

Gauge-invariant and anyonic-symmetric autoregressive neural network for quantum lattice models

Di Luo^{1,2,3,4,*}, Zhuo Chen^{1,3,4,*}, Kaiwen Hu^{1,5,6}, Zhizhen Zhao,⁷ Vera Mikyoung Hur⁶, and Bryan K. Clark^{1,2}

¹Department of Physics, University of Illinois at Urbana-Champaign, Illinois 61801, USA

²IQUIST and Institute for Condensed Matter Theory and NCSA Center for Artificial Intelligence Innovation, University of Illinois at Urbana-Champaign, Illinois 61801, USA

³The NSF AI Institute for Artificial Intelligence and Fundamental Interactions

⁴Center for Theoretical Physics, Massachusetts Institute of Technology, Massachusetts 02139, USA

⁵Department of Physics, University of Michigan, Ann Arbor, Michigan 48109, USA

⁶Department of Mathematics, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA

⁷Department of Electrical and Computer Engineering and CSL, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA



(Received 3 December 2021; revised 26 May 2022; accepted 15 December 2022; published 29 March 2023)

Symmetries such as gauge invariance and anyonic symmetry play a crucial role in quantum many-body physics. We develop a general approach to constructing gauge-invariant or anyonic-symmetric autoregressive neural networks, including a wide range of architectures such as transformer and recurrent neural network, for quantum lattice models. These networks can be efficiently sampled and explicitly obey gauge symmetries or anyonic constraint. We prove that our methods can provide exact representation for the ground and excited states of the two- and three-dimensional toric codes, and the X-cube fracton model. We variationally optimize our symmetry-incorporated autoregressive neural networks for ground states as well as real-time dynamics for a variety of models. We simulate the dynamics and the ground states of the quantum link model of U(1) lattice gauge theory, obtain the phase diagram for the two-dimensional \mathbb{Z}_2 gauge theory, determine the phase transition and the central charge of the SU(2)₃ anyonic chain, and also compute the ground-state energy of the SU(2) invariant Heisenberg spin chain. Our approach provides powerful tools for exploring condensed-matter physics, high-energy physics, and quantum information science.

DOI: [10.1103/PhysRevResearch.5.013216](https://doi.org/10.1103/PhysRevResearch.5.013216)

I. INTRODUCTION

In recent years, there has been a growing interest in machine learning approaches to simulating quantum many-body systems [1–23]. An important step in this direction is the use of neural networks, e.g., restricted Boltzmann machines, to represent variational wave functions. However, many neural networks do not automatically enforce the symmetries of physical models. A considerable amount of work has been devoted to remedy the deficiency for several classes of global symmetries, such as translational symmetry [3], discrete rotational symmetry [3], global U(1) symmetry [4], and antisymmetry [5–7].

In addition to global symmetries, local symmetries can be encoded through gauge invariance. The notion of gauge invariance is crucial in quantum mechanics. In high-energy physics, theory is required to be invariant under the action of gauge symmetry groups [24]. Gauge invariance appears naturally in various condensed matter physics models. For

example, topological states of toric code and double-semion models arise as the ground states of their gauge-invariant Hamiltonians [25,26]. Also, novel quantum matter such as fracton is the ground state of a Hamiltonian where the subsystem symmetry is gauged [27]. In quantum information, various quantum error correction codes can be viewed as eigenstates in a certain gauge-invariant code space [28]. Aside from gauge symmetries, anyonic symmetry is another important local constraint that arises in exotic phases of matter [29,30] and topological quantum computation [31]. The study of quantum lattice models with gauge or anyonic symmetries is significant to enhance our understanding of high-energy physics, condensed-matter physics, and quantum information science.

Simulating quantum many-body gauge theory is exponentially costly. There has been much effort to efficiently simulate quantum lattice gauge theory with both digital and analog quantum computers [32], but more effort is required experimentally to achieve good fidelity. Two standard approaches to simulating gauge theory classically are stochastic, integrating an effective Lagrangian by sampling, and variational. When simulating gauge theory, the stochastic approach naturally obeys gauge invariance but is plagued with exponential costs associated with the sign problem in models with finite density of fermions or involving quantum dynamics [33]. The variational approach overcomes the difficulty by being constrained

*Co-first authors.

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.

to an approximate variational space. Imposing gauge symmetries in the variational approach is particularly important and challenging as, otherwise, lower-energy states can exist in the gauge-violating part of a Hilbert space. Therefore, gauge symmetries must be explicitly constrained. While the stochastic approach has been well studied, there have been limited attempts at using the variational approach for gauge theory. Tensor networks can be readily applied to gauge theory in one dimension and ongoing efforts are required to work with challenges in higher dimensions [32,34]. A variational approach based on gauge-equivariant networks has been introduced very recently [35–37].

We develop a general approach to constructing gauge-invariant or anyonic-symmetric (such as the fusion rule for anyons) autoregressive neural networks (AR-NN) for quantum lattice models. Autoregressive neural networks, such as recurrent neural networks (RNN) [38,39], pixel convolutional neural networks (pixelCNN) [40], and transformers [41], have revolutionized the fields of computer vision and language translation and generation, among many others. Autoregressive neural networks quantum states have recently been introduced in quantum many-body physics [4,42,43] and shown to be capable of representing volume law states (as one generically needs in dynamics) with a number of parameters that scale sublinearly [44]. A central feature of AR-NN is their capability of exactly sampling configurations from them. This is to be contrasted with the standard approach of sampling configurations by doing a random walk over a Markov chain, which is often plagued with long equilibration times and nonergodic behaviors. We construct gauge-invariant AR-NN for the quantum link model of U(1) lattice gauge theory [24], \mathbb{Z}_N gauge theory, and anyonic-symmetric AR-NN for $SU(2)_k$ anyons. We demonstrate the exact representation of gauge-invariant AR-NN for the ground and excited states of the two-dimensional (2D) [26] and three-dimensional (3D) [45] toric codes, and the X-cube fracton model [46]. We optimize our symmetry-incorporated AR-NN for the quantum link model, the 2D toric code in a transverse field, the one-dimensional (1D) Heisenberg chain with $SU(2)$ symmetry, and the $SU(2)_3$ anyonic chain [29,30,47], to obtain ground states accurately and extract phase diagrams and various dynamic properties.

II. CONSTRUCTION FOR GAUGE OR ANYONIC SYMMETRIES

Our goal in this work is to generate autoregressive neural networks (AR-NN) which variationally represent wave functions of quantum lattice models and explicitly obey their gauge symmetries, i.e., given a set of gauge symmetry operators $\{G_i\}$ with local support, we would like to construct a wave function $|\psi\rangle$ such that $G_i|\psi\rangle = |\psi\rangle$ for each i . To do this, we will work within the “gauge basis” $\{|x\rangle\}$ which is diagonal in the gauge $\langle x|\psi\rangle = \langle x|G_i|\psi\rangle$. A sufficient condition of gauge invariance of the wave function is to ensure that the gauge-violating basis elements $|x\rangle$ have zero amplitude in $|\psi\rangle$. Throughout this work, we will primarily work with gauges G_i which are local, i.e., $G_i|x\rangle$ only affects a compact range of sites within the vicinity of site i .

While we would typically want our AR-NN to take as input the configuration $\{x_1, x_2, \dots, x_n\}$ and evaluate $\psi(x_1, x_2, \dots, x_n)$, we will find it useful to instead evaluate $\psi(\tilde{x})$ where $\tilde{x} \equiv \{\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n\}$, $\tilde{x}_i \equiv (x_{i_1}, x_{i_2}, \dots, x_{i_v})$ is a composite particle specifying the configuration of not only site i but also some number of nearby sites. The motivation for working with composite particles is that a particular local gauge constraint G_i might only depend on composite particle \tilde{x}_i (and potentially \tilde{x}_{i+1}), making it easier to apply the gauge constraints. Different composite particles can naturally overlap in physical sites and we will simply augment our gauge constraints to require that the configurations of the composite particles agree on the state of a physical site, i.e., basis states of composite particles which map to disagreeing physical states should also have zero amplitude.

AR-NN perform two functions: sampling and evaluation. AR-NN can sample configurations \tilde{x} from $|\psi(\tilde{x})|^2$. This is done sequentially (in some predetermined order) one composite particle \tilde{x}_i at a time; the probability to sample \tilde{x}_i is equal to $a^2(\tilde{x}_i|\tilde{x}_{<i})$ where $a(\tilde{x}_i|\tilde{x}_{<i})$ is a function which returns the conditional amplitude. Evaluation of the AR-NN gives a value $\psi(\tilde{x}) = \prod_{i=1}^n a(\tilde{x}_i|\tilde{x}_{<i})e^{i\theta(\tilde{x}_i|\tilde{x}_{<i})}$ where $\theta(\tilde{x}_i|\tilde{x}_{<i})$ is a function which returns the conditional phase. Both evaluation and sampling rely on the existence of a gauge block which takes $\tilde{x}_1, \dots, \tilde{x}_{k-1}$ and outputs the possible values $\{\tilde{z}^i\}$ of \tilde{x}_k along with their respective amplitudes $a(\tilde{z}^i|\tilde{x}_{<k})$ and phases $\theta(\tilde{z}^i|\tilde{x}_{<k})$, ensuring that the amplitude of any configuration which is going to violate the gauge constraint is set to zero. To build this gauge block, we start with an autoregressive neural network block which returns a list of amplitudes which do not constrain the gauge (such blocks are standard in autoregressive models such as transformers and RNN); we then zero out those partial configurations which break the gauge (on the already established composite particles) and renormalize the probabilities in this list [see Fig. 1(a)]. Given the gauge block it is then straightforward to both sample and evaluate [see Figs. 1(b) and 1(c)]. Note the probability induced by our AR-NN is different from the probability induced by the AR-NN with only the autoregressive neural network block even if one projects out the gauge-breaking configurations from the latter network.

It is worth noticing that the construction is not limited to gauge theory, but can be generalized to wave functions with either local or global constraints which are checked in the same way as gauge constraints are checked. This will be helpful for describing constraints from certain global symmetries or special algebraic structure, such as the $SU(2)$ symmetry for the Heisenberg model and the $SU(2)_k$ fusion rules for non-Abelian anyons.

III. OPTIMIZATION ALGORITHMS

We use AR-NN to calculate both ground states and real-time dynamic properties. In both cases, we need to optimize our AR-NN. For ground states, an AR-NN is optimized with respect to energy and for real-time dynamics, we optimize an AR-NN at time step $t + 2\tau$ given a network at time t . We describe the details of these optimizations. As these optimization approaches are general, we use x to denote a

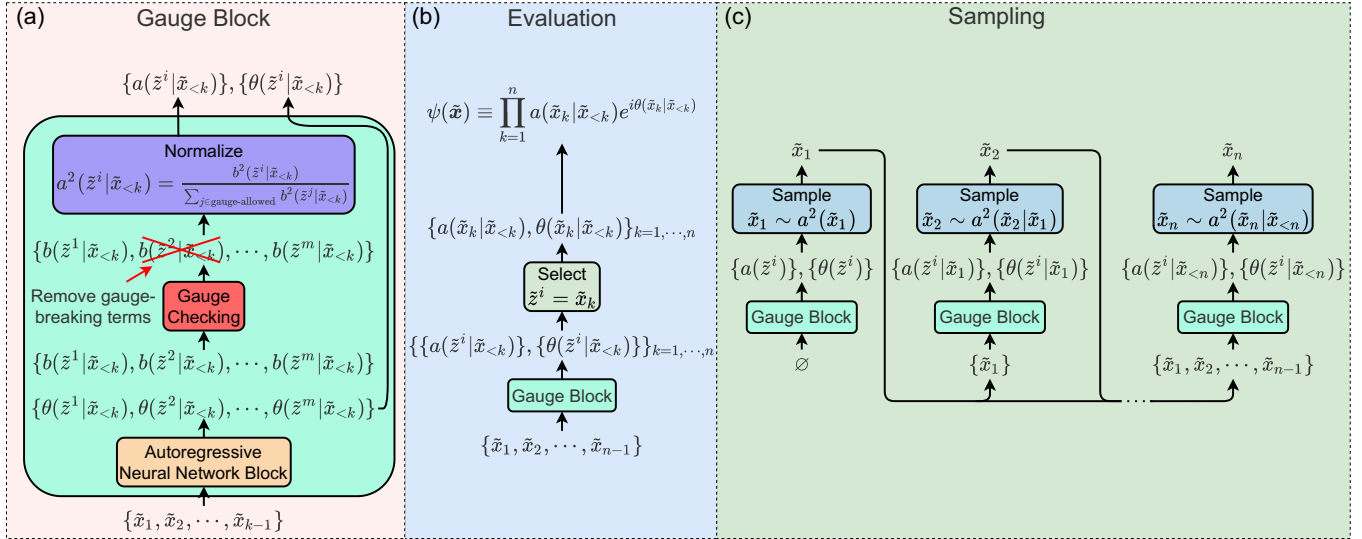


FIG. 1. Autoregressive parametrization of wave function with n composite particles. (a) Gauge block. The input $\{\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_{k-1}\}$ is processed through the autoregressive neural network block (see Appendix E for details) to output amplitude and phase parts. The amplitude part goes through gauge checking, which removes the gauge-breaking terms. Afterwards, the square of the amplitude is normalized. (b) Evaluation process. The evaluation process can be performed in parallel for all the input sites. Given the input $\{\tilde{x}_k\}$, the gauge block simultaneously generates amplitudes and phases for all sites. We then select the correct amplitudes and phases based on the input configuration for each site and construct the wave function from the selected amplitudes and phases. (c) Sampling process. The sampling is done sequentially for each site. We begin with no input and generate the amplitude and phase for the first site. The configuration of the first site is then sampled from the square of the amplitude. Afterwards, we feed the first sample into the gauge block to obtain the second sample. This process continues until we obtain the whole configuration.

configuration, but for the context of the paper, x should be viewed as a composite particle configuration.

For the ground-state optimization, we stochastically minimize the expectation of energy for a Hamiltonian H and a wave function $|\psi_\theta\rangle$ as

$$\langle \psi_\theta | H | \psi_\theta \rangle \approx \frac{1}{N} \sum_{x \sim |\psi_\theta|^2} \frac{H \psi_\theta(x)}{\psi_\theta(x)} \equiv \frac{1}{N} \sum_{x \sim |\psi_\theta|^2} E_{\text{loc}}(x), \quad (1)$$

where N is the batch size and the gradient is given by

$$\frac{\partial}{\partial \theta} \langle \psi_\theta | H | \psi_\theta \rangle \approx \frac{2}{N} \sum_{x \sim |\psi_\theta|^2} \Re \left\{ E_{\text{loc}}(x) \frac{\partial}{\partial \theta} \ln \psi_\theta^*(x) \right\}. \quad (2)$$

We further control the sampling variance [48] by subtracting from $E_{\text{loc}}(x)$ the average over the batch, $E_{\text{avg}} \equiv 1/N \sum_{x \in \text{batch}} E_{\text{loc}}(x)$, and define the stochastic variance reduced loss function as

$$\mathcal{L}_g = \frac{2}{N} \sum_{x \sim |\psi_\theta|^2} \Re \left\{ [E_{\text{loc}}(x) - E_{\text{avg}}] \ln \psi_\theta^*(x) \right\}, \quad (3)$$

where the gradient is taken on $\ln \psi_\theta^*$ using PYTORCH's [49] automatic differentiation.

With this loss function, we also use transfer learning techniques [50,51]. We train our neural networks in smaller systems and use these parameters as the initial starting points for optimizing for larger systems (see Appendix E for details).

For the dynamics optimization, we use a stochastic version of the logarithmic forward-backward trapezoid method [52], which can be viewed as a higher-order generalization of IT-

SWO [53] and the logarithmic version of the loss functions in Refs. [9,42]. We initialize two copies of the neural network $\psi_{\theta(t)}$ and $\psi_{\theta(t+2\tau)}$. At each time step, we train $\psi_{\theta(t+2\tau)}$ to match $(1 + iH\tau)|\psi_{\theta(t+2\tau)}\rangle \equiv |\Psi_\theta\rangle$ and $(1 - iH\tau)|\psi_{\theta(t)}\rangle \equiv |\Phi\rangle$ by minimizing the negative logarithm of the overlap, $-\ln(\langle \Psi_\theta | \Phi \rangle \langle \Phi | \Psi_\theta \rangle) / (\langle \Psi_\theta | \Psi_\theta \rangle \langle \Phi | \Phi \rangle)$. [Since we only take the gradient on $\theta(t+2\tau)$, for simplicity, we write θ for $\theta(t+2\tau)$ and neglect $\theta(t)$.] The inner products related to θ can be evaluated stochastically as

$$\langle \Psi_\theta | \Phi \rangle \approx \frac{1}{N} \sum_{x \sim |\psi_\theta|^2} \frac{\Psi_\theta^*(x) \Phi(x)}{|\psi_\theta(x)|^2} \equiv \frac{1}{N} \sum_{x \sim |\psi_\theta|^2} \alpha(x), \quad (4)$$

$$\langle \Psi_\theta | \Psi_\theta \rangle \approx \frac{1}{N} \sum_{x \sim |\psi_\theta|^2} \frac{|\Psi_\theta(x)|^2}{|\psi_\theta(x)|^2} \equiv \frac{1}{N} \sum_{x \sim |\psi_\theta|^2} \beta(x). \quad (5)$$

The gradient of the negative logarithm of the overlap can be evaluated stochastically as

$$\begin{aligned} & \frac{\partial}{\partial \theta} \left(-\ln \frac{\langle \Psi_\theta | \Phi \rangle \langle \Phi | \Psi_\theta \rangle}{\langle \Psi_\theta | \Psi_\theta \rangle \langle \Phi | \Phi \rangle} \right) \\ & \approx \frac{2}{N} \sum_{x \sim |\psi_\theta|^2} \Re \left\{ \left[\frac{\beta(x)}{\beta_{\text{avg}}} - \frac{\alpha(x)}{\alpha_{\text{avg}}} \right] \frac{\partial}{\partial \theta} \ln \Psi_\theta^*(x) \right\}, \quad (6) \end{aligned}$$

where α_{avg} and β_{avg} are, respectively, the average values of $\alpha(x)$ and $\beta(x)$ over the batch of samples. We can then define the loss function as

$$\mathcal{L}_d \approx \frac{2}{N} \sum_{x \sim |\psi_\theta|^2} \Re \left\{ \left[\frac{\beta(x)}{\beta_{\text{avg}}} - \frac{\alpha(x)}{\alpha_{\text{avg}}} \right] \ln \Psi_\theta^*(x) \right\}, \quad (7)$$

where the gradient is taken on $\ln \Psi_\theta^*$ using PYTORCH's [49] automatic differentiation.

For both optimizations, $\psi_\theta(x)$ is evaluated as described in Fig. 1(a) and x is sampled from $|\psi_\theta|^2$ as described in Fig. 1(b). The full derivations of the stochastic gradients for both optimizations are in Appendix G.

In addition, we extensively use the transfer learning technique by training on small system sizes before moving on to large system sizes. The transfer learning technique provides a good initialization for neural networks that are trained on large system sizes. We observe that the transfer learning technique in general significantly reduces the number of iterations needed. The details of usage of this technique are described in the captions of each figure. (See more details in Appendix E.)

IV. APPLICATIONS IN QUANTUM LATTICE MODELS

A. U(1) quantum link model

The quantum link model (QLM) of U(1) lattice gauge theory in 1 + 1 dimensions in the Hamiltonian formulation with staggered fermions [24] is defined as

$$H_{\text{QLM}} = - \sum_i [\psi_i^\dagger U_{i,i+1} \psi_{i+1} + \psi_{i+1}^\dagger U_{i,i+1}^\dagger \psi_i] + m \sum_i (-1)^i \psi_i^\dagger \psi_i + \frac{g^2}{2} \sum_i E_{i,i+1}^2, \quad (8)$$

where m is the staggered fermion mass, g is the gauge coupling, $i = 1, 2, \dots$ labels the lattice site, ψ_i is the fermion operator, $U_{i,i+1}$ is the link variable, and $E_{i,i+1}$ the electric flux for the U(1) gauge field on link $(i, i+1)$ [24]. We denote by $|q_i\rangle$ the basis state at site i , and by $|e_{i,i+1}\rangle$ the basis at link $(i, i+1)$. Each unit cell is defined to include two sites and two links. The operators $E_{i,i+1}$ and $U_{i,i+1}$ satisfy the following commutation relations: $[E_{i,i+1}, U_{i,i+1}] = U_{i,i+1}$, $[E_{i,i+1}, U_{i,i+1}^\dagger] = -U_{i,i+1}^\dagger$, and $[U_{i,i+1}, U_{i,i+1}^\dagger] = 2E_{i,i+1}$. The gauge constraint is given by the Gauss's law operator $\tilde{G}_i = \psi_i^\dagger \psi_i - E_{i,i+1} + E_{i-1,i} + \frac{1}{2}[(-1)^i - 1]$ such that the ground state $|\psi\rangle$ satisfies $\tilde{G}_i|\psi\rangle = 0$ for each i . The QLM has gained growing interests and been studied in different settings in recent years [32,34,54–56]. We focus on the (1 + 1)-dimensional [(1 + 1)D] QLM with the $S = \frac{1}{2}$ representation for the link operators $U_{i,i+1}$ and $E_{i,i+1}$. Under the Jordan-Wigner transformation, Eq. (8) becomes [57]

$$H = - \sum_i [S_i^+ S_{i,i+1}^+ S_{i+1}^- + \text{H.c.}] + m \sum_i (-1)^i \left(S_i^3 + \frac{1}{2} \right) + \frac{g^2}{2} \sum_i \frac{1}{4}, \quad (9)$$

where $S^\pm \equiv S^1 \pm iS^2$, S^1, S^2, S^3 are the Heisenberg matrices, and the Gauss's law operator becomes $G_i = S_i^3 - S_{i,i+1}^3 + S_{i-1,i}^3 + \frac{1}{2}(-1)^i$. For the $S = \frac{1}{2}$ representation, the last term on the right side of Eq. (9) is constant and, hence, can be discarded.

We define the composite particles of our gauge-invariant AR-NN as in Fig. 2 and choose an order from left to right. Each composite particle $|\sigma_i\rangle$ consists of a fermion $|q_i\rangle$, which can be either $|\bullet\rangle$ or $|\circ\rangle$, and a gauge field in the link $|e_{i,i+1}\rangle$,




FIG. 2. Composite particles for the quantum link model. Each composite particle is defined as $|\sigma_i\rangle \equiv |q_i, e_{i,i+1}\rangle$. We check Gauss's law between $|\sigma_i\rangle$ and $|\sigma_{i+1}\rangle$.

which can be either $|\rightarrow\rangle$ or $|\leftarrow\rangle$. Note that in this case the composite particles do not overlap. The Gauss's law operator G_i acts on $|\sigma_i\rangle$ and $|\sigma_{i+1}\rangle$ to determine allowed configurations and so can be checked in the gauge block which generates the composite particle at site $i+1$. For example, given $|\sigma_i\rangle = |\bullet \rightarrow\rangle$, $|\sigma_{i+1}\rangle$ can only be $|\bullet \rightarrow\rangle$ or $|\circ \leftarrow\rangle$ if i is even, and $|\circ \rightarrow\rangle$ if i is odd.

We implement and variationally optimize this AR-NN for the ground state of Eq. (9). Figure 3 shows the results for 6 unit cells (i.e., 12 particles) which closely match the energy of the exact solution. More importantly, the gauge-invariant construction guarantees that the solution is in the physical space, while the neural network without gauge constraint (i.e., removing the gauge checking from the AR-NN) finds a lower energy but nonphysical state.

We in addition compute the ground state for 40, 80, 120, and 160 unit cells with both transformer and RNN (Fig. 4). The average electric fields are compared with tensor network (TN) results [58]. We find that our results (for matching system sizes) are similar to the TN results for both transformer and RNN. In addition, we extrapolated the ground-state energy for the 160-unit-cell model at $m = 0.7$ (see Fig. 18 in Appendix A) by linearly extrapolating in variance vs energy. We find that the extrapolated ground-state energy is -199.7923 , while our lowest energy is -199.7803 ± 0.0005 , giving us a relative error of only 6×10^{-5} .

We also consider the real-time dynamics for $m = 0.1$ and 2.0 for 6 and 12 unit cells starting with an initial product




FIG. 3. Variational ground-state optimization for the 6-unit-cell (12 sites and 12 links) open-boundary QLM for $m = 0$ with and without gauge-invariant construction. The gauge-invariant autoregressive neural network reaches an accurate ground state while the Ansatz without gauge constraints arrives at a nonphysical state in the optimization. We use the transformer neural network with 1 layer, 32 hidden dimensions, and the real-imaginary parametrization (see Fig. 24). The neural network is randomly initialized and is trained for 1000 iterations with 12 000 samples in each iteration. The neural network architecture and optimization details are discussed in Appendix E.

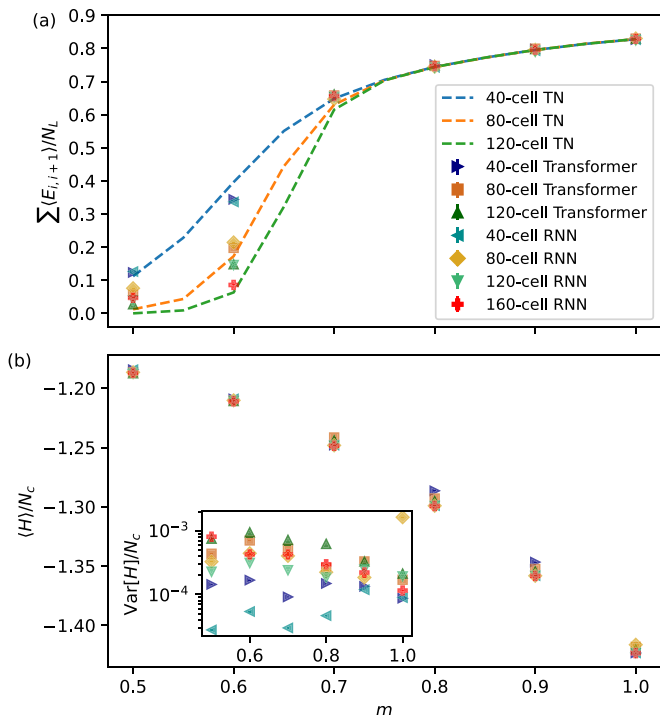


FIG. 4. Variational ground-state optimization for the open-boundary QLM of different system sizes and different m 's with gauge-invariant construction. (a) The expectation value of the electric fields averaged over all links, (b) energy and (inset) energy variance per unit cell. We compare our results with the tensor network (TN) results [dashed lines in (a)] from Ref. [58]. The transformer neural network has 1 layer and 32 hidden dimensions, whereas the RNN has 2 layers and 40 hidden dimensions. For both neural networks, we use the amplitude-phase parametrization (see Fig. 24). The neural networks are randomly initialized. Then they are trained for 3000 iterations with 12 000 samples on 40 unit cells. Then, we use transfer learning technique and train the transformer for 1000 iterations on 80 unit cells and 600 iterations on 120 unit cells. The RNN is then trained for 1000 iterations on 80 and 120 unit cells and 600 iterations on 160 unit cells. The neural network architecture and optimization details are discussed in Appendix E.

state with $|\bullet \rightarrow \circ \rightarrow\rangle$ for each unit cell. Figure 5(a) shows the conservation of energy for different *Ansätze*. The total energy is -1.2 for $m = 0.1$, and -24 for $m = 2.0$. We find that our gauge-invariant AR-NN captures the correct electric field oscillation and has a lower per-step infidelity compared with the nongauge *Ansatz* [see Figs. 5(b) and 5(c)], and, additionally, the anticipated string inversion of the electric flux for small mass (and respectively the static electric flux for large mass) (see Fig. 6).

While this work focuses on the $S = \frac{1}{2}$ representation, our construction can be generalized to an arbitrary S representation. For a higher spin S , composite particles can be defined similarly (see Fig. 2) except that the degree of freedom for each $e_{i,i+1}$ increases to $2S + 1$ as S increases.

B. 2D \mathbb{Z}_N gauge theory

For the 2D toric code [26], consider an $L \times L$ periodic square lattice, where each edge has the basis $\{|0\rangle, |1\rangle\}$. Let




FIG. 5. Dynamics for the 6- and 12-unit-cell (12–24 sites and 12–24 links) open-boundary QLM for $m = 0.1$ and 2.0 with and without gauge-invariant construction. The dashed curves are the exact results from the exact diagonalization for 6 unit cells. (a) The change in the energy during the dynamics. (b) The expectation value of the electric field averaged over all links. (c) The per-step infidelity measure, where $|\Psi\rangle$ and $|\Phi\rangle$ are defined in Sec. III. We use the transformer neural network with 1 layer, 16 hidden dimensions for 6 unit cells, and 32 hidden dimensions for 12 unit cells, and the real-imaginary parametrization (see Fig. 24). The initial state is $|\bullet \rightarrow \circ \rightarrow\rangle$ for each unit cell and we train the neural network using the forward-backward trapezoid method with the time step $\tau = 0.005$, 600 iterations in each time step, and 12 000 samples in each iteration. The neural network architecture, initialization, and optimization details are discussed in Appendix E.

V, P, E denote the sets of vertices, plaquettes, and edges of the lattice, respectively, such that $|V| = L^2$, $|P| = L^2$, $|E| = 2L^2$. Here we consider the toric code with a transverse magnetic field

$$H = H_{\text{TC}} - h \sum_{e \in E} \sigma_e^z, \quad (10)$$




FIG. 6. Dynamics of the gauge-invariant AR-NN for the 12-unit-cell QLM with (a) $m = 0.1$ and (b) $m = 2.0$. The *Ansatz*, initialization, and optimization are the same as in Fig. 5 and are discussed in Appendix E.

where H_{TC} is the toric code Hamiltonian

$$H_{TC} = - \sum_{v \in V} A_v - \sum_{p \in P} B_p, \tag{11}$$

$A_v \equiv \prod_{e \ni v} \sigma_e^z$ (the star operator), $B_p \equiv \prod_{e \ni p} \sigma_e^x$, and h is the strength of the transverse field. Note that A_v is the gauge constraint such that the ground state $|\psi\rangle$ of Eqs. (A1) and (11) satisfies $A_v |\psi\rangle = |\psi\rangle$ for each v .

The composite particle construction is illustrated in Fig. 7. We order our consecutive particles by an ‘‘S’’ shape going up one row and down the next [see Fig. 29(b) in Appendix E]. Two constraints must be checked in the gauge-checking process of a gauge block.

When working on the gauge block associated with composite particle \tilde{x}_v , we check that $A_v |\tilde{x}_v\rangle = |\tilde{x}_v\rangle$ (despite A_v acting on an entire state, this can be checked locally on a single composite particle). In addition, composite particles overlap with their four immediately adjacent composite particles. The gauge block for the composite particle at v therefore checks consistency of the physical sites with the first $v - 1$ compos-

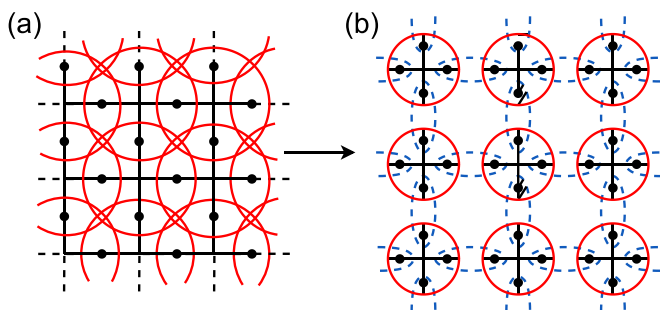


FIG. 7. Composite particles for the 2D toric code. (a) Physical structure of 2D toric code with red circles specifying composite particles. Note multiple composite particles share the same physical sites. (b) Composite particles. We define each star as a composite particle (red circle) and check bond consistency for physical sites shared by adjacent composite particles (blue dashed ovals).

ite particles. For example, given the configuration $|1_0^0\rangle$ for a composite particle, the composite particle to the right can only be $|1_0^1\rangle$, $|1_0^0\rangle$, $|1_1^0\rangle$ or $|1_1^1\rangle$, as the $|1\rangle$ on the right of the left particle must also be on the left of the right particle. In the ‘‘S’’ ordering, there always exist valid choices for each composite particle. For a composite particle that is not the last one, there is an unchosen site which provides freedom of choices to be valid. For the last composite particle, though all sites are fixed, the fixed configuration must be valid because the product of all the Gauss’s law constraints is 1 and all previous Gauss’s law constraints have been satisfied to be 1.

We begin by showing that we can analytically generate an AR-NN for the ground state of H_{TC} . One ground state of Eq. (11) is $|\psi\rangle = \prod_{v \in V} (\mathbb{1} + A_v) |+\rangle^{\otimes n}$ where $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, i.e., an equal superposition of all configurations in the gauge basis which do not violate the gauge constraint. In our construction, if we use an autoregressive neural network block which gives equal weight to all the configurations (this is straightforward to arrange by setting the last linear layer’s weight matrices to zero and bias vectors to equal amplitudes), we exactly achieve this state. Checking the A_v does not affect the relative probabilities because it is not conditional involving only one composite particle. On the other hand, the ‘‘gauge constraints’’ which verify consistency of the underlying state of the sites leave equal probability between all consistent states. To see this, we examine the effect of the gauge constraint on $|a(\tilde{x}_k|\tilde{x}_{<k})|^2$ for any given k , which is the conditional probability of the composite particle \tilde{x}_k . Due to the conditioning from previous composite particles $\{\tilde{x}_{<k}\}$, some sites of the composite particle \tilde{x}_k are fixed. For the Gauss’s law gauge constraints to be 1, the product of all the unchosen site configurations in \tilde{x}_k must be either 1 or -1 , depending on the chosen site configurations. Let $S_1 = \{b_1, \dots, b_j | \prod_{r=1}^j b_r = 1\}$ and $S_{-1} = \{c_1, \dots, c_j | \prod_{r=1}^j c_r = -1\}$ be the two possible sets of unchosen site configurations, where b_r and c_r are the configurations of the unchosen sites in \tilde{x}_k . Consider a function $f : S_1 \rightarrow S_{-1}$ such that $f(b_1) = -c_1$ and $f(b_r) = c_r$ otherwise. Notice that f is bijective and thus S_1 and S_{-1} have the same cardinality, implying that after normalization $|a(\tilde{x}_k|\tilde{x}_{<k})|^2$ will have the same amplitude for any $\{\tilde{x}_{<k}\}$. We can also generate excited states by changing the A_v for a fixed (even) number of vertices to constrain this local eigenvalue to be -1 instead of 1. We provide a numerical verification of this by computing the energy for an exactly represented tower of ground and excited states in Fig. 8.

With a nonzero value of the external field h , the ground state of Eq. (A1) is no longer exactly representable, and we variationally optimize our AR-NN to compute the ground-state energy. Figure 9 shows the minimum energy for Eq. (A1) for different h and the energy derivative, computed using the Hellman-Feynman theorem [59]. The toric code is expected to exhibit a quantum phase transition between the topological and trivial phases at an intermediate value of h , and the sharp change of the energy derivative around $h = 0.34$ is an indicator of this phase transition, which is consistent with the quantum Monte Carlo prediction of $h = 0.328474$




FIG. 8. Energies of the analytical constructions of the ground and excited states of the 11×11 2D and $4 \times 4 \times 4$ 3D toric code, and the $4 \times 4 \times 4$ X-cube fracton model. Here N_{break} is the number of Gauss' law violations. The dashed lines are the exact values for each model. The analytical construction generates the same values as the exact up to stochastic errors from sampling.

in the thermodynamic limit [61]. We can additionally identify the transition by considering the Wilson loop operator $W_C = \prod_{e \in C} \sigma_e^x$ for a closed loop C . It is predicted that the topological order phase follows an area-law decay, $\langle W_C \rangle \sim \exp(-\alpha A_C)$,

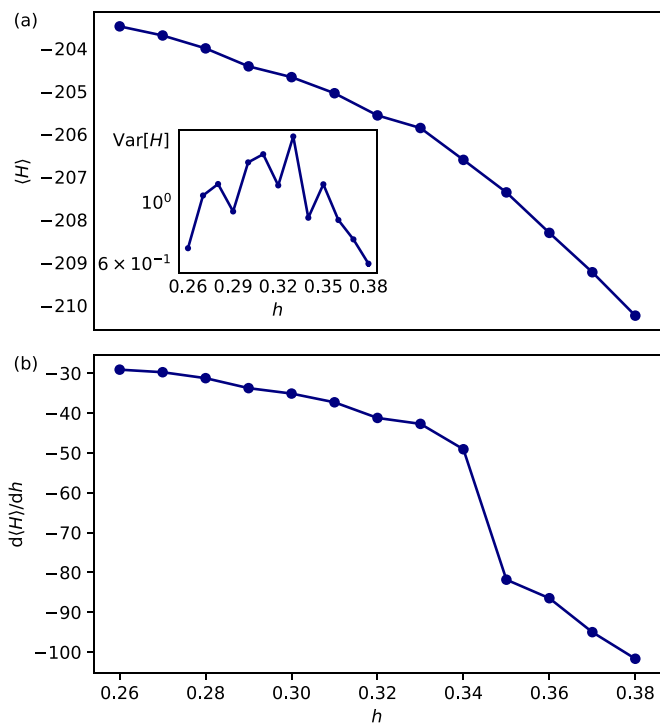


FIG. 9. (a) Energy, (inset) energy variance, and (b) energy derivative (computed by the Hellman-Feynman theorem [59] as $d\langle H \rangle/dh = \langle dH/dh \rangle = -\sum_{e \in E} \langle \sigma_e^z \rangle$) versus h . We use the 2D RNN with 3 layers, 32 hidden dimensions, and the amplitude-phase parametrization (see Fig. 24). We use the transfer learning technique where we first train the neural network on a 6×6 model for 8000 iterations and then we transfer the neural network to the 10×10 model for another 1000 iterations. In each iteration, we use 12000 samples. The neural network architecture, initialization, and optimization details are discussed in Appendix E.

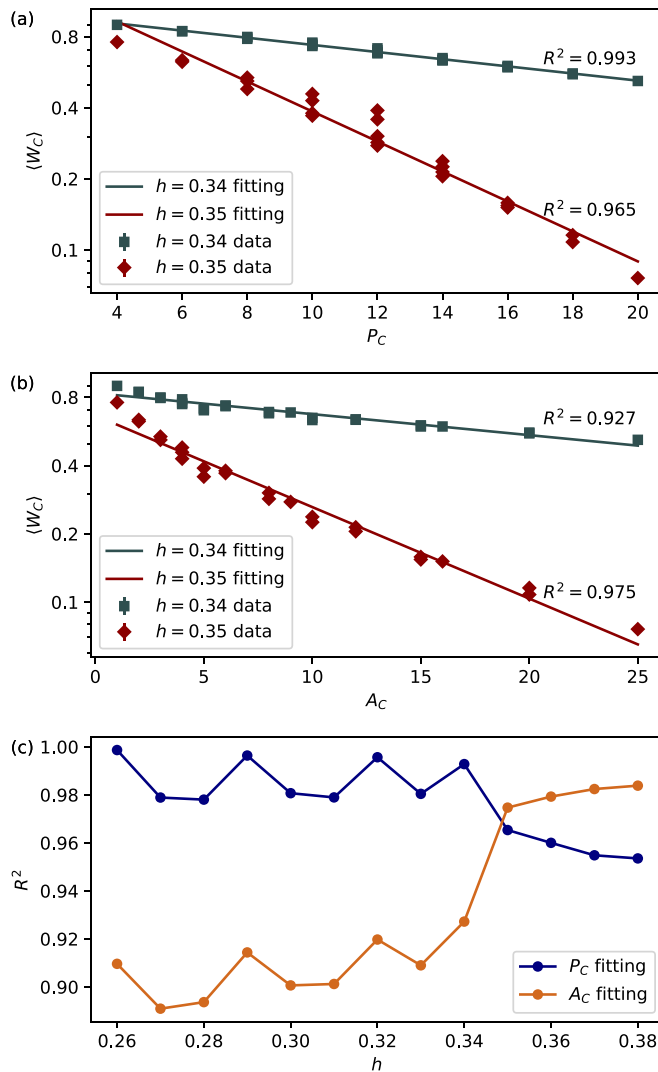


FIG. 10. Perimeter and area laws for the 10×10 2D toric code. The expectation value of the Wilson loop operator with respect to the (a) perimeter and (b) area of the loop in a logarithm-y scale for $h = 0.34$ and 0.35 . (c) The fitting of the correlation coefficient R^2 for the perimeter and area laws for different h . The Ansatz, initialization, and optimization are the same as in Fig. 9 and discussed in Appendix E.

and the trivial phase follows a perimeter-law decay, $\langle W_C \rangle \sim \exp(-\beta P_C)$, where A_C, P_C are the enclosed area and perimeter of the loop C [62]. Figure 10 shows the values of $\langle W_C \rangle$ using our variationally optimized AR-NN. By comparing the respective fits to the area and perimeter laws we again see the transition at $h = 0.34$. Finally, we compare the nonlocal string correlation operators $S_\gamma = \prod_{e \in \gamma} \sigma_e^z$ of our variational states which could be viewed as a measure of the correlation of a pair of excited particle and antiparticle along a path γ . In the topological order phase the nonlocal string operators will decay to zero while they will remain constant at the trivial phase [63]. In Fig. 11, this is seen clearly on both sides of the transition.

At $h = 0.36$, we additionally benchmark our results with the infinite system size infinite projected entangled pair states (iPEPS) results [60] and the gauge-equivariant

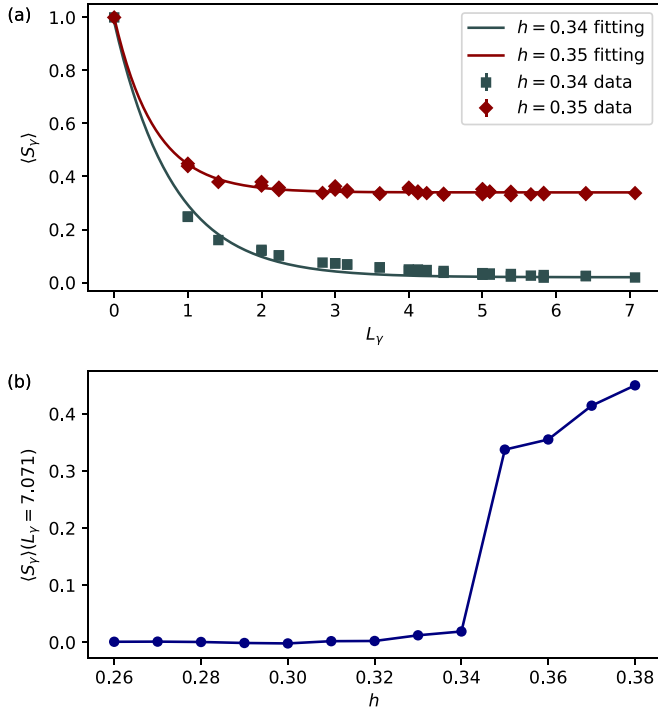


FIG. 11. (a) Nonlocal string correlation function for the 10×10 2D toric code between a pair of particle and antiparticle with a distance of L_y apart. (b) The correlation of a pair of particle and antiparticle at a distance of $L_y = 5\sqrt{2}$ for different h . The *Ansatz*, initialization, and optimization are the same as in Fig. 9 and discussed in Appendix E.

results [35] (Fig. 12). Here we use an improved *Ansatz* with 180° rotation symmetry defined by $\ln \psi_{\text{new}}(x) = \ln[|\psi(x)|^2 + |\psi(R(x))|^2]/2$, where $R(x)$ rotates configuration x by 180° . We find that our results (at least up to $L = 12$) are lower in energy density than the gauge-equivariant results and the iPEPS results, which indicates that our approach is very competitive with the state-of-the-art methods.

Our approach can be naturally generalized to 2D \mathbb{Z}_N gauge theory, which can be described in the language of Kitaev’s $D(G)$ model with group $G = \mathbb{Z}_N$ (see Appendix B). In this case, the basis at each edge becomes a group element in \mathbb{Z}_N . Similarly to Fig. 7, one can define a composite particle over four edges from a vertex and impose gauge invariance. We can also extend our approach to the $(1+1)\text{D}$ \mathbb{Z}_N lattice quantum electrodynamics (QED) model, which is discussed in Appendix C.

C. 3D toric code and fracton model

We turn to gauge-invariant AR-NN for the ground and excited states of the 3D toric code [45] and the fracton model [46,64]. The 3D toric code generalizes the 2D toric code to an $L \times L \times L$ periodic cube where each edge has the basis $\{|0\rangle, |1\rangle\}$. The Hamiltonian takes the same form as the 2D model [see Eq. (11)] except that for each $A_v \equiv \prod_{e \ni v} \sigma_e^z$ there are six edges e associated with each vertex v . A ground state of the 3D toric cube similarly satisfies $A_v |\psi\rangle = B_p |\psi\rangle = |\psi\rangle$ for each v, p . One of the degenerate ground states can also be expressed as $|\psi\rangle = \prod_{v \in V} (\mathbb{1} + A_v) |+\rangle^{\otimes n}$. The excited states




FIG. 12. (a) Energy per site and (b) variance of energy per site for $L \times L$ toric code model with $h = 0.36$. We compare our results (blue squares) with the iPEPS results of infinite system size (dashed lines) from Ref. [60] and the gauge-equivariant neural network results from Ref. [35]. Notice that due to the difference in the definition of h , the h here is twice as large as in Ref. [60]. We use the 2D RNN with 3 layers, 32 hidden dimensions, and the amplitude-phase parametrization (see Fig. 24). The neural network is randomly initialized and trained for 8000 iterations with 12000 samples on 6×6 system. Then, we used the transfer learning technique to train the neural network on 8×8 and 10×10 systems for another 8000 iterations and on 12×12 system for 4000 iterations. The neural network architecture, initialization, and optimization details are discussed in Appendix E.

can be generated by breaking certain constraints from A_v and B_p as in the 2D case.

The X-cube fracton model [46] is also defined on an $L \times L \times L$ periodic cube where each edge has the basis $\{|0\rangle, |1\rangle\}$. The Hamiltonian takes the form

$$H_{\text{fracton}} = - \sum_{v \in V, i} A_v^i - \sum_{c \in C} B_c, \quad (12)$$

where $B_c \equiv \prod_{e \in c} \sigma_e^z$ over the edges in a small cube. The gauge constraint, i.e., Gauss’s law, is $B_c |\psi\rangle = |\psi\rangle$. There are three $A_v^i \equiv \prod_{e \ni v} \sigma_e^x$ for three choices of $i = zy, xy, xz$, depending on which 2D plane A_v^i acts on. The operators are illustrated in Fig. 13. A ground state of the X-cube fracton model satisfies $A_v^i |\psi\rangle = B_c |\psi\rangle = |\psi\rangle$ for each i, v, c . One of the ground states can be expressed as $|\psi\rangle = \prod_C (\mathbb{1} + B_c) |+\rangle^{\otimes n}$. The




FIG. 13. A_v^i and B_c for the X-cube fracton model.

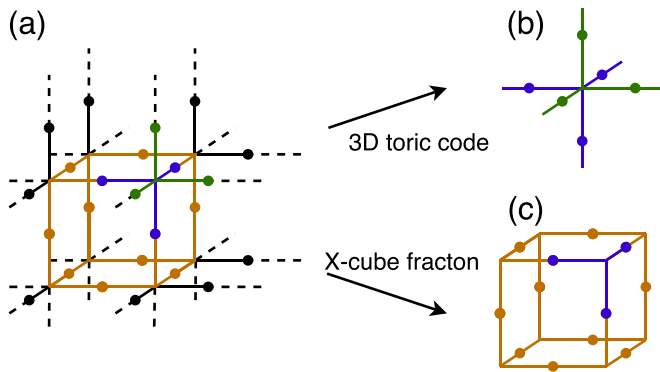


FIG. 14. Composite particles of the 3D toric code and X-cube fracton model. The colors are to help identify how composite particles are defined [i.e., which parts of (a) map to (b) and (c)]. (a) Physical structure of the 3D toric code and X-cube fracton model. (b) Composite particles of the 3D toric code. We define each star as a composite particle and check bond consistency for adjacent particles similarly to the 2D toric code. (c) Composite particles of the X-cube fracton model. We define each cube as a composite particle and check bond consistency on faces of adjacent particles.

excited states break some constraints such that $A_v^i|\psi\rangle = -|\psi\rangle$ or $B_c|\psi\rangle = -|\psi\rangle$ for certain A_v^i, B_c .

The composite particles for the 3D toric code and the X-cube fracton model are illustrated in Fig. 14. For the 3D toric code, a composite particle is made up of six particles associated with a vertex. The ground state can be constructed by initializing the bias of the autoregressive neural network to be all the $|+\rangle$ state and imposing gauge checking on each composite particle to be 1. The excited states can be constructed by forcing even numbers of composite particles to have gauge-checking value -1 . For the X-cube fracton model, a composite particle consists of 12 particles on each mini cube. The ground state comes from initializing all biases of the autoregressive neural network to be $|+\rangle$ and requiring all composite particles to have the gauge-checking value 1. The excited states break the gauge-checking value on sets of four nearby composite particles to be -1 . We numerically verify the exact representations of ground and excited states of the 3D toric code and the X-cube fracton model in Fig. 8, where the energy is shown to be exactly the same as the theoretical predictions.

Our approach can be naturally generalized to the Haah's code fracton [64] and checkerboard fracton [46] models. Similarly to 2D \mathbb{Z}_N gauge theory (see Sec. IV B), one can consider applying gauge-invariant AR-NN to study the 3D \mathbb{Z}_N gauge theory in the context of the 3D toric code or the X-cube fracton model [27] with an external field.

D. $SU(2)_k$ anyonic chain and $SU(2)$ symmetry

Non-Abelian anyons play a crucial role in universal topological quantum computation. Here we consider a chain of Fibonacci anyons, which can be regarded as an $SU(2)_{k=3}$ deformation of the ordinary quantum spin- $\frac{1}{2}$ chain [47]. In this model, there is one type of anyon τ and a trivial vacuum state $\mathbb{1}$ for each site. The constraint from symmetry requires

that τ and $\mathbb{1}$ satisfy the following fusion rule: $\tau \otimes \tau = \tau \oplus \mathbb{1}$, $\tau \otimes \mathbb{1} = \mathbb{1} \otimes \tau = \tau$. We work directly in this basis where each site is either $\mathbb{1}$ or τ , generating an anyonic-symmetric AR-NN. We then proceed to work out the entire phase diagram of the Fibonacci anyons. This can be done particularly efficiently compared with standard Monte Carlo sampling [13] thanks to the exact sampling of the AR-NN.

Our anyonic-symmetric AR-NN is constructed so that it obeys the anyon fusion rule directly by checking two adjacent input configurations and imposing zero amplitude when both are τ . Each anyon is a composite particle and the gauge-checking implements the constraint from the anyon fusion rule.

We consider the Hamiltonian [30]

$$H(\theta) = -\cos\theta \sum_i H_i^{(2)} - \sin\theta \sum_i H_i^{(3)}, \quad (13)$$

where the two-anyon interactions can be described by the golden chain Hamiltonian [29,30]

$$H_i^{(2)} = |\mathbb{1}\tau\mathbb{1}\rangle\langle\mathbb{1}\tau\mathbb{1}| + \phi^{-2}|\tau\mathbb{1}\tau\rangle\langle\tau\mathbb{1}\tau| + \phi^{-1}|\tau\tau\tau\rangle\langle\tau\tau\tau| + \phi^{-3/2}(|\tau\mathbb{1}\tau\rangle\langle\tau\tau\tau| + \text{H.c.}), \quad (14)$$

and the three-anyon interactions can be described by the Majumdar-Gosh chain Hamiltonian [30]

$$H_i^{(3)} = |\mathbb{1}\tau\tau\mathbb{1}\rangle\langle\mathbb{1}\tau\tau\mathbb{1}| + (1 - \phi^{-2})|\tau\tau\tau\tau\rangle\langle\tau\tau\tau\tau| + (1 - \phi^{-1})(|\tau\tau\mathbb{1}\tau\rangle\langle\tau\tau\mathbb{1}\tau| + |\tau\mathbb{1}\tau\tau\rangle\langle\tau\mathbb{1}\tau\tau|) - \phi^{-5/2}(|\tau\mathbb{1}\tau\tau\rangle\langle\tau\tau\tau\tau| + |\tau\tau\mathbb{1}\tau\rangle\langle\tau\tau\tau\tau| + \text{H.c.}) + \phi^{-2}(|\tau\tau\mathbb{1}\tau\rangle\langle\tau\mathbb{1}\tau\tau| + \text{H.c.}), \quad (15)$$

$\phi = (\sqrt{5} + 1)/2$ is the golden ratio.

This model is predicted to exhibit five phases with respect to different θ [30]. Figure 15 shows the optimized energies of the Hamiltonian in Eq. (13) for different θ and the energy derivative computed using the Hellman-Feynman theorem [59]. The nondifferentiable points of the energy derivative indicate the phase transition points, which agree with the conformal field theory prediction. In the special case of $\theta = 0$, the model reduces to the Fibonacci anyons in a golden chain, which has a gapless phase [30]. Using our optimized AR-NN, we compute the second Renyi entropy S_2 [4]. Since the second Renyi entropy S_2 is related to the central charge c under the periodic boundary condition as $S_2 \sim \frac{c}{4} \ln(L)$ with system size L [65], we then extract the central charge finding a value $c = 0.703 \pm 0.005$ very close to the exact result of 0.7 (see Fig. 16).

This can be generalized to the $SU(2)_k$ formulation of anyon theory, for which there are $k + 1$ species of anyons labeled by $j = 0, \frac{1}{2}, 1, \dots, k/2$ with the fusion rule of $SU(2)_k$ [29]. The Hamiltonian can be expressed with operators from the representation of the Temperley-Lieb algebra [29]. To construct an anyonic-symmetric autoregressive neural network for the general $SU(2)_k$ anyonic chain, one works in the angular momentum basis $\{|\dots, j_{i-1}, j_i, j_{i+1}, \dots\rangle\}$ where $j_i \in \{0, \frac{1}{2}, 1, \dots, k/2\}$. Since each j_i is included as the $SU(2)_k$

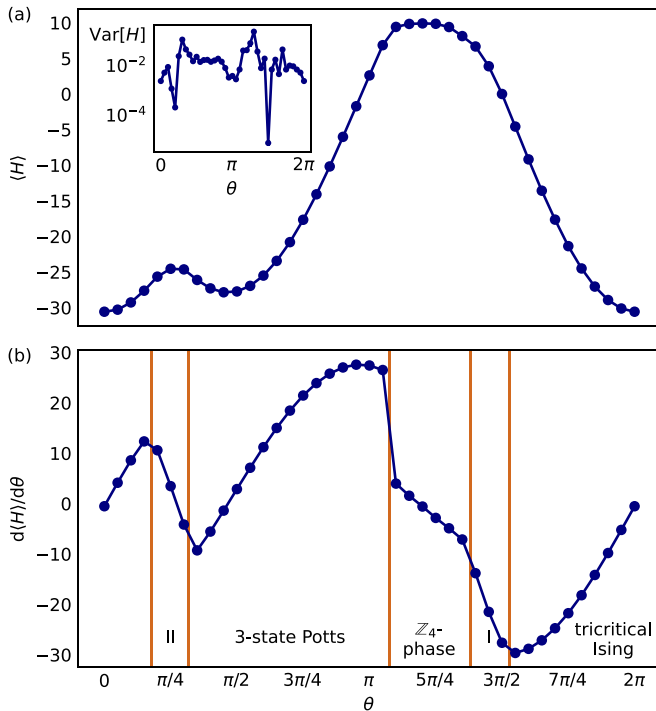


FIG. 15. Phase diagram for 40 anyons with the periodic boundary condition. (a) Energy, (inset) energy variance, and (b) energy derivative, computed using the Hellmann-Feynman theorem [59] as $d\langle H \rangle/d\theta = \langle dH/d\theta \rangle = \sin\theta \sum_i \langle H_i^{(2)} \rangle - \cos\theta \sum_i \langle H_i^{(3)} \rangle$, versus θ . Phase transitions occur when the energy function is not differentiable. Vertical lines in orange are located at the exact phase transition points [30]. We use the 1D RNN with 3 layers, 36 hidden dimensions, and the real-imaginary parametrization. We use the transfer learning technique where we first train the neural network on a 32-anyon model for 3000 iterations and then we transfer the neural network to the 40-anyon model for another 3000 iterations. In each iteration, we use 12 000 samples. The neural network architecture, initialization, and optimization details are discussed in Appendix E.




FIG. 16. The second Renyi entropy S_2 versus the system size L for the optimized AR-NN for the Fibonacci anyons in a golden chain ($\theta = 0$) with the periodic boundary condition. The inset shows the variance of energy of the AR-NN. The slope of the fitted line is the central charge c . The AR-NN is the 1D RNN with 3 layers, L hidden dimensions, and the amplitude-phase parametrization. The neural network is trained for 8000 iterations with a sample size of 12 000. The neural network architecture, initialization, and optimization details are discussed in Appendix E.

fusion rule outcome of j_{i-1} and an extra $\frac{1}{2}$ angular momentum, one can view j_i as a composite particle and gauge checking is the fusion rule.

The Fibonacci anyon is a special case of the $SU(2)_{k=3}$ formulation, considering the mapping $\tau \mapsto j = 1$ and $\mathbb{1} \mapsto j = 0$ and applying $\frac{3}{2} \times j = \frac{3}{2} - j$ from the $SU(2)_3$ fusion rule to the even-number sites [29]. Note that this gives a slightly different AR-NN from what is described above. Aside from the Fibonacci anyon, one can consider the Yang-Lee anyon, which follows the $SU(2)_3$ fusion rule [66].

Using this framework, one can also consider the Heisenberg spin chain with $SU(2)$ symmetry since it can be considered as the $SU(2)_k$ deformation of the ordinary quantum spin- $\frac{1}{2}$ chain [47] as $k \rightarrow \infty$. In Appendix D, we provide the detailed construction of an $SU(2)$ -invariant autoregressive neural network for the Heisenberg model, which can be viewed as the case of $SU(2)_{k=\infty}$, and obtain accurate results for the 1D Heisenberg model.

V. CONCLUSION

We have provided a general approach to constructing gauge-invariant or anyonic-symmetric autoregressive neural network wave functions for various quantum lattice models. These wave functions explicitly satisfy the gauge or algebraic constraints, allow for perfect sampling of configurations, and are capable of explicitly returning the amplitude of a configuration including normalization. To accomplish this, we have upgraded standard AR-NN in such a way that the constraints can be autoregressively satisfied.

We have given explicit constructions of AR-NN which exactly represent the ground and excited states of several models, including the 2D and 3D toric codes as well as the X-cube fracton model. For those models for which exact representations are unknown, we variationally optimize our symmetry-incorporated AR-NN to obtain either high-quality ground states or time-dependent wave functions. This has been done for the $U(1)$ quantum link model, \mathbb{Z}_N gauge theory, the $SU(2)_3$ anyonic chain, and the $SU(2)$ quantum spin- $\frac{1}{2}$ chain. For these systems we are able to measure dynamical properties, produce phase diagrams, and compute observables accurately.

Our approach opens up the possibility of probing a larger variety of models and the physics associated with them. For example, the higher spin representation $S > \frac{1}{2}$ in the $(1+1)$ D QLM models would allow one to probe the quantum chromodynamics related physics of confinement and string breaking [67]. For the $(3+1)$ D QLM models, there is the Coulomb phase which manifests in pyrochlore spin liquids [68]. For \mathbb{Z}_N gauge theory, it will be interesting to consider the general \mathbb{Z}_N toric code with transverse field or disorder, with a goal of understanding its phase diagram. Recently, there have been proposals to understand the \mathbb{Z}_N X-cube fracton model with nontrivial statistical phases [27]. For non-Abelian anyons, the general $SU(2)_k$ formulation exhibits rich physics for different k and one can study the corresponding topological liquids and edge states [47]. Our approach can also be extended to study the phase diagram for the 2D Heisenberg models with $SU(2)$ symmetry.

Aside from exploring various models in condensed-matter physics and high-energy physics, our approach can also be

further applied to quantum information and quantum computation. Fibonacci anyons are known to support universal topological quantum computation, which is robust to local perturbations [69]. It will be interesting to see how well one can approximately simulate topological quantum computation or different braiding operations with anyonic-symmetric autoregressive neural networks. As toric codes are an important example of quantum error correction code, our approach can be used to approximately study the performance of a toric code under different noise conditions. With respect to the recent efforts on simulating lattice gauge theories with quantum computation, our approach also provides an alternative method to compare to and benchmark quantum computers. In summary, the approach we have developed is versatile and powerful for investigating condensed-matter physics, high-energy physics, and quantum information science.

ACKNOWLEDGMENTS

D.L. is grateful for insightful discussion in high-energy physics with J. Stokes and J. Shen. D.L. acknowledges helpful discussion with L. Yeo, O. Dubinkin, R. Levy, P. Xiao, R. Sun, and G. Carleo. This work is supported by the National Science Foundation under Cooperative Agreement No. PHY-2019786 (the NSF AI Institute for Artificial Intelligence and Fundamental Interactions [70]). This work utilizes resources supported by the National Science Foundation’s Major Research Instrumentation program, Grant No. 1725729, as well as the University of Illinois at Urbana-Champaign [71]. The authors acknowledge MIT Satori and MIT SuperCloud [72] for providing HPC resources that have contributed to the research results reported within this paper. Z.Z. is partially supported by NSF DMS-1854791, NSF OAC-1934757, and Alfred P. Sloan Foundation. Vera Mikyoung Hur is partially supported by NSF DMS-1452597 and DMS-2009981. This work is supported in part by the U.S. Department of Energy, Office of Science, Office of High Energy Physics QuantISED program under an award for the Fermilab Theory Consortium “Intersections of QIS and Theoretical Particle Physics”.

APPENDIX A: ADDITIONAL RESULTS FOR QUANTUM LINK MODEL AND TORIC CODE MODEL

In this Appendix, we present additional results. Figure 17 is a 6-unit-cell quantum link model dynamics, which we can exactly diagonalize. We observed that the gauge-invariant AR-NN matches the exact results up to $t = 4$, while the non-gauge *Ansatz* quickly fails to capture the electric fields even with the sign gradient (SG) optimizer [73] [Fig. 17(b)]. In addition, the gauge-invariant AR-NN in general has a lower per-step infidelity [Fig. 17(c)]. In Fig. 18, we extrapolate the ground state energy using variance extrapolation method for the 160 unit cell QLM. We find that our RNN produces variational answer close to the extrapolated result within $\Delta E \sim 0.01$.

In Fig. 19, we measure the local observables of the 12×12 toric code model. We show that even though the neural network does not automatically preserve translational symmetry, the optimization drives the neural network to a translationally




FIG. 17. Dynamics for the 6-unit-cell (12 sites and 12 links) open-boundary QLM for $m = 0.1$ and 2.0 with and without gauge-invariant construction. The dashed curves are the exact results from the exact diagonalization for 6 unit cells. The “6-cell Gauss SG” is the same as the “6-cell Gauss” in Fig. 5. (a) The change in the energy during the dynamics. (b) The expectation value of the electric field averaged over all links. (c) The per-step infidelity measure, where $|\Psi\rangle$ and $|\Phi\rangle$ are defined in Sec. III. We use the transformer neural network with 1 layer, 16 hidden dimensions, and the real-imaginary parametrization (see Fig. 24). The initial state is $|\bullet \rightarrow \circ \rightarrow\rangle$ for each unit cell and we train the neural network using the forward-backward trapezoid method with the time step $\tau = 0.005$, 600 iterations in each time step, and 12000 samples in each iteration. For the results labeled with SG, we used the sign gradient (SG) optimizer [73] for 15–30 iterations (depending on the resulting fidelity) before switching to the regular optimizer. The neural network architecture, initialization, and optimization details are discussed in Appendix E.

symmetric state. In addition, the weights of RNN are translationally invariant, which, although not guaranteed, could be potentially useful for preserving translational symmetry.




FIG. 18. Ground-state energy extrapolation at $m = 0.7$ for 160 unit cells. The energies and variances are obtained from the training output of RNN in Fig. 4. We used a linear fit for variances smaller than 0.2 and obtained a y intercept of -199.7923 . Our variational ground-state energy is -199.7803 ± 0.0005 . The *Ansatz*, initialization, and optimization are the same as in Fig. 4 and are discussed in Appendix E.

Furthermore, we benchmark our method on the following model:

$$H = - \sum_{v \in V} A_v - \sum_{p \in P} B_p - h \sum_{e \in E} \sigma_e^z - j_y \sum_{p \in P} \prod_{e \in p} \sigma_e^y, \quad (\text{A1})$$

With the additional $\sum_{p \in P} \prod_{e \in p} \sigma_e^y$, the Hamiltonian exhibits a sign problem compared to the original toric model, which would be challenging for the Monte Carlo method. We further test our method first on small systems (3×3 to compare against exact diagonalization) and find relative energy differences on the order of 10^{-5} suggesting good agreement (see Fig. 20). We further apply our method on a 10×10 lattice. Although it is not clear what to benchmark exactly against here, we measure the energy difference from a variance extrapolated result, and find our variational answer is close to it within $\Delta E \sim 0.1$ (see Fig. 21).

APPENDIX B: KITAEV'S $D(G)$ MODEL AND EXACT REPRESENTATION OF GROUND STATE

We generalize our gauge-invariant autoregressive construction for the 2D \mathbb{Z}_2 toric code. Kitaev's $D(G)$ model [26] is defined on an $L \times L$ periodic square lattice where each edge has a basis $\{|g\rangle, g \in G\}$ for some group G . Here we focus on finite groups, in particular $G = \mathbb{Z}_N$ for \mathbb{Z}_N theory. Without loss of generality, we attach an upward arrow for each edge in the y direction and a right arrow for each edge in the x direction. We employ the notation of Sec. IV B and introduce operators A_v^g and $B_p^{h_u, h_d, h_l, h_r}$ as in Fig. 22. The Hamiltonian defined on $\mathcal{H}(G)^{\otimes E}$ is

$$H = - \sum_{v \in V} A_v - \sum_{p \in P} B_p, \quad (\text{B1})$$

where $A_v = \frac{1}{|G|} \sum_{g \in G} A_v^g$ is Gauss's law and the gauge constraint, and $B_p = \sum_{h_u, h_r, h_d, h_l = \mathbb{1}_G} B_p^{h_u, h_r, h_d, h_l}$.

Let $|+\rangle = \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle$, and $|\psi\rangle = \prod_{p \in P} B_p |+\rangle^{\otimes E}$ is the ground state. This is because $|\psi\rangle$ is a ground state for each A_v




FIG. 19. Local observables: (a) Expectation value of $\langle \sigma_x \rangle$, (b) $\langle \sigma_z \rangle$, (c) vertex operator $\langle A_v \rangle$, and (d) plaquette operator $\langle B_p \rangle$ [defined in Eq. (11)] for the 12×12 toric code model with $h = 0.36$. The neural network is the same as in Fig. 12. The neural network architecture, initialization, and optimization details are discussed in Appendix E.

and B_p . It is easy to verify that $B_p |\psi\rangle = |\psi\rangle$. To see $A_v |\psi\rangle = |\psi\rangle$, notice that A_v and B_p commute and $A_v |+\rangle^{\otimes E} = |+\rangle^{\otimes E}$. Similarly to the \mathbb{Z}_2 toric code, the ground state can be constructed using gauge-invariant autoregressive neural networks by defining each star as a composite particle and checking Gauss's law and bond consistency.




FIG. 20. (a) Relative error in energy and (b) variance of energy for 3×3 toric code model with an additional term described in Eq. (A1). Here we choose $h = 0.36$ and run the neural network for different J_y 's. We use the 2D RNN neural network with real-imaginary parametrization. The neural network is trained using Adam optimizer for 10000 iterations with 12000 samples in each iteration. The neural network architecture, initialization, and optimization details are discussed in Appendix E.

APPENDIX C: (1 + 1)D \mathbb{Z}_N LATTICE QED MODEL

Our approach in Sec. IV A can be applied to the (1 + 1)D \mathbb{Z}_N lattice quantum electrodynamics (QED) model, which is a discretization of the Schwinger model for the continuous-space QED in 1+1 dimensions [74]. The (1 + 1)D \mathbb{Z}_N model takes a similar form as the (1 + 1)D QLM, which has fermions on sites and electric fields on links between two sites. Let $\{|e_{i,i+1}\rangle\}$ for $1 \leq e_{i,i+1} \leq N$ denote the orthonormal basis on each link $(i, i + 1)$. The (1 + 1)D \mathbb{Z}_N gauge theory can take the form [74]

$$H = - \sum_i [\psi_i^\dagger U_{i,i+1} \psi_{i+1} + \psi_{i+1}^\dagger U_{i,i+1}^\dagger \psi_i] + m \sum_i (-1)^i \psi_i^\dagger \psi_i + \frac{g^2}{8} \sum_i (V_{i,i+1} - \mathbb{1})(V_{i,i+1}^\dagger - \mathbb{1}), \quad (\text{C1})$$

where $U_{i,i+1}|e_{i,i+1}\rangle = |(e_{i,i+1} + 1) \bmod N\rangle$, and $V_{i,i+1}|e_{i,i+1}\rangle = e^{-i2\pi m/N} |e_{i,i+1}\rangle$ for $m = e_{i,i+1}$. The Gauss's law operator G_i of the model can be written as

$$G_i = e^{\frac{i2\pi}{N} [\psi_i^\dagger \psi_i + \frac{1}{2}(-1)^i - \frac{1}{2}]} V_{i,i+1} V_{i-1,i}^\dagger \quad (\text{C2})$$

such that $G_i|\psi\rangle = |\psi\rangle$ for each i [74].

Similarly to the (1 + 1)D QLM, one can construct the gauge-invariant autoregressive neural network as Fig. 2 and perform gauge checking with G_i in Eq. (C2).




FIG. 21. Ground-state energy extrapolation at $h = 0.36$ and $J_y = 0.3$ for 10×10 modified toric code model. The energies and variances are obtained from the training output. We used a linear fit for variances smaller than 0.2 and obtained a y intercept of -213.938 . Our variational ground-state energy is -213.802 ± 0.003 . We use the 2D RNN neural network with real-imaginary parametrization. The neural network is trained using the transfer learning technique for 5000 iterations after the 3×3 result with 12000 samples in each iteration. The neural network architecture, initialization, and optimization details are discussed in Appendix E.

APPENDIX D: SU(2)-INVARIANT AUTOREGRESSIVE NEURAL NETWORK FOR HEISENBERG MODEL

The 1D Heisenberg Model is described as

$$H = \sum_i \sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y + \sigma_i^z \sigma_{i+1}^z. \quad (\text{D1})$$

We work in the angular momentum basis $\{|j_1, j_2, j_3, \dots, j_n\rangle\}$ similarly to [47], instead of the spin basis, to construct an SU(2)-invariant autoregressive wave function. Each j_i is the total angular momentum quantum number for spins from 1 to i and $j_n \equiv J$ is the total angular momentum quantum number for all spins. For the ground state of the Heisenberg model, the total angular momentum is zero, so $j_n = 0$. We define the first composite particle as j_1 and the i th composite particle as the difference $j_i - j_{i-1}$. Note that this uniquely defines a physical state. We then autoregressively enforce $j_{i < n} \geq 0$ and $j_n = 0$ as gauge checking, to achieve the SU(2)-invariant property.

Figure 23 demonstrates the performance of our SU(2)-invariant autoregressive neural network on the Heisenberg model with SU(2) symmetry.

$$A_v^g = \begin{array}{c|c} L_{+,u}^g & L_{+,r}^g \\ \hline v & \\ \hline L_{-,d}^g & \end{array} \quad B_p^{h_u, h_r, h_d, h_l} = \begin{array}{c} T_{-,u}^{h_u} \\ \square \\ T_{+,d}^{h_d} \end{array} \quad T_{+,r}^{h_r}$$

FIG. 22. A_v^g and $B_p^{h_u, h_r, h_d, h_l}$ operators. $A_v^g = L_{+,u}^g L_{+,r}^g L_{-,d}^g L_{-,l}^g$ and $B_p^{h_u, h_r, h_d, h_l} = T_{-,u}^{h_u} T_{+,r}^{h_r} T_{+,d}^{h_d} T_{-,l}^{h_l}$, where $L_+^g|z\rangle = |gz\rangle$, $L_-^g|z\rangle = |zg^{-1}\rangle$, $T_+^h|z\rangle = h\delta_{h,z}|z\rangle$, and $T_-^h|z\rangle = h^{-1}\delta_{h,z}|z\rangle$.




FIG. 23. Relative error of variational ground-state energy for the 1D Heisenberg model with SU(2) symmetry for 22, 24, and 100 spins. We use the 1D RNN with 3 layers, L hidden dimensions, and the real-imaginary parametrization. We train the neural network for 5000 iterations with 12 000 samples in each iteration. The exact solutions for 22 spins and 24 spins are computed with the exact diagonalization, and the exact solution for 100 spins is the density matrix renormalization group (DMRG) result in Ref. [4].

APPENDIX E: NEURAL NETWORK ARCHITECTURE

1. Complex parametrization

Wave functions are complex in general but both the transformer network and 1D/2D RNN are real. We use two approaches [Fig. 24(a) (amplitude phase) and Fig. 24(b) (real-imaginary phase)] to parametrize complex wave functions from real neural networks. In both parametrizations, the input configuration \tilde{x} , together with a default configuration \tilde{x}_0 , is embedded (i.e., each state of a composite particle is mapped

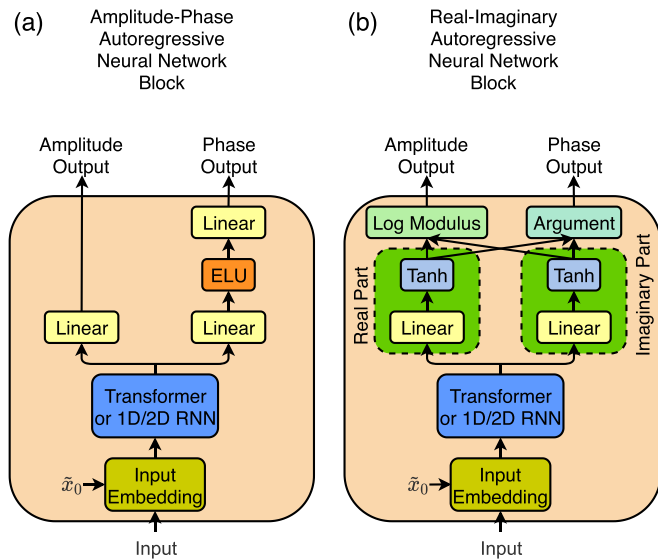


FIG. 24. Two parametrizations of complex wave functions from autoregressive neural networks. (a) The amplitude-phase parametrization. The raw output is used as the input of both the amplitude branch and the phase branch. (b) The real-imaginary parametrization. The raw output is used as the input of both the real branch and the imaginary branch, which later is converted to the amplitude branch and the phase branch.




FIG. 25. A single-layer transformer network. The embedded input is fed into the transformer and the positional encoding is added. After a masked multihead self-attention is applied, a feed-forward layer produces the raw output.

to a unique vector) before fed into the transformer or 1D/2D RNN. Certain gauge blocks in an AR-NN take a default state \tilde{x}_0 as opposed to any state of the composite particles; the embedded vector of this default state has arbitrary parameters which are trained during optimization.

2. Transformer

The transformer used in this work (Fig. 25) is the same as the transformer used in Ref. [8], which can be viewed as the standard transformer encoder with masked multihead attention from Ref. [41] but without an additional add and norm layer. The transformer consists of a standard positional encoding layer, which uses sinusoidal functions to encode the positional information of the embedded input. After positional encoding, the input is fed into the standard masked multihead attention mechanism. The mask here is crucial for autoregressiveness, as it only allows each site to depend on the previous sites. The output of the attention layer is then passed through a standard feed-forward layer. The detailed explanation of the transformer can be found in Refs. [8,41]. This transformer is essentially equivalent to the standard PYTORCH implementation [49], but was implemented independently because that implementation did not exist at the start of our work.

3. RNN cells

For all RNNs in this work, we used the gated recurrent unit (GRU) cell [38] (Fig. 26) in PYTORCH [49], which takes one input vector x_k , the hidden input h_{k-1} , and computes

$$\begin{aligned}
 r &= \sigma(W_{xr}x_k + b_{xr} + W_{hr}h_{k-1} + b_{hr}), \\
 z &= \sigma(W_{xz}x_k + b_{xz} + W_{hz}h_{k-1} + b_{hz}), \\
 n &= \tanh[W_{xn}x_k + b_{xn} + r \odot (W_{hn}h_{k-1} + b_{hn})], \\
 h_k &= (1 - z) \odot n + z \odot h_{k-1}, \\
 y_k &= h_k,
 \end{aligned}
 \tag{E1}$$

where σ is the sigmoid function and \odot means elementwise product.




FIG. 26. The GRU cell [38] on which different RNNs are constructed. This is the same GRU cell as the PYTORCH [49] implementation.

We then build 1D and periodic 2D RNN cells (Fig. 27) based on the GRU cell. The 1D RNN cell computes

$$(y_{\text{raw}}, h_{\text{new}}) = \text{GRU cell}(x_k, h_{\text{old}}), \quad y_k = y_{\text{raw}} + x_k, \quad (\text{E2})$$

whereas the periodic 2D RNN cell computes

$$\begin{aligned} (y_{\text{raw}}, h_{\text{raw}}) &= \text{GRU cell}(x_k, [h_{\text{old}1}, h_{\text{old}2}, h_{\text{old}3}, h_{\text{old}4}]), \\ [h_{\text{new}1}, h_{\text{new}2}, h_{\text{new}3}, h_{\text{new}4}] &= h_{\text{raw}} + x_k, \\ h_{\text{new}} &= \frac{1}{4}(h_{\text{new}1} + h_{\text{new}2} + h_{\text{new}3} + h_{\text{new}4}), \\ y_k &= y_{\text{raw}} + x_k. \end{aligned} \quad (\text{E3})$$

4. RNNs

With the RNN cells, we can build 1D and periodic 2D RNNs.

The 1D RNN (Fig. 28) has a multilayer design and shares the same structure as the PYTORCH [49] GRU [38]. The embedded input configuration is fed into the cells one at a time through multiple layers and produces a raw output. In our work, the cells at different layers share the weight matrices and bias vectors.

The periodic 2D RNN has a more complicated design to capture the most correlations and can be viewed as a periodic extension of the 2D RNN in Ref. [4]. In each layer, the hidden vector h is passed around according to Fig. 29(a), where each

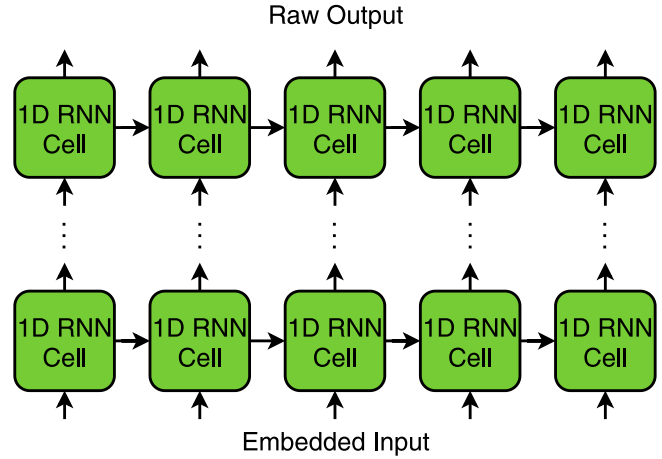


FIG. 28. 1D RNN built from 1D RNN Cells. The neural network has a multilayer design similar to PYTORCH GRU implementation [38,49]. The weight matrices and biases are shared between different layers.

cell receives a maximum number of four hidden vectors and concatenates them according to Fig. 27(b). When the number of hidden vectors received is less than four, zero vectors are used to pad the concatenated hidden vector to the correct length. The configuration is evaluated and sampled in a zigzag S path [Fig. 29(b)] to ensure autoregressiveness.

Before the first layer of the periodic 2D RNN, a special concatenation of the embedded input needs to be performed. At each location, the concatenation layer takes the four surrounding inputs (periodically) and concatenates them into a single vector. If any (or all) of the surrounding inputs lie later in the conditioning order in Fig. 29(b), the corresponding input is replaced with a default input x_0 . For a 4×4 2D input array shown in Fig. 30(a), some concatenation examples are shown in Fig. 30(b). After the first layer, the output of a previous layer can be directly fed into the next layer similar to a regular RNN without any further process.

A multilayer periodic 2D RNN consists of one input concatenation layer and several 2D RNN layers as shown in Fig. 31.




FIG. 27. (a) 1D RNN cell. A ResNet (skip connection) [75] is added between the input x_k and output y_k . To be noted that x_k , y_k , h_{old} , and h_{new} have the same dimension. (b) 2D RNN cell with periodic boundary condition. This cell requires four hidden inputs $h_{\text{old}i}$ ($i = 1, 2, 3, 4$) and generates one hidden output h_{new} . Skip connections are added for both output y_k and hidden output h_{new} . The average pooling reduces the hidden output to have the same dimension as each hidden input. To be noted that the dimension of x_k and y_k is four times the dimension of each $h_{\text{old}i}$ and h_{new} .

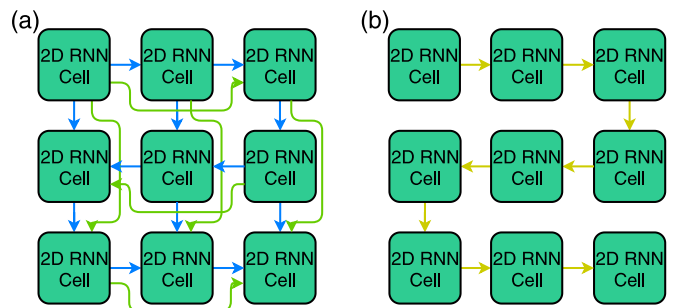


FIG. 29. (a) The hidden information path for 2D RNN with periodic boundary condition for a 3×3 system. Blue arrows show the nonboundary information path whereas the green arrows show the periodic boundary information path. When there are less than four hidden inputs, zero vectors are used for padding. (b) The conditioning and sampling order of the 2D RNN.




FIG. 30. 2D RNN input concatenation layer. For (a) a 3×3 input array with a default input x_0 , (b) the concatenation layer takes the four input vectors surrounding each site with periodic boundary condition and outputs the concatenated vector of the four surrounding vectors. x_0 is used when the surrounding inputs appear later in the conditioning order.

5. Initialization and optimization

We use different initialization techniques for different models. For the QLM model, the initialization is done through tomography, minimizing $-\ln |\psi(x)|^2$ for a desired configuration x ($|\bullet \rightarrow \circ \rightarrow$) for each unit cell in this case). For the 2D toric code model, we set the weight matrix in the last linear layer to be 0 and the bias such that the wave function for each composite particle is $\sqrt{0.23}|0_0^0\rangle_+ + \sqrt{0.11}|1_0^0\rangle_+ + \sqrt{0.11}|1_0^1\rangle_+ + \sqrt{0.11}|0_1^0\rangle_+ + \sqrt{0.11}|0_1^1\rangle_+ + \sqrt{0.11}|1_1^0\rangle_+ + \sqrt{0.11}|1_1^1\rangle_+$, which empirically produces a very low initial energy. We used a transfer learning technique, where we first train our neural network on a 6×6 model before training it on the 10×10 model. When transferred to the larger system, the weight and bias in the last linear layer is dropped and replaced with the initialization scheme described above. In Fig. 32, we show the effect of transfer learning on a 10×10 toric code model. With transfer learning, the energy is clearly lower than the energy without transfer learning.

For the anyon model, we set the weight matrix in the last linear layer to be 0 and the bias to be the state of $1/\sqrt{2}|1\rangle + 1/\sqrt{2}|\tau\rangle$ for each particle. When producing




FIG. 31. The 2D RNN is built from one input concatenation layer and multiple 2D RNN layers. The weight matrices and biases are shared between different layers.

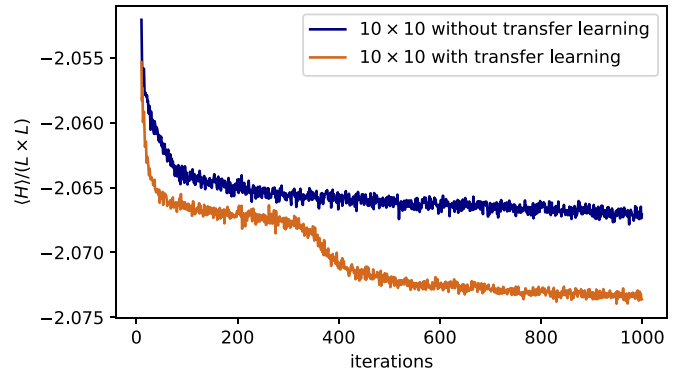


FIG. 32. Per-site energy with and without transfer learning during the training process for the 10×10 toric code model. The first 10 iterations are not included as they are too large and outside the range of the figure. The energy with transfer learning is clearly lower than the energy without transfer learning.

the phase diagram, we used a transfer learning technique, where we first train the neural network on 32 anyons with $\theta = 0, \pi/4, \pi/2, 3\pi/4, \pi, 5\pi/4, 3\pi/2, 7\pi/4,$ and 2π for 3000 iterations, and then transfer the model with θ that is closest to the desired value of θ for 40 anyons for another 3000 iterations. Similar to the toric code case, when transferred to the larger system, the weight and bias in the last linear layer are dropped and replaced by the initialization described above. In all models, except the last layer, the weights and biases are initialized using PYTORCH's [49] default initialization.

For optimization, we used the Adam [76] optimizer with an initial learning rate of 0.01. For the QLM dynamics, the learning rate is halved at iterations 100, 200, 270, 350, and 420, for the QLM ground-state optimization; the learning rate is halved at iterations 300, 600, 900, 1200, 1800, 2400, 3000, 4000, 5000, 6000, and 7000; and for the ground-state optimization of other models, the learning rate is halved at iterations 100, 500, 1000, 1800, 2500, 4000, and 6000. In addition, for the 6-unit-cell cases and 12-unit-cell $m = 0.1$ case with Gauss's law, we use the sign gradient (SG) optimizer [73] for 15–30 iterations (depending on the resulting fidelity) before switching to the regular optimizer, and for the 12-unit-cell $m = 2.0$ case with Gauss's law, we modified the loss function by adding an energy penalty term described in Appendix F.

6. Computational complexity

In this section, we explain the computational complexity of the neural networks used in this work. In terms of scaling, the total cost per sweep is as follows:

(i) transformer: $O(N^2 h^3)$ (evaluation complexity), $O(N^3 h^3)$ (sampling complexity);

(ii) RNN: $O(Nh^2)$ (computational complexity),

where h is the hidden dimension and N is the size of the system. Note that the memory complexity is bounded by the computational complexity. In the training process, we used 2–4 GPUs depending on the availability of GPUs in the cluster and never experienced any memory issue with a total

of 64GB GPU memory. To give a sense for computational difficulty, generating Fig. 12 takes six GPU days using Tesla V100 GPUs.

APPENDIX F: ENERGY PENALTY

While quantum dynamics should exactly preserve the energy, as a practical matter when a variational state cannot fully represent the exact dynamics, there can be a tension between maximizing fidelity per step and preserving the energy. In some cases, it may be desirable to better preserve the total energy at some cost in fidelity per step. Toward that end, we show one can introduce an additional term into the loss function which acts as a penalty toward the drift in energy. We demonstrate this for the $m = 2.0$ 12-unit-cell QLM in Fig. 5. The energy penalty term

$$\mathcal{L}_p = \left| \frac{1}{N} \sum_{x \sim |\psi_\theta|^2} E_{\text{loc}}(x) - E_0 \right|^2, \quad (\text{F1})$$

where $E_{\text{loc}}(x)$ is defined in Eq. (1) and E_0 is the initial energy. This term is added to the dynamics loss function \mathcal{L}_d [Eq. (7)] to obtain the total loss function as

$$\mathcal{L} = \mathcal{L}_d + \alpha \mathcal{L}_p, \quad (\text{F2})$$

with α is a hyperparameter which we choose to be 0.01. We show in Fig. 33 this simulation with and without the energy penalty. We find the dynamics as measured by the observables largely stays the same (and may be better), but the drift in the energy is significantly attenuated.

APPENDIX G: DERIVATION OF STOCHASTIC GRADIENTS FOR VARIATIONAL AND DYNAMICS OPTIMIZATION

In Sec. III, we presented stochastic gradients of the variational and dynamics optimizations. This section includes their derivations. The variational optimization has been widely used and derived many times in other works [1,4]. Here we present the derivation for the sake of completeness:

$$\begin{aligned} \frac{\partial \langle \psi_\theta | H | \psi_\theta \rangle}{\partial \theta} &= \left\langle \frac{\partial \psi_\theta}{\partial \theta} \middle| H \middle| \psi_\theta \right\rangle + \left\langle \psi_\theta \middle| H \middle| \frac{\partial \psi_\theta}{\partial \theta} \right\rangle \\ &= 2 \sum_x \Re \left\{ \frac{\partial \psi_\theta^*(x)}{\partial \theta} H \psi_\theta(x) \right\} \\ &= 2 \sum_x \Re \left\{ \frac{1}{\psi_\theta^*(x)} \frac{\partial \psi_\theta^*(x)}{\partial \theta} \psi_\theta^*(x) H \psi_\theta(x) \right\} \\ &= 2 \sum_x \Re \left\{ \psi_\theta^*(x) H \psi_\theta(x) \frac{\partial}{\partial \theta} \ln \psi_\theta^*(x) \right\} \\ &\approx \frac{2}{N} \sum_{x \sim |\psi_\theta|^2} \Re \left\{ \frac{H \psi_\theta(x)}{\psi_\theta(x)} \frac{\partial}{\partial \theta} \ln \psi_\theta^*(x) \right\} \\ &\equiv \frac{2}{N} \sum_{x \sim |\psi_\theta|^2} \Re \left\{ E_{\text{loc}}(x) \frac{\partial}{\partial \theta} \ln \psi_\theta^*(x) \right\}, \quad (\text{G1}) \end{aligned}$$

where the local energy is $E_{\text{loc}}(x) \equiv H \psi_\theta(x) / \psi_\theta(x)$. We can further control the variance by subtracting from the $E_{\text{loc}}(x)$




FIG. 33. Dynamics for the 6- and 12-unit-cell (12–24 sites and 12–24 links) open-boundary QLM for $m = 2.0$ with and without energy penalty. The dashed curves are the exact results from the exact diagonalization for 6 unit cells. The “12-cell Gauss EP” is the same as the “12-cell Gauss” in Fig. 5. (a) The change in the energy during the dynamics. (b) The expectation value of the electric field averaged over all links. (c) The per-step infidelity measure, where $|\Psi\rangle$ and $|\Phi\rangle$ are defined in Sec. III. We use the transformer neural network with 1 layer, 16 hidden dimensions for 6 unit cells, and 32 hidden dimensions for 12 unit cells, and the real-imaginary parametrization (see Fig. 24). The initial state is $|\bullet \rightarrow \circ \rightarrow\rangle$ for each unit cell and we train the neural network using the forward-backward trapezoid method with the time step $\tau = 0.005$, 600 iterations in each time step, and 12 000 samples in each iteration. The neural network architecture, initialization, and optimization details are discussed in Appendix E.

the average energy $E_{\text{avg}} \equiv \sum_{x \sim |\psi_\theta|^2} E_{\text{loc}}(x) / N$ over the batch of samples [48] as we did in Sec. III and use the stochastic variance reduced gradient as

$$\frac{2}{N} \sum_{x \sim |\psi_\theta|^2} \Re \left\{ [E_{\text{loc}}(x) - E_{\text{avg}}] \frac{\partial}{\partial \theta} \ln \psi_\theta^*(x) \right\}, \quad (\text{G2})$$

which has the same expectation as Eq. (G1).

For the dynamics optimization gradient, as in Sec. III, we define $|\Psi_\theta\rangle = (1 + iH\tau)|\psi_{\theta(t+2\tau)}\rangle$ and $|\Phi\rangle = (1 - iH\tau)|\psi_{\theta(t)}\rangle$, and we drop $\theta(t)$ and name $\theta \equiv \theta(t + 2\tau)$. We start by splitting the negative logarithm overlap:

$$-\ln \frac{\langle \Psi_\theta | \Phi \rangle \langle \Phi | \Psi_\theta \rangle}{\langle \Psi_\theta | \Psi_\theta \rangle \langle \Phi | \Phi \rangle} = -\ln \langle \Psi_\theta | \Phi \rangle - \ln \langle \Phi | \Psi_\theta \rangle + \ln \langle \Psi_\theta | \Psi_\theta \rangle + \ln \langle \Phi | \Phi \rangle. \quad (\text{G3})$$

We then compute the gradient term by term. The first term on the right side of Eq. (G3) becomes

$$\begin{aligned} \frac{\partial}{\partial \theta} \ln \langle \Psi_\theta | \Phi \rangle &= \frac{1}{\langle \Psi_\theta | \Phi \rangle} \left\langle \frac{\partial \Psi_\theta}{\partial \theta} \middle| \Phi \right\rangle = \frac{1}{\sum_x \Psi_\theta^*(x) \Phi(x)} \sum_x \frac{\partial \Psi_\theta^*(x)}{\partial \theta} \Phi(x) = \frac{1}{\sum_x \Psi_\theta^*(x) \Phi(x)} \sum_x \Psi_\theta^*(x) \Phi(x) \frac{\partial}{\partial \theta} \ln \Psi_\theta^*(x) \\ &\approx \frac{1}{\sum_{x \sim |\psi_\theta|^2}^N \frac{\Psi_\theta^*(x) \Phi(x)}{|\psi_\theta(x)|^2}} \sum_{x \sim |\psi_\theta|^2}^N \frac{\Psi_\theta^*(x) \Phi(x)}{|\psi_\theta(x)|^2} \frac{\partial}{\partial \theta} \ln \Psi_\theta^*(x) \equiv \frac{1}{N} \sum_{x \sim |\psi_\theta|^2}^N \frac{\alpha(x)}{\alpha_{\text{avg}}} \frac{\partial}{\partial \theta} \ln \Psi_\theta^*(x), \end{aligned} \quad (\text{G4})$$

where $\alpha(x) = \Psi_\theta^*(x) \Phi(x) / |\psi_\theta(x)|^2$ and $\alpha_{\text{avg}} = \sum_{x \sim |\psi_\theta|^2}^N \alpha(x) / N$ as in Sec. III. The second term on the right side of Eq. (G3) is just the complex conjugate of the first term, whereby

$$\frac{\partial}{\partial \theta} \ln \langle \Phi | \Psi_\theta \rangle = \left[\frac{\partial}{\partial \theta} \ln \langle \Psi_\theta | \Phi \rangle \right]^* \approx \frac{1}{N} \sum_{x \sim |\psi_\theta|^2}^N \left[\frac{\alpha(x)}{\alpha_{\text{avg}}} \frac{\partial}{\partial \theta} \ln \Psi_\theta^*(x) \right]^*. \quad (\text{G5})$$

The third term on the right side of Eq. (G3) becomes

$$\begin{aligned} \frac{\partial}{\partial \theta} \ln \langle \Psi_\theta | \Psi_\theta \rangle &= \frac{1}{\langle \Psi_\theta | \Psi_\theta \rangle} \left(\left\langle \frac{\partial \Psi_\theta}{\partial \theta} \middle| \Psi_\theta \right\rangle + \left\langle \Psi_\theta \middle| \frac{\partial \Psi_\theta}{\partial \theta} \right\rangle \right) = \frac{2}{\sum_x |\Psi_\theta(x)|^2} \sum_x \Re \left\{ \frac{\partial \Psi_\theta^*(x)}{\partial \theta} \Psi_\theta(x) \right\} \\ &= \frac{2}{\sum_x |\Psi_\theta(x)|^2} \sum_x \Re \left\{ |\Psi_\theta(x)|^2 \frac{\partial}{\partial \theta} \ln \Psi_\theta^*(x) \right\} \approx \frac{2}{\sum_{x \sim |\psi_\theta|^2}^N \frac{|\Psi_\theta(x)|^2}{|\psi_\theta(x)|^2}} \sum_{x \sim |\psi_\theta|^2}^N \Re \left\{ \frac{|\Psi_\theta(x)|^2}{|\psi_\theta(x)|^2} \frac{\partial}{\partial \theta} \ln \Psi_\theta^*(x) \right\} \\ &\equiv \frac{2}{N} \sum_{x \sim |\psi_\theta|^2}^N \Re \left\{ \frac{\beta(x)}{\beta_{\text{avg}}} \frac{\partial}{\partial \theta} \ln \Psi_\theta^*(x) \right\}, \end{aligned} \quad (\text{G6})$$

where $\beta(x) = |\Psi_\theta(x)|^2 / |\psi_\theta(x)|^2$ and $\beta_{\text{avg}} = \sum_{x \sim |\psi_\theta|^2}^N \beta(x) / N$ as in Sec. III. The last term on the right side of Eq. (G3) is θ independent such that

$$\frac{\partial}{\partial \theta} \ln \langle \Phi | \Phi \rangle = 0. \quad (\text{G7})$$

Combining all the terms together,

$$\frac{\partial}{\partial \theta} \left(-\ln \frac{\langle \Psi_\theta | \Phi \rangle \langle \Phi | \Psi_\theta \rangle}{\langle \Psi_\theta | \Psi_\theta \rangle \langle \Phi | \Phi \rangle} \right) \approx \frac{2}{N} \sum_{x \sim |\psi_\theta|^2}^N \Re \left\{ \left[\frac{\beta(x)}{\beta_{\text{avg}}} - \frac{\alpha(x)}{\alpha_{\text{avg}}} \right] \frac{\partial}{\partial \theta} \ln \Psi_\theta^*(x) \right\}. \quad (\text{G8})$$

-
- [1] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [2] X. Han and S. A. Hartnoll, Deep Quantum Geometry of Matrices, *Phys. Rev. X* **10**, 011069 (2020).
- [3] K. Choo, T. Neupert, and G. Carleo, Two-dimensional frustrated $j_1 - j_2$ model studied with neural network quantum states, *Phys. Rev. B* **100**, 125124 (2019).
- [4] M. Hibat-Allah, M. Ganahl, L. E. Hayward, R. G. Melko, and J. Carrasquilla, Recurrent neural network wave functions, *Phys. Rev. Res.* **2**, 023358 (2020).
- [5] D. Luo and B. K. Clark, Backflow Transformations via Neural Networks for Quantum Many-Body Wave Functions, *Phys. Rev. Lett.* **122**, 226401 (2019).
- [6] J. Hermann, Z. Schätzle, and F. Noé, Deep neural network solution of the electronic Schrödinger equation, *Nat. Chem.* **12**, 891 (2020).
- [7] D. Pfau, J. S. Spencer, A. G. D. G. Matthews, and W. M. C. Foulkes, Ab initio solution of the many-electron Schrödinger equation with deep neural networks, *Phys. Rev. Res.* **2**, 033429 (2020).
- [8] J. Carrasquilla, D. Luo, F. Pérez, A. Milsted, B. K. Clark, M. Volkovs, and L. Aolita, Probabilistic simulation of quantum circuits with the transformer, *Phys. Rev. A* **104**, 032610 (2021).
- [9] I. L. Gutiérrez and C. B. Mendl, Real time evolution with neural-network quantum states, *Quantum* **6**, 627 (2022).
- [10] S. Lu, X. Gao, and L.-M. Duan, Efficient representation of topologically ordered states with restricted Boltzmann machines, *Phys. Rev. B* **99**, 155136 (2019).
- [11] X. Gao and L.-M. Duan, Efficient representation of quantum many-body states with deep neural networks, *Nat. Commun.* **8**, 662 (2017).

- [12] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, Neural-Network Quantum States, String-Bond States, and Chiral Topological States, *Phys. Rev. X* **8**, 011006 (2018).
- [13] T. Vieijra, C. Casert, J. Nys, W. De Neve, J. Haegeman, J. Ryckebusch, and F. Verstraete, Restricted Boltzmann Machines for Quantum States with Non-Abelian or Anyonic Symmetries, *Phys. Rev. Lett.* **124**, 097201 (2020).
- [14] Y. Nomura, A. S. Darmawan, Y. Yamaji, and M. Imada, Restricted Boltzmann machine learning for solving strongly correlated quantum systems, *Phys. Rev. B* **96**, 205152 (2017).
- [15] M. Schmitt and M. Heyl, Quantum Many-Body Dynamics in Two Dimensions with Artificial Neural Networks, *Phys. Rev. Lett.* **125**, 100503 (2020).
- [16] J. Stokes, J. R. Moreno, E. A. Pnevmatikakis, and G. Carleo, Phases of two-dimensional spinless lattice fermions with first-quantized deep neural-network quantum states, *Phys. Rev. B* **102**, 205122 (2020).
- [17] F. Vicentini, A. Biella, N. Regnault, and C. Ciuti, Variational Neural-Network Ansatz for Steady States in Open Quantum Systems, *Phys. Rev. Lett.* **122**, 250503 (2019).
- [18] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
- [19] K. A. Nicoli, S. Nakajima, N. Strodthoff, W. Samek, K.-R. Müller, and P. Kessel, Asymptotically unbiased estimation of physical observables with neural samplers, *Phys. Rev. E* **101**, 023304 (2020).
- [20] K. A. Nicoli, C. J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel, S. Nakajima, and P. Stornati, Estimation of Thermodynamic Observables in Lattice Field Theories with Deep Generative Models, *Phys. Rev. Lett.* **126**, 032001 (2021).
- [21] N. Yoshioka and R. Hamazaki, Constructing neural stationary states for open quantum many-body systems, *Phys. Rev. B* **99**, 214306 (2019).
- [22] M. J. Hartmann and G. Carleo, Neural-Network Approach to Dissipative Quantum Many-Body Dynamics, *Phys. Rev. Lett.* **122**, 250502 (2019).
- [23] A. Nagy and V. Savona, Variational Quantum Monte Carlo Method with a Neural-Network Ansatz for Open Quantum Systems, *Phys. Rev. Lett.* **122**, 250501 (2019).
- [24] J. Kogut and L. Susskind, Hamiltonian formulation of wilson's lattice gauge theories, *Phys. Rev. D* **11**, 395 (1975).
- [25] M. A. Levin and X.-G. Wen, String-net condensation: A physical mechanism for topological phases, *Phys. Rev. B* **71**, 045110 (2005).
- [26] A. Kitaev, Fault-tolerant quantum computation by anyons, *Ann. Phys.* **303**, 2 (2003).
- [27] W. Shirley, K. Slagle, and X. Chen, Foliated fracton order from gauging subsystem symmetries, *SciPost Phys.* **6**, 041 (2019).
- [28] S. X. Cui, D. Ding, X. Han, G. Penington, D. Ranard, B. C. Rayhaun, and Z. Shangnan, Kitaev's quantum double model as an error correcting code, *Quantum* **4**, 331 (2020).
- [29] A. Feiguin, S. Trebst, A. W. W. Ludwig, M. Troyer, A. Kitaev, Z. Wang, and M. H. Freedman, Interacting Anyons in Topological Quantum Liquids: The Golden Chain, *Phys. Rev. Lett.* **98**, 160409 (2007).
- [30] S. Trebst, E. Ardonne, A. Feiguin, D. A. Huse, A. W. W. Ludwig, and M. Troyer, Collective States of Interacting Fibonacci Anyons, *Phys. Rev. Lett.* **101**, 050401 (2008).
- [31] B. Field and T. Simula, Introduction to topological quantum computation with non-abelian anyons, *Quantum Sci. Technol.* **3**, 045004 (2018).
- [32] M. C. Bañuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero *et al.*, Simulating lattice gauge theories within quantum technologies, *Eur. Phys. J. D* **74**, 165 (2020).
- [33] C. Rebbi, *Lattice Gauge Theories and Monte Carlo Simulations* (World Scientific, Singapore, 1983).
- [34] G. Magnifico, T. Felser, P. Silvi, and S. Montangero, Lattice quantum electrodynamics in (3+1)-dimensions at finite density with tensor networks, *Nat. Commun.* **12**, 3600 (2021).
- [35] D. Luo, G. Carleo, B. K. Clark, and J. Stokes, Gauge Equivariant Neural Networks for Quantum Lattice Gauge Theories, *Phys. Rev. Lett.* **127**, 276402 (2021).
- [36] G. Kanwar, M. S. Albergo, D. Boyda, K. Cranmer, D. C. Hackett, S. Racanière, D. J. Rezende, and P. E. Shanahan, Equivariant Flow-Based Sampling for Lattice Gauge Theory, *Phys. Rev. Lett.* **125**, 121601 (2020).
- [37] D. Boyda, G. Kanwar, S. Racanière, D. J. Rezende, M. S. Albergo, K. Cranmer, D. C. Hackett, and P. E. Shanahan, Sampling using $su(n)$ gauge equivariant flows, *Phys. Rev. D* **103**, 074504 (2021).
- [38] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, Learning phrase representations using RNN encoder-decoder for statistical machine translation, in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Association for Computational Linguistics, Doha, Qatar, 2014), pp. 1724–1734.
- [39] S. Hochreiter and J. Schmidhuber, Long short-term memory, *Neural Comput.* **9**, 1735 (1997).
- [40] A. Van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves *et al.*, Conditional image generation with pixelcnn decoders, in *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, 2016), pp. 4790–4798.
- [41] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, Attention is all you need, *Adv. Neural Inf. Process. Syst.* **30**, 5998 (2017).
- [42] D. Luo, Z. Chen, J. Carrasquilla, and B. K. Clark, Autoregressive neural network for simulating open quantum systems via a probabilistic formulation, *arXiv:2009.05580*.
- [43] O. Sharir, Y. Levine, N. Wies, G. Carleo, and A. Shashua, Deep Autoregressive Models for the Efficient Variational Simulation of Many-Body Quantum Systems, *Phys. Rev. Lett.* **124**, 020503 (2020).
- [44] Y. Levine, O. Sharir, N. Cohen, and A. Shashua, Quantum Entanglement in Deep Learning Architectures, *Phys. Rev. Lett.* **122**, 065301 (2019).
- [45] A. Hamma, P. Zanardi, and X.-G. Wen, String and membrane condensation on three-dimensional lattices, *Phys. Rev. B* **72**, 035307 (2005).
- [46] S. Vijay, J. Haah, and L. Fu, Fracton topological order, generalized lattice gauge theory, and duality, *Phys. Rev. B* **94**, 235157 (2016).
- [47] C. Gils, E. Ardonne, S. Trebst, A. W. W. Ludwig, M. Troyer, and Z. Wang, Collective States of Interacting Anyons, Edge

- States, and the Nucleation of Topological Liquids, *Phys. Rev. Lett.* **103**, 070401 (2009).
- [48] E. Greensmith, P. L. Bartlett, and J. Baxter, Variance reduction techniques for gradient estimates in reinforcement learning, *J. Mach. Learn. Res.* **5**, 1471 (2004).
- [49] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, Pytorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, 2019), pp. 8026–8037.
- [50] C. Roth, Iterative retraining of quantum spin models using recurrent neural networks, [arXiv:2003.06228](https://arxiv.org/abs/2003.06228).
- [51] K. Sprague, J. F. Carrasquilla, S. Whitelam, and I. Tamblyn, Watch and learn—a generalized approach for transferrable learning in deep neural networks via physical principles, *Mach. Learn.: Sci. Technol.* **2**, 02LT02 (2020).
- [52] A. Iserles, Euler’s method and beyond, *A First Course in the Numerical Analysis of Differential Equations*, Cambridge Texts in Applied Mathematics, 2nd ed. (Cambridge University Press, Cambridge, 2008), pp. 8–13.
- [53] D. Kochkov and B. K. Clark, Variational optimization in the AI era: Computational graph states and supervised wave-function optimization, [arXiv:1811.12423](https://arxiv.org/abs/1811.12423).
- [54] Y.-P. Huang, D. Banerjee, and M. Heyl, Dynamical Quantum Phase Transitions in U(1) Quantum Link Models, *Phys. Rev. Lett.* **122**, 250401 (2019).
- [55] P. Karpov, R. Verdel, Y. P. Huang, M. Schmitt, and M. Heyl, Disorder-Free Localization in an Interacting 2D Lattice Gauge Theory, *Phys. Rev. Lett.* **126**, 130401 (2021).
- [56] R. Verdel, M. Schmitt, Y.-P. Huang, P. Karpov, and M. Heyl, Variational classical networks for dynamics in interacting quantum matter, *Phys. Rev. B* **103**, 165103 (2021).
- [57] D. Luo, J. Shen, M. Highman, B. K. Clark, B. DeMarco, A. X. El-Khadra, and B. Gadway, Framework for simulating gauge theories with dipolar spin systems, *Phys. Rev. A* **102**, 032617 (2020).
- [58] E. Rico, T. Pichler, M. Dalmonte, P. Zoller, and S. Montangero, Tensor Networks for Lattice Gauge Theories and Atomic Quantum Simulation, *Phys. Rev. Lett.* **112**, 201601 (2014).
- [59] R. P. Feynman, Forces in molecules, *Phys. Rev.* **56**, 340 (1939).
- [60] S. P. G. Crone and P. Corboz, Detecting a Z_2 topologically ordered phase from unbiased infinite projected entangled-pair state simulations, *Phys. Rev. B* **101**, 115143 (2020).
- [61] F. Wu, Y. Deng, and N. Prokof’ev, Phase diagram of the toric code model in a parallel magnetic field, *Phys. Rev. B* **85**, 195104 (2012).
- [62] K. Gregor, D. A. Huse, R. Moessner, and S. L. Sondhi, Diagnosing deconfinement and topological order, *New J. Phys.* **13**, 025009 (2011).
- [63] M. H. Zarei, Ising order parameter and topological phase transitions: Toric code in a uniform magnetic field, *Phys. Rev. B* **100**, 125159 (2019).
- [64] J. Haah, Local stabilizer codes in three dimensions without string logical operators, *Phys. Rev. A* **83**, 042330 (2011).
- [65] A. Bazavov, Y. Meurice, S.-W. Tsai, J. Unmuth-Yockey, L.-P. Yang, and J. Zhang, Estimating the central charge from the Rényi entanglement entropy, *Phys. Rev. D* **96**, 034514 (2017).
- [66] E. Ardonne, J. Gukelberger, A. W. W. Ludwig, S. Trebst, and M. Troyer, Microscopic models of interacting Yang–Lee anyons, *New J. Phys.* **13**, 045006 (2011).
- [67] D. Banerjee, M. Dalmonte, M. Müller, E. Rico, P. Stebler, U.-J. Wiese, and P. Zoller, Atomic Quantum Simulation of Dynamical Gauge Fields Coupled to Fermionic Matter: From String Breaking to Evolution after a Quench, *Phys. Rev. Lett.* **109**, 175302 (2012).
- [68] U. Wiese, Ultracold quantum gases and lattice systems: Quantum simulation of lattice gauge theories, *Ann. Phys.* **525**, 777 (2013).
- [69] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma, Non-Abelian anyons and topological quantum computation, *Rev. Mod. Phys.* **80**, 1083 (2008).
- [70] <http://iaifi.org/>
- [71] V. Kindratenko, D. Mu, Y. Zhan, J. Maloney, S. H. Hashemi, B. Rabe, K. Xu, R. Campbell, J. Peng, and W. Gropp, Hal: Computer system for scalable deep learning, in *Practice and Experience in Advanced Research Computing*, PEARC ’20 (Association for Computing Machinery, New York, 2020), pp. 41–48.
- [72] A. Reuther, J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron, V. Gadepally, M. Houle, M. Hubbell, M. Jones, A. Klein, L. Milechin, J. Mullen, A. Prout, A. Rosa, C. Yee, and P. Michaleas, Interactive supercomputing on 40,000 cores for machine learning and data analysis, in *2018 IEEE High Performance Extreme Computing Conference (HPEC)* (IEEE, Piscataway, NJ, 2018), pp. 1–6.
- [73] L. Otis and E. Neuscamman, Complementary first and second derivative methods for ansatz optimization in variational Monte Carlo, *Phys. Chem. Chem. Phys.* **21**, 14491 (2019).
- [74] S. Notarnicola, E. Ercolessi, P. Facchi, G. Marmo, S. Pascazio, and F. V. Pepe, Discrete abelian gauge theories for quantum simulations of QED, *J. Phys. A: Math. Theor.* **48**, 30FT01 (2015).
- [75] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2016), pp. 770–778.
- [76] D. Kingma and J. Ba, Adam: A method for stochastic optimization, International Conference on Learning Representations, 2014 (unpublished).