.122.250503).
- [31] J. Eisert, M. Friesdorf and C. Gogolin, *Quantum many-body systems out of equilibrium*, Nature Physics **11**(2), 124 (2015), doi:[10.1038/nphys3215](https://doi.org/10.1038/nphys3215).
- [32] R. J. Lewis-Swan, A. Safavi-Naini, A. M. Kaufman and A. M. Rey, *Dynamics of quantum information*, Nature Reviews Physics **1**(10), 627 (2019), doi:[10.1038/s42254-019-0090-y](https://doi.org/10.1038/s42254-019-0090-y).
- [33] D. Bluvstein, A. Omran, H. Levine, A. Keesling, G. Semeghini, S. Ebadi, T. T. Wang, A. A. Michailidis, N. Maskara, W. W. Ho, S. Choi, M. Serbyn *et al.*, *Controlling quantum many-body dynamics in driven Rydberg atom arrays*, Science **371**(6536), 1355 (2021), doi:[10.1126/science.abg2530](https://doi.org/10.1126/science.abg2530).
- [34] I. Carusotto and C. Ciuti, *Quantum fluids of light*, Reviews of Modern Physics **85**(1), 299 (2013), doi:[10.1103/RevModPhys.85.299](https://doi.org/10.1103/RevModPhys.85.299).
- [35] A. de la Torre, D. M. Kennes, M. Claassen, S. Gerber, J. W. McIver and M. A. Sentef, *Nonthermal pathways to ultrafast control in quantum materials* (2021), [arXiv:2104.10696](https://arxiv.org/abs/2104.10696).
- [36] D.-L. Deng, X. Li and S. Das Sarma, *Quantum entanglement in neural network states*, Phys. Rev. X **7**, 021021 (2017), doi:[10.1103/PhysRevX.7.021021](https://doi.org/10.1103/PhysRevX.7.021021).
- [37] J. Eisert, M. Cramer and M. B. Plenio, *Colloquium: Area laws for the entanglement entropy*, Rev. Mod. Phys. **82**, 277 (2010), doi:[10.1103/RevModPhys.82.277](https://doi.org/10.1103/RevModPhys.82.277).

- [38] G. Carleo, F. Becca, M. Schiró and M. Fabrizio, *Localization and Glassy Dynamics Of Many-Body Quantum Systems*, Scientific Reports **2**, 243 (2012), doi:[10.1038/srep00243](https://doi.org/10.1038/srep00243).
- [39] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems*, Cambridge University Press, first edn., ISBN 978-1-107-12993-1 978-1-316-41704-1, doi:[10.1017/9781316417041](https://doi.org/10.1017/9781316417041) (2017).
- [40] S. Czischek, M. Gärttner and T. Gasenzer, *Quenches near Ising quantum criticality as a challenge for artificial neural networks*, Phys. Rev. B **98**(2), 024311 (2018), doi:[10.1103/PhysRevB.98.024311](https://doi.org/10.1103/PhysRevB.98.024311).
- [41] I. López-Gutiérrez and C. B. Mendl, *Real time evolution with neural-network quantum states*, arXiv:1912.08831 (2019), [1912.08831](https://arxiv.org/abs/1912.08831).
- [42] M. Schmitt and M. Heyl, *Quantum many-body dynamics in two dimensions with artificial neural networks*, Phys. Rev. Lett. **125**, 100503 (2020), doi:[10.1103/PhysRevLett.125.100503](https://doi.org/10.1103/PhysRevLett.125.100503).
- [43] G. Fabiani and J. Mentink, *Investigating ultrafast quantum magnetism with machine learning*, SciPost Physics **7**(1), 004 (2019), doi:[10.21468/SciPostPhys.7.1.004](https://doi.org/10.21468/SciPostPhys.7.1.004).
- [44] G. Fabiani and J. H. Mentink, *Ultrafast entanglement dynamics and spreading of quantum correlations in two-dimensional antiferromagnets* (2019), [arXiv:1912.10845](https://arxiv.org/abs/1912.10845).
- [45] G. Fabiani, M. D. Bouman and J. H. Mentink, *Supermagnonic propagation in two-dimensional antiferromagnets* (2021), [arXiv:2101.10945](https://arxiv.org/abs/2101.10945).
- [46] S.-H. Lin and F. Pollmann, *Scaling of neural-network quantum states for time evolution* (2021), [arXiv:2104.10696](https://arxiv.org/abs/2104.10696).
- [47] C.-Y. Park and M. J. Kastoryano, *Are neural quantum states good at solving non-stoquastic spin hamiltonians?* (2020), [arXiv:2012.08889](https://arxiv.org/abs/2012.08889).
- [48] R. G. Melko, G. Carleo, J. Carrasquilla and J. I. Cirac, *Restricted Boltzmann machines in quantum physics*, Nat. Phys. **15**(9), 887 (2019), doi:[10.1038/s41567-019-0545-1](https://doi.org/10.1038/s41567-019-0545-1).
- [49] D. Bossini, S. Dal Conte, G. Cerullo, O. Gomonay, R. V. Pisarev, M. Borovsak, D. Mihailovic, J. Sinova, J. H. Mentink, T. Rasing and et al., *Laser-driven quantum magnonics and terahertz dynamics of the order parameter in antiferromagnets*, Physical Review B **100**(2) (2019), doi:[10.1103/physrevb.100.024428](https://doi.org/10.1103/physrevb.100.024428).
- [50] T. Barnes, E. Dagotto, J. Riera and E. S. Swanson, *Excitation spectrum of heisenberg spin ladders*, Phys. Rev. B **47**, 3196 (1993), doi:[10.1103/PhysRevB.47.3196](https://doi.org/10.1103/PhysRevB.47.3196).
- [51] S. Gopalan, T. M. Rice and M. Sigrist, *Spin ladders with spin gaps: A description of a class of cuprates*, Phys. Rev. B **49**, 8901 (1994), doi:[10.1103/PhysRevB.49.8901](https://doi.org/10.1103/PhysRevB.49.8901).
- [52] E. Dagotto and T. M. Rice, *Surprises on the way from one- to two-dimensional quantum magnets: The ladder materials*, Science **271**(5249), 618 (1996), doi:[10.1126/science.271.5249.618](https://doi.org/10.1126/science.271.5249.618).

- [53] O. Sharir, Y. Levine, N. Wies, G. Carleo and A. Shashua, *Deep autoregressive models for the efficient variational simulation of many-body quantum systems*, Phys. Rev. Lett. **124**, 020503 (2020), doi:[10.1103/PhysRevLett.124.020503](https://doi.org/10.1103/PhysRevLett.124.020503).
- [54] C.-Y. Park and M. J. Kastoryano, *Geometry of learning neural quantum states*, Phys. Rev. Research **2**, 023232 (2020), doi:[10.1103/PhysRevResearch.2.023232](https://doi.org/10.1103/PhysRevResearch.2.023232).
- [55] I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*, MIT Press, <http://www.deeplearningbook.org> (2016).
- [56] G. Carleo, K. Choo, D. Hofmann, J. E. T. Smith, T. Westerhout, F. Alet, E. J. Davis, S. Efthymiou, I. Glasser, S.-H. Lin, M. Mauri, G. Mazzola *et al.*, *NetKet: A machine learning toolkit for many-body quantum systems*, SoftwareX **10**, 100311 (2019), doi:[10.1016/j.softx.2019.100311](https://doi.org/10.1016/j.softx.2019.100311).
- [57] J. R. Johansson, P. D. Nation and F. Nori, *QuTiP 2: A Python framework for the dynamics of open quantum systems*, Computer Physics Communications **184**(4), 1234 (2013), doi:[10.1016/j.cpc.2012.11.019](https://doi.org/10.1016/j.cpc.2012.11.019).
- [58] D. J. Klein, *Ground-state features for Heisenberg models*, J. Chem. Phys. **77**(6), 3098 (1982), doi:[10.1063/1.444232](https://doi.org/10.1063/1.444232).
- [59] U. Schollwoeck, *Marshall's sign rule and DMRG acceleration*, arXiv:cond-mat/9804231 (1998), [cond-mat/9804231](https://arxiv.org/abs/cond-mat/9804231).
- [60] R. F. Bishop, D. J. J. Farnell and J. B. Parkinson, *Sign rules for anisotropic quantum spin systems*, Phys. Rev. B **61**(10), 6775 (2000), doi:[10.1103/PhysRevB.61.6775](https://doi.org/10.1103/PhysRevB.61.6775).
- [61] W. Marshall, *Antiferromagnetism*, Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences **232**(1188), 48 (1955), doi:[10.1098/rspa.1955.0200](https://doi.org/10.1098/rspa.1955.0200).
- [62] S.-i. Amari, *Natural Gradient Works Efficiently in Learning*, Neural Computation **10**(2), 251 (1998), doi:[10.1162/089976698300017746](https://doi.org/10.1162/089976698300017746).