

Accurate neural quantum states for interacting lattice bosons

Zakari Denis* and Giuseppe Carleo†

*Institute of Physics, École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland and
Center for Quantum Science and Engineering, École Polytechnique
Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*

(Dated: April 14, 2024)

In recent years, neural quantum states have emerged as a powerful variational approach, achieving state-of-the-art accuracy when representing the ground-state wave function of a great variety of quantum many-body systems, including spin lattices, interacting fermions or continuous-variable systems. However, accurate neural representations of the ground state of interacting bosons on a lattice have remained elusive. We introduce a neural backflow Jastrow Ansatz, in which occupation factors are dressed with translationally equivariant many-body features generated by a deep neural network. We show that this neural quantum state is able to faithfully represent the ground state of the 2D Bose-Hubbard Hamiltonian across all values of the interaction strength. We scale our simulations to lattices of dimension up to 20×20 while achieving the best variational energies reported for this model. This enables us to investigate the scaling of the entanglement entropy across the superfluid-to-Mott quantum phase transition, a quantity hard to extract with non-variational approaches.

I. INTRODUCTION

Lattice bosonic systems are important in condensed matter because of their unique properties. They are not subject to the exclusion principle, which allows for the appearance of smooth phases of matter in diluted settings. These phases can often be described by classical fields, which simplifies the understanding of many-body phenomena, including the study of $U(1)$ symmetry breaking, superfluidity, and boson condensation. However, simulating these systems beyond mean-field approximations is a challenge. The quantum statistics of bosons result in an exponential growth of the Hilbert space dimension with system size and particle number, complicating simulations. In driven systems, an infinite-dimensional Hilbert space might be needed for a single bosonic mode, limiting the feasibility of exact diagonalization.

Several methods have been proposed to efficiently simulate interacting lattice bosons at zero temperature. Among earliest approaches, path-integral techniques were proposed at finite temperature [1–4]. In particular, quantum Monte Carlo simulations using the worm algorithm provided remarkably accurate predictions extrapolated to zero temperature [5, 6]. At zero temperature, Green-function Monte Carlo particularly stands out [7–9], while the reptation quantum Monte Carlo algorithm [10] was extended to lattice systems [11]. However, these techniques do not give direct access to the wave function and restrict the structure of operators which may be efficiently evaluated in the chosen computational basis. Furthermore, they suffer from the sign problem [12], even for bosonic systems in cases of geometric frustration [13–15]. This is more generally the case with non-stoquastic Hamiltonians, such as those of

bosonic lattices subjected to artificial magnetic fluxes, which have attracted consistent theoretical interest [16–28] and are of current experimental relevance [29–37].

Variational methods are more flexible in this regard. Variational Monte Carlo (VMC) using a modified Jastrow Ansatz [38–40] reproduced all of the features of the quantum phase transition in all relevant spatial dimensions, though only qualitatively. Tensor networks have been used, in particular, to probe ground-state properties of lattice bosons [41, 42]. However, this is restricted by entanglement. The generic area-law entanglement scaling of lattice bosons [43] and the growth of entanglement upon real-time propagation restrict the efficiency of matrix-product variational representations mainly to the ground state of one-dimensional geometries, such as open chains [44–48] or small-perimeter cylinders [41], and short time evolutions. Moreover, this typically requires to introduce truncations in the local occupation factors.

In recent years, neural quantum states [49] have emerged as a powerful variational method, consistently demonstrating remarkable accuracy in representing the ground-state wave function of a wide range of nontrivial Hamiltonians. In addition to spin problems, where it stands as the state-of-the-art method for frustrated lattices [50–52], properly tailored networks have demonstrated their effectiveness in addressing problems involving other kinds of degrees of freedom, including bosonic [53], fermionic [54–56] and beyond [57, 58]. Despite these successes, highly accurate neural-network representations of the ground state of lattice bosonic systems have remained elusive. Several works [59–62] have recently contributed to these ongoing efforts, where the most notable advances [63, 64] have been achieved by incorporating some of the physical structure of the system in the variational Ansatz.

In this work, we introduce a new bosonic neural Ansatz based on the concept of backflow [65], originally formulated in continuous space [66–69] and later extended to

* Email: zakari.denis@epfl.ch

† Email: giuseppe.carleo@epfl.ch

lattices [70, 71]. Neural-network parametrizations of the backflow transformation [72, 73] have driven the most recent advances in fermionic neural quantum states [54–56]. We provide analytical support to motivate its structure and benchmark it on the celebrated Bose-Hubbard model in two dimensions. We assess its remarkable accuracy and scale our simulations to lattices with up to 20×20 sites without local truncation, allowing us to perform a finite-size scaling analysis around the critical point. We further investigate the entanglement entropy of the system across the phase transition, a matter of theoretical [41, 74, 75] and experimental [76] interest.

The paper is structured as follows: the Bose-Hubbard physics is briefly described in Sec. II; in Sec. III, we introduce the neural backflow-Jastrow architecture and formally derive the structure of some of the correlations it can capture; we then apply it to simulate the ground state of the Bose-Hubbard Hamiltonian in Sec. IV, before finally concluding in Sec. V.

II. BOSE HUBBARD

In this work, we focus on the paradigmatic Bose-Hubbard model, as given by the Hamiltonian

$$\hat{H} = -J \sum_{\langle i,j \rangle} (\hat{a}_i^\dagger \hat{a}_j + \hat{a}_j^\dagger \hat{a}_i) + \frac{U}{2} \sum_i \hat{a}_i^\dagger \hat{a}_i^\dagger \hat{a}_i \hat{a}_i. \quad (1)$$

Here, \hat{a}_i denotes the annihilation operator at site i of the lattice under consideration, J is the rate of the hopping of bosons between nearest-neighbor sites, and U quantifies the strength of on-site repulsive interactions. This system can be realized in a number of platforms, notably with ultra-cold atoms loaded in an optical lattice [76–82], or Josephson-junction arrays [83, 84]. In what follows, we shall restrict ourselves to $L \times L$ two-dimensional lattices at fixed number of particles N and unit density $\bar{n} := N/L^2 = 1$.

This system displays a second-order phase transition resulting from the competition between the kinetic energy and the potential energy. At infinite interaction strength and integer density, the wavefunction is a product of Fock states: the so-called Mott phase. As the ratio J/U increases, the $U(1)$ symmetry spontaneously breaks, and the ground state becomes extensively degenerate. The bosonic population spontaneously condenses into the $\mathbf{k} = 0$ mode, as revealed by a finite value of the condensate fraction:

$$\rho_0 := \langle \hat{n}(\mathbf{k} = 0) \rangle / N, \quad (2)$$

which here plays the role of order parameter. In 2D, this phase transition is known to belong to the 3D XY-model universality class [1, 4, 6, 85, 86]. The critical value of the control parameter, $J_c/U = 0.05974(3)$ [$U_c/J = 16.74(1)$], was first precisely determined by a high-order strong coupling expansion [87, 88] and later refined by means of

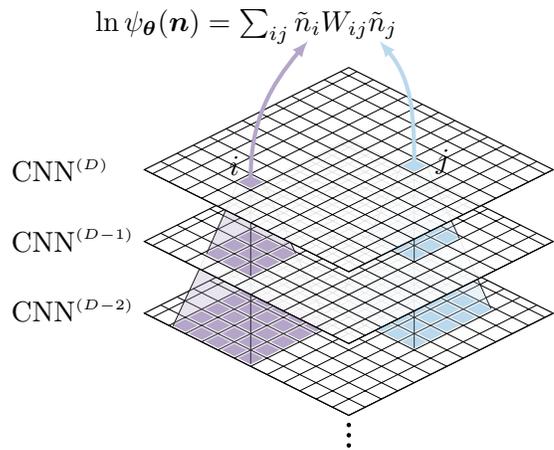


Figure 1. Schematic representation of the neural backflow-Jastrow architecture of Eq. (18). The backflow transformation dresses the bare local quantum numbers entering the Jastrow term with equivariant many-body features. These features bear a dependence on the configuration of all sites within the network’s receptive field. The latter is represented by the shaded region for an exemplary deep convolutional neural network with square filters of width $d_F = 3$.

QMC calculations extrapolated to vanishing temperature [6].

III. VARIATIONAL MODEL

The simplest many-body variational Ansatz capable of qualitatively describing the singular depletion of the condensate while also reproducing the proper behavior of the spatial correlations is the two-body Jastrow wavefunction [89, 90]

$$|\psi_J\rangle = e^{\sum_{ij} \hat{n}_i W_{ij} \hat{n}_j} |\psi_0\rangle, \quad (3)$$

where $|\psi_0\rangle = N^{-1/2} (\sum_i \hat{a}_i^\dagger)^N |\mathbf{0}\rangle$ is the wavefunction of the ideal condensate. This can be equivalently written as

$$\ln \psi_J(\mathbf{n}) = \sum_{ij} n_i W_{ij} n_j + \ln \psi_0(\mathbf{n}), \quad (4)$$

where $\psi(\mathbf{n}) \equiv \langle \mathbf{n} | \psi \rangle$.

In Refs. [38–40], authors showed that the variational accuracy could be greatly enhanced without losing the physical structure and intuition of the two-body Jastrow Ansatz. This is done by adding the following many-body Gutzwiller projector:

$$|\psi_{\text{MBJ}}\rangle = e^{g_{\text{MB}}(\hat{\Pi}_h + \hat{\Pi}_d)} |\psi_J\rangle, \quad (5)$$

where the projectors $\hat{\Pi}_h = \sum_i \hat{h}_i \otimes_{j \in N(i)} (\hat{\mathbf{1}} - \hat{d})_j$ and $\hat{\Pi}_d = \sum_i \hat{d}_i \otimes_{j \in N(i)} (\hat{\mathbf{1}} - \hat{h})_j$, with $\hat{h} = |0\rangle\langle 0|$ and $\hat{d} = |2\rangle\langle 2|$, measure the amount of isolated holons and doublons, respectively. This many-body term serves as a

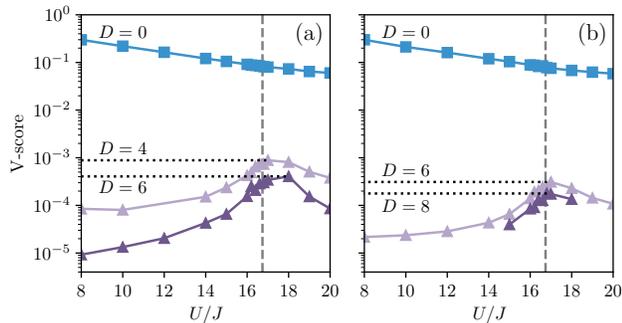


Figure 2. V-score as a function of the depth D of the neural backflow transformation for lattices of size (a) 16×16 and (b) 20×20 at unit filling $\bar{n} = 1$. The gray shaded region indicates the critical value of the control parameter found by QMC [6].

penalty for localized density inhomogeneities, spatially binding holon-doublon pairs, and thereby favoring the onset of a homogeneous Mott insulating phase.

This variational Ansatz, although simple, is able to qualitatively capture the physics of the transition, including the behavior of the correlations and the structure factor [38–40]. However, the quantitative agreement is not excellent, with the ground state variational energy affected by a relative error of the order of 10% and the transition point shifted by 25% to around $U_c/J \sim 21$.

In what follows, we introduce a flexible generalization of the above construction based on the concept of neural backflow transformation [73].

A. Many-body Jastrow

The most straightforward improvement over the Ansatz of Eq. (5) is to lift the Jastrow term to a M -body term. Let us show that the hierarchy of correlations that can be accounted for by a deep convolutional neural network (CNN) encompasses that of such a many-body Jastrow term.

Let us consider a network of depth D built upon the following single-channel convolutional layer with skip connections:

$$h_i^{(n)}(\mathbf{n}) = n_i + \tilde{h}_i^{(n)}(\mathbf{n}), \quad (6)$$

$$\tilde{h}_i^{(n)}(\mathbf{n}) = f(\sum_j w_{\mathbf{r}_{ij}}^{(n)} h_j^{(n-1)}(\mathbf{n}) + b^{(n)}), \quad (7)$$

with $W_{\mathbf{r}} = w_{\mathbf{r}} = 0$ for any \mathbf{r} such that $|\mathbf{r}| > d_F$, where d_F corresponds to the dimension of the convolutional filter. Upon assuming for simplicity that the activation function f belongs to C^2 and possesses a non-vanishing second-order derivative, one has the following recursion

relation [91]:

$$\begin{aligned} \sum_i h_i^{(n)}(\mathbf{n}) &= \dots + \frac{1}{2} \sum_{i_0 i_1} n_{i_0} W_{i_0 i_1}^{(n)} h_{i_1}^{(n-1)}(\mathbf{n}) \\ &= \dots + \sum_{n=1}^D \frac{1}{2^n} \sum_{i_0 \dots i_n} n_{i_0} W_{i_0 i_1}^{(n)} n_{i_1} W_{i_1 i_2}^{(n-1)} \dots W_{i_{n-1} i_n}^{(1)} n_{i_n}, \end{aligned} \quad (8)$$

that is a linear combination of n -body Jastrow terms with $n \leq D+1$, with symmetric two-body translationally equivariant Jastrow weights given by

$$W_{ij}^{(n)} \equiv W_{ji}^{(n)} \propto \sum_k w_{\mathbf{r}_{ki}}^{(n)} w_{\mathbf{r}_{kj}}^{(n)}. \quad (9)$$

Trivially, the two-body Jastrow weights can be made independent for each n -body term at the cost of at most M channels in the CNN.

An analogous result is obtained, replacing skip connections by residual connections, as will be later introduced. Therefore, range- d_F M -body Jastrow factors can be generated with a residual neural-network (ResNet) Ansatz with depth $M-1$ and filters of width d_F .

B. Generalized many-body Gutzwiller projector

In Eq. (5), the many-body term acts as a partial projection of the wavefunction onto a manifold with specific values of the local quantum numbers within some finite-size patch. The choice of the specific values of these quantum numbers is set from *a priori* knowledge about the underlying physics. We may loosen this inductive bias at the cost of a larger number of variational parameters.

Let us consider the following many-body Gutzwiller Ansatz:

$$|\psi_{\text{MB}}\rangle = e^{g_{\text{MB}} \sum_i \otimes_{\mathbf{v} \in \mathcal{V}} |x_{\mathbf{v}}\rangle \langle x_{\mathbf{v}}|_{T_{\mathbf{v}}(i)}} |\psi_0\rangle, \quad (10)$$

where \mathcal{V} denotes a set of relative positions with respect to the i th site, thereby defining a *patch*, and $\mathcal{X} = \{x_{\mathbf{v}}, \mathbf{v} \in \mathcal{V}\}$ some corresponding fixed local configurations given *a priori*. Here, $T_{\mathbf{v}}(i)$ denotes the site corresponding to i translated on the lattice by \mathbf{v} . This Ansatz may be evaluated as

$$\ln \psi_{\text{MB}}(\mathbf{n}) = \ln \Pi_{\text{MB}}(\mathbf{n}) + \ln \psi_0(\mathbf{n}), \quad (11)$$

$$\ln \Pi_{\text{MB}}(\mathbf{n}) = g_{\text{MB}} \sum_i \mathbb{1}[\wedge_{\mathbf{v} \in \mathcal{V}} n_{T_{\mathbf{v}}(i)} = x_{\mathbf{v}}]; \quad (12)$$

where $\mathbb{1}$ denotes the indicator function. The diagonal many-body projector of Eq. (12) may be represented by

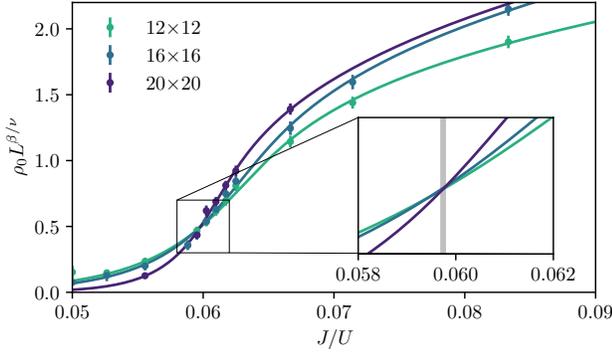


Figure 3. Rescaled condensate fraction as a function of the inverse interaction strength for increasing lattice sizes. Plain lines denote a fit to the scaling function in Eq. (26) to the variational data. The gray shaded region in the inset indicates the confidence interval for the critical value of the control parameter as found by QMC calculations [6]. As expected by scaling theory, all curves intersect for a same abscissa, showing good correspondence with the QMC estimation.

a CNN. Indeed, one has that

$$\ln \Pi_{\text{MB}}(\mathbf{n}) = \sum_i h_{i,1}^{(2)}(\mathbf{n}), \quad (13)$$

$$h_{i,1}^{(2)}(\mathbf{n}) = \text{ReLU}\left(-2g_{\text{MB}}\Delta x \sum_{\mu=1}^2 h_{i,\mu}^{(1)}(\mathbf{n}) + g_{\text{MB}}\right), \quad (14)$$

$$h_{i,\mu}^{(1)}(\mathbf{n}) = \text{ReLU}\left(\sum_{\mathbf{v} \in \mathcal{V}} \frac{(-1)^\mu}{x_{T_{\mathbf{v}}(i)}} n_{T_{\mathbf{v}}(i)} - (-1)^\mu\right), \quad (15)$$

where $h_{i,\mu}^{(n)}$ denotes the μ th channel of the i th output of the n th layer of the CNN, and $\Delta x > \max_{x \in \mathcal{X}}(x)$ [92]. This corresponds to a two-layer CNN with parameters:

$$\begin{aligned} K_{\mathbf{v},\mu,1}^{(1)} &= (-1)^\mu / x_{T_{\mathbf{v}}(i)} \in \mathbb{R}^{|\mathcal{V}| \times 2 \times 1}, & b_\mu^{(1)} &= (-1)^{\mu+1} \in \mathbb{R}^2, \\ K_{\mathbf{0},1,\mu'}^{(2)} &= -2g_{\text{MB}}\Delta x \in \mathbb{R}^{1 \times 1 \times 2}, & b_1^{(2)} &= g_{\text{MB}} \in \mathbb{R}. \end{aligned} \quad (16)$$

Therefore, a many-body Gutzwiller Ansatz with K many-body projectors may be represented by a two-layer CNN with a filter geometry matching \mathcal{V} and K channels. Relaxing the conditions in Eq. (16) generalizes the projection to the optimum quantum numbers. In App. A, we show how this architecture can also account for an attractive confining potential for holons and doublons close apart, effectively stabilizing the Mott phase.

C. Neural backflow Jastrow Ansatz

This motivates the use of a NQS Ansatz of the form [63]

$$\ln \psi_{\boldsymbol{\theta}}(\mathbf{n}) = \sum_{i,\mu} h_{i,\mu}^{(D)}(\mathbf{n}) + \ln \psi_0(\mathbf{n}), \quad (17)$$

where $\boldsymbol{\theta}$ denotes the set of all variational parameters, and $h_{i,\mu}$ is the channel μ of site i of a depth- D residual network with stride equal to one.

Provided the number of channels is large enough, such a network should be able to encompass both the M -body Jastrow term and the many-body Gutzwiller projector. However, the range of the Jastrow terms in such a parametrization is limited by the width of the filters d_F . This is in contrast with the initial Jastrow Ansatz of Eq. (3) which could capture the crucial long-range correlations responsible of the singular behavior of the density structure factor at zero momentum [39].

To circumvent this, we instead introduce the many-body features of the ResNet as an equivariant backflow transformation of a translationally invariant two-body Jastrow with maximum range. This leads to the variational Ansatz that we shall consider throughout this article:

$$\ln \psi_{\boldsymbol{\theta}}(\mathbf{n}) = \sum_{i,j=1}^{L^d} \tilde{n}_i W_{d_{ij}} \tilde{n}_j, \quad (18)$$

where the weights only depend on the L^1 minimum-image distance d_{ij} between all pairs (i,j) . We also dropped the original mean-field prior ψ_0 which proves useful only for small system sizes. Fed to this Jastrow factor are translation-equivariant many-body features obtained according to the following backflow transformation of the original local occupation factors:

$$\tilde{n}_i = n_i + \sum_{\mu=1}^{\alpha_D} a_\mu h_{i,\mu}^{(D)}. \quad (19)$$

Here, $h_{i,\mu}^{(D)}$ are translation-covariant features extracted from the original occupations through a residual neural network of depth D with N_c channels, and the mixing weights a_μ both combine all channels and control the magnitude of the backflow transformation. The former are obtained as $h_{i,\mu}^{(2\ell)} = \tilde{h}_{i,\mu}^{(2\ell-1)}$ and $h_{i,\mu}^{(2\ell+1)} = \text{LayerNorm}(h_{i,\mu}^{(2\ell-1)} + \tilde{h}_{i,\mu}^{(2\ell)})$, with

$$\tilde{h}_{i,\mu}^{(\ell)} = \sigma\left(\sum_{\mu'=1}^{\alpha_\ell} \sum_{\mathbf{v}: v_\infty \leq d_K} K_{\mathbf{v},\mu,\mu'}^{(\ell)} h_{T_{\mathbf{v}}(i),\mu'}^{(\ell-1)} + b_\mu^{(\ell)}\right), \quad (20)$$

where, σ is an activation function (GELU), α_ℓ the number of channels of the ℓ th layer, d_K the kernel size, $T_{\mathbf{v}}(i)$ corresponds to the site i translated by \mathbf{v} , and $(K_{\mathbf{v},\mu,\mu'}^{(\ell)}, b_\mu^{(\ell)})$ the filters and bias of the ℓ th layer. To further improve stability, the input occupations are rescaled as $n_i \mapsto n_i/\bar{n} - 1, \forall i$, with $\bar{n} = N_p/L^d$. The architecture of the model is illustrated in Fig. 1.

Let us note that this construction, in presence of a bias term in the last layer of the backflow transformation, is at least as expressive as the bare ResNet. Furthermore it allows one to perform VMC more efficiently by initializing the weights of the Jastrow with those obtained by first optimizing a bare two-body Jastrow Ansatz of the

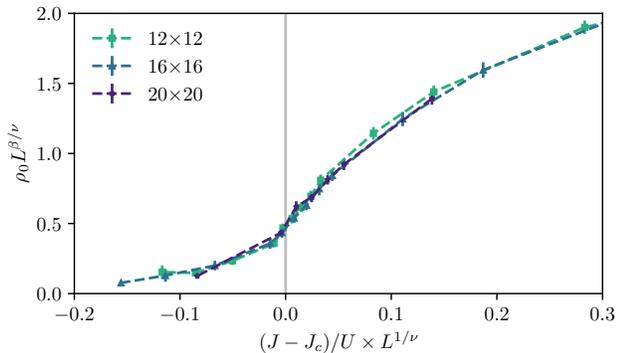


Figure 4. Data collapse of the data of Fig. 3 when plotted against the rescaled centered control parameter, showing the universal behavior of observables close to the critical point.

form of Eq. (3), and then optimizing the entire network altogether. Indeed, provided the backflow weights are initialized such that they initially act perturbatively, this allows to start the optimization closer from the ground state and thus with a much lower variance on all Monte Carlo estimates involved in the process.

Splitting the layout of the wavefunction into a simple mean-field-like long-range behaviour and a complex short-range structure seems a sensible design choice, especially for gapped Hamiltonians whose ground states exhibit some finite-range quantum correlations.

Very recent variational Ansätze based upon the visual-transformer architecture [52, 93, 94] adopt a reminiscent structure, rooted in representation learning [95], where a strong scission of the model into a deep embedding and a shallow final correlator proves optimal.

IV. NUMERICAL RESULTS

In what follows, we optimize the parameters θ of the model defined by Eq. (18) by minimizing the variational expectation value of the energy

$$E_{\theta} = \frac{\langle \psi_{\theta} | \hat{H} | \psi_{\theta} \rangle}{\langle \psi_{\theta} | \psi_{\theta} \rangle}, \quad (21)$$

with \hat{H} as given by Eq. (1). This is achieved with guarantees of exponential convergence via the stochastic-reconfiguration prescription [96], where the parameters are evolved in time according to

$$\dot{\theta} = -\mathbf{S}^{-1} \mathbf{F}, \quad (22)$$

where \mathbf{S} and \mathbf{F} respectively denote the quantum geometric tensor [97] and the vector of forces, as given by

$$S_{kk'} = \mathbb{E}_{\mathbf{n}} [O_k^*(\mathbf{n}) O_{k'}(\mathbf{n})] - \mathbb{E}_{\mathbf{n}} [O_k^*(\mathbf{n})] \mathbb{E}_{\mathbf{n}} [O_{k'}(\mathbf{n})], \quad (23)$$

$$F_k = \mathbb{E}_{\mathbf{n}} [O_k^*(\mathbf{n}) E_{\text{loc}}(\mathbf{n})] - \mathbb{E}_{\mathbf{n}} [O_k^*(\mathbf{n})] \mathbb{E}_{\mathbf{n}} [E_{\text{loc}}(\mathbf{n})], \quad (24)$$

Table I. Comparison of the variational energy obtained by the neural backflow Jastrow (NBFJ) Ansatz with depth $D = 6$ and those of Green-function Monte Carlo (GFMC) [99] and recent NQS simulations [64].

Lattice size	U/J	Method	# parameters	E/JL^2
8×8	8.5	GFMC [99]	—	$-1.544669(8)$
		NBFJ	6681	$-1.544668(3)$
	17.0	GFMC [99]	—	$-0.5375(4)$
		NBFJ	6681	$-0.53804(2)$
10×10	16.0	NQS-OH [64]	426	$-0.5721(3)$
		NBFJ	6683	$-0.58972(3)$
	20.0	NQS-B [64]	206	$-0.4068(9)$
		NBFJ	6683	$-0.43329(3)$

with $O_k(\mathbf{n}) := \partial_{\theta_k} \ln \psi_{\theta}(\mathbf{n})$ the log-derivative of the variational Ansatz and $E_{\text{loc}}(\mathbf{n}) := \langle \mathbf{n} | \hat{H} | \psi_{\theta} \rangle / \langle \mathbf{n} | \psi_{\theta} \rangle$ the so-called *local energy*. In the above, expectation values are implicitly taken with respect to the Born distribution $|\psi_{\theta}|^2$.

Importantly, all of these quantities can be efficiently estimated by drawing samples from this probability density function thanks to, e.g., Markov chain Monte Carlo and the Metropolis-Hastings algorithm. More details about the variational optimization are given in App. B.

In all following simulations, ResNet-based backflow transformations with square filters of width $d_F = 3$ and $\alpha = 12$ channels per layer were used. A depth of $D = 8$ was used for 20×20 lattices and $D = 6$ for the rest. All models were optimized through NetKet 3.0 [98] on a single GPU.

A. Benchmarking

The scaling of the Hilbert-space dimension with system size is rather adverse for bosonic lattices. Therefore, exact-diagonalization reference quantities are restricted to very small systems which are not very representative of the physics at work in the thermodynamic limit.

At moderate system size, ground-state predictions may be efficiently obtained via Green-function Monte Carlo (GFMC) [7–9]. These provide a valuable reference value since they are exact up to a controllable systematic bias. In Table I, we compare the variational energies achieved by our Ansatz for a 8×8 lattice against a GFMC reference [99], showing perfect agreement. Therein, we further compare our variational energies to the best previously achieved NQS results [64] for a 10×10 system, obtaining values 3 to 6% lower, although with a higher number of parameters.

For larger system sizes, instead, we assess the accuracy of our wavefunction by means of the recently introduced variational score (V-score) [100], which allows to systematically compare Ansätze. For bosonic systems, this score

can be computed as

$$\text{V-score} = \frac{L^d \text{Var}[E_\theta]}{(E_\theta - E_{\text{MF}})}, \quad (25)$$

Where L^d is the number of sites of the system, $\text{Var}[E_\theta]$ is the variance of the variational energy, and E_{MF} is the mean-field energy. The latter is given by $E_{\text{MF}}/L^d = \bar{n}(U\bar{n}/2 - zJ)$ for Bose Hubbard with coordination number z ($z = 4$ here).

This quantity is intensive. Therefore, it allows one to systematically compare the accuracy of a simulations at varying system sizes. It vanishes when the variational Ansatz exactly matches the ground state by virtue of the zero-variance property. Furthermore, it was empirically verified [100] that the relative error on the variational energy of the ground state positively correlates to the V-score of the Ansatz following a universal trend. This can be used to roughly infer the error, which proves very useful whenever exact diagonalization is not available.

A V-score in the range 10^{-4} – 10^{-6} signals a relative error roughly below 10^{-5} , indicating that the problem has been—for all practical purposes—solved variationally. This is the case, for instance, of spin models in the absence of magnetic frustration.

In Fig. 2, we show the V-score achieved by the backflow-Jastrow Ansatz for Bose Hubbard at unit filling on a periodic 2D lattice with $N = 16 \times 16$ and $N = 20 \times 20$ sites. The problem is clearly harder in the vicinity of the critical point, where the gap closes in the thermodynamic limit. A few ResNet layers in the backflow transformation dramatically improve the accuracy over bare Jastrow, lowering the V-score from 10^{-1} down to about 2×10^{-4} for a depth of $D = 8$ layers. We observe that increasing the depth D of the backflow transformation consistently improves the accuracy. Interestingly, the depth required to reach a given accuracy does not seem to depend on the lattice size. Variational energies along with corresponding V-scores for all tested neural-backflow depths at $U/J = 16.8$ ($\sim U_c/J$) are provided in App. C.

B. Finite-size scaling

We can exploit the ability of our approach to scale to large system sizes to investigate the universal finite-size scaling of physical quantities of the system. According to the usual finite-size scaling argument, the order parameter ρ_0/N for a lattice of linear size L should depend on the control parameter $(J - J_c)/U$ as $\rho_0/N|_L \sim L^{-\beta/\nu} \tilde{f}(L^{1/\nu}(J - J_c)/U)$ where \tilde{f} is some scaling function. Hence, when plotting $\rho_0/N|_L L^{\beta/\nu}$ against J/U for various values of L , all curves should intersect at the critical value J_c/U , provided the ratio of critical exponents β/ν is correct.

In Fig. 3, we perform the above analysis using the critical exponents of the 3D XY universality class [101]. Upon

fitting to our variational data the fitting function

$$f(J/U) = [\text{softplus}_a(J/U - b)]^c, \quad (26)$$

with $\text{softplus}_\alpha(x) := \ln(1 + e^{\alpha x})/\alpha$, and parameters a , $b \sim J_c/U$ and c , all curves intersect within the confidence intervals of the best available estimation of the critical value of J/U [6]. In Fig. 4, we verify that all of the variational data collapse into a universal curve upon rescaling and centering the abscissa of Fig. 3.

C. Entanglement entropy

One of the strengths of variational approaches is the direct access to the wavefunction, in contrast with other techniques such as QMC. This allows us to compute quantities beyond usual linear forms on the wavefunction. In particular, a quantity of interest is the entanglement entropy, which quantizes the degree of entanglement of a subsystem composed of a subset of the lattice A and its complement \bar{A} . It is defined as the entropy of the reduced density $\hat{\rho}_A = \text{Tr}_{\bar{A}}[|\psi_\theta\rangle\langle\psi_\theta|]$. While the von Neumann entropy is typically used, the Rényi-2 entanglement entropy, as given by

$$S_2(\hat{\rho}_A) = -\ln(\text{Tr} \hat{\rho}_A^2), \quad (27)$$

can be easily estimated via Monte Carlo sampling of two independent replicas through the estimator [102–104]

$$S_2(\hat{\rho}_A) = -\ln \mathbb{E}_{\mathbf{n}, \mathbf{n}' \sim |\psi_\theta|^2} \left[\frac{\psi_\theta(\mathbf{n}'_A, \mathbf{n}_{\bar{A}}) \psi_\theta(\mathbf{n}_A, \mathbf{n}'_{\bar{A}})}{\psi_\theta(\mathbf{n}_A, \mathbf{n}_{\bar{A}}) \psi_\theta(\mathbf{n}'_A, \mathbf{n}'_{\bar{A}})} \right]. \quad (28)$$

In this estimator, samples \mathbf{n} and \mathbf{n}' are first drawn from two independent Markov chains and split into configurations of either complementary subspace $\mathbf{n}^{(\prime)} = (\mathbf{n}_A^{(\prime)}, \mathbf{n}_{\bar{A}}^{(\prime)})$. In the numerator, the wavefunction is evaluated on configurations generated by partially swapping the configurations of the replicas according to $\mathbf{n}_A \leftrightarrow \mathbf{n}'_A$. Importantly, when evaluating the above estimator, we enforce that $\psi_\theta(\mathbf{n}) = 0$ whenever $\sum_i n_i \neq N$. In what follows, we shall consider the first half of the lattice, of dimension $L \times L/2$ and boundary length of $|\partial A| = 2L$, as the subsystem A .

Such a system is expected to exhibit an area-law scaling of the entanglement entropy. However, at finite size, the degeneracy of the ground state in the superfluid is slightly lifted. This *tower-of-states* mechanism [74] induces a correction in the superfluid phase scaling as the logarithm of the size of the boundary, resulting into the law [41]

$$S_2(\hat{\rho}_A) = aL + b\mathbf{1}[U > U_c] \ln L + c + O(1/L), \quad (29)$$

where $b = 1/2$ is a universal coefficient related to the number of Goldstone modes (1 for the 2D Bose-Hubbard and 3D XY models).

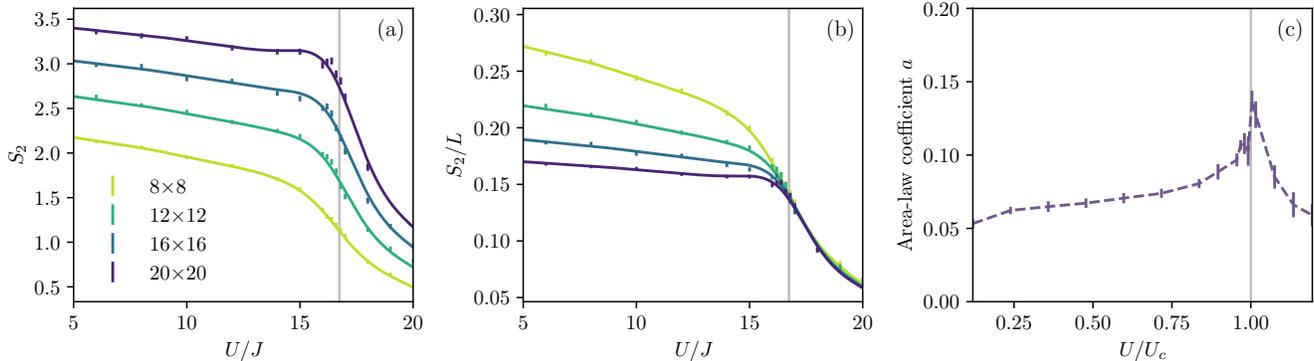


Figure 5. (a) Rényi-2 entropy $S_2(\hat{\rho}_A)$ evaluated on the reduced density matrix of the first half (A) of the lattice. (b) Entanglement entropy rescaled by the typical dimension of the subsystem’s boundary length $|\partial A| \sim L$. In the Mott insulating phase, a strict area law is observed while the spontaneous breaking of the continuous symmetry induces a logarithmic correction in the superfluid phase. (c) Area-law coefficient a obtained by fitting Eq. (29) to the variational data for linear sizes $L \in \{8, 10, 12, 14, 16, 20\}$. In panels (a) and (b), errors are estimated by bootstrap; a cubic spline interpolation of the S_2 fit is added as a guide to the eye.

In Fig. 5(a), we show the Rényi-2 entropy $S_2(\hat{\rho}_A)$ as a function of the strength of the interactions. Upon increasing the lattice size, S_2 approaches a step-like singular behavior. This singular behavior is a signature of bosonic criticality [75]. In Fig. 5(b), we rescale this quantity by the the dimension of the boundary L . As in Ref. [41], we observe a strict area-law dependence of the entanglement entropy in the Mott phase along with a logarithmic departure in the superfluid phase. In Fig. 5(c), we further extract area-law coefficient a of Eq. (29). There, as predicted by Ref. [75], one observes a pronounced cusp at the critical point, a signature of the contribution of the Higgs mode to the entanglement entropy.

V. CONCLUSION

In this paper, we introduced the neural backflow-Jastrow architecture. We physically motivated its relevance for the study of interacting lattice bosons and analytically proved that it encompasses the structure of various standard variational architectures while providing a natural generalization. We benchmarked our Ansatz on the Bose-Hubbard problem and systematically assessed its variational accuracy, achieving a V-score as low as 2×10^{-4} on a 20×20 lattice. Its great scalability allowed us to scale up VMC simulations to lattices with up to 20×20 sites on a single GPU with no local Hilbert-space truncation. Thanks to this, we performed a finite-size scaling analysis at zero temperature, showing remarkable agreement with the results obtained in QMC calculations by extrapolating finite-temperature results. Furthermore, we were able to investigate the entanglement properties of the system. We observed the logarithmic correction of the entanglement entropy due to the spon-

aneous breaking of the $U(1)$ continuous symmetry and were able to extract its universal scaling prefactor, which we found to display a second-order singularity at the critical point.

This work clears the way for the simulation of many bosonic lattice systems which remain out of reach for other techniques. A prime example is the simulation of bosons in the presence of a synthetic magnetic field. Indeed, such systems are intrinsically plagued with the sign problem, thereby ruling out QMC, and set in two spatial dimensions, which severely challenges tensor-network approaches, hitherto constrained to narrow-ladder geometries. Another interesting prospect is the extension of this architecture to driven-dissipative scenarios, where accurate neural representations of the density matrix are still unavailable. This would enlarge the scope of the NQS approach to the field of quantum optics.

ACKNOWLEDGMENTS

We thank Federico Becca for his valuable early input and for conducting the GFMC simulations presented in this study. Our gratitude extends to Stephen R Clark for sharing data from Ref. [64]. We also appreciate the helpful discussions with Markus Holzmann, David Clément, and Tommaso Roscilde. This work was supported by SEFRI through Grant No. MB22.00051 (NEQS - Neural Quantum Simulation).

Appendix A: Holon-doublon confinement

Close to the Mott-to-superfluid transition, the state of the system can be expressed as a perturbative expansion around the $U/J \sim +\infty$ limit. To leading order, this

yields the following correction to the Mott state $|\psi_\infty\rangle$:

$$|\psi_U\rangle - |\psi_\infty\rangle \propto \frac{J}{U} \sum_{\langle i,j \rangle} (\hat{a}_i^\dagger \hat{a}_j + \hat{a}_j^\dagger \hat{a}_i) |\psi_\infty\rangle + O((J/U)^2). \quad (\text{A1})$$

At unit filling this corresponds to the emergence of isolated holon-doublon bound pairs at any neighboring pair of sites. Similarly, higher-order contributions involve additional pairs, or pairs separated by a larger distance. Within the Mott phase, such an expansion converges and holon-doublon pairs must thus be spatially confined exponentially in order to stabilize the phase.

Such a confinement can be encoded in the variational wavefunction in the form of an attractive potential [38]

$$|\psi'\rangle = e^{-\sum_{ij} \hat{h}_i V_{d_{ij}} \hat{d}_j} |\psi\rangle, \quad (\text{A2})$$

where $V_d \leq 0$, for $0 < d \leq R$, and $V_d = 0$ otherwise. Upon assuming any holon (doublon) is not surrounded by more than one doublon (holon), the many-body projector of Eq. (5) is a particular case of the above with range $R = 1$ and $V_1 = -g_{\text{MB}}$. The $R > 1$ case can be easily represented by a convolutional network. Indeed, upon neglecting occupations above 2 bosons per site and assuming that pairs are distant from each other by more than R , one has

$$-\sum_{ij} V_{d_{ij}} \langle \mathbf{n} | \hat{h}_i \hat{d}_j | \mathbf{n} \rangle = \sum_i h_{i,1}^{(2)}(\mathbf{n}), \quad (\text{A3})$$

$$h_{i,1}^{(2)}(\mathbf{n}) = \text{ReLU}\left(-V_0 \sum_{\mu=1}^2 h_{i,\mu}^{(1)}(\mathbf{n}) - \sum_{d=1}^R V_d h_{i,2+d}^{(1)}(\mathbf{n})\right), \quad (\text{A4})$$

$$h_{i,\mu \leq 2}^{(1)}(\mathbf{n}) = \text{ReLU}\left((-1)^\mu (n_i + 1)\right), \quad (\text{A5})$$

$$h_{i,\mu=2+d}^{(1)}(\mathbf{n}) = \text{ReLU}\left(\sum_{\mathbf{v} \in \mathcal{V}} \delta_{v,d} x_{T_v(i)}\right), \quad (\text{A6})$$

with $V_0 = \sum_{d>0} |V_d|$ and where input occupations are shifted by the unit density such that configurations $n = -1$ and $n = 1$ correspond respectively to a hole and a doublon, such as in the main text. This bears the form of a two-layer convolutional network with parameters

$$\begin{aligned} K_{\mathbf{v},\mu \leq 2,1}^{(1)} &= (-1)^\mu \delta_{v,0} \in \mathbb{R}^{|\mathcal{V}| \times 2 \times 1}, & b_\mu^{(1)} &= (-1)^\mu \in \mathbb{R}^2, \\ K_{\mathbf{v},\mu=2+d,1}^{(1)} &= \delta_{v,d} \in \mathbb{R}^{|\mathcal{V}| \times R \times 1}, & b_\mu^{(1)} &= 0 \in \mathbb{R}^R, \\ K_{\mathbf{0},1,\mu'}^{(2)} &= -V_\mu \in \mathbb{R}^{1 \times 1 \times 2}, & b_1^{(2)} &= 0 \in \mathbb{R}. \end{aligned} \quad (\text{A7})$$

Appendix B: Optimization procedure

In all simulations the learning rate was set to $\eta = 1 \times 10^{-3}$, the diagonal-shift regularization of the quantum geometric tensor was set to $\lambda = 5 \times 10^{-4}$ (1×10^{-3}) for the (neural backflow) Jastrow. 8192 samples were used for most of the VMC calculations, and 12288 for backflow transformations of depth $D = 8$ for the largest lattice, while the number of samples was set to 32768 for our simulations in Table I.

1. Metropolis-Hastings transition rule

In all simulations, Markov-chain Monte Carlo (MCMC) was used in order to sample from the Born probability distribution $p_\theta(\mathbf{n}) := |\psi_\theta(\mathbf{n})|^2 / \|\psi_\theta\|^2$. At each MCMC step, an update from the current configuration \mathbf{n} into a new configuration \mathbf{n}' is proposed with probability given by the proposition distribution $g(\mathbf{n}'|\mathbf{n})$, and accepted with probability

$$p_{\text{acc}}(\mathbf{n} \rightarrow \mathbf{n}') = \min\left(1, \frac{p_\theta(\mathbf{n}') g(\mathbf{n}|\mathbf{n}')}{p_\theta(\mathbf{n}) g(\mathbf{n}'|\mathbf{n})}\right). \quad (\text{B1})$$

Our proposition distribution is induced by the first-quantization local transition rule $g^{(1)}(\mathbf{x}'|\mathbf{x}) = \prod_{\mu}^N g_{\text{loc}}^{(1)}(\mathbf{x}'_\mu|\mathbf{x}_\mu)$ with

$$g_{\text{loc}}^{(1)}(\mathbf{x}'|\mathbf{x}) \propto \mathbb{1}[\|\mathbf{x} - \mathbf{x}'\| = a], \quad (\text{B2})$$

where a denotes the lattice parameter. This proposal is clearly ergodic and corresponds to the Hamiltonian rule for a non-interacting system. It simply moves a particle picked at random into any of its neighboring sites. In second quantization, this can be achieved through the update kernel $g^{(2)}(\mathbf{n}'|\mathbf{n}) = \prod_{\langle i,j \rangle} g_{\text{loc}}^{(2)}(n'_i, n'_j|n_i, n_j)$ with

$$g_{\text{loc}}^{(2)}(n'_i, n'_j|n_i, n_j) \propto n_i \delta_{n'_i, n_i-1} \delta_{n'_j, n_j+1}. \quad (\text{B3})$$

This corresponds to choosing a site with a probability proportional to its occupation number and then transferring one of the occupations to any neighboring site. This kernel is not symmetric, hence, a correction is needed to enforce detail balance. This is accounted for in Eq. (B1) thanks to the following ratio of proposal densities:

$$\frac{g_{\text{loc}}^{(2)}(n_i, n_j|n'_i, n'_j)}{g_{\text{loc}}^{(2)}(n'_i, n'_j|n_i, n_j)} = \frac{n'_j}{n_i}. \quad (\text{B4})$$

Appendix C: Worst-case variational figures of merit

Variational Monte Carlo proves the most demanding in the Mott phase at the vicinity of the critical point U_c^+/J . In Table II, we compile the variational energies and corresponding variational scores in that challenging region ($U/J = 16.8$) as a function of system size and depth D of the neural backflow transformation.

Table II. Variational figures of merit at $U/J = 16.8$. 12 288 samples were used in optimizing the networks of depth $D = 6$ and $D = 8$ for the largest system size, 8192 otherwise.

Lattice size	depth	E_{θ}/JL^2	V-score
8×8	0	-0.452(2)	8.4×10^{-2}
	6	-0.54783(6)	5.6×10^{-5}
12×12	0	-0.444(1)	7.9×10^{-2}
	6	-0.54320(7)	1.6×10^{-4}
16×16	0	-0.446(1)	8.1×10^{-2}
	2	-0.5375(2)	3.7×10^{-3}
	4	-0.5411(1)	7.3×10^{-4}
	6	-0.54191(7)	3.0×10^{-4}
20×20	0	-0.4414(9)	7.8×10^{-2}
	2	-0.5373(1)	3.1×10^{-3}
	4	-0.54119(5)	6.8×10^{-4}
	6	-0.54162(5)	3.1×10^{-4}
	8	-0.54182(4)	1.8×10^{-4}

- [1] G. G. Batrouni, R. T. Scalettar, and G. T. Zimanyi, Quantum critical phenomena in one-dimensional Bose systems, *Physical Review Letters* **65**, 1765 (1990).
- [2] W. Krauth, N. Trivedi, and D. Ceperley, Superfluid-insulator transition in disordered boson systems, *Physical Review Letters* **67**, 2307 (1991).
- [3] R. T. Scalettar, G. G. Batrouni, and G. T. Zimanyi, Localization in interacting, disordered, Bose systems, *Physical Review Letters* **66**, 3144 (1991).
- [4] G. G. Batrouni and R. T. Scalettar, World-line quantum Monte Carlo algorithm for a one-dimensional Bose model, *Physical Review B* **46**, 9051 (1992).
- [5] B. Capogrosso-Sansone, N. V. Prokof'ev, and B. V. Svistunov, Phase diagram and thermodynamics of the three-dimensional Bose-Hubbard model, *Physical Review B* **75**, 134302 (2007).
- [6] B. Capogrosso-Sansone, Ş. G. Söyler, N. Prokof'ev, and B. Svistunov, Monte Carlo study of the two-dimensional Bose-Hubbard model, *Physical Review A* **77**, 015602 (2008).
- [7] N. Trivedi and D. M. Ceperley, Ground-state correlations of quantum antiferromagnets: A Green-function Monte Carlo study, *Physical Review B* **41**, 4552 (1990).
- [8] M. Calandra Buonaura and S. Sorella, Numerical study of the two-dimensional Heisenberg model using a Green function Monte Carlo technique with a fixed number of walkers, *Physical Review B* **57**, 11446 (1998).
- [9] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems*, 1st ed. (2017).
- [10] S. Baroni and S. Moroni, Reptation Quantum Monte Carlo: A Method for Unbiased Ground-State Averages and Imaginary-Time Correlations, *Physical Review Letters* **82**, 4745 (1999).
- [11] G. Carleo, F. Becca, S. Moroni, and S. Baroni, Reptation quantum Monte Carlo algorithm for lattice Hamiltonians with a directed-update scheme, *Physical Review E* **82**, 046710 (2010).
- [12] M. Troyer and U.-J. Wiese, Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations, *Physical Review Letters* **94**, 170201 (2005).
- [13] P. Henelius and A. W. Sandvik, Sign problem in Monte Carlo simulations of frustrated quantum spin systems, *Physical Review B* **62**, 1102 (2000).
- [14] P. Hauke, T. Roscilde, V. Murg, J. I. Cirac, and R. Schmied, Modified spin-wave theory with ordering vector optimization: Frustrated bosons on the spatially anisotropic triangular lattice, *New Journal of Physics* **12**, 053036 (2010).
- [15] G. Pan and Z. Y. Meng, The sign problem in quantum Monte Carlo simulations, in *Encyclopedia of Condensed Matter Physics (Second Edition)*, edited by T. Chakraborty (Oxford, 2024) pp. 879–893.
- [16] M. Ö. Oktel, M. Niță, and B. Tanatar, Mean-field theory for Bose-Hubbard model under a magnetic field, *Physical Review B* **75**, 045133 (2007).
- [17] G. Möller and N. R. Cooper, Condensed ground states of frustrated Bose-Hubbard models, *Physical Review A* **82**, 063625 (2010).
- [18] S. D. Huber and N. H. Lindner, Topological transitions for lattice bosons in a magnetic field, *Proceedings of the National Academy of Sciences* **108**, 19925 (2011).
- [19] A. Dhar, M. Maji, T. Mishra, R. V. Pai, S. Mukerjee, and A. Paramekanti, Bose-Hubbard model in a strong effective magnetic field: Emergence of a chiral Mott insulator ground state, *Physical Review A* **85**, 041602 (2012).
- [20] A. Dhar, T. Mishra, M. Maji, R. V. Pai, S. Mukerjee, and A. Paramekanti, Chiral Mott insulator with staggered loop currents in the fully frustrated Bose-Hubbard model, *Physical Review B* **87**, 174501 (2013).
- [21] A. Tokuno and A. Georges, Ground states of a Bose-Hubbard ladder in an artificial magnetic field: Field-theoretical approach, *New Journal of Physics* **16**, 073005 (2014).
- [22] F. Kolley, M. Piraud, I. P. McCulloch, U. Schollwöck, and F. Heidrich-Meisner, Strongly interacting bosons on a three-leg ladder in the presence of a homogeneous flux,

- New Journal of Physics* **17**, 092001 (2015).
- [23] Y.-C. He, S. Bhattacharjee, R. Moessner, and F. Pollmann, Bosonic Integer Quantum Hall Effect in an Interacting Lattice Model, *Physical Review Letters* **115**, 116803 (2015).
- [24] C. Romen and A. M. Läuchli, Chiral Mott insulators in frustrated Bose-Hubbard models on ladders and two-dimensional lattices: A combined perturbative and density matrix renormalization group study, *Physical Review B* **98**, 054519 (2018).
- [25] T.-S. Zeng, D. N. Sheng, and W. Zhu, Continuous phase transition between bosonic integer quantum Hall liquid and a trivial insulator: Evidence for deconfined quantum criticality, *Physical Review B* **101**, 035138 (2020).
- [26] Y.-F. Song and S.-J. Yang, Quantum phases for bosons in a magnetic lattice with a harmonic trap, *New Journal of Physics* **22**, 073001 (2020).
- [27] M. Buser, C. Hubig, U. Schollwöck, L. Tarruell, and F. Heidrich-Meisner, Interacting bosonic flux ladders with a synthetic dimension: Ground-state phases and quantum quench dynamics, *Physical Review A* **102**, 053314 (2020).
- [28] C.-M. Halati and T. Giamarchi, Bose-Hubbard triangular ladder in an artificial gauge field, *Physical Review Research* **5**, 013126 (2023).
- [29] D. Jaksch and P. Zoller, Creation of effective magnetic fields in optical lattices: The Hofstadter butterfly for cold neutral atoms, *New Journal of Physics* **5**, 56 (2003).
- [30] F. Gerbier and J. Dalibard, Gauge fields for ultracold atoms in optical superlattices, *New Journal of Physics* **12**, 033007 (2010).
- [31] M. Aidelsburger, M. Atala, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch, Realization of the Hofstadter Hamiltonian with Ultracold Atoms in Optical Lattices, *Physical Review Letters* **111**, 185301 (2013).
- [32] H. Miyake, G. A. Siviloglou, C. J. Kennedy, W. C. Burton, and W. Ketterle, Realizing the Harper Hamiltonian with Laser-Assisted Tunneling in Optical Lattices, *Physical Review Letters* **111**, 185302 (2013).
- [33] J. Struck, M. Weinberg, C. Ölschläger, P. Windpassinger, J. Simonet, K. Sengstock, R. Höppner, P. Hauke, A. Eckardt, M. Lewenstein, and L. Mathey, Engineering Ising-XY spin-models in a triangular lattice using tunable artificial gauge fields, *Nature Physics* **9**, 738 (2013).
- [34] M. Atala, M. Aidelsburger, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch, Observation of chiral currents with ultracold atoms in bosonic ladders, *Nature Physics* **10**, 588 (2014).
- [35] N. Goldman, G. Juzeliūnas, P. Öhberg, and I. B. Spielman, Light-induced gauge fields for ultracold atoms, *Reports on Progress in Physics* **77**, 126401 (2014).
- [36] A. Sterdyniak, N. R. Cooper, and N. Regnault, Bosonic Integer Quantum Hall Effect in Optical Flux Lattices, *Physical Review Letters* **115**, 116802 (2015).
- [37] L. Barbiero, J. Cabedo, M. Lewenstein, L. Tarruell, and A. Celi, Frustrated magnets without geometrical frustration in bosonic flux ladders, *Physical Review Research* **5**, L042008 (2023).
- [38] M. Capello, F. Becca, M. Fabrizio, S. Sorella, and E. Tosatti, Variational Description of Mott Insulators, *Physical Review Letters* **94**, 026406 (2005).
- [39] M. Capello, F. Becca, M. Fabrizio, and S. Sorella, Superfluid to Mott-Insulator Transition in Bose-Hubbard Models, *Physical Review Letters* **99**, 056402 (2007).
- [40] M. Capello, F. Becca, M. Fabrizio, and S. Sorella, Mott transition in bosonic systems: Insights from the variational approach, *Physical Review B* **77**, 144517 (2008).
- [41] V. Alba, M. Haque, and A. M. Läuchli, Entanglement Spectrum of the Two-Dimensional Bose-Hubbard Model, *Physical Review Letters* **110**, 260403 (2013).
- [42] M. M. Rams, P. Czarnik, and L. Cincio, Precise Extrapolation of the Correlation Function Asymptotics in Uniform Tensor Network States with Application to the Bose-Hubbard and XXZ Models, *Physical Review X* **8**, 041033 (2018).
- [43] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, *Reviews of Modern Physics* **82**, 277 (2010).
- [44] R. V. Pai, R. Pandit, H. R. Krishnamurthy, and S. Ramasesha, One-Dimensional Disordered Bosonic Hubbard Model: A Density-Matrix Renormalization Group Study, *Physical Review Letters* **76**, 2937 (1996).
- [45] T. D. Kühner and H. Monien, Phases of the one-dimensional lattice Bose-Hubbard model, *Physical Review B* **58**, R14741 (1998).
- [46] S. Rapsch, U. Schollwöck, and W. Zwerger, Density matrix renormalization group for disordered bosons in one dimension, *Europhysics Letters* **46**, 559 (1999).
- [47] C. Kollath, A. M. Läuchli, and E. Altman, Quench Dynamics and Nonequilibrium Phase Diagram of the Bose-Hubbard Model, *Physical Review Letters* **98**, 180601 (2007).
- [48] A. M. Läuchli and C. Kollath, Spreading of correlations and entanglement after a quench in the one-dimensional Bose-Hubbard model, *Journal of Statistical Mechanics: Theory and Experiment* **2008**, P05018 (2008).
- [49] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [50] Y. Nomura and M. Imada, Dirac-Type Nodal Spin Liquid Revealed by Refined Quantum Many-Body Solver Using Neural-Network Wave Function, Correlation Ratio, and Level Spectroscopy, *Physical Review X* **11**, 031034 (2021).
- [51] A. Chen and M. Heyl, Efficient optimization of deep neural quantum states toward machine precision (2023), [arxiv:2302.01941](https://arxiv.org/abs/2302.01941).
- [52] R. Rende, L. L. Viteritti, L. Bardone, F. Becca, and S. Goldt, A simple linear algebra identity to optimize Large-Scale Neural Network Quantum States (2023), [arxiv:2310.05715](https://arxiv.org/abs/2310.05715).
- [53] G. Pescia, J. Han, A. Lovato, J. Lu, and G. Carleo, Neural-network quantum states for periodic systems in continuous space, *Physical Review Research* **4**, 023138 (2022).
- [54] J. Hermann, Z. Schätzle, and F. Noé, Deep-neural-network solution of the electronic Schrödinger equation, *Nature Chemistry* **12**, 891 (2020).
- [55] 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, *Physical Review Research* **2**, 033429 (2020).
- [56] G. Pescia, J. Nys, J. Kim, A. Lovato, and G. Carleo, Message-Passing Neural Quantum States for the Homogeneous Electron Gas (2023), [arxiv:2305.07240](https://arxiv.org/abs/2305.07240).

- [57] D. Luo, G. Carleo, B. K. Clark, and J. Stokes, Gauge Equivariant Neural Networks for Quantum Lattice Gauge Theories, *Physical Review Letters* **127**, 276402 (2021).
- [58] M. Medvidović and D. Sels, Variational Quantum Dynamics of Two-Dimensional Rotor Models, *PRX Quantum* **4**, 040302 (2023).
- [59] H. Saito, Solving the Bose–Hubbard Model with Machine Learning, *Journal of the Physical Society of Japan* **86**, 093001 (2017).
- [60] K. McBrien, G. Carleo, and E. Khatami, Ground state phase diagram of the one-dimensional Bose-Hubbard model from restricted Boltzmann machines, *Journal of Physics: Conference Series* **1290**, 012005 (2019).
- [61] V. Vargas-Calderón, H. Vinck-Posada, and F. A. González, Phase Diagram Reconstruction of the Bose–Hubbard Model with a Restricted Boltzmann Machine Wavefunction, *Journal of the Physical Society of Japan* **89**, 094002 (2020).
- [62] Z. Zhu, M. Mattheakis, W. Pan, and E. Kaxiras, HubbardNet: Efficient predictions of the Bose-Hubbard model spectrum with deep neural networks, *Physical Review Research* **5**, 043084 (2023).
- [63] K. Choo, G. Carleo, N. Regnault, and T. Neupert, Symmetries and Many-Body Excitations with Neural-Network Quantum States, *Physical Review Letters* **121**, 167204 (2018).
- [64] M. Y. Pei and S. R. Clark, *Specialising Neural-network Quantum States for the Bose Hubbard Model* (2024), arxiv:2402.15424.
- [65] R. P. Feynman and M. Cohen, Energy Spectrum of the Excitations in Liquid Helium, *Physical Review* **102**, 1189 (1956).
- [66] M. A. Lee, K. E. Schmidt, M. H. Kalos, and G. V. Chester, Green’s Function Monte Carlo Method for Liquid ^3He , *Physical Review Letters* **46**, 728 (1981).
- [67] Y. Kwon, D. M. Ceperley, and R. M. Martin, Effects of three-body and backflow correlations in the two-dimensional electron gas, *Physical Review B* **48**, 12037 (1993).
- [68] Y. Kwon, D. M. Ceperley, and R. M. Martin, Effects of backflow correlation in the three-dimensional electron gas: Quantum Monte Carlo study, *Physical Review B* **58**, 6800 (1998).
- [69] M. Holzmann, D. M. Ceperley, C. Pierleoni, and K. Esler, Backflow correlations for the electron gas and metallic hydrogen, *Physical Review E* **68**, 046707 (2003).
- [70] L. F. Tocchio, F. Becca, A. Parola, and S. Sorella, Role of backflow correlations for the nonmagnetic phase of the t - t' Hubbard model, *Physical Review B* **78**, 041101 (2008).
- [71] L. F. Tocchio, F. Becca, and C. Gros, Backflow correlations in the Hubbard model: An efficient tool for the study of the metal-insulator transition and the large- U limit, *Physical Review B* **83**, 195138 (2011).
- [72] M. Ruggeri, S. Moroni, and M. Holzmann, Nonlinear Network Description for Many-Body Quantum Systems in Continuous Space, *Physical Review Letters* **120**, 205302 (2018).
- [73] D. Luo and B. K. Clark, Backflow Transformations via Neural Networks for Quantum Many-Body Wave Functions, *Physical Review Letters* **122**, 226401 (2019).
- [74] M. A. Metlitski and T. Grover, Entanglement Entropy of Systems with Spontaneously Broken Continuous Symmetry (2015), arxiv:1112.5166.
- [75] I. Frérot and T. Roscilde, Entanglement Entropy across the Superfluid-Insulator Transition: A Signature of Bosonic Criticality, *Physical Review Letters* **116**, 190401 (2016).
- [76] R. Islam, R. Ma, P. M. Preiss, M. Eric Tai, A. Lukin, M. Rispoli, and M. Greiner, Measuring entanglement entropy in a quantum many-body system, *Nature* **528**, 77 (2015).
- [77] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, Cold Bosonic Atoms in Optical Lattices, *Physical Review Letters* **81**, 3108 (1998).
- [78] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms, *Nature* **415**, 39 (2002).
- [79] I. B. Spielman, W. D. Phillips, and J. V. Porto, Mott-Insulator Transition in a Two-Dimensional Atomic Bose Gas, *Physical Review Letters* **98**, 080404 (2007).
- [80] I. B. Spielman, W. D. Phillips, and J. V. Porto, Condensate Fraction in a 2D Bose Gas Measured across the Mott-Insulator Transition, *Physical Review Letters* **100**, 120402 (2008).
- [81] I. Bloch, J. Dalibard, and W. Zwerger, Many-body physics with ultracold gases, *Reviews of Modern Physics* **80**, 885 (2008).
- [82] M. Endres, T. Fukuhara, D. Pekker, M. Cheneau, P. Schauß, C. Gross, E. Demler, S. Kuhr, and I. Bloch, The ‘Higgs’ amplitude mode at the two-dimensional superfluid/Mott insulator transition, *Nature* **487**, 454 (2012).
- [83] A. van Otterlo, K.-H. Wagenblast, R. Fazio, and G. Schön, Response of Josephson-junction arrays near the quantum phase transition, *Physical Review B* **48**, 3316 (1993).
- [84] C. Bruder, R. Fazio, and G. Schön, The Bose-Hubbard model: from Josephson junction arrays to optical lattices, *Annalen der Physik* **517**, 566 (2005).
- [85] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher, Boson localization and the superfluid-insulator transition, *Physical Review B* **40**, 546 (1989).
- [86] W. Krauth and N. Trivedi, Mott and Superfluid Transitions in a Strongly Interacting Lattice Boson System, *Europhysics Letters* **14**, 627 (1991).
- [87] J. K. Freericks and H. Monien, Strong-coupling expansions for the pure and disordered Bose-Hubbard model, *Physical Review B* **53**, 2691 (1996).
- [88] N. Elstner and H. Monien, Dynamics and thermodynamics of the Bose-Hubbard model, *Physical Review B* **59**, 12184 (1999).
- [89] A. Bijl, The lowest wave function of the symmetrical many particles system, *Physica* **7**, 869 (1940).
- [90] R. Jastrow, Many-Body Problem with Strong Forces, *Physical Review* **98**, 1479 (1955).
- [91] Here, the Lagrange form of the first-order Taylor expansion of f around the bias b was used: $\tilde{h}_i^{(n)} = f(b) + \sum_j [f'(b)w_{r_{ij}}]h_j^{(n-1)} + \frac{1}{2}\sum_{jk} [f''(b_\star)w_{r_{ij}}w_{r_{ik}}]h_j^{(n-1)}h_k^{(n-1)}$. Interestingly, the mean-value theorem, at the core of the Lagrange remainder, was also used to show that the ground state of a system of interacting particles in continuous space is exactly captured by a general backflow transformation [56].

- [92] The case where the value 0 belongs to \mathcal{X} can be simply worked out by a fixed initial linear transformation of all inputs as $n_i \mapsto n_i + \xi$, $\forall \xi \in \mathbb{R}_+^*$. In particular, Eq. (5) is recovered using local fluctuations $n_i - \bar{n}$ as an input and identifying holes (doublons) with the value $x = -1$ ($x = +1$).
- [93] L. L. Viteritti, R. Rende, and F. Becca, Transformer Variational Wave Functions for Frustrated Quantum Spin Systems, *Physical Review Letters* **130**, 236401 (2023).
- [94] L. L. Viteritti, R. Rende, A. Parola, S. Goldt, and F. Becca, *Transformer Wave Function for the Shastry-Sutherland Model: Emergence of a Spin-Liquid Phase* (2024), [arxiv:2311.16889](https://arxiv.org/abs/2311.16889).
- [95] Y. Bengio, A. Courville, and P. Vincent, Representation Learning: A Review and New Perspectives, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35**, 1798 (2013).
- [96] S. Sorella, Wave function optimization in the variational Monte Carlo method, *Physical Review B* **71**, 241103 (2005).
- [97] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, Quantum Natural Gradient, *Quantum* **4**, 269 (2020).
- [98] F. Vicentini, D. Hofmann, A. Szabó, D. Wu, C. Roth, C. Giuliani, G. Pescia, J. Nys, V. Vargas-Calderón, N. Astrakhantsev, and G. Carleo, NetKet 3: Machine Learning Toolbox for Many-Body Quantum Systems, *SciPost Physics Codebases* , 007 (2022).
- [99] Federico Becca (private communication, 2023).
- [100] D. Wu, R. Rossi, F. Vicentini, N. Astrakhantsev, F. Becca, X. Cao, J. Carrasquilla, F. Ferrari, A. Georges, M. Hibat-Allah, M. Imada, A. M. Läuchli, G. Mazzola, A. Mezzacapo, A. Millis, J. R. Moreno, T. Neupert, Y. Nomura, J. Nys, O. Parcollet, R. Pohle, I. Romero, M. Schmid, J. M. Silvester, S. Sorella, L. F. Tocchio, L. Wang, S. R. White, A. Wietek, Q. Yang, Y. Yang, S. Zhang, and G. Carleo, *Variational Benchmarks for Quantum Many-Body Problems* (2023), [arxiv:2302.04919](https://arxiv.org/abs/2302.04919).
- [101] M. Campostrini, M. Hasenbusch, A. Pelissetto, P. Rossi, and E. Vicari, Critical behavior of the three-dimensional XY universality class, *Physical Review B* **63**, 214503 (2001).
- [102] M. B. Hastings, I. González, A. B. Kallin, and R. G. Melko, Measuring Rényi Entanglement Entropy in Quantum Monte Carlo Simulations, *Physical Review Letters* **104**, 157201 (2010).
- [103] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nature Physics* **14**, 447 (2018).
- [104] J. Zhao, B.-B. Chen, Y.-C. Wang, Z. Yan, M. Cheng, and Z. Y. Meng, Measuring Rényi entanglement entropy with high efficiency and precision in quantum Monte Carlo simulations, *npj Quantum Materials* **7**, 1 (2022).