

Neural network–augmented Pfaffian wave-functions for scalable simulations of interacting fermions

Ao Chen, Zhou-Quan Wan, Anirvan Sengupta, Antoine Georges, Christopher Roth

Angaben zur Veröffentlichung / Publication details:

Chen, Ao, Zhou-Quan Wan, Anirvan Sengupta, Antoine Georges, and Christopher Roth. 2026. "Neural network–augmented Pfaffian wave-functions for scalable simulations of interacting fermions." *Proceedings of the National Academy of Sciences (PNAS)* 123 (35): e2535288123. <https://doi.org/10.1073/pnas.2535288123>.

Nutzungsbedingungen / Terms of use:

CC BY-NC-ND 4.0





Neural network–augmented Pfaffian wave-functions for scalable simulations of interacting fermions

Ao Chen^{a,b,c,1}, Zhou-Quan Wan^a, Anirvan Sengupta^{a,d,e}, Antoine Georges^{a,f,g,h,1} , and Christopher Roth^a

Affiliations are included on p. 9.

Contributed by Antoine Georges; received December 11, 2025; accepted July 17, 2026; reviewed by Di Luo and Markus Schmitt

Developing accurate numerical methods for strongly interacting fermions is crucial for improving our understanding of various quantum many-body phenomena, especially unconventional superconductivity. Recently, neural quantum states have emerged as a promising approach for studying correlated fermions, highlighted by the hidden fermion and backflow methods, which use neural networks to model corrections to fermionic quasiparticle orbitals. In this work, we expand these ideas to the space of Pfaffians, a wave-function that naturally expresses superconducting pairings, and propose the hidden fermion Pfaffian state (HFPS), which flexibly represents both unpaired and superconducting phases and scales to large systems with favorable asymptotic complexity. In our numerical experiments, HFPS provides state-of-the-art variational accuracy in different regimes of both the attractive and repulsive Hubbard models. We show that the HFPS is able to capture general phases with and without pairing, and therefore may be a useful tool for modeling phases with unconventional superconductivity.

strong correlations | superconductivity | machine learning | Pfaffians | neural networks

Understanding the behavior of strongly interacting fermions is a prominent challenge in condensed matter physics. A remarkable feature of strongly correlated fermionic systems is that they host phases with distinct physical properties but very small differences in total energy. This is notoriously the case for quantum materials with strong electronic correlations, such as copper-oxide (cuprate), high critical temperature superconductors (1, 2). But, remarkably, it is also a feature of even the most idealized lattice model of interacting fermions, the Hubbard model. It is indeed becoming increasingly apparent that the phase diagram of the Hubbard model hosts a rich diversity of possible phases, the actual ground state depending sensitively on details of the Hamiltonian parameters (3, 4).

This is both an experimental opportunity, allowing for the switching between phases with different properties by slightly altering the environment, and a major numerical challenge. Computational methods must indeed reach exquisite accuracy in order to reliably predict phase diagrams, often on the scale of only a few thousandths of the bare electronic energy scales defining the system.

Variational wave-function methods are a tool of choice to address this problem at zero temperature (5, 6). Because the dimension of the Hilbert space depends exponentially on the size of the system, the full many-body wave-function is an intractable object for moderately sized systems. As a result, the wave-function must be represented in a compressed form, using parameterized variational *ansätze*.

One of the most efficient compression methods is the fermionic matrix product state representation (fMPS) which, combined with the density matrix renormalization group algorithm (DMRG) has transformed our ability to compute the properties of (quasi-) one-dimensional systems (7–12). Generalization of fMPS to higher dimensions such as fermionic projected entangled-pair states (fPEPS), or generally fermionic tensor networks, provides asymptotically exact solutions for fermionic systems with arbitrary dimensions, but is bottlenecked by huge optimization complexity (13–17).

Quantum Monte Carlo (QMC) methods, including determinant QMC (DQMC) (18–22) and auxiliary-field QMC (AFQMC) (23–25), provide accurate solutions for systems without sign problems but are limited by either exponential costs or systematic bias when the sign problem exists (26). Variational Monte Carlo (VMC) optimization of many-body fermionic wave-functions (27–38) is a powerful method to overcome the limitations of the aforementioned methods, especially within the strongly correlated regime, but its accuracy depends strongly on the specific parameterization of the wave-function and the possible physical biases in choosing this parameterization.

Significance

Understanding how electrons interact in two-dimensional materials is essential for explaining phenomena like high-temperature superconductivity, but existing numerical methods are often limited by systematic biases. Neural quantum states offer a promising alternative, yet their success has largely been restricted to spin systems, since representing fermion pairing and controlling computational cost have remained major obstacles. Here, we introduce the hidden fermion Pfaffian state, a neural-network-augmented variational wave function that explicitly encodes fermion pairing. Combined with sublattice structures and low-rank updates, this reduces computational scaling from $\mathcal{O}(N^5)$ to $\mathcal{O}(N^3)$, enabling accurate, large-scale simulations of the Fermi-Hubbard model in superconducting and stripe-ordered phases, opening a path toward scalable, unbiased simulations of strongly correlated electron systems and high-temperature superconductivity.

The authors declare no competing interest.

Copyright © 2026 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](#).

¹To whom correspondence may be addressed. Email: chenao.phys@gmail.com or ageorges@flatironinstitute.org.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2535288123/-DCSupplemental>.

Published August 28, 2026.

Recently, a more “agnostic” parameterization of many-body variational wave-functions has been introduced, based on artificial neural networks (ANN) (39). Neural quantum states (NQS) are a much more expressive class of variational wave-functions which encode quantum correlations in the “synapses” of an ANN. This approach has become a mainstream method for modeling spin systems and has provided great insights into quantum spin liquids (40–43). On the other hand, the development of NQS for fermionic systems has progressed somewhat more slowly. Although great progress has been made in continuous space for quantum chemistry problems (44–48) and the electron gas (49, 50), fermionic lattice problems are still challenging due to strong interactions and correlations in lattice models. Several NQS architectures have been proposed for fermions, including a combination of a Pfaffian and a restricted Boltzmann machine (PP+RBM) (51), neural network backflow (NNBF) (52), and hidden fermion determinant states (HFDS) (53). These architectures, nevertheless, either use an ANN as a generalized Jastrow factor, which only provides a weak correction to the mean-field wave-function, or have unfavorable scaling complexity.

Since superconductivity is a focal point of research on strongly correlated systems, it is natural to consider variational wave-functions which explicitly embody the possibility for fermions to form pairs, a classic example being the mean-field BCS ansatz (54). More generally, a mean-field wave-function with pairing, called a Thouless state (55), forms a class of wave-functions with the structure of a Pfaffian (56, 57) when it is projected onto a fixed number of particles. In quantum chemistry calculations, this is often called the antisymmetric geminal power (AGP) (58). In addition to superconductors, Pfaffians have been shown to describe states with nonabelian anyons (59, 60), and to provide a good representation of the 5/2 fractional quantum Hall state. Most generally, the Pfaffian parameterizes a broad class of antisymmetric objects which encompasses Slater determinant states (37) and paired BCS wave-functions. While Pfaffians are in general expensive objects, taking $\mathcal{O}(n^3)$ time to compute, Pfaffian-based VMC can be accelerated by imposing sublattice structure (38) and applying low-rank update (LRU) techniques (61, 62), to efficiently push simulations to large scales. Some early works have considered the use of Pfaffians in the framework of neural quantum states (51, 63–66), but this line of research remains rather uncharted, especially in combination with the most expressive ANN architectures and efficient optimization methods.

In this manuscript, we propose an ANN-augmented Pfaffian variational wave-function, which we call hidden fermion Pfaffian states (HFPS), for the simulation of interacting fermions. This class of wave-functions, which generalizes the hidden fermion formalism (53), provides an architecture which is compatible with deep ANNs and large-scale VMC simulations. We show that the HFPS can successfully approximate ground states of the Hubbard model in various phases, including s-wave superconductors, Mott insulators, and the “stripe” spin and charge-density wave phase. In small systems where exact diagonalization (ED) is feasible, or for systems where QMC is sign-free, HFPS provides variational energies close to those of benchmark methods. In all of the studied cases, HFPS gives state-of-the-art variational results, often surpassing other methods by more than an order of magnitude in variational accuracy. This gives promise that HFPS can resolve the small energy differences between competing phases and map out precise phase diagrams of strongly interacting models. Furthermore, HFPS scales rather well with system

size and offers the promise of accessing large enough systems to understand fermionic phases with long-range correlations, including unconventional superconductors.

1. Methods

1.1. Pfaffian Wave-Function. The Pfaffian wave-function is a projected solution of the general bilinear Hamiltonian

$$\hat{\mathcal{H}}_0 = \sum_{p,q} \left(t_{pq} \hat{c}_p^\dagger \hat{c}_q + \frac{1}{2} \Delta_{pq} \hat{c}_p^\dagger \hat{c}_q^\dagger + \frac{1}{2} \Delta_{qp}^* \hat{c}_p \hat{c}_q \right), \quad [1]$$

where the indices p and q run over all orbitals of the system including spatial and spinful degrees of freedom, and M is the total number of orbitals. In [SI Appendix](#), we discuss how to obtain the ground state of $\hat{\mathcal{H}}_0$ using the Bogoliubov transformation. A common form of the ground state, often known as the Thouless state, can be expressed as

$$|\psi_0\rangle = \exp \left(\frac{1}{2} \sum_{p,q} F_{pq} \hat{c}_p^\dagger \hat{c}_q^\dagger \right) |0\rangle, \quad [2]$$

where $\mathbf{F}$ is an $M \times M$ antisymmetric matrix depending on t_{pq} and Δ_{pq} , and $|0\rangle$ denotes the vacuum state. Since the BCS Hamiltonian is a special case of $\hat{\mathcal{H}}_0$, $|\psi_0\rangle$ can express the BCS state by

$$\begin{aligned} |\psi_{\text{BCS}}\rangle &= \prod_{\mathbf{k}} (u_{\mathbf{k}} + v_{\mathbf{k}} \hat{c}_{\mathbf{k},\uparrow}^\dagger \hat{c}_{-\mathbf{k},\downarrow}^\dagger) |0\rangle \\ &= \left(\prod_{\mathbf{k}} u_{\mathbf{k}} \right) \times \exp \left(\sum_{ij} F_{ij} \hat{c}_{i,\uparrow}^\dagger \hat{c}_{j,\downarrow}^\dagger \right) |0\rangle, \end{aligned} \quad [3]$$

where $\hat{c}_{\mathbf{k},\sigma}^\dagger = \sum_j \hat{c}_{j,\sigma}^\dagger e^{i\mathbf{k}\cdot\mathbf{r}_j} / \sqrt{N_{\text{site}}}$ is the quasi-particle in the momentum space of a lattice with N_{site} unit cells, and $F_{ij} = \sum_{\mathbf{k}} e^{i\mathbf{k}\cdot(\mathbf{r}_i - \mathbf{r}_j)} v_{\mathbf{k}} / u_{\mathbf{k}} N_{\text{site}}$.

For closed systems where the number of particles is conserved, it is convenient to project $|\psi_0\rangle$ onto a sector of fixed particle number

$$|\psi_{\text{pf}}\rangle = \hat{\mathcal{P}}_N |\psi_0\rangle = \frac{1}{(N/2)!} \left(\sum_{p<q} F_{pq} \hat{c}_p^\dagger \hat{c}_q^\dagger \right)^{N/2} |0\rangle, \quad [4]$$

where $\hat{\mathcal{P}}_N$ is the projector onto the sector with N fermions. Given a Fock state

$$|\mathbf{n}\rangle = \prod_p (\hat{c}_p^\dagger)^{n_p} |0\rangle, \quad [5]$$

where n_p is 0/1 for an empty/occupied orbital and $\sum_p n_p = N$, the wave-function of $|\psi_{\text{pf}}\rangle$ is

$$\psi_{\text{pf}}(\mathbf{n}) = \langle \mathbf{n} | \psi_{\text{pf}} \rangle = \text{pf}(\mathbf{n} \star \mathbf{F} \star \mathbf{n}), \quad [6]$$

where $\text{pf}(\mathbf{A})$ represents the Pfaffian of a matrix $\mathbf{A}$. A brief introduction to Pfaffians and to some of their mathematical properties is provided in [SI Appendix](#). The $\star$ symbol represents a slicing operation selecting the rows of $\mathbf{F}$ by $\mathbf{n} \star \mathbf{F}$ and columns by $\mathbf{F} \star \mathbf{n}$ according to the occupied orbitals of $\mathbf{n}$. The sliced matrix size becomes $N \times N$ if the Fock state $|\mathbf{n}\rangle$ has N particles. Eq. 6 is the most common form of the Pfaffian wave-function in VMC, where independent elements of $\mathbf{F}$ are treated as trainable variational parameters.

The Pfaffian wave-function can also express a Slater determinant state,

$$\psi_{\text{det}}(\mathbf{n}) = \det(\mathbf{n} \star \mathbf{u}) = \text{pf}(\mathbf{n} \star \mathbf{u} \mathbf{J} \mathbf{u}^T \star \mathbf{n}), \quad [7]$$

where $\mathbf{u}$ is an $M \times N$ matrix denoting single-particle orbitals in a Slater determinant state, and $\mathbf{J}$ is any antisymmetric matrix satisfying $\text{pf} \mathbf{J} = 1$.




Fig. 1. Illustration of HFPS architecture in Eq. 14. Here we show a one-dimensional (1D) spinless fermion system for simplicity. In this example, there are four orbitals, $M = 4$, two visible fermions, $N = 2$, and two hidden fermions, $\tilde{N} = 2$. The CNN outputs are separated into two parts along the channel dimension, one for $\tilde{\mathbf{F}}^{vh}(\mathbf{n})$ and the other for $\mathbf{J}(\mathbf{n})$. The 4×4 Pfaffian matrix contains four 2×2 parts, $\mathbf{n} \star \mathbf{F}^{vv} \star \mathbf{n}$, $\mathbf{n} \star \tilde{\mathbf{F}}^{vh}(\mathbf{n})$, $-\tilde{\mathbf{F}}^{vh}(\mathbf{n})^T \star \mathbf{n}$, and $\tilde{\mathbf{F}}^{hh}$. The orange blocks represent the matrix elements of $\mathbf{F}^{vv}$, and the purple ones represent the matrix elements of $\tilde{\mathbf{F}}^{hh}$. The remaining colored blocks represent $\tilde{\mathbf{F}}^{vh}(\mathbf{n})$ generated by a CNN which are equivariant to a translation of the input Fock state $\mathbf{n}$.

Therefore, a Pfaffian wave-function with $\mathbf{F} = \mathbf{u}\mathbf{J}\mathbf{u}^T$ expresses a Slater determinant state.

In summary, the Pfaffian wave-function in Eq. 6 is a general projected solution of the bilinear Hamiltonian, where the Slater determinant state in Eq. 7 with noninteracting quasi-particles, and the projected BCS state in Eq. 3 with Cooper pairs are special cases. It can also express superconducting states with p-wave and d-wave pairings given suitable $\hat{\mathcal{H}}_0$. Furthermore, the structure of the Pfaffian wave-function allows for fast low-rank updates (61) and sublattice translational symmetry (38). These techniques greatly accelerate simulations in big systems and make large-scale VMC optimization possible.

1.2. Hidden Fermion Pfaffian State. In ref. 53, it was shown how correlations between particles can be taken into account by enlarging the physical Hilbert space with “hidden” degrees of freedom, while keeping a Slater determinant form of the wave-function. Higher-order correlations between electrons emerge from a projection onto the physical Hilbert space.

Here, we generalize this concept to Pfaffian wave-functions. We expand the Hilbert space with M orbitals and N visible fermions to include additional $\tilde{M}$ hidden orbitals and $\tilde{N}$ hidden fermions. The Pfaffian state in this expanded Hilbert space, which we name the hidden fermion Pfaffian state (HFPS), becomes

$$|\psi\rangle = \frac{1}{(N_{\text{tot}}/2)!} \left(\sum_{p < q} F_{pq}^{vv} \hat{c}_p^\dagger \hat{c}_q^\dagger + \sum_{\tilde{p} < \tilde{q}} F_{\tilde{p}\tilde{q}}^{hh} \hat{a}_{\tilde{p}}^\dagger \hat{a}_{\tilde{q}}^\dagger + \sum_{p, \tilde{p}} F_{p\tilde{p}}^{vh} \hat{c}_p^\dagger \hat{a}_{\tilde{p}}^\dagger \right)^{N_{\text{tot}}/2} |0\rangle, \quad [8]$$

where $\mathbf{F}^{vv}$, $\mathbf{F}^{vh}$, and $\mathbf{F}^{hh}$ are matrices representing the various couplings between visible and hidden fermions, $\hat{c}_p^\dagger$ and $\hat{a}_{\tilde{p}}^\dagger$ denote respectively visible and hidden fermion operators, and $N_{\text{tot}} = N + \tilde{N}$ is the total number of fermions. The wave-function component, similar to Eq. 6, is given by

$$\begin{aligned} \psi(\mathbf{n}, \tilde{\mathbf{n}}) &= \langle \mathbf{n}, \tilde{\mathbf{n}} | \psi \rangle \\ &= \text{pf} \left[(\mathbf{n}, \tilde{\mathbf{n}}) \star \begin{pmatrix} \mathbf{F}^{vv} & \mathbf{F}^{vh} \\ -(\mathbf{F}^{vh})^T & \mathbf{F}^{hh} \end{pmatrix} \star (\mathbf{n}, \tilde{\mathbf{n}}) \right] \\ &= \text{pf} \begin{pmatrix} \mathbf{n} \star \mathbf{F}^{vv} \star \mathbf{n} & \mathbf{n} \star \mathbf{F}^{vh} \star \tilde{\mathbf{n}} \\ -\tilde{\mathbf{n}} \star (\mathbf{F}^{vh})^T \star \mathbf{n} & \tilde{\mathbf{n}} \star \mathbf{F}^{hh} \star \tilde{\mathbf{n}} \end{pmatrix}. \end{aligned} \quad [9]$$

where $\mathbf{n}$ and $\tilde{\mathbf{n}}$ denote the occupation number of visible and hidden fermions, respectively.

In order to describe a physical wave-function, we need to project $|\psi\rangle$ back onto the Hilbert space of the real fermions. A natural choice is to make the hidden fermion occupations $\tilde{\mathbf{n}}$ depend on visible fermions $\mathbf{n}$, i.e. $\tilde{\mathbf{n}} \rightarrow \tilde{\mathbf{n}}(\mathbf{n})$. Then

$$\psi(\mathbf{n}) = \text{pf} \begin{pmatrix} \mathbf{n} \star \mathbf{F}^{vv} \star \mathbf{n} & \mathbf{n} \star \mathbf{F}^{vh} \star \tilde{\mathbf{n}}(\mathbf{n}) \\ -\tilde{\mathbf{n}}(\mathbf{n}) \star (\mathbf{F}^{vh})^T \star \mathbf{n} & \tilde{\mathbf{n}}(\mathbf{n}) \star \mathbf{F}^{hh} \star \tilde{\mathbf{n}}(\mathbf{n}) \end{pmatrix}, \quad [10]$$

In practice, one cannot easily parameterize the discrete distribution $\tilde{\mathbf{n}}(\mathbf{n})$ by ANNs. Alternatively, we define an $M \times \tilde{N}$ matrix $\mathbf{F}^{vh}(\mathbf{n}) = \mathbf{F}^{vh} \star \tilde{\mathbf{n}}(\mathbf{n})$ and an $\tilde{N} \times \tilde{N}$ matrix $\mathbf{F}^{hh}(\mathbf{n}) = \tilde{\mathbf{n}}(\mathbf{n}) \star \mathbf{F}^{hh} \star \tilde{\mathbf{n}}(\mathbf{n})$, such that

$$\psi(\mathbf{n}) = \text{pf} \begin{pmatrix} \mathbf{n} \star \mathbf{F}^{vv} \star \mathbf{n} & \mathbf{n} \star \mathbf{F}^{vh}(\mathbf{n}) \\ -\mathbf{F}^{vh}(\mathbf{n})^T \star \mathbf{n} & \mathbf{F}^{hh}(\mathbf{n}) \end{pmatrix}. \quad [11]$$

Then one can use ANNs to parameterize $\mathbf{F}^{vh}(\mathbf{n})$ and $\mathbf{F}^{hh}(\mathbf{n})$.

In SI Appendix, we review the HFDS (53), an earlier work combining the idea of hidden fermions with Slater determinants. The HFDS wave-function, can be written as

$$\psi_{\text{HFDS}}(\mathbf{n}) = \det \begin{pmatrix} \mathbf{n} \star \mathbf{u}^v \\ \mathbf{u}^h(\mathbf{n}) \end{pmatrix}, \quad [12]$$

where $\mathbf{u}^v$ is the visible orbitals, and $\mathbf{u}^h(\mathbf{n})$ is the hidden orbitals parameterized by ANNs.

As with the fermionic mean field, we can rewrite ψ_{HFDS} in the Pfaffian form,

$$\psi_{\text{HFDS}}(\mathbf{n}) = \text{pf} \left[\begin{pmatrix} \mathbf{n} \star \mathbf{u}^v \\ \mathbf{u}^h(\mathbf{n}) \end{pmatrix} \mathbf{J} \begin{pmatrix} \mathbf{n} \star \mathbf{u}^v \\ \mathbf{u}^h(\mathbf{n}) \end{pmatrix}^T \right], \quad [13]$$

where $\mathbf{J}$ is an arbitrary $(N + \tilde{N}) \times (N + \tilde{N})$ antisymmetric matrix with $\text{pf} \mathbf{J} = 1$. The HFPS wave-function in Eq. 11 can express any HFDS wave-function in Eq. 13 by letting $\mathbf{F}^{vv} = \mathbf{u}^v \mathbf{J} (\mathbf{u}^v)^T$, $\mathbf{F}^{vh}(\mathbf{n}) = \mathbf{u}^v \mathbf{J} \mathbf{u}^h(\mathbf{n})^T$, and $\mathbf{F}^{hh}(\mathbf{n}) = \mathbf{u}^h(\mathbf{n}) \mathbf{J} \mathbf{u}^h(\mathbf{n})^T$. Conversely, not all HFPS can be expressed by HFDS. For instance, a full-rank $\mathbf{F}^{vv}$ cannot be decomposed into $\mathbf{u}^v \mathbf{J} (\mathbf{u}^v)^T$. Therefore, the HFPS can be viewed as a generalization of HFDS.

In most modern ANN architectures, especially those utilized in NQS including the convolutional neural network (CNN) (42, 67, 68), group CNN (GCNN) (41), and transformers (69, 70), the shape of output matrices is usually $N_{\text{site}} \times C$, where N_{site} is the number of sites in the system and C is the number of output channels in the network. As shown by Fig. 1, these outputs can be directly reshaped into the required shape $M \times \tilde{N}$ for the matrix $\mathbf{F}^{vh}(\mathbf{n})$ while keeping translation equivariance. Nevertheless, the network output shape is not compatible with the $\tilde{N} \times \tilde{N}$ shape of $\mathbf{F}^{hh}(\mathbf{n})$. To solve this issue, we utilize the




Fig. 2. Illustration of HFPS architecture with sublattice and translation symmetry as an extension of the architecture in Fig. 1. The $\mathbf{F}^{vv}$ matrix elements have a banded structure due to the sublattice symmetry discussed in *SI Appendix*. The sliced matrices therefore only differ by a permutation applied to both the rows and columns, and the Pfaffian values are equivalent up to a permutation sign, as indicated by the “ $\approx$ ” symbol in the plot. In practice, the Pfaffian is only computed once to reduce the computational cost. Finally, the translation symmetries that are broken by the sublattice are applied to obtain the HFPS with translation symmetry.

spectral theory of antisymmetric matrices to decompose $\mathbf{F}^{hh}(\mathbf{n})$ into $\mathbf{F}^{hh}(\mathbf{n}) = \mathbf{V}^T(\mathbf{n})\tilde{\mathbf{F}}^{hh}\mathbf{V}(\mathbf{n})$, where $\tilde{\mathbf{F}}^{hh}$ is an $\tilde{N} \times \tilde{N}$ antisymmetric matrix, $\mathbf{V}(\mathbf{n})$ is an $\tilde{N} \times \tilde{N}$ unitary matrix, and we attribute all $\mathbf{n}$ dependence to $\mathbf{V}(\mathbf{n})$ without loss of generality. By defining $\tilde{\mathbf{F}}^{vh}(\mathbf{n}) = \mathbf{F}^{vh}(\mathbf{n})\mathbf{V}^T(\mathbf{n})$ and $\mathbf{U}(\mathbf{n}) = \text{diag}(\mathbf{1}, \mathbf{V}(\mathbf{n}))$, we rewrite Eq. 11 into

$$\begin{aligned} \psi(\mathbf{n}) &= \text{pf} \left[\mathbf{U}^T(\mathbf{n}) \begin{pmatrix} \mathbf{n} \star \mathbf{F}^{vv} \star \mathbf{n} & \mathbf{n} \star \tilde{\mathbf{F}}^{vh}(\mathbf{n}) \\ -\tilde{\mathbf{F}}^{vh}(\mathbf{n})^T \star \mathbf{n} & \tilde{\mathbf{F}}^{hh} \end{pmatrix} \mathbf{U}(\mathbf{n}) \right] \\ &= J(\mathbf{n}) \times \text{pf} \begin{pmatrix} \mathbf{n} \star \mathbf{F}^{vv} \star \mathbf{n} & \mathbf{n} \star \tilde{\mathbf{F}}^{vh}(\mathbf{n}) \\ -\tilde{\mathbf{F}}^{vh}(\mathbf{n})^T \star \mathbf{n} & \tilde{\mathbf{F}}^{hh} \end{pmatrix}, \end{aligned} \quad [14]$$

where we have defined $J(\mathbf{n}) = \det(\mathbf{U}(\mathbf{n}))$. Therefore, we drop the dependence of hidden-hidden pairing on $\mathbf{n}$ by introducing an additional Jastrow factor $J(\mathbf{n})$. In practice, $J(\mathbf{n})$ and $\tilde{\mathbf{F}}^{vh}(\mathbf{n})$ are computed by the ANN, while $\mathbf{F}^{vv}$ and $\tilde{\mathbf{F}}^{hh}$ are directly treated as variational parameters. In Fig. 1, we show how to obtain $J(\mathbf{n})$ and $\tilde{\mathbf{F}}^{vh}(\mathbf{n})$ from a CNN and combine them with the Pfaffian. In *SI Appendix*, we further show that the HFPS can be viewed as a Pfaffian NNBF with controlled rank, which allows us to perform low-rank updates (LRU) to greatly accelerate the computation of Pfaffian matrices. The LRU technique can also be utilized in HFDS and low-rank NNBF to reduce the computational cost.

In Fig. 2, we show an illustration of HFPS with sublattice structure, which reduces the number of Pfaffian computations required to build a translationally symmetric wave-function.

1.3. Optimization. Typically, NQS are trained using stochastic reconfiguration (SR), which implements imaginary time evolution. Neglecting errors that arise from sampling bias and model expressivity, the variational state $|\psi_\tau\rangle$ at training iteration τ can be expressed in terms of the initial state $|\psi_0\rangle$ as $|\psi_\tau\rangle = \sum_i e^{-\tau E_i} \langle i | \psi_0 \rangle |i\rangle$, where $|i\rangle$'s are the eigenstates of the Hamiltonian with energies E_i . This implies that convergence time strongly depends on the overlap between the initial state and the ground state. As an additional consequence, a variational state with a pathological initialization, such that the overlap is very large with a different eigenstate than the ground state, is very likely to get stuck in a local minimum. Therefore, we improve training efficiency by starting the mean-field part $\mathbf{F}^{vv}$ of HFPS from a mean-field state that we expect to have sizable overlap with the ground state. We either optimize a Slater determinant state in Eq. 7, a BCS state in Eq. 3, or a full Thouless state in Eq. 2,

to obtain an optimal mean-field state $|\psi_0\rangle$ and a suitable initial $\mathbf{F}^{vv}$. We train these models using exact gradient descent of variational energy with automatic differentiation, taking advantage of the fact that observables of mean-field states can be computed in $\mathcal{O}(M^2)$ time without Monte-Carlo sampling. A well-known problem with mean-field methods is that they tend to overestimate long-range order. On the square lattice Hubbard, it has been shown that a mean-field Slater determinant trained with a small interaction, $U \approx 2.7$ for 1/8 doping, has the largest overlap with the ground state at $U = 8$ (71). Therefore, we initialize our time evolution with a mean-field state trained at $|U| = 3$. In some cases, we also find it is best to add a d-wave pairing during mean-field optimization, as detailed in Section 2.4.

After properly initializing $\mathbf{F}^{vv}$, we train all variational parameters in Eq. 14, including $\mathbf{F}^{vv}$, $\tilde{\mathbf{F}}^{hh}$, and ANN weights θ , to search for the ground state of interacting fermion systems by employing SR (72–74) to perform imaginary-time evolution. To solve the SR equation, we utilize minimum-norm SR (MinSR) (42) and subsampled projected-increment natural gradient descent (SPRING) (75) with momentum $\mu = 0.5$.

In most simulations, we utilize a CNN similar to the ResNet in refs. 42 and 70 with 16 layers (8 residual blocks), kernel size 3×3 , channel number $C = 32$, and in total around 150 K parameters. Among 32 channels, 28 of them are transformed into $\tilde{\mathbf{F}}^{vh}$ with $\tilde{N} = 14$ hidden fermions in the studied spinful system, and the remaining channels are used for the generalized Jastrow factor $J(\mathbf{n})$. The symmetry projection explained in *SI Appendix* is also implemented to improve the accuracy. For simulations in Section 2.4, we utilize a GCNN (41). Unless otherwise specified, the total number of samples in each VMC iteration is fixed at $N_s = 16,384$, and the training takes around 10^4 VMC iterations.

1.4. Reducing the Asymptotic Complexity. For large fermionic NQS with either determinants or Pfaffians, the time complexity of every VMC iteration is dominated by the following parts.

- Monte Carlo sampling or local energy computation in VMC: $\mathcal{O}(N)$;
- Computing determinant or Pfaffian of matrices: $\mathcal{O}((N + \tilde{N})^3)$;
- Imposing translation symmetry: $\mathcal{O}(M)$.

Then the total complexity is $\mathcal{O}(NM(N + \tilde{N})^3)$ for each VMC step. This complexity exceeds that of a CNN or GCNN, which costs $\mathcal{O}(M)$ per forward pass and $\mathcal{O}(NM)$ in each VMC step. The particle density n and hidden fermion number $\tilde{N}$ are usually kept unchanged when one increases the system size,




Fig. 3. Comparison of the time cost of an HFPS forward pass with and without the acceleration given by low-rank updates (LRU) and sublattice symmetry.

so the complexity can be simplified as $\mathcal{O}(N^5)$. This complexity becomes the main bottleneck in the NQS simulation of large fermion systems, limiting most of the previous fermionic NQS attempts to small systems or low accuracy (52, 53, 76).

Fortunately, there are several techniques one can employ in HFPS to greatly reduce the complexity. In Monte Carlo sampling or local energy computation, the sliced Pfaffian orbitals are often generated by replacing a few rows and columns from previous Pfaffian orbitals if the Hamiltonian only contains local interactions. In this case, one does not have to recompute the full Pfaffian, as the low-rank update (LRU) can be utilized to reduce the complexity from $\mathcal{O}((N + \tilde{N})^3)$ to $\mathcal{O}(N^2\tilde{N})$. In this work, we utilize our open-source package *lrux* (62) for efficient LRU. Although LRU can be applied to NNBF as well, it in general imposes a strong constraint on network architectures and is not as natural as HFDS, which we explain in *SI Appendix* in detail. Therefore, the direct utilization of LRU provides a great advantage of the hidden fermion formalism, including HFDS and HFPS.

On the other hand, the $\mathcal{O}(M)$ complexity stemming from imposing translation symmetry can also be reduced. The CNN and GCNN utilized in this work are translation covariant by definition, and the translation symmetry breaking only happens in the pairing orbitals between visible fermions. Naively, one can apply constraints to make the visible pairing orbitals fully translation invariant, but this design choice impedes the ability to represent symmetry-broken phases. Alternatively, one can enforce Pfaffians to only be invariant under translation over multiple lattice sites, which is often described as a sublattice symmetry (38). In this construction, the wave-function has an associated unit cell, spanned by the sublattice translation vectors, in which translation symmetry is broken. In *SI Appendix*, we show that the computational cost of translation symmetry projection does not scale up with system size when sublattice symmetry is utilized. As illustrated in Fig. 2, the computation of the translation-symmetrized wave-function involves M terms, but they are equivalent up to a fermion permutation sign due to sublattice structure, so one only needs to

compute S terms, where S is the size of the sublattice unit cell independent of the full system size. With these techniques, the time complexity of the HFPS forward pass is reduced to $\mathcal{O}(N^2)$ and the complexity of each VMC step is reduced to $\mathcal{O}(N^3)$. We note that sublattice symmetry can also be imposed on determinants in principle (*SI Appendix*), but in practice we find they are more difficult to train than in the Pfaffian case.

In Fig. 3, we show the time cost of 10^3 parallel HFPS forward passes on an A6000 GPU. As illustrated in the figure, the complexity is greatly reduced from $\mathcal{O}(N^4)$ to $\mathcal{O}(N^2)$ by utilizing LRU and sublattice. For the largest lattice with 400 sites, the forward pass is accelerated by a factor of 10^3 , making it possible to simulate large fermionic systems.

Although the LRU and sublattice techniques can be utilized in HFPS to provide great acceleration, they have not been utilized in previous fermionic NQS work. As shown in Table 1, only PP+RBM (51) has utilized these techniques, while NNBF (52, 77) and HFDS (53) with multilayer perceptrons (MLP) have not. However, the PP+RBM only utilizes a multiplicative Jastrow and not an orbital correction, therefore it cannot easily represent interaction-induced modifications to the nodal structure. Therefore, we expect HFPS to be an ideal NQS architecture for large fermionic systems with strongly correlated electrons.

2. Results

To illustrate the performance of HFPS, we employ it to study the two-dimensional Hubbard model on the square lattice

$$\hat{\mathcal{H}} = -t \sum_{\langle ij \rangle, \sigma} \hat{c}_{i, \sigma}^\dagger \hat{c}_{j, \sigma} + U \sum_i \hat{n}_{i, \uparrow} \hat{n}_{i, \downarrow}. \quad [15]$$

We will show the performance of HFPS in different phases of the Hubbard model. To quantify the accuracy of variational states, we measure the variational energy E in VMC and compute the relative error of the variational energy defined as

$$\epsilon_{\text{rel}} = \frac{E - E_0}{E_\infty - E_0}, \quad [16]$$

where E_0 is the exact ground-state energy, usually provided by ED in small systems or QMC in sign-free cases, and E_∞ is the energy zero point given by the energy at infinite temperature. The reason for introducing E_∞ is to remove the dependency of ϵ_{rel} on a constant shift of the Hamiltonian (6).

The raw data of our numerical experiments are included in *SI Appendix*.

2.1. Small-Scale Benchmark. We first consider a case in which the density is low (close to quarter filling). Electrons are then relatively free to move around in the Hubbard model, and the system is in a rather weakly correlated regime even if U is large. In the thermodynamic limit, the system is expected to be in a Fermi liquid state with well-defined quasiparticle excitations. As a starting point, we study the performance of HFPS in this simple case to demonstrate its ability to encode correlations.

In Fig. 4, we show the results of various fermionic NQs in the 4×4 Hubbard model with periodic boundary conditions

Table 1. Scaling complexity of different fermionic NQS wave-functions that have been used in the literature

NQS	Low-rank update (LRU)	Sublattice	Local update with symmetry	VMC with symmetry
PP+RBM (51)	Yes	Yes	$\mathcal{O}(N^2)$	$\mathcal{O}(N^3)$
MLP-NNBF (52, 77)	No	No	$\mathcal{O}(N^4)$	$\mathcal{O}(N^5)$
HFDS (53, 78)	Yes*	No	$\mathcal{O}(N^3)$	$\mathcal{O}(N^4)$
HFPS (this work)	Yes	Yes	$\mathcal{O}(N^2)$	$\mathcal{O}(N^3)$

The PP+RBM only uses a neural network Jastrow, while the other three also use a neural network to modify the fermion orbitals. *Low-rank updates were not used for these implementations of HFDS but they could be used with no architectural modifications.




Fig. 4. Relative error of energy ϵ_{rel} in the 4×4 Hubbard model with particle density $n = 5/8$. The definition of ϵ_{rel} is given in Eq. 16, and the error is computed with respect to the ED energy. The results given by the determinant with RBM (Det+RBM) and hidden fermion determinant state (HFDS) (53) are provided to compare with our HFPS energy.

(PBC), visible particle number $N = 10$, and particle density $n = 5/8$, within the Fermi liquid regime. The HFPS provides a great improvement in variational energy compared with the neural Jastrow method and HFDS (53). Compared with the best previous NQS result given by HFDS, the HFPS error is roughly reduced by a factor of 100. We attribute this improvement to the utilization of a modern deep network architecture in HFPS as compared to the shallow fully connected network in HFDS.

2.2. Superconductivity with Attractive Interactions. When the on-site interaction of the Hubbard model is attractive with $U < 0$, the electrons form Cooper pairs and exhibit s-wave superconductivity. This superconducting state evolves from a BCS regime (54) at smaller $|U|$ to a strong coupling regime at larger $|U|$ where local pairs form with a large binding energy but become phase coherent at a smaller energy scale, forming a Bose–Einstein condensate (BEC) (79). Although the nature of

the off-diagonal long-range order (ODLRO) in the ground-state is qualitatively captured by the BCS wave-function in Eq. 3, the BCS state becomes quantitatively inaccurate at moderate to large $|U|$. To quantify the superconductivity in a system with a conserved particle number, the usual order parameter $\langle \hat{\Delta}(\mathbf{r}) \rangle = \langle \hat{c}_{\mathbf{r},\uparrow} \hat{c}_{\mathbf{r},\downarrow} \rangle$ in the BCS wave-function is not directly accessible. Instead, we define the pair–pair correlation

$$C_p(\mathbf{r}) = \langle \hat{\Delta}^\dagger(\mathbf{r}) \hat{\Delta}(\mathbf{0}) \rangle. \quad [17]$$

Although $C_p(\mathbf{r})$ remains finite even without Cooper pairs when $|\mathbf{r}|$ is small, it can serve as an order parameter of BCS superconductivity when $|\mathbf{r}| \rightarrow \infty$, corresponding to ODLRO. Therefore, we can use the long-range pair correlations as a proxy for the pairing amplitude in the thermodynamic limit

$$|\langle \hat{\Delta} \rangle| \sim \left| \lim_{|\mathbf{r}| \rightarrow \infty} C_p(\mathbf{r}) \right|^{1/2}. \quad [18]$$

In Fig. 5, we show $C_p(\mathbf{r})$ given by several methods in the 8×8 attractive Hubbard model with particle density $n = 7/8$ and PBC. The BCS state is defined in Eq. 3 and optimized by exact gradient descent of energy. It is a qualitatively correct description, as we expect for attractive electrons, and provides good accuracy under weak interaction. However, the BCS wave-function tends to overestimate the long-range correlation when the interaction is strong because of the underestimation of quantum fluctuations. Consequently, it is not a quantitatively accurate description for electrons with strong attractions. In HFPS, although the mean-field part $\mathbf{F}^{vv}$ is initialized by the BCS wave-function, it is trained to encode correlations through the ANN and achieves great accuracy in predicting pair–pair correlations, within the error bar of the exact result given by DQMC. The relative error of energy ϵ_{rel} given by HFPS also reaches the level of 10^{-4} , greatly outperforming the 10^{-2} level provided by the BCS wave-function. Therefore, HFPS provides a reliable approach in studying the behavior of systems with BCS superconductivity and strong interactions, for instance, the crossover to BEC.

2.3. Mott Insulator. At half-filling (one particle per site on average) and for strong repulsive interactions $U > 0$, the kinetic energy of electrons is suppressed, and the ground-state of the




Fig. 5. The pair–pair correlation $C_p(\mathbf{r})$ in the 8×8 Hubbard model with attractive interaction $U < 0$ and particle density $n = 7/8$. The definition of $C_p(\mathbf{r})$ is given in Eq. 17. The methods shown here include the BCS wave-function in Eq. 3 with minimized variational energy, our HFPS, and the determinant quantum Monte Carlo (DQMC) for an exact reference in this sign-free system. (A) $C_p(\mathbf{r})$ is measured at $\mathbf{r} = (x, x)$. The dashed and solid lines represent $U = -2$ and $U = -8$, respectively. (B) $C_p(\mathbf{r})$ is measured at $\mathbf{r} = (4, 4)$, the furthest distance in the 8×8 system with PBC. The uncertainty is smaller than the marker size in both panels.

Hubbard model on the square lattice is a Mott insulator with antiferromagnetic long-range order. On a bipartite lattice, the half-filling case is special since it does not display a sign problem in Monte Carlo methods and projective auxiliary field quantum Monte Carlo (AFQMC) becomes exact. Hence, AFQMC provides a benchmark to which variational wave-functions can be compared. Furthermore, the spin-balanced Hubbard model with U and $-U$ can be considered equivalent at half filling. This can be shown from a partial particle-hole transformation on spin-down operators $\hat{c}_{i,\downarrow} \rightarrow \text{sign}(i) \hat{c}_{i,\downarrow}^\dagger$, where $\text{sign}(i) = +1$ for one bipartite lattice and -1 for the other. Then the original Hubbard Hamiltonian in Eq. 15 is transformed into

$$\hat{\mathcal{H}} = -t \sum_{\langle ij \rangle, \sigma} \hat{c}_{i,\sigma}^\dagger \hat{c}_{j,\sigma} - U \sum_i \hat{n}_{i,\uparrow} \hat{n}_{i,\downarrow} + U \sum_i \hat{n}_{i,\uparrow}. \quad [19]$$

As the last term is a constant for a system with conserved spin-up and spin-down particles, this transformation shows that the attractive and repulsive Hubbard models are equivalent at half filling. In this work, we utilize this freedom of basis choice to simulate an attractive Hubbard model for the Mott insulator phase in HFPS, which allows us to use the same technique as the simulations in Section 2.2 and initialize HFPS using a BCS state in Eq. 3 trained with exact gradient.

In Fig. 6, we show the energy error of different NQs in the 8×8 lattice at half-filling. The HFDS result is provided in the system with PBC along one side and antiperiodic boundary condition (APBC) along the other, while other results are provided with PBC along both sides. The AFQMC reference energy is selected with respective boundary conditions for each method to compute the relative error ϵ_{rel} . The existing numerical methods, including PP+RBM with an exact Lanczos step (LS), HFDS (53), and fPEPS with bond dimension $D = 16$, provide ϵ_{rel} at the level of 10^{-3} , while HFPS provides better accuracy. The fPEPS result presented here is optimized by the simple update method and not

pushed to its limiting expressive power. We expect the gradient-based optimization (17, 81, 82) will reduce the error of fPEPS by a factor of 2 to 10, but HFPS still achieves a similar or better accuracy. Compared to a recent result provided by Transformer-NNBF (80), our HFPS achieves similar accuracy. It is worth noting that the uncertainty of the benchmark AFQMC energy is around 10^{-4} , which leads to a big uncertainty of ϵ_{rel} when the accuracy of HFPS also reaches this level, as shown in Fig. 6.

As also presented in Fig. 6, the HFPS wave-functions with no sublattice and 4×4 sublattice show similar error because a 4×4 unit cell is large enough to encode the translational symmetry breaking of the best mean-field state. The sublattice symmetry provides great acceleration in VMC, allowing us to train with more iterations given a similar total time cost. Therefore, the HFPS with sublattice reaches slightly better accuracy compared to the one without sublattice, although the former in principle has lower expressive power.

2.4. Stripe Phase. In order to test our method on a computationally challenging task, we study the Hubbard model with small $\delta = 1/8$ hole doping (particle density $n = 7/8$) and strong repulsive interactions $U = 8$. In this regime, the Hubbard model displays stripe order, where one-dimensional hole-rich domain walls lie between patches of antiferromagnetic order. The sublattice polarization switches across the domain walls so that when electrons hop across, the magnetism is not frustrated. This manifests itself as a density wave pattern of charge and spin ordering, with an associated structure factor peak. At the same time, the doped square lattice Hubbard model has a tendency toward d-wave pairing, and states with off-diagonal long-range order (ODLRO) are energetically competitive (71, 83–85). Determining whether density-wave order, superconducting order, or both, survive to the thermodynamic limit is an ongoing problem for condensed matter physics that requires accurate variational methods. In existing studies, it is widely believed that the ground state of the Hubbard model with $U = 8$ and $\delta = 1/8$ is a density wave phase without ODLRO (84, 86, 87). While our simulations corroborate these findings, we find that d-wave pairing still plays an essential role in the short-range physics and energetics.

As explained in Section 1.3, suitable initialization with a mean-field state is crucial for avoiding local minima, especially in the stripe regime where there are many competing low-energy states. For a Hubbard model with $U \geq 0$, the mean-field state with the lowest energy is a Slater determinant state without pairing. For the 4×4 and 8×4 lattices, we indeed achieve good accuracy starting from a Slater determinant trained at $U = 3$. Nevertheless, we find that on the 16×4 lattice we are able to improve our final energies by training the mean-field state with an additional d-wave pairing field, i.e.

$$\hat{\mathcal{H}}' = \hat{\mathcal{H}} + \Delta \sum_{\langle ij \rangle} \text{sign}(i, j) (\hat{c}_{i,\uparrow} \hat{c}_{j,\downarrow} - \hat{c}_{i,\downarrow} \hat{c}_{j,\uparrow} + H.C.), \quad [20]$$

where $\hat{\mathcal{H}}$ is the original Hubbard Hamiltonian in Eq. 15 with a weaker interaction $U = 3$, Δ is the strength of the d-wave field, and $\text{sign}(i, j)$ denotes the sign for d-wave pairing, which is $+1$ for horizontal neighbors and -1 for vertical neighbors. The lowest energy state of this Hamiltonian is a general Pfaffian state with both stripe order and d-wave pairing. For the 16×4 stripe, we train our mean-field state with $\Delta = 0.2$. In the following, we utilize a GCNN in our HFPS with roughly 600 K parameters.

The variational energies of different methods on $L \times 4$ PBC lattices are presented in Fig. 7. In this system, HFPS and the recent Transformer-NNBF (80) greatly outperform earlier works based on shallow ANNs. These results showcase the importance




Fig. 6. Relative error of energy ϵ_{rel} in the 8×8 Hubbard model at half-filling (particle density $n = 1$). Here we include the results provided by PP+RBM with a Lanczos step (LS) (51), HFDS (53), Transformer-NNBF (80), and fPEPS with bond dimension $D = 16$ (provided by J. Gray, private communication) for comparison. The ground-state energy is estimated by AFQMC at $U = -4$ and -8 (6) for the computation of ϵ_{rel} . The uncertainty of the relative error of HFPS and Transformer-NNBF mainly comes from the uncertainty of AFQMC, and in other methods the uncertainty is smaller than the marker size.




Fig. 7. Variational energy in the spin- and charge-ordered (stripe) ground state of the $L \times 4$ Hubbard model with repulsive interaction $U = 8$ and particle density $n = 7/8$. The presented variational results include multilayer perceptron neural network backflow (MLP-NNBF) (52), hidden fermion determinant state (HFDS) (53), tensor backflow (TBF) (88), TBF with a Lanczos step (LS) (88), Transformer-NNBF (80), and HFPS.

of large-scale simulations in fermionic NQS, as also demonstrated in previous simulations on spin systems (42, 89). Compared to the ground state energy marked by the black star in the small 4×4 system, HFPS provides nearly exact energy with $\epsilon_{\text{rel}} = 1.2 \times 10^{-5}$, while HFDS has $\epsilon_{\text{rel}} = 3 \times 10^{-4}$. In larger systems, we also present tensor backflow (TBF) (88), a simple backflow with coefficients directly represented by tensors instead of ANNs, for reference. As shown in the plot, the previous MLP-NNBF (52) and HFDS (53) provide worse energy than the trivial tensor representation of backflow, indicating that the ANNs in these methods probably do not encode correlations between electrons properly. Our HFPS method, in contrast to these NQs, outperforms TBF and even the TBF with a Lanczos step (LS), demonstrating that the scalable HFPS architecture captures the essential degrees of freedom in strongly correlated electrons. The energy we obtain in the 16×4 lattice is $E/N = -0.76413$, outperforming a recent result $E/N = -0.76298$ produced by Transformer-NNBF (80). We believe that this energy discrepancy is likely due to the presence of strong short-ranged d-wave pairing in the ground state (86), as we find a fairly similar energy, $E/N = -0.76256$, to Transformer-NNBF when we initialize without a pairing field, which drives us to a meta-stable state without pairing. This is evidence that using Pfaffian-based wave-functions, which are better equipped to model pairing, may be useful even for phases without ODLRO, as long as there is short-range pairing.

Utilizing the accurate numerical solution in the 16×4 lattice, we compute the correlations in the system to verify the existence of stripe patterns and short-range d-wave pairings, including the spin-spin correlation

$$C_s(\mathbf{r}) = \langle \hat{\mathbf{S}}_{\mathbf{r}} \cdot \hat{\mathbf{S}}_{\mathbf{0}} \rangle - \langle \hat{\mathbf{S}}_{\mathbf{r}} \rangle \cdot \langle \hat{\mathbf{S}}_{\mathbf{0}} \rangle, \quad [21]$$

the charge-charge correlation

$$C_c(\mathbf{r}) = \langle \hat{n}_{\mathbf{r}} \cdot \hat{n}_{\mathbf{0}} \rangle - \langle \hat{n}_{\mathbf{r}} \rangle \cdot \langle \hat{n}_{\mathbf{0}} \rangle, \quad [22]$$

and the pair-pair correlation

$$C_p(\mathbf{r}) = \langle \hat{\Delta}^\dagger(\mathbf{r}) \hat{\Delta}(\mathbf{0}) \rangle \quad [23]$$

for d-wave symmetrized pairs

$$\hat{\Delta}(\mathbf{r}) = \frac{1}{4} \sum_{\delta} \frac{1}{\sqrt{2}} \text{sign}(\delta) (\hat{c}_{\mathbf{r},\uparrow} \hat{c}_{\mathbf{r}+\delta,\downarrow} - \hat{c}_{\mathbf{r},\downarrow} \hat{c}_{\mathbf{r}+\delta,\uparrow}), \quad [24]$$

where $\sum_{\delta}$ iterates over nearest neighbors, and $\text{sign}(\delta) = +1$ for horizontal neighbors and -1 for vertical neighbors. The correlations are presented in Fig. 8 A–C.

To show a more clear tendency, in Fig. 8D we choose $\mathbf{r} = (x, 2)$ and plot the staggered spin correlation $(-1)^x C_s(\mathbf{r})$ and the charge correlation $C_c(\mathbf{r})$ as a function of x . The spin correlation shows a strong spin density wave (SDW) with spatial period 16, which matches the $1/8$ doping in the system. The charge correlation indicates a superposition of short-range repulsion and the charge density wave (CDW) with spatial period 8, also consistent with the $1/8$ doping. In Fig. 8E, we show the spin and charge structure factor

$$C_{\gamma}(\mathbf{k}) = \sum_{\mathbf{r}} C_{\gamma}(\mathbf{r}) e^{-i\mathbf{k} \cdot \mathbf{r}}, \quad [25]$$

where γ represents either the spin structure s or the charge structure c . The momentum cuts are chosen to be $\mathbf{k} = (k_x, \pi)$ for spin structures and $\mathbf{k} = (k_x, 0)$ for charge structures so that they pass through their respective peaks at $\mathbf{k} = (7\pi/8, \pi)$ and $\mathbf{k} = (\pi/4, 0)$, reflecting the existence of SDW with period 16 and CDW with period 8. These results confirm the existence of density wave order in the doped Hubbard model. In Fig. 8F, we plot $C_p(\mathbf{r})$ for $\mathbf{r} = (x, 2)$, which shows finite short-range d-wave pairings coexisting with SDW and CDW. Since the pairing decays over long distances, our results at $1/8$ -doping do not display long-range superconducting order. However, some mechanisms, e.g., the next-nearest-neighbor hopping t' (85), might enhance the pairing to turn the system into an unconventional superconductor.

3. Discussion

In this work, we introduced an ANN-augmented Pfaffian wavefunction, HFPS, for performing variational Monte Carlo on interacting fermion problems. This construction generalizes the hidden fermion formalism (53) to an antisymmetrized wavefunction of fermionic pairs, the Pfaffian. The motivation for this switch is to accurately model superconductors, as the Pfaffian can represent them at the mean-field level.

We describe how to initialize the HFPS with a good fermionic mean-field, so that the training converges well. Additionally, we utilize various numerical techniques that reduce the overall VMC complexity from $\mathcal{O}(N^5)$ to $\mathcal{O}(N^3)$, hence providing a systematic approach to large-scale simulations of strongly correlated electrons.

We demonstrate the performance of HFPS in various phases of the Hubbard model, including the superconducting phase of the attractive case in both weak coupling (BCS) and strong coupling (BEC), the antiferromagnetic Mott insulator, and the stripe phase. The HFPS is able to achieve significantly higher accuracy than existing variational methods and produces accurate energies and correlation functions in sign-free cases when benchmark comparisons with QMC are possible.

The progress allowed by HFPS pushes fermionic NQS studies from proof-of-principle results on small systems to a new stage where it can be used to study challenging problems in a systematic fashion. We believe that this will allow us to shed light on fermion problems with strong and long-range correlations




Fig. 8. The spin, charge, and pair correlations in the 16×4 Hubbard model with 1/8 doping and $U = 8$. (A–C) The correlations C_s , C_c , and C_p defined by Eqs. 21–23 in the full real space. For visibility of weak correlations, their strength is truncated to ceiling values $|C_s| \leq 0.1$ and $|C_c|, |C_p| \leq 0.01$. The star marks the central site with $\mathbf{r} = (0, 0)$. (D) The staggered spin-spin correlation $(-1)^x C_s(\mathbf{r})$ and the charge-charge correlation $C_c(\mathbf{r})$ with $\mathbf{r} = (x, 2)$. The two curves are given different scales to make the oscillations prominent. (E) The spin structure factor $C_s(\mathbf{k})$ with $\mathbf{k} = (k_x, \pi)$ and the charge structure factor $C_c(\mathbf{k})$ with $\mathbf{k} = (k_x, 0)$ in Eq. 25. The arrows indicate the expected structure factor peaks for SDW and CDW. (F) The d-wave pair-pair correlation $C_p(\mathbf{r})$ with $\mathbf{r} = (x, 2)$. The uncertainty is smaller than the marker size in (D–F).

in 2D, similar to how NQS has improved our understanding of quantum spin systems.

Finally, we list several possible future directions for HFPS and more generally for fermionic NQS.

- In Fig. 5, we show that HFPS correctly captures s-wave superconductivity in the attractive Hubbard model and provides systematic improvement on the BCS wave-function. As the Pfaffian wave-function represents a general Gaussian state, it can also capture unconventional pairing (e.g., p-wave and d-wave superconductivity) given a suitable initialization specified by the mean-field Hamiltonian $\hat{H}_0$ in Eq. 1. Furthermore, as shown by Figs. 7 and 8, the HFPS reaches high accuracy in the stripe phase with strong correlations, which has significant d-wave short-range pairing, and may be closely related to superconducting phases. These results reveal the great potential of HFPS for studying high-temperature superconductivity in the t - t' square lattice Hubbard model. While there is strong evidence for d-wave superconductivity, especially with $|t'| > 0$ (84, 85, 87), NQS could provide detailed descriptions of these phases, including determining whether there are residual density wave orders, and the spatial extent of the pair binding (90).
- HFPS could be used to study multiband/multiorbital models of strongly correlated electrons, as relevant to the description of many quantum materials such as nickelate superconductors in the two-band case. It could also provide more accurate computational descriptions of cuprates by including the oxygen orbitals. Additionally, HFPS could be used to study Hofstadter-Hubbard models and effective models for moiré materials.

- By measuring particle correlations in the many-body NQS, one might be able to construct the effective low-energy Hamiltonian governing the behavior of quasi-particles. This would provide insights into the mechanism of d-wave pairing in high-temperature superconductors and help us to better understand the rich phases of cuprates and other quantum materials with strong electronic correlations.

The ED is performed with QuSpin (91, 92). The NQS simulations are performed by Quantax (93), in which the scalable ANNs are implemented by JAX (94) and Equinox (95), and the LRU is implemented by lux (62).

Data, Materials, and Software Availability. Code for training machine learning models have been deposited in Quantax (<https://github.com/chrisrothUT/quantax/tree/custom/quantax>) (96). All other data are included in the manuscript and/or SI Appendix.

ACKNOWLEDGMENTS. A.C. acknowledges Markus Heyl for his support and Johnnie Gray for providing fPEPS data. We are grateful to Miguel Morales, Javier Robledo-Moreno, Conor Smith, Shiwei Zhang, and Garnet Chan for useful discussions. The Flatiron Institute is a division of the Simons Foundation.

Author affiliations: ^aCenter for Computational Quantum Physics, Flatiron Institute, New York, NY 10010; ^bDivision of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, CA 91125; ^cTheoretical Physics III, Center for Electronic Correlations and Magnetism, Institute of Physics, University of Augsburg, Augsburg 86135, Germany; ^dCenter for Computational Mathematics, Flatiron Institute, New York, NY 10010; ^eDepartment of Physics and Astronomy, Rutgers University, Piscataway, NJ 08854; ^fCollège de France, Paris 75005, France; ^gCentre de Physique Théorique, CNRS, École Polytechnique, IP Paris, Palaiseau F-91128, France; and ^hDepartment of Quantum Matter Physics, Université de Genève, Genève CH-1211, Suisse

Author contributions: A.C., A.S., A.G., and C.R. designed research; A.C. and C.R. performed research; A.C. and C.R. contributed new reagents/analytic tools; A.C. and Z.-Q.W. analyzed data; and A.C., A.S., A.G., and C.R. wrote the paper.

Reviewers: D.L., University of California, Los Angeles; and M.S., Universität Regensburg.

1. J. G. Bednorz, K. A. Müller, Possible high- T_c superconductivity in the Ba-La-Cu-O system. *Z. Phys. B Condens. Matter* **64**, 189–193 (1986).
2. M. K. Wu *et al.*, Superconductivity at 93 K in a new mixed-phase Y-Ba-Cu-O compound system at ambient pressure. *Phys. Rev. Lett.* **58**, 908–910 (1987).
3. M. Qin, T. Schäfer, S. Andergassen, P. Corboz, E. Gull, The Hubbard model: A computational perspective. *Ann. Rev. Condens. Matter Phys.* **13**, 275–302 (2022).
4. D. P. Arovas, E. Berg, S. A. Kivelson, S. Raghu, The Hubbard model. *Ann. Rev. Condens. Matter Phys.* **13**, 239–274 (2022).
5. F. Becca, S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, 2017).
6. D. Wu *et al.*, Variational benchmarks for quantum many-body problems. *Science* **386**, 296–301 (2024).
7. S. R. White, Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **69**, 2863–2866 (1992).
8. S. R. White, Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B* **48**, 10345–10356 (1993).
9. U. Schollwöck, The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.* **326**, 96–192 (2011).
10. S. R. White, D. J. Scalapino, Density matrix renormalization group study of the striped phase in the $2d - t - J$ model. *Phys. Rev. Lett.* **80**, 1272–1275 (1998).
11. H. C. Jiang, T. P. Devereaux, Superconductivity in the doped Hubbard model and its interplay with next-nearest hopping. *Science* **365**, 1424–1428 (2019).
12. H. C. Jiang, S. A. Kivelson, Stripe order enhanced superconductivity in the Hubbard model. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2109406119 (2022).
13. P. Corboz, G. Vidal, Fermionic multiscale entanglement renormalization ansatz. *Phys. Rev. B* **80**, 165129 (2009).
14. P. Corboz, R. Orús, B. Bauer, G. Vidal, Simulation of strongly correlated fermions in two spatial dimensions with fermionic projected entangled-pair states. *Phys. Rev. B* **81**, 165104 (2010).
15. C. V. Kraus, N. Schuch, F. Verstraete, J. I. Cirac, Fermionic projected entangled pair states. *Phys. Rev. A* **81**, 052338 (2010).
16. Q. Mortier *et al.*, Fermionic tensor network methods. *SciPost Phys.* **18**, 012 (2025).
17. W. Y. Liu, H. Zhai, R. Peng, Z. C. Gu, G. K. L. Chan, Accurate simulation of the Hubbard model with finite fermionic projected entangled pair states. *Phys. Rev. Lett.* **134**, 256502 (2025).
18. R. Blankenbecler, D. J. Scalapino, R. L. Sugar, Monte Carlo calculations of coupled boson-fermion systems. I. *Phys. Rev. D* **24**, 2278–2286 (1981).
19. J. E. Hirsch, Two-dimensional Hubbard model: Numerical simulation study. *Phys. Rev. B* **31**, 4403–4419 (1985).
20. S. R. White *et al.*, Numerical study of the two-dimensional Hubbard model. *Phys. Rev. B* **40**, 506–516 (1989).
21. D. J. Scalapino, S. R. White, S. Zhang, Insulator, metal, or superconductor: The criteria. *Phys. Rev. B* **47**, 7995–8007 (1993).
22. F. Assaad, H. Evertz, *World-line and Determinantal Quantum Monte Carlo Methods for Spins, Phonons and Electrons*, H. Fehske, R. Schneider, A. Weiße, Eds. (Springer, Berlin Heidelberg, Berlin, Heidelberg, 2008), pp. 277–356.
23. G. Sugiyama, S. Koonin, Auxiliary field Monte-Carlo for quantum many-body ground states. *Ann. Phys.* **168**, 1–26 (1986).
24. S. Zhang, J. Carlson, J. E. Gubernatis, Constrained path Monte Carlo method for fermion ground states. *Phys. Rev. B* **55**, 7464–7477 (1997).
25. H. Shi, S. Zhang, Symmetry in auxiliary-field quantum Monte Carlo calculations. *Phys. Rev. B* **88**, 125132 (2013).
26. M. Troyer, U. J. Wiese, Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations. *Phys. Rev. Lett.* **94**, 170201 (2005).
27. D. Ceperley, G. V. Chester, M. H. Kalos, Monte Carlo simulation of a many-fermion study. *Phys. Rev. B* **16**, 3081–3099 (1977).
28. H. Yokoyama, H. Shiba, Variational Monte-carlo studies of Hubbard model. I. *J. Phys. Soc. Jpn.* **56**, 1490–1506 (1987a).
29. H. Yokoyama, H. Shiba, Variational Monte-Carlo studies of Hubbard model. II. *J. Phys. Soc. Jpn.* **56**, 3582–3592 (1987b).
30. H. Yokoyama, H. Shiba, Variational Monte-Carlo studies of Hubbard model. III. Intersite correlation effects. *J. Phys. Soc. Jpn.* **59**, 3669–3686 (1990).
31. C. Gros, Superconductivity in correlated wave functions. *Phys. Rev. B* **38**, 931–934 (1988).
32. C. Gros, Physics of projected wavefunctions. *Ann. Phys.* **189**, 53–88 (1989).
33. A. Paramekanti, M. Randeria, N. Trivedi, Projected wave functions and high temperature superconductivity. *Phys. Rev. Lett.* **87**, 217002 (2001).
34. S. Sorella *et al.*, Superconductivity in the two-dimensional $t - J$ model. *Phys. Rev. Lett.* **88**, 117002 (2002).
35. A. Paramekanti, M. Randeria, N. Trivedi, High- T_c superconductors: A variational theory of the superconducting state. *Phys. Rev. B* **70**, 054504 (2004).
36. L. F. Tocchio, F. Becca, A. Parola, S. Sorella, Role of backflow correlations for the nonmagnetic phase of the $t - t'$ Hubbard model. *Phys. Rev. B* **78**, 041101 (2008).
37. D. Tahara, M. Imada, Variational Monte Carlo method combined with quantum-number projection and multi-variable optimization. *J. Phys. Soc. Jpn.* **77**, 114701 (2008).
38. T. Misawa *et al.*, mvmc-open-source software for many-variable variational Monte Carlo method. *Comput. Phys. Commun.* **235**, 447–462 (2019).
39. G. Carleo, M. Troyer, Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).
40. Y. Nomura, M. Imada, Dirac-type nodal spin liquid revealed by refined quantum many-body solver using neural-network wave function, correlation ratio, and level spectroscopy. *Phys. Rev. X* **11**, 031034 (2021).
41. C. Roth, A. Szabó, A. H. MacDonald, High-accuracy variational Monte Carlo for frustrated magnets with deep neural networks. *Phys. Rev. B* **108**, 054410 (2023).
42. A. Chen, M. Heyl, Empowering deep neural quantum states through efficient optimization. *Nat. Phys.* **20**, 1476–1481 (2024).
43. L. L. Viteritti, R. Rende, A. Parola, S. Goldt, F. Becca, Transformer wave function for two dimensional frustrated magnets: Emergence of a spin-liquid phase in the Shastry-Sutherland model. *Phys. Rev. B* **111**, 134411 (2025).
44. K. Choo, A. Mezzacapo, G. Carleo, Fermionic neural-network states for ab-initio electronic structure. *Nat. Commun.* **11**, 2368 (2020).
45. D. Pfau, J. S. Spencer, A. G. D. G. Matthews, W. M. C. Foulkes, Ab initio solution of the many-electron Schrödinger equation with deep neural networks. *Phys. Rev. Res.* **2**, 033429 (2020).
46. J. Hermann, Z. Schätzle, F. Noé, Deep-neural-network solution of the electronic Schrödinger equation. *Nat. Chem.* **12**, 891–897 (2020).
47. I. von Glehn, J. S. Spencer, D. Pfau, "A self-attention ansatz for ab-initio quantum chemistry" in *The Eleventh International Conference on Learning Representations (ICLR)* (OpenReview.net, Kigali, Rwanda, 2023). <https://openreview.net/forum?id=xvTeHVIF7j>.
48. D. Pfau, S. Axelrod, H. Sutterud, I. von Glehn, J. S. Spencer, Accurate computation of quantum excited states with neural networks. *Science* **385**, eadn0137 (2024).
49. G. Pescia, J. Nys, J. Kim, A. Lovato, G. Carleo, Message-passing neural quantum states for the homogeneous electron gas. *Phys. Rev. B* **110**, 035108 (2024).
50. C. Smith *et al.*, Unified variational approach description of ground-state phases of the two-dimensional electron gas. *Phys. Rev. Lett.* **133**, 266504 (2024).
51. Y. Nomura, A. S. Darmawan, Y. Yamaji, M. Imada, Restricted Boltzmann machine learning for solving strongly correlated quantum systems. *Phys. Rev. B* **96**, 205152 (2017).
52. D. Luo, B. K. Clark, Backflow transformations via neural networks for quantum many-body wave functions. *Phys. Rev. Lett.* **122**, 226401 (2019).
53. J. R. Moreno, G. Carleo, A. Georges, J. Stokes, Fermionic wave functions from neural-network constrained hidden states. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2122059119 (2022).
54. J. Bardeen, L. N. Cooper, J. R. Schrieffer, Microscopic theory of superconductivity. *Phys. Rev.* **106**, 162–164 (1957).
55. D. Thouless, Vibrational states of nuclei in the random phase approximation. *Nucl. Phys.* **22**, 78–95 (1961).
56. A. J. Coleman, Structure of fermion density matrices. II. Antisymmetrized geminal powers. *J. Math. Phys.* **6**, 1425–1431 (1965).
57. J. P. Bouchaud, A. Georges, C. Lhuillier, Pair wave functions for strongly correlated fermions and their determinantal representation. *J. Phys. Fr.* **49**, 553–559 (1988).
58. P. R. Surján, *An Introduction to the Theory of Geminals*, P. R. Surján *et al.*, Eds. (Springer, Berlin Heidelberg, Berlin, Heidelberg, 1999), pp. 63–88.
59. G. Moore, N. Read, Nonabelions in the fractional quantum hall effect. *Nucl. Phys. B* **360**, 362–396 (1991).
60. M. Storni, R. H. Morf, S. Das Sarma, Fractional quantum hall state at $\nu = \frac{5}{2}$ and the Moore-read Pfaffian. *Phys. Rev. Lett.* **104**, 076803 (2010).
61. R. G. Xu, T. Okubo, S. Todo, M. Imada, Optimized implementation for calculation and fast-update of Pfaffians installed to the open-source fermionic variational solver mvmc. *Comput. Phys. Commun.* **277**, 108375 (2022).
62. A. Chen, C. Roth, IruX: Fast low-rank updates of determinants and Pfaffians in JAX. *arXiv [Preprint]* (2026). <https://doi.org/10.48550/arXiv.2602.05255> (Accessed 5 February 2026).
63. B. Fore *et al.*, Dilute neutron star matter from neural-network quantum states. *Phys. Rev. Res.* **5**, 033062 (2023).
64. J. Kim *et al.*, Neural-network quantum states for ultra-cold fermi gases. *Commun. Phys.* **7**, 148 (2024).
65. N. Gao, S. Günnemann, "Neural Pfaffians: Solving many many-electron Schrödinger equations" in *Advances in Neural Information Processing Systems*, A. Globerson, *et al.*, (Curran Associates, Inc., 2024), vol. 37, pp. 125336–125369.
66. B. Fore, J. Kim, M. Hjorth-Jensen, A. Lovato, Investigating the crust of neutron stars with neural-network quantum states. *Commun. Phys.* **8**, 108 (2025).
67. K. Choo, T. Neupert, G. Carleo, Two-dimensional frustrated $J_1 - J_2$ model studied with neural network quantum states. *Phys. Rev. B* **100**, 125124 (2019).
68. X. Liang *et al.*, Deep learning representations for quantum many-body systems on heterogeneous hardware. *Mach. Learn. Sci. Technol.* **4**, 015035 (2023).
69. L. L. Viteritti, R. Rende, F. Becca, Transformer variational wave functions for frustrated quantum spin systems. *Phys. Rev. Lett.* **130**, 236401 (2023).
70. A. Chen, V. D. Naik, M. Heyl, Convolutional transformer wave functions. *Phys. Rev. Res.* **8**, L022040 (2026).
71. B. X. Zheng *et al.*, Stripe order in the underdoped region of the two-dimensional Hubbard model. *Science* **358**, 1155–1160 (2017).
72. S. Sorella, Green function Monte Carlo with stochastic reconfiguration. *Phys. Rev. Lett.* **80**, 4558–4561 (1998).
73. S. Sorella, M. Casula, D. Rocca, Weak binding between two aromatic rings: Feeling the van der Waals attraction by quantum Monte Carlo methods. *J. Chem. Phys.* **127**, 014105 (2007).
74. G. Mazzola, A. Zen, S. Sorella, Finite-temperature electronic simulations without the Born-Oppenheimer constraint. *J. Chem. Phys.* **137**, 134112 (2012).
75. G. Goldshlager, N. Abrahamsen, L. Lin, A kaczmarz-inspired approach to accelerate the optimization of neural network wavefunctions. *J. Comput. Phys.* **516**, 113351 (2024).
76. I. Romero, J. Nys, G. Carleo, Spectroscopy of two-dimensional interacting lattice electrons using symmetry-aware neural backflow transformations. *Commun. Phys.* **8**, 46 (2025).
77. K. Loehr, B. K. Clark, Enhancing neural network backflow. *arXiv [Preprint]* (2025). <https://doi.org/10.48550/arXiv.2510.26906> (Accessed 30 October 2025).
78. H. Lange *et al.*, Simulating superconductivity in mixed-dimensional $t||-J||-J_{\perp}$ bilayers with neural quantum states. *arXiv [Preprint]* (2026). <https://doi.org/10.48550/arXiv.2602.10091> (Accessed 10 February 2026).
79. Q. Chen, J. Stajic, S. Tan, K. Levin, BCS-BEC crossover: From high temperature superconductors to ultracold superfluids. *Phys. Rep.* **412**, 1–88 (2005).
80. Y. Gu *et al.*, Solving the Hubbard model with neural quantum states. *Nat. Commun.* **17**, 7838 (2026).
81. W. Y. Liu, S. J. Dong, Y. J. Han, G. C. Guo, L. He, Gradient optimization of finite projected entangled pair states. *Phys. Rev. B* **95**, 195154 (2017).
82. W. Y. Liu, Y. Z. Huang, S. S. Gong, Z. C. Gu, Accurate simulation for finite projected entangled pair states in two dimensions. *Phys. Rev. B* **103**, 235155 (2021).
83. J. P. F. LeBlanc *et al.*, Solutions of the two-dimensional Hubbard model: Benchmarks and results from a wide range of numerical algorithms. *Phys. Rev. X* **5**, 041041 (2015).
84. M. Qin *et al.*, Absence of superconductivity in the pure two-dimensional Hubbard model. *Phys. Rev. X* **10**, 031016 (2020).

85. H. Xu *et al.*, Coexistence of superconductivity with partially filled stripes in the Hubbard model. *Science* **384**, eadh7691 (2024).
86. A. Wietek, Fragmented cooper pair condensation in striped superconductors. *Phys. Rev. Lett.* **129**, 177001 (2022).
87. S. Sorella, Systematically improvable mean-field variational ansatz for strongly correlated systems: Application to the Hubbard model. *Phys. Rev. B* **107**, 115133 (2023).
88. Y. T. Zhou, Z. W. Zhou, X. Liang, Solving fermi-Hubbard-type models by tensor representations of backflow corrections. *Phys. Rev. B* **109**, 245107 (2024).
89. R. Rende, L. L. Viteritti, L. Bardone, F. Becca, S. Goldt, A simple linear algebra identity to optimize large-scale neural network quantum states. *Commun. Phys.* **7**, 260 (2024).
90. C. Roth, A. Chen, A. Sengupta, A. Georges, Superconductivity in the two-dimensional Hubbard model revealed by neural quantum states. *arXiv [Preprint]* (2025). <https://doi.org/10.48550/arXiv.2511.07566> (Accessed 10 November 2025).
91. P. Weinberg, M. Bukov, QuSpin: A Python package for dynamics and exact diagonalisation of quantum many body systems part I: Spin chains. *SciPost Phys.* **2**, 003 (2017).
92. P. Weinberg, M. Bukov, QuSpin: A Python package for dynamics and exact diagonalisation of quantum many body systems. Part II: Bosons, fermions and higher spins. *SciPost Phys.* **7**, 020 (2019).
93. A. Chen, C. Roth, Quantax: Flexible neural quantum states in JAX. GitHub. <https://github.com/ChenAo-Phys/quantax>. Accessed 26 June 2026.
94. J. Bradbury *et al.*, JAX: Composable transformations of Python+NumPy programs. GitHub. <https://github.com/google/jax>. Accessed 10 August 2026.
95. P. Kidger, C. Garcia, Equinox: Neural networks in Jax via callable Pytrees and filtered transformations. *arXiv [Preprint]* (2021). <https://doi.org/10.48550/arXiv.2111.00254> (Accessed 30 October 2021).
96. A. Chen, C. Roth, Quantax custom branch for HFPS (2025). GitHub. <https://github.com/chrisrothUT/quantax/tree/custom>. Deposited 4 February 2026.