Variational quantum dynamics of two-dimensional rotor models

Matija Medvidović
Department of Physics, Columbia University
Center for Computational Quantum Physics, Flatiron Institute
matija.medvidovic@columbia.edu

Dries Sels
Center for Quantum Phenomena, Department of Physics, New York University
Center for Computational Quantum Physics, Flatiron Institute
dsels@flatironinstitute.org

Abstract

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 systems with 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 to perform previously unexplored non-equilibrium simulations bridging the gap between simulation and experiment.

1 Introduction

Non-equilibrium quantum many-body physics has been at the center of physics and chemistry research for over a decade [1, 2]. The field is driven by remarkable progress in our ability to control matter at the atomic scale [3–5]. Such quantum control [6–9] of modern experiments and hardware is becoming increasingly limited by numerical simulation of the real-time evolution of quantum systems. Fast entanglement growth out of equilibrium forces one to keep track of many-body correlations. Challenges remain despite recent progress [10–19], especially in cases of realistic systems with continuous-variables. Practically, calculations are roadblocked by numerical instabilities resulting from a combination of Monte Carlo noise and flatness of the quantum geometry of modern neural network wave functions [20, 17, 21–23].

In this work, we present an approach for capturing long-time dynamics of continuous-variable 2D lattice models, using a combination of methods that were previously unexplored in the field – the Hamiltonian Monte Carlo sampler, a tailored variational ansatz and novel regularization of 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.

2 Model and Methods

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

Machine Learning and the Physical Sciences Workshop, NeurIPS 2023.

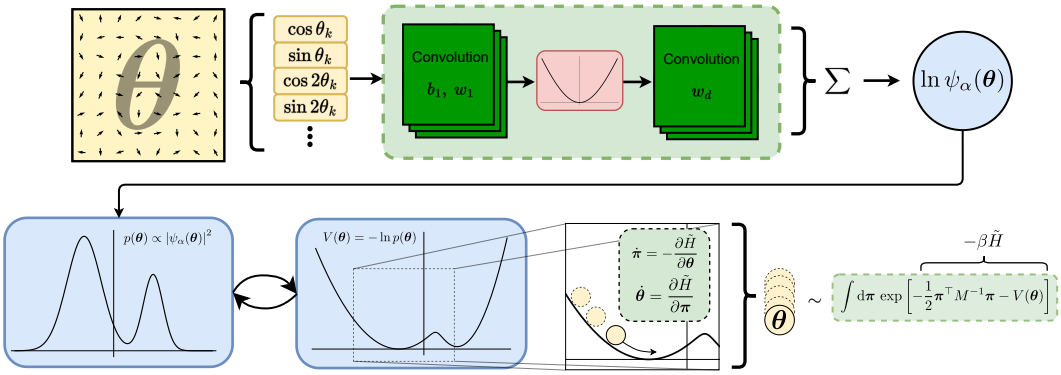

Figure 1: **Top**: The (complex) two-layer convolutional neural network $\psi_\alpha(\boldsymbol{\theta})$ used for simulations of two-dimensional QRM systems. The output is a single complex number we interpret as $\ln\psi_\alpha(\boldsymbol{\theta})$. **Bottom**: The Hamiltonian Monte Carlo algorithm. Samples are collected as snapshots of solutions of Hamilton's equations.

that angles θ_k represent superconducting phases of adjacent Josephson junctions:

$$H = -\frac{gJ}{2}\sum_k \frac{\partial^2}{\partial\theta_k^2} - J\sum_{\langle k,l\rangle}\cos(\theta_k - \theta_l)\ , \tag{1}$$

where $L_k = -i\,\partial_{\theta_k}$ is the 2D angular momentum operator and $\hat{\mathbf{n}}_k = (\cos\theta_k, \sin\theta_k)$ are the rotors in the continuous basis $|\boldsymbol{\theta}\rangle$ of choice. The Hamiltonian in Eq. 1 is defines the quantum rotor model (QRM). Its equilibrium properties [27] have been studied using Quantum Monte Carlo (QMC) [28, 29] methods. However, QRM real-time dynamics have not been explored despite prospects of practical applications in the study of dynamics of arrays of coupled Josephson junctions [30]. Direct solution of the relevant time-dependent Schrödinger equation (TDSE) $i\,\partial_t\psi = H\,\psi$ is prohibitively expensive even for a handful of interacting rotors. The continuous nature of the $|\boldsymbol{\theta}\rangle$ basis exacerbates the problem.

We represent a quantum state using a convolutional neural network (CNN) wavefunction $\psi_\alpha(\boldsymbol{\theta})$ where $\alpha \in \mathbb{C}^P$ is a set of P complex variational parameters. The full neural-network quantum state (NQS) then reads $|\psi_\alpha\rangle = \int \mathrm{d}\boldsymbol{\theta}\ \psi_\alpha(\boldsymbol{\theta})\,|\boldsymbol{\theta}\rangle$ where $\mathrm{d}\boldsymbol{\theta} \equiv \mathrm{d}\theta_1\cdots\mathrm{d}\theta_N$. The integral is performed over the cube $[-\pi,\pi]^N$.

Our simulation of the real-time dynamics of NQS is based on the time-dependent Variational Monte Carlo (t-VMC) method [31, 16, 15]. The core assumption that allows us to approximately solve the TDSE is that of time dependence of neural-network parameters $\alpha = \alpha(t)$. Trajectories $\alpha(t)$ can be derived by substituting $\partial_t\psi_\alpha = \dot{\alpha}\cdot\nabla_\alpha\psi_\alpha\cdot$ into the TDSE. Evolution equations [22] then read $i\,S\dot{\alpha} = g$, where

$$S_{\mu\nu} = \langle\mathcal{O}_\mu^\dagger\mathcal{O}_\nu\rangle - \langle\mathcal{O}_\mu^\dagger\rangle\langle\mathcal{O}_\nu\rangle \quad\text{and}\quad g_\mu = \langle\mathcal{O}_\mu^\dagger H\rangle - \langle\mathcal{O}_\mu^\dagger\rangle\langle H\rangle\ . \tag{2}$$

Averages $\langle\cdot\rangle \equiv {}^{\langle\psi_\alpha|\cdot|\psi_\alpha\rangle}\!/_{\langle\psi_\alpha|\psi_\alpha\rangle}$ are performed at time t (i.e. for $\alpha = \alpha(t)$). Parameters α are indexed by $\{\mu,\nu,\ldots\}$ in our notation. Operator $\mathcal{O}_\mu$ is defined by $\partial_{\alpha_\mu}|\psi_\alpha\rangle = \mathcal{O}_\mu\,|\psi_\alpha\rangle$. The matrix S is commonly called the *quantum geometric tensor* (QGT) [32, 21, 23] and corresponds to the metric tensor of the neural-network parameter manifold.

Since quantum averages over an exponentially large Hilbert space $\mathcal{H}$ in Eq. 2 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. 2 using the *local operator* trick. We refer the interested reader to an excellent summary in Refs. [22, 31].

$$\langle H\rangle = \frac{\langle\psi_\alpha|\,H\,|\psi_\alpha\rangle}{\langle\psi_\alpha|\psi_\alpha\rangle} = \frac{\int \mathrm{d}\boldsymbol{\theta}\,\psi_\alpha^*(\boldsymbol{\theta})H\psi_\alpha(\boldsymbol{\theta})}{\int \mathrm{d}\boldsymbol{\theta}\,|\psi_\alpha(\boldsymbol{\theta})|^2} = \mathop{\mathbb{E}}_{\boldsymbol{\theta}\sim|\psi_\alpha|^2}\left[\frac{H\psi_\alpha(\boldsymbol{\theta})}{\psi_\alpha(\boldsymbol{\theta})}\right]\ . \tag{3}$$

After computing the matrix S and the vector g at time t, one can formally define $\dot{\alpha} = -i\,S^{-1}g$ and use any ordinary differential equation (ODE) integrator to obtain the next set of parameters at time

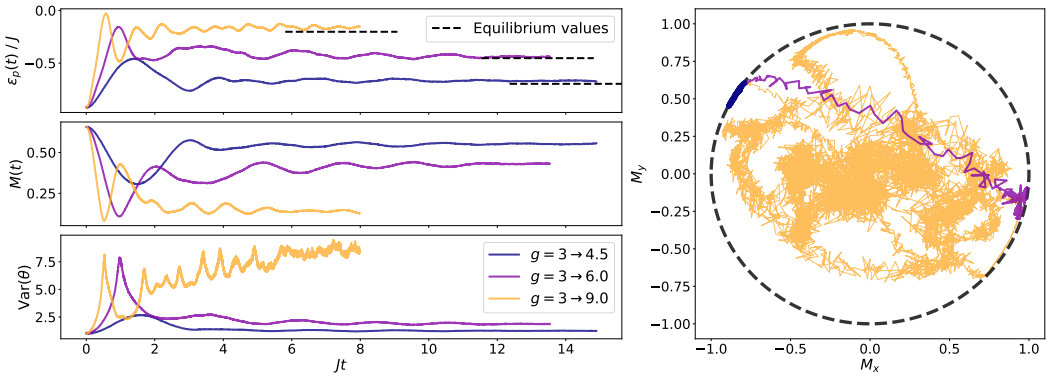

Figure 2: Results for different quenches from initial value $g_i = 3$ on a two-dimensional 8×8 square lattice. **Left**: Potential energy, magnetization and angular variance as functions of real time. **Right**: A parametric plot of the mean rotor direction, parametrized by the real time t.

$t + \delta t$. This update rule strongly resembles the quantum natural gradient optimization scheme [21, 23]. We use the Bogacki-Shampine Runge Kutta 3(2) solver. [35–37].

However, the inverse S^{-1} is often ill-defined. One reason is noisy Monte Carlo estimates of matrix elements make small eigenvalues vanish. Therefore, quickly and efficiently obtaining many uncorrelated samples from $|\psi_{\alpha(t)}(\boldsymbol{\theta})|^2$ is crucial. In addition, overparametrizing ψ_α introduces redundancy, producing linearly dependent or vanishing rows/columns in S. Therefore, choosing an smaller CNN trial wavefunction is equally important. In practice, adding more parameters to the wavefunction can sometimes unexpectedly reduce accuracy by making S ill-conditioned.

We introduce a heuristic regularization (pseudoinverse) that allows us to propagate the TDSE longer in time. The S matrix is diagonalized $S = U\Sigma U^\dagger$ at each time step. Having obtained eigenvalues σ_μ^2 such that $\Sigma = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_P^2)$, we define the pseudoinverse as $S^{-1} \approx U\tilde{\Sigma}^{-1}U^\dagger$ with $\tilde{\Sigma}^{-1}_{\mu\nu} = \frac{1/\sigma_\mu^2}{1+\left(\lambda^2/\sigma_\mu^2\right)^6}\,\delta_{\mu\nu}$. We heuristically find that such smooth cutoff is superior to traditional pseudoinverses when using adaptive integrators for updating parameters α.

We employ Hamiltonian Monte Carlo (HMC) [38, 39] to estimate Hilbert space averages in Eq. 2 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. This results in a Markov chain with considerably lower autocorrelation times.

For a generic probability distribution $p(\boldsymbol{\theta})$, HMC augments the configuration space with artificial momentum variables $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_N)$. Hamiltonian dynamics in the resulting $2N$-dimensional space conserves energy (probability). Monte Carlo updates can be defined through numerical integration of Hamilton's equations. Given initial conditions $\boldsymbol{\theta}(0)$, $\boldsymbol{\pi}(0)$ and a small step size, a common choice is the symplectic leapfrog integrator [38, 37] because it conserves energy (probability) exactly, allowing for large jumps in the $\boldsymbol{\theta}$-space while keeping high acceptance probabilities. Randomness is injected by sampling the normal distribution $\boldsymbol{\pi}(0) \sim \mathcal{N}(0, M)$. Due to the space limitation, we refer interested readers to the excellent review of HMC in Ref. [39].

We use the CNN architecture [40, 41] to model $\psi_\alpha(\boldsymbol{\theta})$, similar to Refs. [17, 42] and employ automatic differentiation (AD) techniques to obtain all derivatives $\mathcal{O}_\mu$ in Eqs. 2. We concatenate $h_0 = \{(\cos n\theta_k, \sin n\theta_k) \mid n = 1, \ldots, K\}$ along the input channel axis (Fig. 1) to enforce periodicity in θ. We choose a simple 2-layer CNN model to control the number of parameters P. In addition nontrivially affecting the QGT inverse, the diagonalization cost grows as $\mathcal{O}(P^3)$. Heuristically, introducing more parameters α requires more Monte Carlo samples to correctly resolve the averages in Eq. 2 and does not significantly contribute to simulation accuracy in our case.

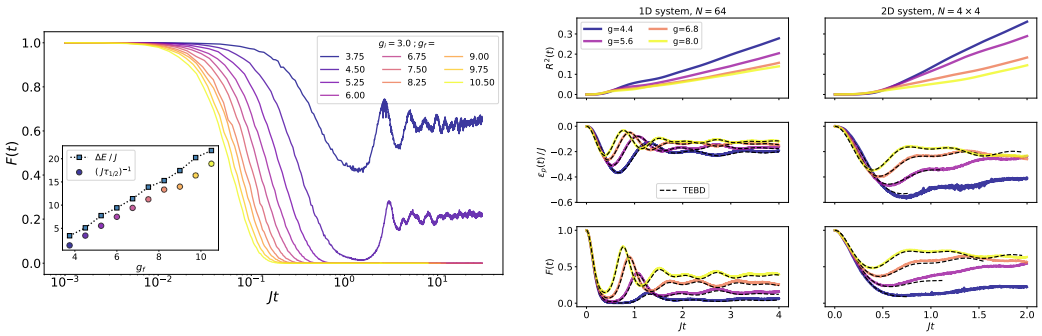

Figure 3: Fidelity and vorticity as functions of time. **Left**: Time-dependent many-body Loschmidt echo $F(t)$ for a number of quenches. **Right**: One- and two-dimensional benchmarks on a chain with $N = 64$ rotors and 4×4 lattice (respectively) with open boundary conditions.

3 Results

We simulate the effects instantaneous *quenches* – we approximate a ground state of the Hamiltonian in Eq. 1 using standard VMC methods and then preform time evolution under a different value of g. In Fig. 2, we choose a square 8×8 lattice, tracking the dynamics of the potential energy density $\epsilon_{\mathrm{p}}(t) = -\frac{J}{N} \sum_{\langle k,l \rangle} \langle \hat{\mathbf{n}}_k \cdot \hat{\mathbf{n}}_l \rangle_t$ and the average magnetization magnitude $M(t) = \frac{1}{N} \langle | \sum_k \hat{\mathbf{n}}_k | \rangle_t$. Averages $\langle \cdot \rangle_t$ are performed with respect to the trial state at time t.

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 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 slower equilibriation with only fluctuations around the magnetization direction. For moderate to large quenches in Fig. 2, we observe a (transient) demagnetization of the sample and convergence to a new equilibrium state.

One also has access to global observables such as the Loschmidt echo with applications in the context of dynamical phase transitions [43] and quantum chaos [44]. It is defined as the (normalized) overlap $F(t) = |\langle \Psi(0) | \Psi(t) \rangle|^2$. Following Refs. [45, 46], we evaluate Monte Carlo estimators of $F(t)$. Expectedly, we find that that F decays with time, (see 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.

We introduce another time scale $\tau_{1/2}$ defined as the time needed for the fidelity to decrease by 50%. We observe that $\tau_{1/2}$ increases linearly with the quench g_f. This result matches basic estimates given by the 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 $\psi_{\alpha(0)}$ [47]. The comparison in Fig. 3 (left, inset) demonstrates that the t-VMC method can be used to estimate quantities of experimental interest for system sizes unreachable by other wavefunction methods.

We compare our results to tensor-network time-evolving block decimation (TEBD) [48, 49] simulations for a one- and two-dimensional versions of the model. For all benchmarks, states were initialized to the coherent superposition of all basis states $|\psi(0)\rangle \propto \int \mathrm{d}\boldsymbol{\theta} \, |\boldsymbol{\theta}\rangle$. Following Refs. [17, 50], we use $r(t) = \mathcal{D}\big(\psi(t+\delta t), e^{-iH\delta t}\,\psi(t)\big) / \mathcal{D}\big(\psi(t), e^{-iH\delta t}\,\psi(t)\big)$ as a figure of merit where $\mathcal{D}(\cdot,\cdot)$ is the Fubini-Study distance on the Hilbert space $\mathcal{H}$. 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 $\big|\psi_{\alpha(t+\delta t)}\big\rangle$. We plot the integrated error $R^2(t) = \int_0^t r^2(s) \, \mathrm{d}s$ as an upper bound on the square of the integrated error $R(t) = \int_0^t r(s) \, \mathrm{d}s$ due to the triangle inequality.

In Fig. 3 (right), 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. The integrated residual $R^2(t)$ grows more rapidly for lower values of g because the initial state $\psi(0)$ representing a more typical state in the disordered phase. In contrast to the 1D case, we observe that the TEBD method

exponentially grows the MPS bond dimension χ past the cutoff $\chi_{\max} = 1000$ at relatively short times. We see qualitative agreement between the two methods for early times, before χ grows to the point where further simulation is numerically prohibitively expensive.

Overall, both t-VMC and TEBD algorithms predict similar dynamical behavior where the comparison is possible. However, the number of parameters in the MPS grows exponentially due to entropy build-up, blocking tensor-network algorithms [51–53, 12] from scaling to higher dimsnsions and longer times. Continuous degrees of freedom amplify the problem due to infinite local basis.

4 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. Our calculations are able to access nontrivial local and global observables. Good agreement was found with tensor-network-based TEBD simulations for the case of one-dimensional systems of comparable size. Our approach paves the way for accurate non-equilibrium simulations of continuous systems at previously unexplored system sizes and evolution times, bridging the gap between simulation and experiment.

5 Acknowledgements

MM acknowledges insightful discussions with Filippo Vicentini about t-VMC 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 ansatze were very helpful for fine-tuning simulations. MM also acknowledges support from the CCQ graduate fellowship in computational quantum physics. The Flatiron Institute is a division of the Simons Foundation. DS was supported by AFOSR: Grant FA9550-21-1-0236 and NSF: Grant OAC-2118310.

Software libraries

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

It was built on JAX [54] for array manipulations, automatic differentiation for sampling and optimization and GPU support, Flax [55] for neural-network construction and manipulation and NumPy [56] and SciPy [57] for CPU array manipulations. Matplotlib [58] was used to produce figures.

References

[1] Warren, W. S., Rabitz, H. & Dahleh, M. Coherent control of quantum dynamics: The dream is alive. *Science* **259**, 1581–1589 (1993).

[2] Polkovnikov, A., Sengupta, K., Silva, A. & Vengalattore, M. Colloquium: Nonequilibrium dynamics of closed interacting quantum systems. *Rev. Mod. Phys.* **83**, 863–883 (2011). URL https://link.aps.org/doi/10.1103/RevModPhys.83.863.

[3] Budden, M. *et al.* Evidence for metastable photo-induced superconductivity in k3c60. *Nature Physics* **17**, 611–618 (2021). URL https://doi.org/10.1038/s41567-020-01148-1.

[4] Thomas, A. *et al.* Tilting a ground-state reactivity landscape by vibrational strong coupling. *Science* **363**, 615–619 (2019). URL https://www.science.org/doi/abs/10.1126/science.aau7742. https://www.science.org/doi/pdf/10.1126/science.aau7742.

[5] Zhang, J. *et al.* Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator. *Nature* **551**, 601–604 (2017). URL https://doi.org/10.1038/nature24654.

[6] Sivak, V. V. *et al.* Model-free quantum control with reinforcement learning. *Physical Review X* **12**, 011059 (2022). URL https://link.aps.org/doi/10.1103/PhysRevX.12.011059.

[7] Porotti, R., Essig, A., Huard, B. & Marquardt, F. Deep reinforcement learning for quantum state preparation with weak nonlinear measurements. *Quantum* **6**, 747 (2022).

[8] Metz, F. & Bukov, M. Self-correcting quantum many-body control using reinforcement learning with tensor networks (2022). URL http://arxiv.org/abs/2201.11790.

[9] Bukov, M. *et al.* Reinforcement learning in different phases of quantum control. *Physical Review X* **8**, 031086 (2018). URL https://link.aps.org/doi/10.1103/PhysRevX.8.031086.

[10] Zaletel, M. P., Mong, R. S. K., Karrasch, C., Moore, J. E. & Pollmann, F. Time-evolving a matrix product state with long-ranged interactions. *Phys. Rev. B* **91**, 165112 (2015). URL https://link.aps.org/doi/10.1103/PhysRevB.91.165112.

[11] Wurtz, J., Polkovnikov, A. & Sels, D. Cluster truncated wigner approximation in strongly interacting systems. *Annals of Physics* **395**, 341–365 (2018). URL https://www.sciencedirect.com/science/article/pii/S0003491618301647.

[12] Czarnik, P., Dziarmaga, J. & Corboz, P. Time evolution of an infinite projected entangled pair state: An efficient algorithm. *Physical Review B* **99**, 035115 (2019). URL https://link.aps.org/doi/10.1103/PhysRevB.99.035115.

[13] Hubig, C., Bohrdt, A., Knap, M., Grusdt, F. & Cirac, I. Evaluation of time-dependent correlators after a local quench in ipeps: hole motion in the t-j model. *SciPost Physics* **8**, 021 (2020). URL https://scipost.org/10.21468/SciPostPhys.8.2.021.

[14] Zhou, Y., Stoudenmire, E. M. & Waintal, X. What limits the simulation of quantum computers? *Phys. Rev. X* **10**, 041038 (2020). URL https://link.aps.org/doi/10.1103/PhysRevX.10.041038.

[15] Carleo, G., Becca, F., Schiró, M. & Fabrizio, M. Localization and glassy dynamics of many-body quantum systems. *Scientific Reports* **2**, 243 (2012). URL http://www.nature.com/articles/srep00243.

[16] Carleo, G., Cevolani, L., Sanchez-Palencia, L. & Holzmann, M. Unitary dynamics of strongly interacting bose gases with the time-dependent variational monte carlo method in continuous space. *Physical Review X* **7**, 031026 (2017). URL https://journals.aps.org/prx/abstract/10.1103/PhysRevX.7.031026.

[17] Schmitt, M. & Heyl, M. Quantum many-body dynamics in two dimensions with artificial neural networks. *Physical Review Letters* **125**, 100503 (2020). URL https://link.aps.org/doi/10.1103/PhysRevLett.125.100503.

[18] Hofmann, D., Fabiani, G., Mentink, J., Carleo, G. & Sentef, M. Role of stochastic noise and generalization error in the time propagation of neural-network quantum states. *SciPost Physics* **12**, 165 (2022). URL https://scipost.org/10.21468/SciPostPhys.12.5.165.

[19] Barison, S., Vicentini, F. & Carleo, G. An efficient quantum algorithm for the time evolution of parameterized circuits. *Quantum* **5**, 512 (2021).

[20] Czischek, S., Gärttner, M. & Gasenzer, T. Quenches near ising quantum criticality as a challenge for artificial neural networks. *Phys. Rev. B* **98**, 024311 (2018). URL https://link.aps.org/doi/10.1103/PhysRevB.98.024311.

[21] Amari, S. I. Natural gradient works efficiently in learning. *Neural Computation* **10**, 251–276 (1998). URL http://www.mitpressjournals.org/doi/10.1162/089976698300017746.

[22] Yuan, X., Endo, S., Zhao, Q., Li, Y. & Benjamin, S. C. Theory of variational quantum simulation. *Quantum* **3**, 191 (2019). URL https://quantum-journal.org/papers/q-2019-10-07-191/.

[23] Stokes, J., Izaac, J., Killoran, N. & Carleo, G. Quantum natural gradient. *Quantum* **4**, 269 (2020). URL https://quantum-journal.org/papers/q-2020-05-25-269/.

[24] Vogt, N. *et al.* One-dimensional Josephson junction arrays: Lifting the Coulomb blockade by depinning. *Physical Review B* **92**, 045435 (2015). URL http://dx.doi.org/10.1103/PhysRevB.92.045435.

[25] Martinoli, P. & Leemann, C. Two Dimensional Josephson Junction Arrays. *Journal of Low Temperature Physics* **118**, 699–731 (2000). URL https://link.springer.com/article/10.1023/A:1004651730459.

[26] Kockum, A. F. & Nori, F. Quantum bits with josephson junctions. *Springer Series in Materials Science* **286**, 703–741 (2019). URL http://dx.doi.org/10.1007/978-3-030-20726-7_17. 1908.09558.

[27] José, J. V., Kadanoff, L. P., Kirkpatrick, S. & Nelson, D. R. Renormalization, vortices, and symmetry-breaking perturbations in the two-dimensional planar model. *Physical Review B* **16**, 1217–1241 (1977). URL https://link.aps.org/doi/10.1103/PhysRevB.16.1217.

[28] Jiang, W., Pan, G., Liu, Y. & Meng, Z. Y. Solving quantum rotor model with different monte carlo techniques. *Chinese Physics B* **31** (2019). URL http://arxiv.org/abs/1912.08229.

[29] Stokes, J., De, S., Veerapaneni, S. & Carleo, G. Continuous-variable neural-network quantum states and the quantum rotor model (2021). URL http://arxiv.org/abs/2107.07105.

[30] Berke, C., Varvelis, E., Trebst, S., Altland, A. & DiVincenzo, D. P. Transmon platform for quantum computing challenged by chaotic fluctuations. *Nature Communications* **13**, 2495 (2022). URL https://doi.org/10.1038/s41467-022-29940-y.

[31] Becca, F. & Sorella, S. *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, 2017).

[32] Sorella, S. Green function monte carlo with stochastic reconfiguration. *Physical Review Letters* **80**, 4558–4561 (1998). URL https://link.aps.org/doi/10.1103/PhysRevLett.80.4558.

[33] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics* **21**, 1087–1092 (1953).

[34] Hastings, W. K. Monte carlo sampling methods using markov chains and their applications. *Biometrika* **57**, 97–109 (1970).

[35] Bogacki, P. & Shampine, L. A 3(2) pair of runge - kutta formulas. *Applied Mathematics Letters* **2**, 321–325 (1989).

[36] Butcher, J. C. *Numerical Methods for Ordinary Differential Equations* (John Wiley & Sons, Ltd, 2008).

[37] Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. *Numerical Recipes in C* (Cambridge University Press, Cambridge, USA, 1992), second edn.

[38] Neal, R. M. *Handbook of Markov Chain Monte Carlo* (Chapman and Hall/CRC, 2011). URL https://www.taylorfrancis.com/books/9781420079425.

[39] Betancourt, M. A conceptual introduction to hamiltonian monte carlo (2017). URL http://arxiv.org/abs/1701.02434.

[40] LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–44 (2015). URL http://www.ncbi.nlm.nih.gov/pubmed/26017442.

[41] Carleo, G. *et al.* Machine learning and the physical sciences. *Reviews of Modern Physics* **91**, 045002 (2019). URL https://link.aps.org/doi/10.1103/RevModPhys.91.045002.

[42] Pescia, G., Han, J., Lovato, A., Lu, J. & Carleo, G. Neural-network quantum states for periodic systems in continuous space. *Physical Review Research* **4**, 023138 (2022). URL https://link.aps.org/doi/10.1103/PhysRevResearch.4.023138.

[43] Heyl, M., Polkovnikov, A. & Kehrein, S. Dynamical quantum phase transitions in the transverse-field ising model. *Phys. Rev. Lett.* **110**, 135704 (2013). URL https://link.aps.org/doi/10.1103/PhysRevLett.110.135704.

[44] Schmitt, M., Sels, D., Kehrein, S. & Polkovnikov, A. Semiclassical echo dynamics in the sachdev-ye-kitaev model. *Phys. Rev. B* **99**, 134301 (2019). URL https://link.aps.org/doi/10.1103/PhysRevB.99.134301.

[45] Medvidović, M. & Carleo, G. Classical variational simulation of the quantum approximate optimization algorithm. *npj Quantum Information* **7**, 101 (2021). URL https://www.nature.com/articles/s41534-021-00440-z.

[46] Jónsson, B., Bauer, B. & Carleo, G. Neural-network states for the classical simulation of quantum computing (2018). URL http://arxiv.org/abs/1808.05232.

[47] Mandelstam, L. & Tamm, I. *The Uncertainty Relation Between Energy and Time in Non-relativistic Quantum Mechanics*, 115–123 (Springer Berlin Heidelberg, Berlin, Heidelberg, 1991). URL https://doi.org/10.1007/978-3-642-74626-0_8.

[48] Vidal, G. Efficient classical simulation of slightly entangled quantum computations. *Physical Review Letters* **91**, 147902 (2003). URL https://link.aps.org/doi/10.1103/PhysRevLett.91.147902. 0301063.

[49] Vidal, G. Efficient simulation of one-dimensional quantum many-body systems. *Physical Review Letters* **93**, 040502 (2004). URL https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.93.040502. 0310089.

[50] Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).

[51] Haegeman, J. *et al.* Time-dependent variational principle for quantum lattices. *Physical Review Letters* **107**, 070601 (2011).

[52] Haegeman, J., Lubich, C., Oseledets, I., Vandereycken, B. & Verstraete, F. Unifying time evolution and optimization with matrix product states. *Physical Review B* **94**, 165116 (2016).

[53] Verstraete, F. & Cirac, J. I. Renormalization algorithms for quantum-many body systems in two and higher dimensions (2004). URL http://arxiv.org/abs/cond-mat/0407066.

[54] Bradbury, J. *et al.* JAX: composable transformations of Python+NumPy programs (2018). URL http://github.com/google/jax.

[55] Heek, J. *et al.* Flax: A neural network library and ecosystem for JAX (2020). URL http://github.com/google/flax.

[56] Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585**, 357–362 (2020). 2006.10256.

[57] Virtanen, P. *et al.* Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature Methods* **17**, 261–272 (2020). URL http://www.nature.com/articles/s41592-019-0686-2.

[58] Hunter, J. D. Matplotlib: A 2D graphics environment. *Comput. Sci. Eng.* **9**, 99–104 (2007).