


Variational Quantum Dynamics of Two-Dimensional Rotor Models

Matija Medvidović^{1,2,*} and Dries Sels^{1,3}

¹Center for Computational Quantum Physics, Flatiron Institute, 162 5th Avenue, New York, New York 10010, USA

²Department of Physics, Columbia University, New York 10027, USA

³Center for Quantum Phenomena, Department of Physics, New York University, 726 Broadway, New York, NY, 10003, USA

 (Received 12 January 2023; revised 18 August 2023; accepted 11 September 2023; published 4 October 2023)

We present a numerical method to simulate the dynamics of continuous-variable quantum many-body systems. Our approach is based on custom neural-network many-body quantum states. We focus on dynamics of two-dimensional quantum rotors and simulate large experimentally relevant system sizes by representing a trial state in a continuous basis and using state-of-the-art sampling approaches based on Hamiltonian Monte Carlo. We demonstrate the method can access quantities like the return probability and vorticity oscillations after a quantum quench in two-dimensional systems of up to 64 (8×8) coupled rotors. Our approach can be used for accurate nonequilibrium simulations of continuous systems at previously unexplored system sizes and evolution times, bridging the gap between simulation and experiment.

DOI: [10.1103/PRXQuantum.4.040302](https://doi.org/10.1103/PRXQuantum.4.040302)

I. INTRODUCTION

Nonequilibrium quantum many-body physics has been at the forefront of condensed matter, atomic physics, and chemistry research for over a decade [1,2]. The field is driven by remarkable progress in our ability to coherently control matter at the atomic scale. This control has resulted in the creation of novel phases of matter, including observations of light-induced superconductivity [3], cavity-enhanced chemical reactions [4], and dynamical phase transitions [5].

The capacity to precisely control [6–9] modern quantum experiments and hardware is becoming increasingly limited by numerical simulation of the real-time evolution of quantum systems. At its core, the problem is related to fast entanglement growth in systems out of equilibrium, which forces one to keep track of all the intricate correlations that build up in the system. While there has been considerable progress [10–14], challenges remain, in particular, if one moves away from one-dimensional spin models.

Recently, it has been proposed that methods inspired by classical and quantum machine learning might alleviate some of these problems [15–19]. In practice, however, it has been difficult to achieve reliable results due to numerical instabilities resulting from a combination of Monte Carlo noise and flatness of the quantum geometry of modern neural-network wave functions [17,20–23].

In this work, we present an approach for capturing long-time dynamics of two-dimensional (2D) lattice models with continuous degrees of freedom, using a combination of methods that were previously unexplored in the field of variational simulations—the Hamiltonian Monte Carlo sampler, a tailored variational ansatz and proper regularization of the projected dynamics. We focus on the quantum rotor model with direct applications to arrays of coupled Josephson junctions and explore previously unreachable system sizes and evolution times, up to 8×8 square lattices.

The paper is organized as follows. First, we introduce the physics of the quantum rotor model and the variational wave function. Then, we outline the Hamiltonian Monte Carlo sampler and its connection to the time-dependent variational Monte Carlo algorithm. Finally, we present results for the two-dimensional model, showing magnetization, vorticity, and the Loschmidt echo converging to appropriate equilibrium values. Our Monte Carlo results are substantiated by self-consistency checks when key hyperparameters are changed and by comparing our approach to tensor-network calculations in one and two spatial dimensions.

*matija.medvidovic@columbia.edu

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

II. MODEL AND METHODS

Consider a system of continuous planar rotors, whose angles θ_k (with respect to an arbitrary axis) could, for example, represent superconducting phases of adjacent Josephson junctions on a lattice Λ with N sites. We use the basis $|\boldsymbol{\theta}\rangle \equiv |\theta_1, \dots, \theta_N\rangle$ for the Hilbert space \mathcal{H} . We start with an effective Hamiltonian that captures the relevant physics of superconducting Josephson junctions [24–26]:

$$H = \frac{gJ}{2} \sum_k L_k^2 - J \sum_{\langle k,l \rangle} \hat{\mathbf{n}}_k \cdot \hat{\mathbf{n}}_l, \quad (1)$$

where $L_k = -i\partial_{\theta_k}$ and $\hat{\mathbf{n}}_k = (\cos \theta_k, \sin \theta_k)$ in the continuous basis $|\boldsymbol{\theta}\rangle$ of choice. The Hamiltonian in Eq. (1) is often called the quantum rotor model (QRM). Its equilibrium properties [27] have been studied using variational Monte Carlo (VMC) [28] and other quantum Monte Carlo (QMC) [29] methods. Perhaps most notably, the quantum critical point separating the disordered and O(2) broken phase has been predicted at $g_c \approx 4.25$.

However, as noted in the introduction, real-time evolution properties of the QRM have barely been explored. This is mainly due to the lack of suitable methods that can access experimentally relevant times $t \gg J^{-1}$, at large system sizes in two dimensions. The ability to simulate relatively large system sizes is not only of theoretical interests but has technological applications in the study of dynamics of arrays of coupled Josephson junctions [30].

The evolution equation for a state $\Psi = \Psi(\boldsymbol{\theta})$, in the continuous basis $|\boldsymbol{\theta}\rangle$, reads

$$i \frac{\partial \Psi}{\partial t} = -\frac{gJ}{2} \sum_k \frac{\partial^2 \Psi}{\partial \theta_k^2} - J \sum_{\langle k,l \rangle} \cos(\theta_k - \theta_l) \Psi, \quad (2)$$

with appropriate periodic boundary conditions

$$\Psi(\theta_1, \dots, \theta_k + 2\pi, \dots, \theta_N) = \Psi(\theta_1, \dots, \theta_k, \dots, \theta_N) \quad (3)$$

for each rotor k . Equation (2) is prohibitively expensive to solve exactly even for a handful of interacting rotors. The continuous nature of the $|\boldsymbol{\theta}\rangle$ basis exacerbates the problem.

A. Variational simulation

We represent a quantum state using a wave function $\psi_\alpha(\boldsymbol{\theta})$ where $\alpha \in \mathbb{C}^P$ is a set of P real or complex variational parameters. Since any $|\psi\rangle \in \mathcal{H}$ admits an expansion in terms of $|\boldsymbol{\theta}\rangle$, we define the following un-normalized variational quantum state (VQS):

$$|\psi_\alpha\rangle = \int d\boldsymbol{\theta} \psi_\alpha(\boldsymbol{\theta}) |\boldsymbol{\theta}\rangle, \quad (4)$$

where $d\boldsymbol{\theta} \equiv d\theta_1 \cdots d\theta_N$. The integral is performed over the cube $[-\pi, \pi]^N$.

Building on previous work on continuous systems [16], our simulation of the real-time dynamics of the state given in Eq. (4) is based on the time-dependent variational Monte Carlo (t-MVC) method [15,31]. The core assumption that allows us to approximately solve Eq. (2) is that of time dependence of parameters $\alpha = \alpha(t)$.

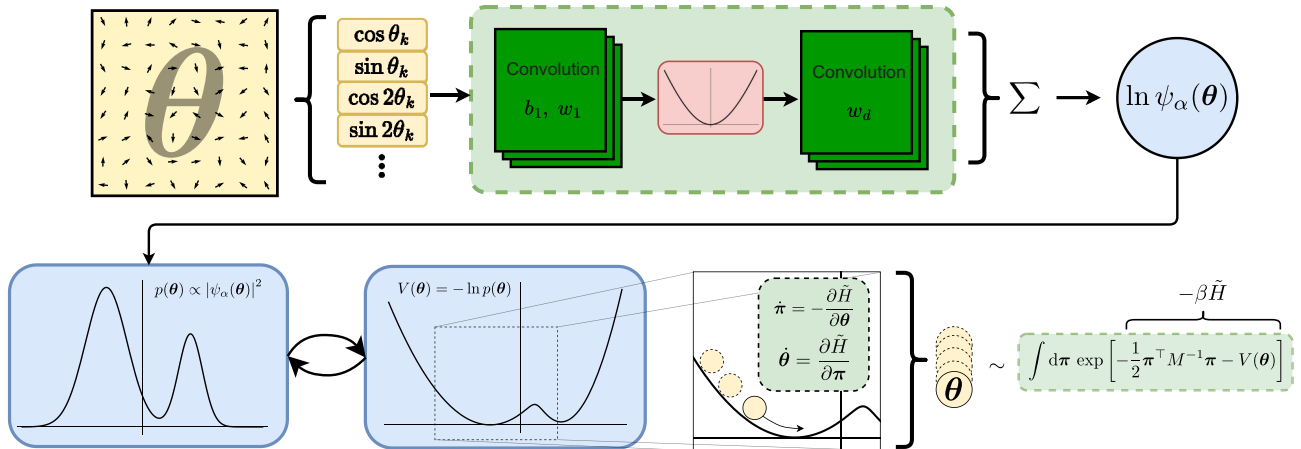


FIG. 1. Top: the ansatz $\psi_\alpha(\boldsymbol{\theta})$ architecture used for simulations of two-dimensional QRM systems. It amounts to a two-layer convolutional neural network with an activation function given by Eq. (15). To enforce periodicity and improve expressivity, we precalculate sines and cosines of input angles, which are treated as different input channels by the CNN. The final layer outputs a single channel and all of its components are summed into a single complex number (because of complex parameters $\alpha \in \mathbb{C}^P$) we then interpret as $\ln \psi_\alpha(\boldsymbol{\theta})$. Bottom: an illustration of the Hamiltonian Monte Carlo algorithm. Dummy momentum variables are introduced and sampling the given N -dimensional probability distribution is rewritten in $2N$ -dimensional phase space with an artificial effective Hamiltonian \tilde{H} . Samples are collected as snapshots of solutions of Hamilton's equations of motion.

Optimal trajectories $\alpha(t)$ induced by unitary hamiltonian evolution $e^{-iHt} |\psi_\alpha\rangle$ can conveniently be found by extremizing the time-dependent variational principle (TDVP) [22] action

$$\mathcal{C}[\alpha] = \int dt \langle \Psi_{\alpha(t)} | \left(i \frac{d}{dt} - H \right) | \Psi_{\alpha(t)} \rangle, \quad (5)$$

where $|\Psi_\alpha\rangle$ is a normalized version of state $|\psi_\alpha\rangle$. Optimal evolution equations read $iS\dot{\alpha} = g$, where

$$\begin{aligned} S_{\mu\nu} &= \langle \mathcal{O}_\mu^\dagger \mathcal{O}_\nu \rangle - \langle \mathcal{O}_\mu^\dagger \rangle \langle \mathcal{O}_\nu \rangle, \\ g_\mu &= \langle \mathcal{O}_\mu^\dagger H \rangle - \langle \mathcal{O}_\mu^\dagger \rangle \langle H \rangle, \end{aligned} \quad (6)$$

with averages $\langle \cdot \rangle \equiv \langle \psi_\alpha | \cdot | \psi_\alpha \rangle / \langle \psi_\alpha | \psi_\alpha \rangle$ being performed at time t [i.e., for $\alpha = \alpha(t)$]. Operator \mathcal{O}_μ is defined by $\partial_{\alpha_\mu} |\psi_\alpha\rangle = \mathcal{O}_\mu |\psi_\alpha\rangle$. We note that the matrix S is commonly called the *quantum geometric tensor* (QGT) [21,23,32] and corresponds to the metric tensor of the parameter manifold induced by the distance in \mathcal{H} between un-normalized states defined in Eq. (4). In Eqs. 6, we have chosen our ansatz ψ_α such that it is a holomorphic function of complex parameters α .

Since quantum averages over an exponentially large Hilbert space \mathcal{H} in the TDVP Eq. (6) cannot be computed exactly, Markov chain Monte Carlo (MCMC) sampling methods are often employed [33,34]. In VMC calculations, it is common to rewrite quantum averages, such as those in Eq. (6), as expressions amenable to estimation through sampling. For example, in the case of the Hamiltonian H , we obtain the *local energy* E_L :

$$\langle H \rangle = \frac{\langle \psi_\alpha | H | \psi_\alpha \rangle}{\langle \psi_\alpha | \psi_\alpha \rangle} = \int d\theta p_\alpha(\theta) E_L(\theta) \quad (7)$$

where

$$p_\alpha(\theta) \propto |\psi_\alpha(\theta)|^2 \quad \text{and} \quad E_L(\theta) = \frac{\langle \theta | H | \psi_\alpha \rangle}{\langle \theta | \psi_\alpha \rangle}. \quad (8)$$

For more details about the specific sampling algorithm employed in this work, we refer the reader to Sec. II B and Appendix A 1.

After computing the matrix S and the vector g at time t , one can formally define $\dot{\alpha} = -iS^{-1}g$ and use any ordinary differential equation (ODE) integrator (see Appendix A 3) to obtain the next set of parameters, at time $t + \delta t$. However, the inverse is often ill defined.

One reason is that Monte Carlo estimates of matrix elements are noisy. Noise accumulates to render the matrix singular by making a small eigenvalues vanish. Therefore, quickly and efficiently obtaining many uncorrelated samples from $p(\theta, t) \propto |\psi_{\alpha(t)}(\theta)|^2$ is crucial. The other reason is that the specific choice of ψ_α introduces redundancy between different parameters, producing linearly

dependent or vanishing rows and columns in S . Therefore, choosing an efficiently parameterized trial wave function is equally important. In practice, adding more parameters to the wave function can sometimes unexpectedly reduce accuracy by making S ill conditioned.

In order to move forward with the algorithm, regularization schemes must be used. For ground-state optimization tasks, simply replacing $S \rightarrow S + \epsilon \mathbb{1}$, for some small positive constant ϵ , often suffices to diminish the effect of small eigenvalues.

However, in this work, we regularize the S matrix by diagonalization $S = U\Sigma U^\dagger$ at each time step. Having obtained eigenvalues σ_μ^2 such that $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_P^2)$, we define the pseudoinverse as $S^{-1} \approx U\tilde{\Sigma}^{-1}U^\dagger$ with

$$\tilde{\Sigma}_{\mu\nu}^{-1} = \frac{1/\sigma_\mu^2}{1 + (\lambda^2/\sigma_\mu^2)^6} \delta_{\mu\nu}. \quad (9)$$

We heuristically find that the smooth cutoff with a hyperparameter λ^2 in Eq. (9) is superior to traditional pseudoinverses when using adaptive integrators for updating parameters α . For more details on regularization, see Appendix A 2.

After calculating averages in Eq. (6) and appropriately regularizing the QGT inverse S^{-1} , one can use any external ODE integrator to perform time stepping in the top-level equation $\dot{\alpha} = -iS^{-1}g$. In this work, we use the embedded Bogacki-Shampine adaptive solver RK3(2) from the Runge-Kutta family [35–37].

B. Hamiltonian Monte Carlo

Hilbert-space averages defined in Eq. (6) cannot be evaluated analytically for an arbitrary ψ_α . To perform this task in an efficient and scalable way, we employ Hamiltonian Monte Carlo (HMC) [38,39] to obtain samples from the distribution $p(\theta, t)$ at each time step t . We make this choice because HMC offers a systematic way of making large steps in MCMC proposals while still keeping acceptance probabilities high, unlike more conventional approaches like random-walk Metropolis. This results in a Markov chain with considerably lower autocorrelation times, allowing for treatments of larger systems with less overall runtime spent on sampling.

For a generic probability distribution $p(\theta)$, HMC augments the configuration space with artificial momentum variables $\boldsymbol{\pi} = (\pi_1, \dots, \pi_N) \sim \mathcal{N}(0, M)$:

$$p(\theta) \propto \int d\boldsymbol{\pi} \exp \left\{ -\frac{1}{2} \boldsymbol{\pi}^\top M^{-1} \boldsymbol{\pi} + \ln p(\theta) \right\} \quad (10)$$

for some choice of a positive-definite *mass matrix* M . Interpreting the exponent in Eq. (10) as an effective classical Hamiltonian $\beta\tilde{H}(\theta, \boldsymbol{\pi})$ inducing a Boltzmann weight

$e^{-\beta\tilde{H}}$, Monte Carlo updates can be defined through numerical integration of relevant Hamilton's equations. Owing to insights from statistical physics, we know that a large number of particles in equilibrium following classical equations of motion have precisely this desired Boltzmann distribution.

Given $\boldsymbol{\theta}(0)$, $\boldsymbol{\pi}(0)$ and a small step size ε , a common choice is the leapfrog integrator:

$$\begin{aligned}\boldsymbol{\pi}(\tau + \varepsilon/2) &= \boldsymbol{\pi}(\tau) - \frac{\varepsilon}{2} \frac{\partial V}{\partial \boldsymbol{\theta}}(\boldsymbol{\theta}(\tau)), \\ \boldsymbol{\theta}(\tau + \varepsilon) &= \boldsymbol{\theta}(\tau) + \varepsilon M^{-1} \boldsymbol{\pi}(\tau + \varepsilon/2), \\ \boldsymbol{\pi}(\tau + \varepsilon) &= \boldsymbol{\pi}(\tau + \varepsilon/2) - \frac{\varepsilon}{2} \frac{\partial V}{\partial \boldsymbol{\theta}}(\boldsymbol{\theta}(\tau + \varepsilon)),\end{aligned}\quad (11)$$

where $V(\boldsymbol{\theta}) = -\ln p(\boldsymbol{\theta})$ and τ is the fictitious HMC time variable, unrelated to t in Eq. (5). This specific integrator is chosen because of its symplectic [37,38] property—it conserves energy and probability exactly, allowing for large jumps in the $\boldsymbol{\theta}$ space while keeping high acceptance probabilities. We note that higher-order symplectic integrators can be used as well.

After integrating for L steps, the new configuration $(\boldsymbol{\theta}(L\varepsilon), \boldsymbol{\pi}(L\varepsilon))$ is proposed as the next sample in the Markov chain. It is common to apply the Metropolis-Hastings accept-reject step [33,34] despite the fact that the new configuration has the same energy (probability) as the initial one. This is done to offset the effects of unwanted numerical errors in the leapfrog scheme, usually improving overall performance for many samples [38,39].

Equations (11) simulate a swarm of effective classical particles whose positions and momenta follow the desired joint Boltzmann distribution in Eq. (10). Discarding all $\boldsymbol{\pi}$ samples is equivalent to marginalizing the distribution in Eq. (10). In practice, randomness is injected by sampling the normal distribution $\boldsymbol{\pi}(0) \sim \mathcal{N}(0, M)$ each time initial conditions are required for numerical integration.

Choosing the mass matrix M , the time step ε and the integration length L carefully is crucial for efficient exploration of the configuration space. In this work, we chose to set M and ε automatically, by using heuristically proven [39–41] algorithms operating samples from an extended warmup phase for each Markov chain individually. Integration length L was treated as a hyperparameter. For more details and specific values, see Appendix A 1.

C. The trial wave function

In this work, we use a variant of the standard convolutional neural-network (CNN) architecture [42,43] to model $\psi_\alpha(\boldsymbol{\theta})$. Our approach is built on those of Refs [17,44].

Specifically, we set

$$\ln \psi_\alpha(\boldsymbol{\theta}) = \frac{1}{\sqrt{2KN}} \sum_{c=1}^{2K} \sum_k [w_D^c * h_{D-1}^c(\boldsymbol{\theta})]_k, \quad (12)$$

where $*$ denotes a convolution over lattice indices k and $c = 1, \dots, 2K$ is the channel index. Features $h_{D-1}^c(\boldsymbol{\theta})$ are the output of $D - 1$ -layer CNN defined by

$$h_d^c(\boldsymbol{\theta}) = f_d \left(b_d^c + \sum_{c'=1}^{2K} w_d^{cc'} * h_{d-1}^{c'}(\boldsymbol{\theta}), \right) \quad (13)$$

with an elementwise nonlinear activation function f_d , biases b_d , and weights w_d at layer d . We include all weights and biases into the set of trainable parameters α and use automatic differentiation (AD) techniques to obtain all derivatives \mathcal{O}_μ required for evaluation of Eqs. (6). For CNN inputs h_0 , we concatenate the following features:

$$h_0 = \left\{ (\cos n\theta_k, \sin n\theta_k) \mid n = 1, \dots, K \right\} \quad (14)$$

along the channel axis, as illustrated on Fig. 1. This construction allows us to include a limited number of higher Fourier modes *a priori*, improving ansatz expressivity in a controlled way. In this work, we set $D = 2$, $K = 4$ for larger two-dimensional (8×8) experiments and $K = 1$ for smaller systems.

To maintain analytic dependence on parameters α , we restrict the CNN nonlinearities f_d to polynomial functions. The Taylor expansion of the logarithm of the zeroth-order modified Bessel function of the first kind is used:

$$\ln I_0(z) = \frac{z^2}{4} - \frac{z^4}{64} + \frac{z^6}{576} + \mathcal{O}(z^8). \quad (15)$$

This particular activation function choice is motivated by the appearance of I_0 in the version of the restricted Boltzmann machine (RBM) adapted to the QRM in Ref. [28]. This approach has the advantage of maintaining the holomorphic dependence of ψ_α on α and preserving the form of Eqs. (6).

In this work, we focus on a simple two-layer CNN ansatz to control the number of parameters P . In addition nontrivially affecting the QGT inverse (see Sec. II A), the cost to diagonalize the QGT in order to regularize the inverse in Eq. (9) grows as $\mathcal{O}(P^3)$. Heuristically, we also find that introducing more parameters α requires more Monte Carlo samples to correctly resolve the relevant averages in Eq. (6) and does not significantly contribute to simulation accuracy in our case. A systematic investigation of larger neural-network architecture details is left for future work.

III. RESULTS

In this section, we study dynamical properties of several observables of the QRM, focusing on the two-dimensional model. A series of benchmarks in one and two dimensions can be found in Sec. III A.

We simulate the effects instantaneous *quenches* of the coupling constant g in Eq. (1). Specifically, we initialize parameters α of the ansatz ψ_α illustrated on Fig. 1 to the ground state of the QRM Hamiltonian with $g = g_i$ using imaginary-time variational Monte Carlo (VMC) [31,43] methods. We then simulate real-time dynamics under $g = g_f$. In this work, we focus on quenches from the ordered phase to the disordered: $g_i < g_c < g_f$.

In Fig. 2, we choose a square 8×8 lattice, tracking the dynamics of the potential energy density

$$\epsilon_p(t) = -\frac{J}{N} \left\langle \sum_{\langle k,l \rangle} \hat{\mathbf{n}}_k \cdot \hat{\mathbf{n}}_l \right\rangle_t \quad (16)$$

and the average magnetization magnitude M

$$M(t) = \frac{1}{N} \left\langle \left| \sum_k \hat{\mathbf{n}}_k \right| \right\rangle_t, \quad (17)$$

along with its x , y components defined by $\mathbf{M} = N^{-1} \sum_k \langle \hat{\mathbf{n}}_k \rangle_t$. Averages $\langle \cdot \rangle_t$ are performed with respect to the ansatz state at time t . In addition, corresponding circular variances were defined as $\text{Var}(\theta_k) = -2 \ln |\langle \hat{\mathbf{n}}_k \rangle_t|$ and averaged over the lattice index k .

These observables were chosen as a proxy for thermalization. Across a wide range of quenches we observe

convergence to their respective equilibrium values at $g = g_f$, see Fig. 2. We observe two distinct dynamical regimes in relation to the quantum critical point $g_c \approx 4.25$, when $g_i < g_c$. For small quenches (left column of Fig. 2) we see the expected outcome—slower equilibration with only small fluctuations in the direction of the magnetization. However, for moderate to large quenches in Fig. 2, we observe a (transient) demagnetization of the sample and convergence to a new equilibrium state.

In addition, we define a measure of average vorticity

$$v(A) = \frac{1}{|A|} \int_A d\mathbf{a} \cdot \nabla \times \hat{\mathbf{n}} = \frac{1}{|A|} \oint_{\partial A} d\ell \cdot \hat{\mathbf{n}} \quad (18)$$

over a surface A with edge ∂A on the lattice. Using Stokes' theorem, we rewrite the expression as a contour integral over ∂A in the positive direction. On Fig. 3 (right panel), we plot $v(A)$ averaged over all n_ℓ square $\ell \times \ell$ surfaces: $v_\ell = n_\ell^{-1} \sum_{|A|=\ell^2} v(A)$. As expected, we find almost zero vorticity for quenches in the ordered phase, while larger fluctuations are generated for quenches across the critical point. We postpone a detailed analysis for future work.

Aside from local observables, such as energy and magnetization, one also has access to global observables such as the Loschmidt echo. The latter has some interesting properties in the context of dynamical phase transitions [45] and quantum chaos [46]. The Loschmidt echo expresses the quantum state overlap between the initial state and some time-evolved state. In general, the fidelity $F(\Psi, \Phi)$ between two generic normalized quantum states Ψ and Φ is defined as $F(\Psi, \Phi) = |\langle \Psi | \Phi \rangle|^2$. For real-time evolution, we expect the fidelity $F(\Psi(t=0), \Psi(t))$




FIG. 2. Results for different quenches from initial value $g_i = 3$ on a two-dimensional 8×8 square lattice. Left: potential energy, magnetization, and angular variance as functions of real time. For the small quench to $g_f = 4.5$, we observe the expected behavior with slower approach to the new ordered equilibrium state. Convergence is similar to adiabatic change. The moderate quench to $g_f = 6.0$ exhibits a sharp increase in rotor angle variance is accompanied by a single flip (right panel) in the average magnetization at $t \approx J^{-1}$. For the large quench to $g_f = 9.0$, many rotor flips occur after the first one, indicating much more detailed exploration of the underlying Hilbert space. Convergence to the new equilibrium starts taking place only for $t \gtrsim 5J^{-1}$. Right: a parametric plot of the mean rotor direction. We observe a more thorough exploration of the magnetization sphere for larger quenches.

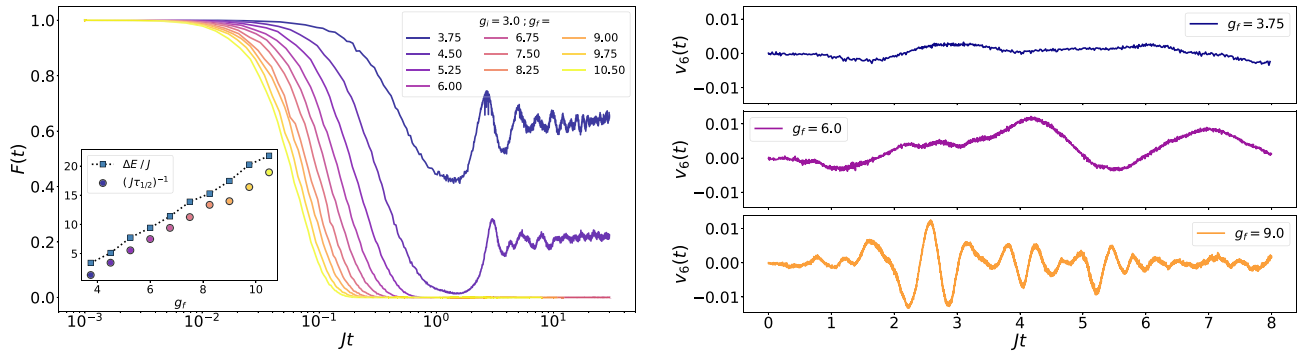


FIG. 3. Fidelity and vorticity as functions of time. Left: time-dependent many-body fidelity $F(t)$ defined in Eq. (19), for a number of quenches. For trajectories quenching to values of g_f in the same equilibrium phase, we see convergence to nonzero values at late times. Conversely, trajectories with $g_f > g_c$ converge to $F(t \rightarrow \infty) = 0$. Additionally, $\tau_{1/2}$ (the time it takes for fidelity to decrease by 50%) is shown to scale linearly with g in agreement with the appropriate uncertainty relation $\Delta E \Delta t \geq 1/2$. Right: the onset of vorticity [defined in Eq. (18)] for three quenches of increasing magnitude.

to decay as a function of time t , for any given initial state $|\Psi(t=0)\rangle$.

To evaluate this quantity using Monte Carlo sampling of unnormalized ansatz wave functions $\psi(\boldsymbol{\theta}, t) = \psi_{\alpha(t)}(\boldsymbol{\theta})$, we rewrite the fidelity definition as

$$F(t) = \left\langle \frac{\psi(\boldsymbol{\theta}, t)}{\psi(\boldsymbol{\theta}, 0)} \right\rangle_{\boldsymbol{\theta} \sim |\psi(\cdot, 0)|^2} \left\langle \frac{\psi(\boldsymbol{\theta}, 0)}{\psi(\boldsymbol{\theta}, t)} \right\rangle_{\boldsymbol{\theta} \sim |\psi(\cdot, t)|^2}, \quad (19)$$

following Refs [47,48]. The expression in Eq. (19) is manifestly independent of the normalization factor. In practice, we take the real part of Eq. (19) to discard the small nonzero imaginary part coming from finite-sample estimates of the two factors. In addition, we calculate and store both factors in log space to preserve accuracy and maintain numerical stability.

As expected, we find that the return probability (or fidelity in short) decays quickly with time, as illustrated in Fig. 3 (left panel). For smaller quenches, the fidelity shoots back up to a nonzero value suggesting a finite overlap between the initial state the long time “equilibrium” state after the quench. The latter may be interpreted as a signature of quenching between two Hamiltonians in the ordered phase.

As a measure of the fidelity decay, we introduce another time scale $\tau_{1/2}$ defined as the time needed for the fidelity to decrease by 50%. We observe that $\tau_{1/2}$ has increased linearly with the quench g_f . This result matches basic estimates given by the second-order short-time expansion of $F(t)$ and uncertainty relation $\Delta E \Delta t \geq 1/2$. Therefore, fidelity decay time can be lower bounded by ΔE^{-1} , estimated using samples from the initial state $\psi_{\alpha(0)}$ [49]. Reference points from this calculation are presented in Fig. 3 (left, inset). This comparison demonstrates that the t-MVC method can be used to estimate quantities of experimental interest for system sizes unreachable by other wave-function-based methods.

A. Benchmarks

To substantiate our results, we perform a series of benchmarks and compare results to tensor-network simulations for a one- and two-dimensional versions of the model. In particular, we benchmark the results with the time-evolving block decimation (TEBD) [50,51] algorithm. For all benchmarks, states were initialized to the coherent superposition of all basis states $|\psi(0)\rangle \propto \int d\boldsymbol{\theta} |\boldsymbol{\theta}\rangle$ by explicitly setting the final convolution kernel w_D^c [Eq. (12)] to zero. All presented tensor-network simulations have been performed with a fixed singular value cutoff. Convergence within the matrix product state (MPS) variational manifold has been confirmed by repeating simulations with larger cutoff values.

We organize numerical benchmarks as follows. First, we compare t-MVC results with TEBD for an extended one-dimensional and a smaller two-dimensional system. Practical error estimates are defined. Then, we turn to examining effects of key hyperparameters in the t-MVC approach and show evidence of self-consistent convergence.

Following Refs [17,52], we use the following figure of merit:

$$r(t) = \frac{\mathcal{D}(\psi(t + \delta t), e^{-iH\delta t}\psi(t))}{\mathcal{D}(\psi(t), e^{-iH\delta t}\psi(t))}, \quad (20)$$

where $|\psi(t)\rangle = |\psi_{\alpha(t)}\rangle$. In Eq. (20), $\mathcal{D}(\cdot, \cdot)$ represents the Fubini-Study distance on the Hilbert space \mathcal{H} . We estimate $r^2(t)$ at each time t using HMC samples from the ansatz (see Ref. [17] and Appendix D). Intuitively, $r^2(t)$ measures an appropriately normalized measure of deviation between the full state $e^{-iH\delta t} |\psi(t)\rangle$ after one time step δt and its projection onto the variational manifold $|\psi_{\alpha(t+\delta t)}\rangle$. We plot the integrated error

$$R^2(t) = \int_0^t r^2(s) ds \quad (21)$$




FIG. 4. One- and two-dimensional benchmarks and comparison with tensor-network data. Evolution was performed starting from a coherent superposition state $|\psi(0)\rangle \propto \int d\theta |\theta\rangle$. Results are compared with the TEBD tensor-network algorithm evolving a matrix product state (MPS) in the conjugate angular-momentum eigenbasis (see Sec. III A and Appendix B). Left: a one-dimensional benchmark on a chain with $N = 64$ rotors and open boundary conditions. Center: a two-dimensional benchmark of the t-MVC method on a 4×4 lattice and open boundary conditions. We note that disagreement between t-MVC and TEBD results appears as the maximum bond dimension χ_{\max} is reached. Singular value cutoff of 10^{-12} was used. Right: the growing number of MPS parameters P_{MPS} associated with the increasing bond dimension χ is plotted in units of the number of the CNN parameter count P_{CNN} as a function of time. One- and two-dimensional cases are compared. A cutoff of $\chi_{\max} = 1000$ was reached in the 2D system for the singular value cutoff of 10^{-9} .

to reflect error propagation through time as accurately as possible. We remark that the integrated-squared error in Eq. (21) should be interpreted an upper bound on the square of the integrated error $R(t) = \int_0^t r(s) ds$ due to the triangle inequality.

In Fig. 4 (left), we show that this algorithm performs well on a one-dimensional system of $N = 64$ rotors where the growth of the so-called *bond dimension* χ is limited. Convergence to appropriate equilibrium values is reached for both methods with good agreement at intermediate times for the dynamics of potential energy density $\epsilon_p(t)$ and the Loschmidt echo $F(t)$. The *integrated residual* $R^2(t)$ grows more rapidly for lower values of g . This is expected because the initial state $\psi(0)$ has lower energy for larger values of g in the QRM hamiltonian, Eq. (1), representing a more typical state in the disordered phase.

In contrast to the one-dimensional (1D) case, in Fig. 4 (center), we observe that the TEBD method exponentially grows the MPS bond dimension χ past the cutoff $\chi_{\max} = 1000$ at relatively short times. We plot the number of parameters P_{MPS} in the MPS as a function of time in the right panel of Fig. 4, in units of the number of parameters P_{CNN} in the CNN ansatz presented in this work. We see qualitative agreement between the two methods for early times, before χ grows to the point where further simulation is numerically prohibitively expensive.

In Fig. 5 we show evidence that the variance of observables is controllable through the most important Monte Carlo (HMC) hyperparameters while the bias is mostly controlled by different regularizations of the S -matrix inverse [Eq. (6)]. In the top panel of Fig. 5, we see that the standard deviation of the estimator for total magnetization $M(t)$ scales with the number of HMC samples N_s in an expected way: $\sigma_M \propto N_s^{-1/2}$ for three different times during the evolution.

In addition, we report that heuristically varying the number of leapfrog integration steps L increases estimator variances the most around segments of trajectories with higher curvature, as evidenced by the middle panel of Fig. 5. Intuitively, in the limit of $L \rightarrow 1$ and small leapfrog step sizes ε , HMC approaches random-walk Metropolis sampling (see Ref. [39] and Appendix A 1), which suffers from lower acceptance rates and longer mixing times in cases of sharply peaked target distributions. We observe that even a moderate increase to $L \approx 10$ accompanied by automatic hyperparameter tuning described in Sec. II B considerably reduces variance.

Finally, we explore the effects of S -matrix regularization [Eq. (9)]. In practice, we set λ^2 itself in an adaptive manner each iteration:

$$\lambda^2 = \lambda^2(\sigma_1^2, \dots, \sigma_P^2) = \max \left(a_c, r_c \times \max_{\mu} (\sigma_{\mu}^2) \right) \quad (22)$$

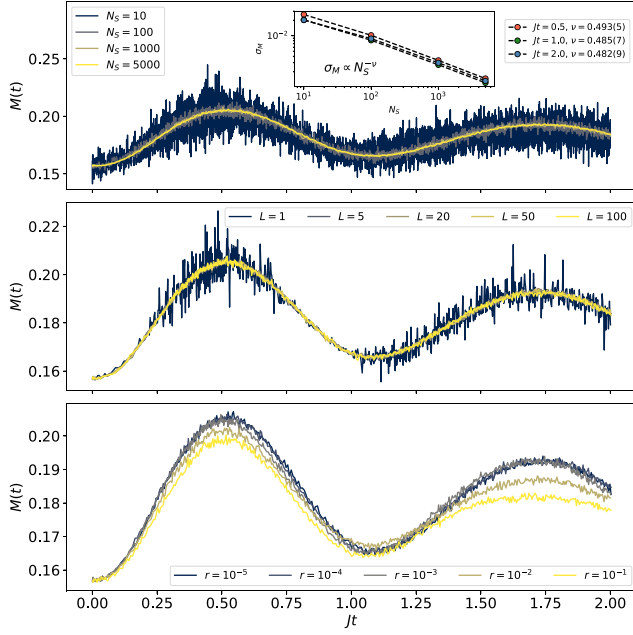


FIG. 5. Effects of key hyperparameters on magnetization measurements. All experiments were performed on a one-dimensional chain with $N = 32$. Top: effects on magnetization estimates by varying the number of HMC samples N_s . Errors were estimated using bootstrap resampling independently at different times show expected scaling $\sigma_M \propto N_s^{-1/2}$ in all cases. Middle: variance change in magnetization estimates by varying the number of leapfrog integrator steps L between HMC proposals. In the $L \rightarrow 1$ limit, HMC approaches the random-walk Metropolis sampler. Bottom: bias increase associated with changing the λ^2 cutoff parameter in Eq. (9).

depending on the S -matrix spectrum. In the bottom panel of Fig. 5, we see that, for a fixed $a_c = 10^{-5}$, increasing r_c leads to increasing the estimator bias. Excluding relevant eigenvalues from participating in time evolution through Eq. (9) can lead to a failure to capture parts of relevant physics.

Overall, both t-MVC and TEBD algorithms predict similar dynamical behavior of the potential energy density [Eq. (16)] and the fidelity [Eq. (19)], as shown on Fig. 4. However, the number of parameters in the MPS grows exponentially due to entropy buildup during time evolution. Tensor-network real-time evolution algorithms [53,54] based on MPS or two-dimensional architectures such as projected entangled pair states (PEPS) [12,55] face several challenges to extend to late times and higher dimensions. Incorporating continuous degrees of freedom exacerbates the problem—tensor network algorithms are limited to using the locally truncated eigenbasis of the angular-momentum operator L_k in the QRM Hamiltonian in Eq. (1), in contrast to the t-MVC method (see Appendix B).

IV. CONCLUSION

We present a method to approximate unitary dynamics of continuous-variable quantum many-body systems, based on custom neural-network quantum states. The approach employs Hamiltonian Monte Carlo sampling and custom regularization of the quantum geometric tensor. The method was benchmarked on quench dynamics of two-dimensional quantum rotors. We indicated that our calculations are able to access nonlocal quantities like the return probability. Good agreement was found with tensor-network-based TEBD simulations for the case of one-dimensional systems of comparable size. Finally, we showed evidence that the method is controlled by a handful of key hyperparameters. Our approach paves the way for accurate nonequilibrium simulations of continuous systems at previously unexplored system sizes and evolution times, bridging the gap between simulation and experiment.

ACKNOWLEDGMENTS

M.M. acknowledges insightful discussions with Filippo Vicentini about t-MVC regularization, Bob Carpenter about the role of circular geometry in Monte Carlo sampling and Hamiltonian Monte Carlo details. In addition, discussions with Sandro Sorella about the infinite variance problem and James Stokes about different ansatz were very helpful for fine tuning simulations. M.M. also acknowledges support from the CCQ graduate fellowship in computational quantum physics. The Flatiron Institute is a division of the Simons Foundation. D.S. was supported by AFOSR Grant No. FA9550-21-1-0236 and NSF Grant No. OAC-2118310.

SOFTWARE LIBRARIES

The code used in this work has been packaged into an installable library and is publicly available to reproduce any results in this work or explore new ones: github.com/Matematija/continuous-vmc.

It was built on JAX [56] for array manipulations, automatic differentiation for sampling and optimization and GPU support, Flax [57] for neural-network construction and manipulation and NumPy [58] and SciPy [59] for CPU array manipulations. Matplotlib [60] was used to produce figures.

APPENDIX A: SIMULATION DETAILS

In this Appendix, we mention some of the details of numerical simulations performed in this work that have not been discussed in the main text. We also clearly state different hyperparameters and their observed effect on performance and numerical stability.

1. Hamiltonian Monte Carlo details

As noted in the main text, the Hamiltonian Monte Carlo (HMC) algorithm used in this work has many important hyperparameters. To define the proposal, we must specify the following: the leapfrog integration length L , leapfrog step size ε , and the mass matrix M . We fix L heuristically and adaptively set M and ε during an extended *warmup* phase for each Markov chain independently. We assume that the mass matrix is diagonal $M = \text{diag}(m_1, \dots, m_N)$.

Before any samples are collected for evaluation of Eq. (6), each chain is run for N_w steps. Following the popular software package *Stan* [61], we subdivide the warmup period into $N_p + 2$ phases (*windows*), each of which is one of two types:

- (a) Fast: samples are collected and only step size ε is adapted using the online optimization algorithm in Ref. [40]. Mass matrix remains unchanged. Fast windows are used to efficiently *initialize* the chain by moving it towards a typical set of highly probable samples.
- (b) Slow: samples are collected and both step size ε and the mass matrix M are estimated. Step size is estimated the same way as in the fast window. Mass matrix elements are estimated as the variance of corresponding variables: $m_k = \text{Var}(\theta_k)$ using the appropriate formula for the variance of periodic random variables presented in the main text, Sec. III.

After initializing each $\theta_k \sim \text{Uniform}(-\pi, \pi)$, we begin the warmup phase with a single fast window of length $N_w/12$, followed by five fast windows. The first fast window is $N_w/36$ steps long with each subsequent slow window doubling in size. Finally, we end the warmup

by running an additional fast window for the remaining $N_w/18$ steps. After each window, the HMC transition kernel (the leapfrog ODE solver) is updated with adapted values for ε and M (for slow windows). After the final fast window, all hyperparameters are locked in and actual collection of the N_s for Eq. (6) begins. The full list of relevant hyperparameters can be found in Table I.

We use automatic differentiation (using JAX [56]) to obtain numerically exact gradients $\nabla_{\theta} \ln p(\theta, t)$ of needed to run the leapfrog integrator. To avoid loss of accuracy or numerical instabilities through exponentiation, we employ the following identity:

$$\ln p(\theta, t) = \ln |\psi_{\alpha(t)}(\theta)|^2 = 2\text{Re} \{ \ln \psi_{\alpha(t)}(\theta) \}, \quad (\text{A1})$$

when the logarithm of the wave function is parameterized instead of the wave function itself.

For completeness, we note that a common precaution against leapfrog integration getting stuck in regions of high curvature used in this work. Instead of fixing the integration length to a specific value $L = L_0$, it is randomly chosen between $(1 - \gamma)L_0$ and $(1 + \gamma)L_0$ each time the integrator is called, with a new hyperparameter $0 \leq \gamma < 1$. This *jittering* of trajectory lengths can help HMC walkers move away from regions of high curvature if they get stuck [38,39,61].

Finally, to collect more independent samples by utilizing modern massively parallel GPU hardware, we run N_c such chains in parallel, each one warmed up independently.

Finally, we note that the HMC proposal outlined in Eq. (11) approaches the RWM update:

$$\theta' = \theta + \sqrt{\Sigma} \mathbf{z}; \quad \mathbf{z} \sim \mathcal{N}(0, \mathbb{1}), \quad (\text{A2})$$

TABLE I. The list of relevant hyperparameters for the Hamiltonian Monte Carlo algorithm with their values used in this work.

Symbol	Name	Value	Domain	Description
ε	Step size	Dynamically adapted	\mathbb{R}	The leapfrog integrator step size.
M	Mass matrix	Dynamically adapted	\mathbb{R}^{N^2}	The covariance (metric) tensor of the dummy momentum variables π .
L	(Average) integration length	20	\mathbb{N}	The number of leapfrog steps taken before proposing a sample. (If $\gamma > 0$, we relabel $L \rightarrow L_0$.)
γ	Jitter	0.2	$[0, 1)$	Randomness for L during sampling—it is drawn uniformly between $(1 - \gamma)L_0$ and $(1 + \gamma)L_0$.
ε_0	Initial step size	0.1	\mathbb{R}	A guess for the value of ε to refine during the warmup phase.
δ	Target acceptance rate	0.8	$[0, 1]$	Target acceptance rate used for optimization of ε by algorithm in Ref. [40].
N_w	Length of warmup phase	800	\mathbb{N}	Total number of MC samples used for extended warmup.
N_p	Number of slow windows	5	\mathbb{N}	Total number of slow adaptation windows during warmup.
N_s	Number of samples	2000	\mathbb{N}	Total number of samples (per chain).
N_c	Number of chains	20	\mathbb{N}	Total number of independent Markov chains.

in the limit of few leapfrog integrator steps: $L \rightarrow 1$. Indeed, for $L = 1$ and small step sizes ε , Eq. (11) becomes

$$\begin{aligned} \theta' &= \theta(\varepsilon) = \theta(0) + \varepsilon M^{-1} \pi(\varepsilon/2) = \theta(0) + \varepsilon M^{-1} \pi(0) \\ &\quad - \frac{\varepsilon^2}{2} M^{-1} \frac{\partial V}{\partial \theta}(\theta(0)) \\ &= \theta(0) + \varepsilon M^{-1} \pi(0) + \mathcal{O}(\varepsilon^2), \end{aligned} \quad (\text{A3})$$

where M^{-1} is equivalent in effect to the $\sqrt{\Sigma}$ matrix and $\pi(0) \sim \mathcal{N}(0, M^{-1})$ by construction in Eq. (10).

2. Numerical regularization schemes

After evaluating the averages in Eq. (6) at time t , one needs to solve the linear system $iS\dot{\alpha} = g$ to obtain $\dot{\alpha}$ needed to progress to time $t + \delta t$. Since the S matrix is singular in most cases of interest, a robust regularization scheme is needed. As pointed out in the main text, replacing $S \rightarrow S + \varepsilon \mathbb{1}$ is often enough in the case of ground-state searches (imaginary-time evolution). We remark that this is equivalent to the L2-regularized least-squares solution of $iL\dot{\alpha} = h$.

$$\dot{\alpha} = \operatorname{argmin}_{\dot{\alpha} \in \mathbb{C}^P} \left\{ \|iL\dot{\alpha} - h\|_2^2 + \varepsilon \|\dot{\alpha}\|_2^2 \right\}, \quad (\text{A4})$$

where $L^\dagger L = S$ is the Cholesky decomposition of the S matrix (assuming S is positive definite), $L^\dagger h = g$, and $\|\cdot\|_2$ is the standard euclidean 2-norm on \mathbb{C}^P .

As outlined in the main text, we instead adopt a regularization scheme based on the spectrum of the S matrix, $S = U\Sigma U^\dagger$, where $\Sigma = \operatorname{diag}(\sigma_1^2, \dots, \sigma_P^2)$. Our definition of the pseudoinverse is $S^{-1} \approx U\tilde{\Sigma}^{-1}U^\dagger$ with

$$\tilde{\Sigma}_{\mu\nu}^{-1} = f(\sigma_\mu^2) \frac{\delta_{\mu\nu}}{\sigma_\mu^2} \quad \text{and} \quad f(\sigma^2) = \frac{1}{1 + (\lambda^2/\sigma^2)^6}. \quad (\text{A5})$$

In the limit of $\lambda^2 \rightarrow 0$, we recover the actual matrix inverse. As opposed to the more traditional choice of the step function $f(\sigma^2) = \theta(\sigma^2 - \lambda^2)$, we find that choosing

a smooth functional form for $f(\sigma^2)$ in Eq. (A5) makes the adaptive time stepping in the top-level integration routine (see Appendix A3) more stable.

As noted in the main text, we set λ^2 to

$$\lambda^2 = \lambda^2(\sigma_1^2, \dots, \sigma_P^2) = \max \left(a_c, r_c \times \max_\mu(\sigma_\mu^2) \right), \quad (\text{A6})$$

each iteration, with $a_c = 10^{-4}$ and $r_c = 10^{-2}$ chosen for 2D calculations and $a_c = 10^{-5}$ and $r_c = 10^{-4}$ for 1D benchmarks. To track potential over regularization and as a measure of ansatz expressivity, we define the *effective rank* $\rho(S) = \sum_\mu f(\sigma_\mu^2)$. Intuitively, since $0 < f(\sigma^2) < 1$ for all eigenvalues σ^2 , $\rho(S)$ can be interpreted as the effective number of eigenvalues that have not been set to zero by the regularization function f . In other words, it corresponds to the number of parameters in α that get updated at time t .

We plot $\rho(S)$ as a function of time on Fig. 6 for some simulated quenches. In all cases, we see that the effective rank increases rapidly to $\rho \sim 1$ at intermediate times that, for larger quenches, correspond to rapid oscillations and onset of vorticity. In those cases, it is natural to interpret this regime as almost all parameters α being important to capture the relevant physics. At later times, ρ converges to values below 10%, indicating equilibration and less oscillatory behavior.

For completeness, we mention that alternative regularization techniques have been explored as well. For example, the method of Schmitt and Heyl in Ref. [17], based on the SNR for each eigenvalue in σ^2 , represents a computationally and physically well-motivated approach. However, it did not bring any measurable performance improvement in our case.

3. The time-dependent variational principle and ODE integrators

To make use of the TDVP action in Eq. (5) in the main text, to propagate the variational parameters forward




FIG. 6. Some ODE integrator and regularization details. Left: the *effective rank* of the S matrix $\rho(S)$ as a function of time, reflecting the internal dimensionality of the parameter manifold α , as discussed in Appendix A2. Center: time steps δt taken by the adaptive ODE integrator of choice. Right: approximate conservation of energy as a function of time, for different quenches.

in time, one must construct the corresponding Euler-Lagrange equations. To this end, we first manipulate the action into a more transparent form:

$$\delta\mathcal{C}[\alpha] = \delta \int dt \langle \Psi_{\alpha(t)} | \left(i \frac{d}{dt} - H \right) | \Psi_{\alpha(t)} \rangle \quad (\text{A7})$$

$$= \delta \int dt \left\{ \frac{i}{2} \frac{\langle \Psi_{\alpha(t)} | \dot{\Psi}_{\alpha(t)} \rangle - \langle \dot{\Psi}_{\alpha(t)} | \Psi_{\alpha(t)} \rangle}{\langle \Psi_{\alpha(t)} | \Psi_{\alpha(t)} \rangle} - \frac{\langle \Psi_{\alpha(t)} | H | \Psi_{\alpha(t)} \rangle}{\langle \Psi_{\alpha(t)} | \Psi_{\alpha(t)} \rangle} \right\} \quad (\text{A8})$$

$$= -\delta \int dt \int d\theta \frac{|\langle \theta | \Psi_{\alpha(t)} \rangle|^2}{\langle \Psi_{\alpha(t)} | \Psi_{\alpha(t)} \rangle} \left\{ \text{Im} \sum_{\mu} \mathcal{O}_{\mu}(\theta, t) \dot{\alpha}_{\mu}(t) + E_L(\theta, t) \right\} \quad (\text{A9})$$

$$\propto \int dt \sum_{\mu} \left\{ i \sum_v S_{\mu v}(t) \dot{\alpha}_v(t) - g_{\mu}(t) \right\} \delta \alpha_{\mu}(t) - \text{c.c.}, \quad (\text{A10})$$

where we used definitions of $S_{\mu\nu}$, g_{μ} , and E_L from Eqs. (7) and (6), respectively, as well as $\dot{\alpha} = d\alpha/dt$. The explicit form of log-derivative operators \mathcal{O}_{μ} introduced in the main text is

$$\mathcal{O}_{\mu}(\theta, t) = \frac{\partial}{\partial \alpha_{\mu}} \ln \psi_{\alpha(t)}(\theta) \quad (\text{A11})$$

in the $|\theta\rangle$ basis. We note that the expression in Eq. (A9) and the definitions of S and g in Eq. (A10) change form if $\psi_{\alpha}(\theta)$ cannot be interpreted as a holomorphic function of α . The reader is referred to Ref. [22] for detailed derivations.

We implement and experiment with a number of different Runge-Kutta [37] (RK) ODE solvers. Heuristically, we notice that higher-order adaptive embedded solvers do help offset the effects of imperfect Monte Carlo estimates of S and g in two ways:

- (a) Adaptive solvers are naturally higher-order because of an embedded lower-order method. Using an adaptive solver can locally adjust the time step δt , usually drastically reducing the overall number of time steps required.
- (b) The solution at $t + \delta t$ is constructed as a linear combination of solutions estimated on a fixed set of points within the interval $[t, t + \delta t]$. Any leftover errors in these intermediate estimates have a higher probability of canceling out.

In this work, we choose the adaptive third-order method with an embedded second order method using the Bogacki-Shampine [35–37,62] pair of formulas. It balances being low-enough order to avoid wasting computational resources with still being high-enough order to allow for adaptive time-stepping.

For quenches shown in Fig. 2, we show variations in δt in Fig. 6 (center panel). We note that the time stepping varies more in regions of higher curvature (kinetic energy,

shorter times) and for larger quenches, successfully adjusting to conserve energy in all cases (Fig. 6 right). For longer times, on the order of thermalization in observed quenches, variations are reduced and δt approximately converges to a constant value.

APPENDIX B: TENSOR-NETWORK CALCULATIONS

In this Appendix we present the details of tensor-network calculations performed in this work. First, we lay out some of the conventions and formalism involved with treating the QRM in the discrete eigenbasis of the angular momentum operator L_k . This basis is useful for attacking the model with tensor-network methods or perturbation theory. Finally, we report the dependence of results reported in the main text on two different cutoffs—one in singular value magnitude and one in local basis size.

1. Angular momentum basis

In the $|\theta\rangle$ basis, we have $L_k = -i\partial_k$ (where we adopt the convention $\partial_k \equiv \partial_{\theta_k}$). Therefore,

$$-i \frac{\partial}{\partial \theta} \langle \theta | m \rangle = m \langle \theta | m \rangle \implies \langle \theta | m \rangle = \frac{e^{-im\theta}}{\sqrt{2\pi}}; \quad m \in \mathbb{Z}, \quad (\text{B1})$$

which is identical to eigenfunctions for a particle on a circle at each lattice site—we have a product basis $|\mathbf{m}\rangle = |m_1, \dots, m_N\rangle$. The Hamiltonian given in Eq. (1) then reads

$$\langle \mathbf{m}' | H | \mathbf{m} \rangle = \frac{gJ}{2} \sum_k m_k^2 - J \sum_{\langle k,l \rangle} \langle \mathbf{m}' | \hat{\mathbf{n}}_k \cdot \hat{\mathbf{n}}_l | \mathbf{m} \rangle. \quad (\text{B2})$$

After inserting the identity $\mathbb{1} = \int d\theta |\theta\rangle \langle\theta|$ into the second term and simple integration, we obtain

$$\begin{aligned} & \langle m'_k, m'_l | \hat{\mathbf{n}}_k \cdot \hat{\mathbf{n}}_l | m_k, m_l \rangle \\ &= \frac{1}{2} \left(\delta_{m'_k, m_k+1} \delta_{m'_l, m_l-1} + \delta_{m'_k, m_k-1} \delta_{m'_l, m_l+1} \right), \end{aligned} \quad (\text{B3})$$

where $\delta_{..}$ is the Kronecker δ symbol. The structure of Eq. (B3) suggests rewriting the original Hamiltonian as

$$H = \frac{gJ}{2} \sum_k L_k^2 - \frac{J}{2} \sum_{\langle k,l \rangle} (L_k^+ L_l^- + L_l^+ L_k^-), \quad (\text{B4})$$

where

$$L_k^+ \equiv \sum_{m_k \in \mathbb{Z}} |m_k + 1\rangle \langle m_k|$$

$$\text{and } L_k^- \equiv \sum_{m_k \in \mathbb{Z}} |m_k - 1\rangle \langle m_k| \quad \text{so that } [L_k^+, L_k^-] = 0, \quad (\text{B5})$$

for all $k \in \Lambda$. To perform tensor-network calculations, we truncate local basis states to $\{|-M\rangle, \dots, |M\rangle\}$ so that $L^+ |M\rangle = L^- | -M\rangle = 0$. Namely, we set,

$$\begin{aligned} L^+ &\rightarrow \begin{pmatrix} 0 & 1 & & & & \\ & 0 & 1 & & & \\ & & & \ddots & & \\ & & & & 0 & 1 \\ & & & & & & 0 & 1 \\ & & & & & & & & 0 & 1 \end{pmatrix}, \\ L^- &\rightarrow \begin{pmatrix} 0 & & & & & & & & & \\ 1 & 0 & & & & & & & & \\ & 1 & & & & & & & & \\ & & & \ddots & & & & & & \\ & & & & 0 & & & & & \\ & & & & & & 1 & & & \\ & & & & & & & & 1 & 0 \end{pmatrix} \end{aligned} \quad (\text{B6})$$

to build a matrix product operator (MPO) representation of the Hamiltonian. We use $M = 5$ throughout. For use in real-time evolution through the TEBD algorithm in Sec. III A in the main text, we initialize the trial wave function as a matrix product state (MPS) [63] (see Appendix C 2). We exploit

$$|\psi(t=0)\rangle \propto \int d\theta |\theta\rangle = \prod_{i \in \Lambda} |m_i = 0\rangle \quad (\text{B7})$$

to match the initial state given in the main text in Sec. III A by initializing the corresponding MPS to have bond dimension $\chi = 1$.

2. Convergence and cutoffs

To reinforce TEBD results as a benchmark in two dimensions, we further study the dependence of obtained results on the singular value cutoff c and the local basis truncation parameter M defined in Appendix B 1. In both cases, we perform a series of calculations on a 4×4 system for a set of different cutoffs and for two different values of the coupling constant g in the QRM Hamiltonian [Eq. (1)].

In the case of the singular value cutoff c , we look at a range of values between 10^{-4} and 10^{-11} in Fig. 7. We observe that plotted curves seem to converge only for $c \leq 10^{-9}$. Satisfactory energy conservation is reached for $c \approx 10^{-9}$ as well as convergence of measured observables, except for the highest value $c = 10^{-4}$. In that case, almost all singular values are discarded, extinguishing any nontrivial dynamics. Any further decrease in c has the undesirable side-effect of rapidly growing the MPS bond dimension χ , making longer simulations prohibitively expensive.

We explore a range of values for the local basis truncation parameter M as well (see Fig. 8). For both

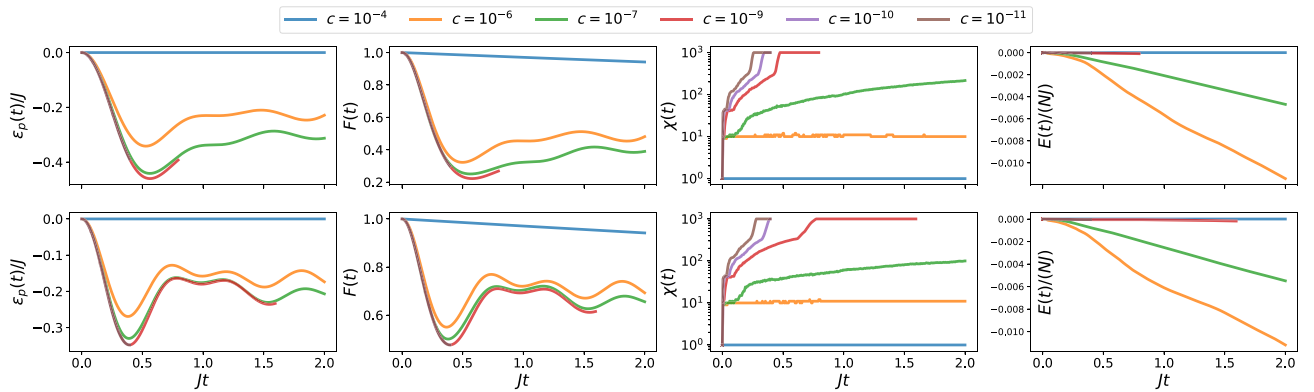


FIG. 7. TEBD results for a range of different values of singular value cutoffs c . We plot time evolution of potential energy density ϵ_p , fidelity F as well as bond dimension χ growth. Additionally, we show energy conservation (or lack thereof). The initial state is set to the one described in Eq. (B7) in all cases. Two independent sets of calculations were performed under the QRM Hamiltonian [Eq. (1)] with $g = 5.6$ (top row) and $g = 8.0$ (bottom row).

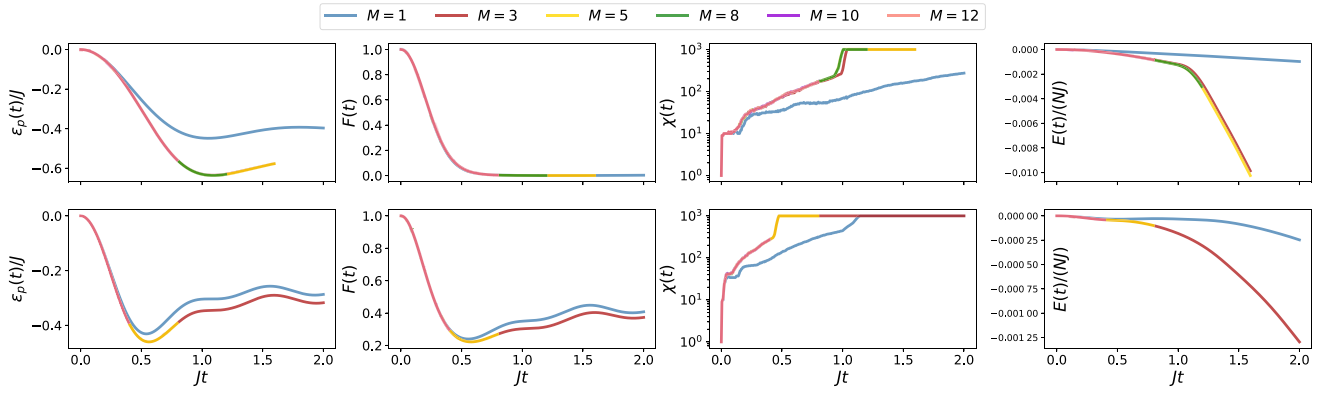


FIG. 8. TEBD results for a range of different choices of truncated local bases (M). We plot the same observables for the same initial state as in Fig. 7. Similarly, two independent sets of calculations were performed with $g = 2.0$ and $c = 10^{-7}$ (top row) as well as $g = 5.6$ and $c = 10^{-9}$ (bottom row).

reference values of the coupling constant g , we observe fast convergence towards self-consistent time evolution. These results indicate lower sensitivity to values of M , as long as they lie above a threshold of $M \gtrsim 3$ ($M = 5$ was used in the main text). Therefore, for short to intermediate time evolution, $Jt \sim 1 - 10$, we expect the numerical cost of increased bond dimensions to remain dominated by singular value cutoff c .

In all cases, bond dimensions χ increase quickly and saturate at the cutoff value χ_{\max} , as is characteristic of MPS-based calculations in two dimensions. Therefore, to estimate any additional errors coming from the choice of χ_{\max} (set to 10^3 in the main text), we perform an additional set of independent calculations for a range of values, at $g = 8.0$. Results are presented in Fig. 9 where we plot estimated errors in reference observables. We discover that the relevant errors can be neglected for values as low as $\chi_{\max} \approx 600$, at least for intermediate times. Reference observables quickly converge to the desired precision in this case.

APPENDIX C: VARIATIONAL QUANTUM STATES ON A CIRCLE

In order to perform numerically efficient t-MVC iterations described in Sec. II A, we need to keep the

parameter count in ψ_α relatively low (order $10^2 - 10^4$ on modern GPU hardware). The reason for this constraint is the regularization scheme we employ to stabilize the QGT inverse in Eq. (6)—each iteration requires us to diagonalize the $P \times P$ Hermitian matrix S , which becomes prohibitively expensive and memory consuming for large parameter counts P . We note that iterative solvers such as the *conjugate gradient* can formally help push the limit of tractable P by several orders of magnitude. However, that speedup comes at the cost of having to rely on weaker regularization schemes that do not require the full QGT spectrum. In this Appendix, we describe different ansatzes (trial wave functions) considered in our simulations.

1. The Jastrow wave function

The Jastrow wave function is defined as

$$\ln \psi_\alpha(\theta) = \frac{1}{2} \sum_{ij} w_{ij} \hat{\mathbf{n}}_i \cdot \hat{\mathbf{n}}_j = \frac{1}{2} \sum_{ij} w_{ij} \cos(\theta_i - \theta_j), \quad (\text{C1})$$

where $\alpha = \{w_{ij}\}$ is a symmetric matrix of variational parameters. The advantage of a simple Jastrow ansatz is the fact that the QGT given in Eq. (6) is never




FIG. 9. Convergence of TEBD calculations with respect to the cutoff bond dimension value χ_{\max} . We perform independent calculations for a fixed set of times, plot reference observable values obtained at different saturated bond dimensions. Plotted errors are defined as the difference between the value at bond dimension χ and the value at bond dimension $\chi = 1000$: $\delta A = A(\chi) - A(\chi = 1000)$ where $A = \epsilon_p$ in the left panel and $A = F$ in the middle panel.

ill conditioned. However, Jastrow expressivity is limited compared to deeper neural-network quantum state parametrizations.

2. Matrix product states

Using notation and conventions laid out in Appendix B, one can write down a traditional MPS ansatz in the discrete angular momentum basis $|\mathbf{m}\rangle$:

$$\begin{aligned} \psi_\alpha(\boldsymbol{\theta}) &= \sum_{m_1=-M}^M \cdots \sum_{m_N=-M}^M c_{m_1 \cdots m_N} e^{-i \sum_i m_i \theta_i}; \\ c_{m_1 \cdots m_N} &= \sum_{l_1=1}^{\chi_1} \cdots \sum_{l_N=1}^{\chi_N} A_{l_1}^{m_1} A_{l_1 l_2}^{m_2} \cdots A_{l_{N-2} l_{N-1}}^{m_{N-1}} A_{l_{N-1}}^{m_N}. \end{aligned} \quad (\text{C2})$$

In one spatial dimension and for ground-state searches (imaginary time evolution), this trial wave-function form is the most accurate due to the density matrix renormalization group algorithm (DMRG) [63] algorithm. Expressivity is controlled by *bond dimensions* χ_i and the basis truncation parameter M .

3. The (circular) restricted Boltzmann machine

The circular restricted Boltzmann machine (RBM) [28] is defined as

$$\begin{aligned} \psi_\alpha(\boldsymbol{\theta}) \propto \int d\mu(\hat{\mathbf{h}}) \exp \left\{ \sum_j \mathbf{a}_j \cdot \hat{\mathbf{n}}_j + \sum_k \mathbf{b}_k \cdot \hat{\mathbf{h}}_k \right. \\ \left. + \sum_{jk} w_{jk} \hat{\mathbf{n}}_j \cdot \hat{\mathbf{h}}_k \right\} \end{aligned} \quad (\text{C3})$$

where $\alpha = \{\mathbf{a}_j, \mathbf{b}_k, w_{jk}\}$ are variational parameters and $d\mu(\hat{\mathbf{h}})$ is the relevant measure for *hidden units* $\hat{\mathbf{h}}_k$. It is natural to choose hidden units to have the same intrinsic Hilbert space as visible rotors $\hat{\mathbf{n}}_j$. Therefore, for the $O(2)$ quantum rotor model, we choose

$$\hat{\mathbf{h}}_k = (\cos \phi_k, \sin \phi_k)$$

$$\text{so that } d\mu(\hat{\mathbf{h}}) = d\boldsymbol{\phi} = d\phi_1 d\phi_2 \dots d\phi_{N_h}, \quad (\text{C4})$$

up to an overall multiplicative constant. We note that the number hidden units N_h is a hyperparameter and can be increased to control ansatz expressivity.

After performing the integrals in Eq. (C3), one obtains the following closed-form expression:

$$\ln \psi_\alpha(\boldsymbol{\theta}) = \sum_{j=1}^N \mathbf{a}_j \cdot \hat{\mathbf{n}}_j + \sum_{k=1}^{N_h} \ln I_0 \left(\sqrt{\sum_l (\mathbf{x}_k)_l^2} \right), \quad (\text{C5})$$

where $(\mathbf{x}_k)_l$ stands for the l th component of vector $\mathbf{x}_k = \mathbf{b}_k + \sum_j w_{jk} \hat{\mathbf{n}}_j$ and I_0 is the zeroth-order modified Bessel function of the first kind.

Instead of a full dense matrix, we can restrict the general linear map $\hat{\mathbf{n}}_j \mapsto \sum_j w_{jk} \hat{\mathbf{n}}_j$ to a convolution, assuming that underlying rotors are arranged in a square lattice. This restriction cuts P down by approximately an order of magnitude while not sacrificing any measurable accuracy in ground-state optimization tasks.

4. Activation functions

In order to define an analytic ansatz ψ_α with no hidden singularities, care must be taken when choosing activation functions for complex-valued inputs. Informally, singularities often appear in one of the following two ways, when using holomorphic activations:

- A well-behaved function (or its derivatives) on the real axis has singularities on the imaginary axis. This is the case for \tanh and $\ln I_0$ from Eq. (C3), for example.
- An otherwise well-behaved function has a branch cut that is crossed during time evolution. Side effects include sudden jumps in conserved quantities during real-time evolution. This is the case for $\ln I_0$ from Eq. (C3) and similar functions involving logarithms and/or roots.

There are two solutions to this problem. As noted in the main text, one can restrict themselves to (higher-order) polynomial activations, which are analytic everywhere and have no branch cuts. Inspired by Eq. (C3), we use Taylor expansions of $\ln I_0$ and its gradient:

$$\begin{aligned} \ln I_0(z) &= \frac{z^2}{4} - \frac{z^4}{64} + \frac{z^6}{576} + \mathcal{O}(z^8) \quad \text{and} \\ \frac{I_1(z)}{I_0(z)} &= \frac{z}{2} - \frac{z^3}{16} + \frac{z^5}{96} + \mathcal{O}(z^7). \end{aligned} \quad (\text{C6})$$

This approach has the advantage of maintaining the holomorphic dependence of ψ_α on α and preserving the form of Eq. (6). We note that if the effect of isolated singularities is not as important, Padé approximants often provide better approximations of target functions (and better ground-state energies) while still eliminating branch cuts.

The second option is abandoning holomorphicity in parameters α and applying well-behaved real activations to real and imaginary parts of the input separately. In that case, Eq. (6) must be corrected. We refer interested readers to the excellent overview of subtleties associated with complex parameters in Ref. [22].

APPENDIX D: THE R^2 PERFORMANCE METRIC

In the main text, we define a measure of time-dependent integration error in Eq (20), following Refs [17,52]. The Fubini-Study distance $\mathcal{D}(\cdot, \cdot)$ on the Hilbert space \mathcal{H} defined as

$$\mathcal{D}(\psi, \phi) = \cos^{-1}\left(\sqrt{F(\psi, \phi)}\right) = \cos^{-1}\left(\sqrt{\frac{|\langle \psi | \phi \rangle|^2}{\langle \psi | \psi \rangle \langle \phi | \phi \rangle}}\right). \quad (\text{D1})$$

Using a consistent Taylor expansion in the limit of $\delta t \ll J^{-1}$,

$$e^{-iH\delta t} = \mathbb{1} - iH\delta t + \mathcal{O}(\delta t^2) \quad (\text{D2})$$

$$|\psi(t + \delta t)\rangle = \left(1 + \delta t \sum_{\mu} \dot{\alpha}_{\mu} \mathcal{O}_{\mu}\right) |\psi_{\alpha(t)}\rangle + \mathcal{O}(\delta t^2), \quad (\text{D3})$$

authors in Ref. [17] rewrite Eq. (20) as

$$r^2(t) = 1 - \frac{1}{\text{Var}_t H} \sum_{\mu\nu} S_{\mu\nu}^{-1} g_{\mu}^* g_{\nu} + \mathcal{O}(\delta t^2). \quad (\text{D4})$$

In Eq. (D4), we used notation from Eq. (6) in the main text with

$$\text{Var}_t H = \langle H^2 \rangle_t - \langle H \rangle_t^2 \approx \langle |E_L - \langle E_L \rangle|^2 \rangle_t. \quad (\text{D5})$$

Finally, the R^2 figure of merit is constructed as a time integral of r^2 : $R^2(t) = \int_0^t r^2(s) ds$. We note that $r^2(t)$ and $R^2(t)$ are readily available for estimation through Monte Carlo sampling.

-
- [1] W. S. Warren, H. Rabitz, and M. Dahleh, Coherent control of quantum dynamics: The dream is alive, *Science* **259**, 1581 (1993).
 - [2] A. Polkovnikov, K. Sengupta, A. Silva, and M. Vengalattore, Colloquium: Nonequilibrium dynamics of closed interacting quantum systems, *Rev. Mod. Phys.* **83**, 863 (2011).
 - [3] M. Budden, T. Gebert, M. Buzzi, G. Jotzu, E. Wang, T. Matsuyama, G. Meier, Y. Laplace, D. Pontiroli, M. Riccò, F. Schlawin, D. Jaksch, and A. Cavalleri, Evidence for metastable photo-induced superconductivity in K_3C_{60} , *Nat. Phys.* **17**, 611 (2021).
 - [4] A. Thomas, L. Lethuillier-Karl, K. Nagarajan, R. M. A. Vergauwe, J. George, T. Chervy, A. Shalabney, E. Devaux, C. Genet, J. Moran, and T. W. Ebbesen, Tilting a ground-state reactivity landscape by vibrational strong coupling, *Science* **363**, 615 (2019).
 - [5] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe,

Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator, *Nature* **551**, 601 (2017).

- [6] V. V. Sivak, A. Eickbusch, H. Liu, B. Royer, I. Tsioutsios, and M. H. Devoret, Model-free quantum control with reinforcement learning, *Phys. Rev. X* **12**, 011059 (2022).
- [7] R. Porotti, A. Essig, B. Huard, and F. Marquardt, Deep reinforcement learning for quantum state preparation with weak nonlinear measurements, *Quantum* **6**, 747 (2022).
- [8] F. Metz and M. Bukov, Self-correcting quantum many-body control using reinforcement learning with tensor networks, *ArXiv:2201.11790* (2022).
- [9] M. Bukov, A. G. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, Reinforcement learning in different phases of quantum control, *Phys. Rev. X* **8**, 031086 (2018).
- [10] M. P. Zaletel, R. S. K. Mong, C. Karrasch, J. E. Moore, and F. Pollmann, Time-evolving a matrix product state with long-ranged interactions, *Phys. Rev. B* **91**, 165112 (2015).
- [11] J. Wurtz, A. Polkovnikov, and D. Sels, Cluster truncated Wigner approximation in strongly interacting systems, *Ann. Phys. (N. Y.)* **395**, 341 (2018).
- [12] P. Czarnik, J. Dziarmaga, and P. Corboz, Time evolution of an infinite projected entangled pair state: An efficient algorithm, *Phys. Rev. B* **99**, 035115 (2019).
- [13] C. Hubig, A. Bohrdt, M. Knap, F. Grusdt, and I. Cirac, Evaluation of time-dependent correlators after a local quench in iPEPS: Hole motion in the t-j model, *SciPost Phys.* **8**, 021 (2020).
- [14] Y. Zhou, E. M. Stoudenmire, and X. Waintal, What limits the simulation of quantum computers?, *Phys. Rev. X* **10**, 041038 (2020).
- [15] G. Carleo, F. Becca, M. Schiró, and M. Fabrizio, Localization and glassy dynamics of many-body quantum systems, *Sci. Rep.* **2**, 243 (2012).
- [16] G. Carleo, L. Cevolani, L. Sanchez-Palencia, and M. Holzmann, Unitary dynamics of strongly interacting Bose gases with the time-dependent variational Monte Carlo method in continuous space, *Phys. Rev. X* **7**, 031026 (2017).
- [17] M. Schmitt and M. Heyl, Quantum many-body dynamics in two dimensions with artificial neural networks, *Phys. Rev. Lett.* **125**, 100503 (2020).
- [18] D. Hofmann, G. Fabiani, J. Mentink, G. Carleo, and M. Sentef, Role of stochastic noise and generalization error in the time propagation of neural-network quantum states, *SciPost Phys.* **12**, 165 (2022).
- [19] S. Barison, F. Vicentini, and G. Carleo, An efficient quantum algorithm for the time evolution of parameterized circuits, *Quantum* **5**, 512 (2021).
- [20] S. Czischek, M. Gärtner, and T. Gasenzer, Quenches near Ising quantum criticality as a challenge for artificial neural networks, *Phys. Rev. B* **98**, 024311 (2018).
- [21] S. I. Amari, Natural gradient works efficiently in learning, *Neural. Comput.* **10**, 251 (1998).
- [22] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, Theory of variational quantum simulation, *Quantum* **3**, 191 (2019).
- [23] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, Quantum natural gradient, *Quantum* **4**, 269 (2020).

- [24] N. Vogt, R. Schäfer, H. Rotzinger, W. Cui, A. Fiebig, A. Shnirman, and A. V. Ustinov, One-dimensional Josephson junction arrays: Lifting the Coulomb blockade by depinning, *Phys. Rev. B* **92**, 045435 (2015).
- [25] P. Martinoli and C. Leemann, Two dimensional Josephson junction arrays, *J. Low. Temp. Phys.* **118**, 699 (2000).
- [26] A. F. Kockum and F. Nori, Quantum bits with Josephson junctions, *Springer Ser. Mater. Sci.* **286**, 703 (2019).
- [27] J. V. José, L. P. Kadanoff, S. Kirkpatrick, and D. R. Nelson, Renormalization, vortices, and symmetry-breaking perturbations in the two-dimensional planar model, *Phys. Rev. B* **16**, 1217 (1977).
- [28] J. Stokes, S. De, S. Veerapaneni, and G. Carleo, Continuous-variable neural-network quantum states and the quantum rotor model, *ArXiv:2107.07105* (2021).
- [29] W. Jiang, G. Pan, Y. Liu, and Z. Y. Meng, Solving quantum rotor model with different Monte Carlo techniques, *Chin. Phys. B* **31**, 00 (2019).
- [30] C. Berke, E. Varvelis, S. Trebst, A. Altland, and D. P. DiVincenzo, Transmon platform for quantum computing challenged by chaotic fluctuations, *Nat. Commun.* **13**, 2495 (2022).
- [31] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, Cambridge, United Kingdom, 2017).
- [32] S. Sorella, Green function Monte Carlo with stochastic reconfiguration, *Phys. Rev. Lett.* **80**, 4558 (1998).
- [33] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.* **21**, 1087 (1953).
- [34] W. K. Hastings, Monte Carlo sampling methods using Markov chains and their applications, *Biometrika* **57**, 97 (1970).
- [35] P. Bogacki and L. Shampine, A 3(2) pair of Runge-Kutta formulas, *Appl. Math. Lett.* **2**, 321 (1989).
- [36] J. C. Butcher, *Numerical Methods for Ordinary Differential Equations* (John Wiley & Sons, Ltd, Chichester, United Kingdom, 2008).
- [37] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C* (Cambridge University Press, Cambridge, USA, 1992), 2nd ed.
- [38] R. M. Neal, *Handbook of Markov Chain Monte Carlo* (Chapman and Hall/CRC, New York, USA, 2011).
- [39] M. Betancourt, A conceptual introduction to Hamiltonian Monte Carlo, *ArXiv:1701.02434* (2017).
- [40] Y. Nesterov, Primal-dual subgradient methods for convex problems, *Math. Program.* **120**, 221 (2009).
- [41] M. D. Hoffman and A. Gelman, The no-u-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo, *J. Mach. Learn. Res.* **15**, 1593 (2011).
- [42] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature* **521**, 436 (2015).
- [43] 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**, 045002 (2019).
- [44] G. Pescia, J. Han, A. Lovato, J. Lu, and G. Carleo, Neural-network quantum states for periodic systems in continuous space, *Phys. Rev. Res.* **4**, 023138 (2022).
- [45] M. Heyl, A. Polkovnikov, and S. Kehrein, Dynamical quantum phase transitions in the transverse-field Ising Model, *Phys. Rev. Lett.* **110**, 135704 (2013).
- [46] M. Schmitt, D. Sels, S. Kehrein, and A. Polkovnikov, Semiclassical echo dynamics in the Sachdev-Ye-Kitaev model, *Phys. Rev. B* **99**, 134301 (2019).
- [47] M. Medvidović and G. Carleo, Classical variational simulation of the quantum approximate optimization algorithm, *npj Quantum Inf.* **7**, 101 (2021).
- [48] B. Jónsson, B. Bauer, and G. Carleo, Neural-network states for the classical simulation of quantum computing, *ArXiv:1808.05232* (2018).
- [49] L. Mandelstam and I. Tamm, *The Uncertainty Relation Between Energy and Time in Non-Relativistic Quantum Mechanics* (Springer Berlin Heidelberg, Berlin, Heidelberg, 1991), p. 115.
- [50] G. Vidal, Efficient classical simulation of slightly entangled quantum computations, *Phys. Rev. Lett.* **91**, 147902 (2003).
- [51] G. Vidal, Efficient simulation of one-dimensional quantum many-body systems, *Phys. Rev. Lett.* **93**, 040502 (2004).
- [52] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [53] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Time-dependent variational principle for quantum lattices, *Phys. Rev. Lett.* **107**, 070601 (2011).
- [54] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Unifying time evolution and optimization with matrix product states, *Phys. Rev. B* **94**, 165116 (2016).
- [55] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions, *ArXiv:cond-mat/0407066* (2004).
- [56] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of Python+NumPy programs, <http://github.com/google/jax> (2018).
- [57] J. Heek, A. Levskaya, A. Oliver, M. Ritter, B. Rondepierre, A. Steiner, and M. van Zee, Flax: A neural network library and ecosystem for JAX, <http://github.com/google/flax> (2020).
- [58] C. R. Harris, *et al.*, Array programming with NumPy, *Nature* **585**, 357 (2020).
- [59] P. Virtanen, *et al.*, Scipy 1.0: fundamental algorithms for scientific computing in Python, *Nat. Methods.* **17**, 261 (2020).
- [60] J. D. Hunter, Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 99 (2007).
- [61] B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. A. Brubaker, J. Guo, P. Li, and A. Riddell, Stan: A probabilistic programming language, *J. Stat. Softw.* **76**, 1 (2017).
- [62] J. C. Butcher, Coefficients for the study of Runge-Kutta integration processes, *J. Aust. Math. Soc.* **3**, 185 (1963).
- [63] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).