

Self-attention neural network for solving correlated electron problems in solids

Max Geier¹,* Khachatur Nazaryan¹,* Timothy Zaklama¹, and Liang Fu¹

Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA



(Received 24 March 2025; revised 29 May 2025; accepted 16 June 2025; published 14 July 2025)

The attention mechanism has transformed artificial intelligence research by its ability to learn relations between objects. In this work, we explore how a many-body wave-function *Ansatz* constructed from a large-parameter self-attention neural network can be used to solve the interacting electron problem in solids. By a systematic neural-network variational Monte Carlo study on a moiré quantum material, we demonstrate that the self-attention *Ansatz* provides an accurate and efficient solution without human bias. Moreover, our numerical study finds that the required number of variational parameters scales roughly as N^2 with the number of electrons, which opens a path towards efficient large-scale simulations.

DOI: [10.1103/qxc3-bkc7](https://doi.org/10.1103/qxc3-bkc7)

I. INTRODUCTION

Solving the many-electron Schrödinger equation for solids is an exceedingly difficult problem due to the exponential growth of the Hilbert space dimension. Various techniques based on the variational principle have long been developed to approximate the ground state of interacting electron systems using trial wave functions. The Hartree-Fock method based on independent electron approximation [1,2] usually captures 99% of the total energy [3], but misses electron correlation effects, the driving force behind fascinating quantum phenomena such as high-temperature superconductivity and the fractional quantum Hall effect.

Recently, neural network (NN) based variational Monte Carlo (VMC) has been developed for solving the many-electron problem with high accuracy [4–12]. Compared to human-designed trial wave functions, neural network wave functions contain a large number of parameters, have enormous representation power, and can be optimized efficiently. NN-VMC has been shown to be highly accurate in calculating the ground-state energy of interacting electrons in atoms and molecules [9–11,13,14], lattice models [6,12,15–17], uniform electron gas [18–24], moiré semiconductors [25,26], and fractional quantum Hall liquids [27,28].

Despite the rapid progress, two important questions remain open. First, a number of NN architectures have so far been introduced and used to study different many-electron problems. Is there any hope of finding a unifying architecture that applies to a wide range of interacting electron systems? Second, it is essential to assess the finite-size effect in numerical simulations of solid-state systems. How does the performance of the neural *Ansatz* change as the system size increases?

In this work, we present a general NN-VMC method to solve many-electron problems in solids, where electron correlations in the NN *Ansatz* are *entirely* produced from the self-attention mechanism. The attention mechanism was

originally introduced in the context of large language models to learn relations between words [29]. Here, the attention mechanism is employed to identify and quantify how electrons influence each other and how such influence affects their individual orbitals. This enables the construction of NN wave functions from Slater determinants of generalized orbitals that depend on the configuration of all electrons [11]. By minimizing the energy using Monte Carlo and neural network techniques [3], variational energy and wave function for the many-electron ground state are obtained.

The performance of our attention based NN *Ansatz* is evaluated for interacting electrons in semiconductor moiré heterobilayers, such as WSe₂/WS₂. This moiré platform hosts a fascinating variety of correlated electronic states, including Mott insulators [30–32], generalized Wigner crystals [33], and strongly correlated Fermi liquids [34]. Here, doped electrons reside on one semiconductor layer and experience a moiré potential. The Hamiltonian thus takes the form of a two-dimensional Coulomb electron gas in a periodic potential [35]:

$$H = H_0 + H_{ee} = \sum_i \left(-\frac{1}{2} \nabla_i^2 + V(\mathbf{r}_i) \right) + \frac{1}{2} \sum_i \sum_{i \neq j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}, \quad (1)$$

where $V(\mathbf{r}) = -2V_0 \sum_{j=1}^3 \cos(\mathbf{g}_j \cdot \mathbf{r} + \varphi)$ is the moiré potential with reciprocal lattice vectors $\mathbf{g}_j = \frac{4\pi}{\sqrt{3}a_M}(\cos \frac{2\pi j}{3}, \sin \frac{2\pi j}{3})$, moiré lattice constant a_M , and φ controls the shape of the moiré potential. Despite its conceptual simplicity, this Hamiltonian exhibits a variety of electron phases that emerge from the interplay between kinetic energy, moiré potential, and Coulomb interaction [36,37].

First, for small system size, we benchmark our NN results with band-projected exact diagonalization. Remarkably, the NN energies are found to be lower even when five bands are included in the exact diagonalization. Next, we assess the performance of the NN wave function as the system size increases. Specifically, we study how the required number of

*These authors contributed equally to this work.

parameters N_{par} scales with the number of electrons N , and numerically discover a scaling law $N_{\text{par}} \propto N^\alpha$ with a small exponent $\alpha \approx 2$. These findings suggests that self-attention NN *Ansatz* is an accurate and efficient method in solving large-scale interacting electron problems.

We note that self-attention NN wave functions were first introduced in Ref. [11] to solve quantum chemistry problems, reaching state-of-the-art accuracy. The recent work Ref. [27] demonstrated that the self-attention *Ansatz* accurately describes the fractional quantum Hall ground state, achieving a lower ground-state energy than Landau-level projected exact diagonalization. The remarkable success of the self-attention NN wave function in atoms, molecules, electron gas, and now moiré materials, taken as a whole, suggests that self-attention is a key ingredient for a unifying solution to the correlated electron problem.

In the remainder of the article, we detail the self-attention neural network wave-function *Ansatz* and present its benchmark with self-consistent Hartree-Fock and band-projected exact diagonalization. First, in Sec. II we summarize the traditional methods of Hartree-Fock and band-projected exact diagonalization for comparison. Then, Sec. III describes the construction of the self-attention neural network *Ansatz*. Section IV summarizes principal ideas of the variational Monte Carlo method applied to optimize the wave-function *Ansatz*. Finally, our numerical results and benchmarks are presented in Sec. V.

II. TRADITIONAL NUMERICAL METHODS

In this section, we present Hartree-Fock and band-projected exact diagonalization solutions of the interacting Hamiltonian for transition metal dichalcogenides (TMD) heterobilayers. Results from these standard numerical methods will provide the benchmark for NN-VMC in the following sections.

Throughout this work we study periodic solids. Numerical simulation of a periodic solid requires to divide the space into periodic supercells of finite size and require periodic boundary conditions on the wave function. With periodic boundary conditions, the Coulomb interaction needs to be carefully calculated to account for interactions of particles within the supercell with all their images in the other supercells (see Appendix A for details).

A. Hartree-Fock

The Hartree-Fock approximation to the ground state of an interacting fermionic system is obtained by minimizing the energy over the space of single Slater determinant states:

$$\Psi^{\text{HF}}(\mathbf{r}_1, \dots, \mathbf{r}_N) = \frac{1}{\sqrt{N!}} \det \begin{pmatrix} \phi_1(\mathbf{r}_1) & \dots & \phi_N(\mathbf{r}_1) \\ \vdots & \ddots & \vdots \\ \phi_1(\mathbf{r}_N) & \dots & \phi_N(\mathbf{r}_N) \end{pmatrix}. \quad (2)$$

This wave-function *Ansatz* captures the quantum-mechanical “exchange” effect by incorporating the Pauli exclusion principle, but neglects correlations arising from electron-electron interactions. Single Slater determinant states have a simple structure that enables efficient computation of observables

using Wick’s theorem. At the same time, this structure enables a self-consistent determination of the optimal one-body orbitals $\phi_j(\mathbf{r})$ underlying the Slater determinant wave function.

By construction, the energy of the Hartree-Fock ground state is the energy of the best wave function constructed from independent orbitals. This motivates the definition of *correlation energy* as the energy difference between the true ground state and the Hartree-Fock ground state:

$$E_{\text{corr}} = E_{\text{GS}} - \langle \Psi^{\text{HF}} | \hat{H} | \Psi^{\text{HF}} \rangle. \quad (3)$$

In molecular systems and most solid-state systems, the correlation energy typically is only about 1% of the ground-state energy [3]. However, it plays an essential role in driving various quantum phases of matter, ranging from superconductivity to fractional quantum Hall states.

B. Band-projected exact diagonalization

One approach to capture electron-electron correlations is to construct a finite set of Slater determinant wave functions and exactly diagonalize the full Hamiltonian projected onto this finite subspace of the Hilbert space. This approach is variational because the restricted choice of Slater determinants forms a variational wave-function *Ansatz*. The obtained ground-state energy is an upper bound on the true ground-state energy because inclusion of additional Slater determinants can further lower the energy.

A good choice of Slater determinants spanning the variational subspace is obtained by diagonalizing the noninteracting part of the Hamiltonian H_0 , i.e., the first term in Eq. (1), and keeping only the lowest-energy eigenstates up to a cutoff. For a periodic system, the single-particle eigenstates are Bloch states labeled by the Bloch momentum \mathbf{k} and band index n . Projecting onto the subspace of lowest N_{bands} bands, the Hamiltonian matrix elements between Slater determinants are written in second quantization

$$\begin{aligned} \tilde{\mathbf{H}} = & \sum_{n=1}^{N_{\text{bands}}} \sum_{\mathbf{k}, \sigma} \varepsilon_{\mathbf{k}\sigma}^n c_{\mathbf{k}\sigma n}^\dagger c_{\mathbf{k}\sigma n} + \frac{1}{2} \sum_i \xi_{\text{M}} \\ & + \frac{1}{2} \sum_{\substack{n, m, \\ n', m'=1}}^{N_{\text{bands}}} \sum_{\substack{\mathbf{k}' \mathbf{p}' \mathbf{k} \mathbf{p} \\ \sigma \sigma'}} V_{\mathbf{k}' \mathbf{p}' \mathbf{k} \mathbf{p}; \sigma \sigma'}^{nmn'm'} c_{\mathbf{k}' \sigma' n'}^\dagger c_{\mathbf{k} \sigma n}^\dagger c_{\mathbf{p}' \sigma' m} c_{\mathbf{p} \sigma m}, \quad (4) \end{aligned}$$

where $c_{\mathbf{k}\sigma n}^\dagger$ creates a Bloch state in the n th band at crystal momentum \mathbf{k} and spin σ with corresponding single-particle energy $\varepsilon_{\mathbf{k}\sigma}^n$, and $V_{\mathbf{k}' \mathbf{p}' \mathbf{k} \mathbf{p}; \sigma \sigma'}^{nmn'm'} \equiv \langle \mathbf{k}', \sigma', n; \mathbf{p}', \sigma', m | \hat{V} | \mathbf{k}, \sigma, n; \mathbf{p}, \sigma, m \rangle$ are the corresponding matrix elements of the Coulomb interaction projected onto the band basis. Equation (4) contains the Madelung energy $N\xi_{\text{M}}/2$ that describes the interaction of charges in the periodic supercell with their own images in other supercells (see Appendix A for details).

For a small size system, the band-projected Hamiltonian $\tilde{\mathbf{H}}$ can be numerically diagonalized to yield the ground-state energy and wave function. The band-projected exact diagonalization (BP-ED) method is accurate only if the interaction strength is small compared to the gap to excitations involving high-energy bands. However, as we shall show below, interaction-induced band mixing is substantial for realistic

material parameters, and therefore it is necessary to include *multiple* low-lying bands to obtain quantitatively accurate results, which limits the practical application of BP-ED to very small system size.

III. NEURAL NETWORK WAVE FUNCTION

While Hartree-Fock and BP-ED studies have provided useful results on moiré semiconductors, they suffer from respective limitations. The Hartree-Fock method fails to capture electron correlation effect and cannot describe entangled states such as fractional Chern insulators, while BP-ED is severely limited by finite-size effect and band truncation.

Recently, a new type of variational method based on neural network wave functions has been developed to approximate the ground state of many-electron systems. Traditional variational wave functions for correlated systems are artfully designed by humans, tailor made for a given quantum phase, and usually contain only a few variational parameters. In contrast, thanks to their universal approximation capability, deep neural networks can be used to generate powerful wave-function *Ansatz*, which has a large number of variational parameters and can accurately represent distinct phases of matter in a unified way [10,38,39].

In the following, we show how an expressive neural network wave-function *Ansatz* capturing correlations between particles can be constructed. First, in Sec. III A, a universal approximation of single-particle orbitals is constructed by a deep feed-forward neural network. Second, in Sec. III B, by generating a set of single-particle orbitals from the feed-forward neural network, a single Slater determinant wave function is obtained, which enables neural network implementation of unrestricted Hartree-Fock calculations. Finally, in Sec. III C, electron correlations are captured by mixing different electron streams in a permutation-equivariant way with self-attention mechanism. The resulting neural network wave function takes the form of generalized Slater determinants that are built from “correlated orbitals,” and by construction is antisymmetric in electron coordinates [10]. This representation is in principle universal as all many-body fermionic wave functions can be written as a generalized Slater determinant when allowing for discontinuous functions [10]. Figure 1 shows schematically and compares the network architectures representing general single Slater determinant wave functions, and its extension by self-attention capturing correlations.

A. Deep neural network for one-body orbitals

A single particle orbital describes the motion of a single electron in an external potential. Mathematically, it is represented by a function $\phi(\mathbf{r})$ from real space coordinates \mathbb{R}^d to the space of complex numbers \mathbb{C} describing amplitude and phase of the electronic wave at the given position $\mathbf{r} \in \mathbb{R}^d$. Due to the universal approximation theorem for deep feed forward neural networks [38], any such function can be approximated to arbitrary accuracy when sufficiently many hidden layers are included. Therefore, deep feed-forward neural networks are the main structural element that we employ to construct general single-particle orbitals.

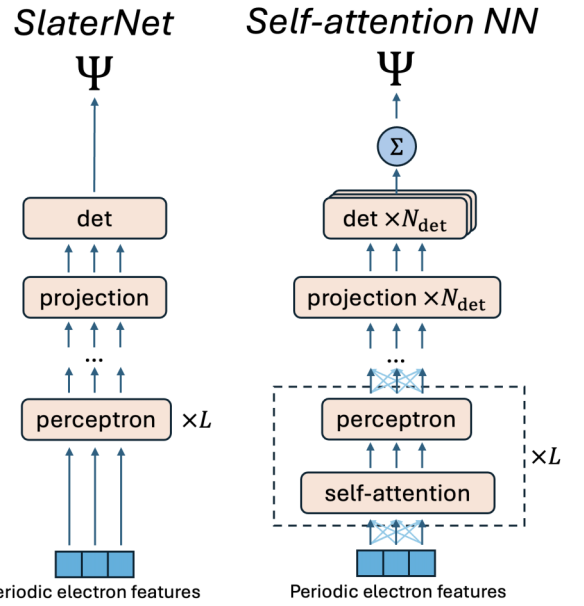


FIG. 1. Architecture of the neural network wave-function *Ansatz*. *SlaterNet*: Multilayer perceptron neural network generates one-body orbitals to approximate general single Slater determinant wave functions. *Self-attention NN*: Self-attention neural network for solids based on Psiformer [11] capturing correlations by communicating mutual information between individual particle streams.

The structure of the compositional wave-function *Ansatz* consists of three sections with distinct functionality. First, the function input, i.e., the electron coordinate \mathbf{r} in coordinate space in our case, is transformed into a representative *feature*

$$\mathbf{f}^0 = \text{feature}(\mathbf{r}). \quad (5)$$

For systems with open boundary conditions, it is sufficient for the function $\text{feature}(\mathbf{r})$ to represent \mathbf{r} in dimensionless units. Systems with periodic boundary conditions satisfy

$$\Psi(\mathbf{r}_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_N) = \Psi(\mathbf{r}_1, \dots, \mathbf{r}_i + \mathbf{L}, \dots, \mathbf{r}_N) \quad (6)$$

for any particle i , where the two vectors $\mathbf{L} = n\mathbf{L}_1 + m\mathbf{L}_2$ with $n, m \in \mathbb{Z}$ specify the supercell size and geometry. For periodic systems, we choose the function $\text{feature}(\mathbf{r})$ to express the coordinate \mathbf{r} in terms of periodic coordinates [18,40]

$$\text{feature}(\mathbf{r}) = \begin{pmatrix} \sin(\mathbf{G}_1^T \mathbf{r}) \\ \sin(\mathbf{G}_2^T \mathbf{r}) \\ \cos(\mathbf{G}_1^T \mathbf{r}) \\ \cos(\mathbf{G}_2^T \mathbf{r}) \end{pmatrix}, \quad (7)$$

where $\mathbf{G}_i^T \mathbf{L}_j = 2\pi\delta_{ij}$ are the reciprocal supercell vectors and δ_{ij} is the Kronecker delta. This periodic feature representation uniquely specifies the electron coordinate on the torus and ensures that proximate positions on the torus are passed as proximate features to the network, enabling efficient representation of general periodic functions by the neural network. This modification adapts the quantum chemistry implementation Psiformer of Ref. [11] to periodic solids.

Second, the featured coordinate is embedded in the internal representation of the neural network by a linear

transformation

$$\mathbf{h}^0 = W^0 \mathbf{f}^0, \quad (8)$$

where $W_0 \in \mathbb{R}^{d_L} \times \mathbb{R}^{2d_{\text{dim}}}$ is a rectangular matrix with d_L the width of the internal network layers. This embedded feature is passed through the multilayer perceptron neural network with functional form

$$\mathbf{h}^{l+1} = \mathbf{h}^l + \tanh(W^{l+1} \mathbf{h}^l + \mathbf{b}^{l+1}), \quad (9)$$

where $W^{l+1} \in \mathbb{R}^{d_L} \times \mathbb{R}^{d_L}$ is the linear transformation between layers, $\mathbf{b}^{l+1} \in \mathbb{R}^{d_L}$ is an activation bias vector, and \tanh is the nonlinear activation function of the nodes, and $l = 0, 1, \dots, L-1$ enumerates the layers.

Finally, the wave function is obtained by projecting the output of the neural network \mathbf{h}_L onto real and imaginary parts of a complex number,

$$\phi(\mathbf{r}) = \mathbf{w}_0 \cdot \mathbf{h}^L + i \mathbf{w}_1 \cdot \mathbf{h}^L, \quad (10)$$

where $\mathbf{w}_0, \mathbf{w}_1 \in \mathbb{R}^{d_L}$ are learnable projection vectors and the center dot denotes the scalar product.

B. SlaterNet: Neural network for unrestricted Hartree-Fock

The ability to construct general single-particle wave functions directly enables the construction of general Slater determinant wave functions for N electrons. This is achieved by passing each electron coordinate \mathbf{r}_i through the feed-forward neural network,

$$\mathbf{f}_i^0 = \text{feature}(\mathbf{r}_i), \quad (11)$$

$$\mathbf{h}_i^0 = W^0 \mathbf{f}_i^0, \quad (12)$$

$$\mathbf{h}_i^{l+1} = \mathbf{h}_i^l + \tanh(W^{l+1} \mathbf{h}_i^l + \mathbf{b}^{l+1}), \quad (13)$$

where the same transformations W^l and shifts \mathbf{b}^l are applied to the *single-particle stream* \mathbf{h}_i^l generated by the coordinates of each particle. Thereby, this compositional function generates a single function in a high-dimensional space before projecting onto the space of N single-particle orbitals \mathbb{C}^N . Then, N distinct single-particle orbitals are constructed as

$$\phi_j(\mathbf{r}_i) = \mathbf{w}_{2j} \cdot \mathbf{h}_i^L + i \mathbf{w}_{2j+1} \cdot \mathbf{h}_i^L \quad (14)$$

with individually learnable projection vectors $\mathbf{w}_{2j}, \mathbf{w}_{2j+1}$. The projection vectors are not required to satisfy an orthogonality criterion because nonorthogonal directions are projected out upon forming the Slater determinant as in Eq. (2).

Because all single-electron wave functions are universally approximated by the deep neural network, this *Ansatz* universally approximates all single Slater determinant wave functions. By minimizing the energy over this variational space, the best set of single-particle orbitals is identified. This is equivalent to an unrestricted Hartree-Fock calculation. We refer to this neural network structure as “SlaterNet” throughout the paper.

C. Self-attention neural network for electron correlations

Correlations due to interactions among particles can be described by specifying how the state of individual particles i is modified by interactions with all remaining particles

$j \neq i$. To capture electron correlation effects, it is necessary to go beyond Hartree-Fock *Ansatz* as described earlier. One approach is to promote the single-particle orbitals $\phi_j(\mathbf{r}_i)$ to the *correlated orbitals* $\phi_j(\mathbf{r}_i; \{\mathbf{r}_{j/i}\})$ with explicit dependence on the other particles' position $\{\mathbf{r}_{j/i}\}$. The idea of correlated orbitals dates back to the backflow transformation used in variational study of superfluid helium and uniform electron gas [41,42], and has recently been applied to neural network wave functions [7,10,11,43].

In this work, inspired by [11], we use the *self-attention* mechanism [29] to learn how particles influence each other and how such influence affects their individual orbitals. Specifically, the self-attention SELFATTN_{*i*} operates on the entire set of outputs $\{\mathbf{h}_i^l\}$ from the previous layer and generates the set of intermediate states $\{\mathbf{f}_i^l\}$ that is passed as input to the next perceptron layer that generates the next state $\{\mathbf{h}_i^{l+1}\}$:

$$\dots \rightarrow \{\mathbf{h}_i^l\} \xrightarrow{\text{SELFATTN}} \{\mathbf{f}_i^l\} \xrightarrow{\text{PERCEPTRON}} \{\mathbf{h}_i^{l+1}\} \rightarrow \dots$$

The following explains the self-attention mechanism and the mathematical expression of the compositional function is summarized in Eqs. (19) and (20) below.

The first step of the self-attention operation is to define three distinct features for each element of the set $\{\mathbf{h}_i^l\}$. The three features, called “keys,” “queries,” and “values,” are vectors obtained by the linear transformations

$$\mathbf{k}_i^{lh} = W_k^{lh} \mathbf{h}_i^l, \quad W_k^{lh} : \mathbb{R}^{d_L} \rightarrow \mathbb{R}^{d_{\text{Attn}}}, \quad (15)$$

$$\mathbf{q}_i^{lh} = W_q^{lh} \mathbf{h}_i^l, \quad W_q^{lh} : \mathbb{R}^{d_L} \rightarrow \mathbb{R}^{d_{\text{Attn}}}, \quad (16)$$

$$\mathbf{v}_i^{lh} = W_v^{lh} \mathbf{h}_i^l, \quad W_v^{lh} : \mathbb{R}^{d_L} \rightarrow \mathbb{R}^{d_{\text{AttnValues}}}, \quad (17)$$

where the transformations W_k^{lh} , W_q^{lh} , and W_v^{lh} are learned and are independent of the particle index i . The vector space dimensions d_{Attn} and $d_{\text{AttnValues}}$ are typically much smaller than the dimension d_L of \mathbf{h}_i^l because they only represent individual features of the state \mathbf{h}_i^l . Here, we have used multihead attention mechanism, with the index h labeling independent applications of distinct transformations $W_k^{lh}, W_q^{lh}, W_v^{lh}$ to extract multiple sets of keys, queries, and values.

The features “key” \mathbf{k}_i^{lh} and “query” \mathbf{q}_j^{lh} are elements of the same vector space $\mathbb{R}^{d_{\text{Attn}}}$ to allow direct comparison between different electron streams i and j . Specifically, in the attention mechanism the relevance of streams j for a selected “key” streams i is quantified by the similarity measure $\exp(\mathbf{k}_i^{lh} \cdot \mathbf{q}_j^{lh})$. By weighting the scalar product $\mathbf{k}_i^{lh} \cdot \mathbf{q}_j^{lh}$ with the exponential function, the most relevant streams j for each “key” stream i are singled out. This procedure can be interpreted as approximating the adjacency matrix of a graph describing the most relevant relations [44]. The “value” feature \mathbf{v}_j^{lh} quantifies the influence the stream j can exert on other streams $i \neq j$.

Altogether, the self-attention operation identifies for each key stream i the most relevant stream j according to the feature representations W_k^{lh}, W_q^{lh} and returns the value \mathbf{v}_j^{lh} corresponding to the most relevant stream j , up to exponentially suppressed contributions from less relevant

streams:

$$\begin{aligned} \text{SELFATTN}_i(\{\mathbf{h}_j^l\}; W_k^{lh}, W_q^{lh}, W_v^{lh}) \\ = \frac{1}{\mathcal{N}} \sum_{j=1}^N \exp(\mathbf{q}_j^{lh} \cdot \mathbf{k}_i^{lh}) \mathbf{v}_j^{lh}, \end{aligned} \quad (18)$$

where a normalization factor is included:

$$\mathcal{N} = \sqrt{d_{\text{AttnValues}}} \sum_{j=1}^N \exp(\mathbf{q}_j^{lh} \cdot \mathbf{k}_i^{lh}).$$

The output of all attention heads h is then accumulated and used to generate the input for the next perceptron layer:

$$\begin{aligned} \mathbf{f}_i^{l+1} &= \mathbf{h}_i^l + W_o^l \text{concat}_h \\ &\times [\text{SELFATTN}_i(\{\mathbf{h}_j^l\}; W_k^{lh}, W_q^{lh}, W_v^{lh})] \end{aligned} \quad (19)$$

$$\mathbf{h}_i^{l+1} = \mathbf{f}_i^{l+1} + \tanh(W^{l+1} \mathbf{f}_i^{l+1} + \mathbf{b}^{l+1}), \quad (20)$$

where the accumulation is performed by the learned transformation $W_o^l \in \mathbb{R}^{d_L \times \mathbb{R}^{N_{\text{heads}} d_{\text{AttnValues}}}}$ projecting the output of the attention heads onto the dimension of \mathbf{h}_i^l . Equations (19) and (20) replace the compositional relation (13) in the construction of one-body orbitals as in SlaterNet. Optionally, multiple perceptron layers can be applied between the attention layers, which is realized by multiple subsequent applications of (20) with individually learnable parameters. By mixing different single-particle streams with the self-attention mechanism, correlations between the particle states are captured.

The generalized single-particle orbitals $\phi_j(\mathbf{r}_i, \{\mathbf{r}_{/i}\})$ constructed by the self-attention operation are permutation equivariant in the coordinates of the remaining electrons $\{\mathbf{r}_{/i}\}$:

$$\phi_j(\mathbf{r}_i; \{\mathbf{r}_l, \dots, \mathbf{r}_k, \dots\}) = \phi_j(\mathbf{r}_i; \{\dots \mathbf{r}_k, \dots, \mathbf{r}_l, \dots\}) \quad (21)$$

for any $k, l \neq i$. The permutation equivariance is necessary for the generalized Slater determinant to represent a fermionic wave function that is antisymmetric under permutation of particle coordinates. This property is inherent to the functional form of the self-attention operation (18).

To further enhance the expressive power of our neural network wave function, it is useful to construct multiple Slater determinants from the neural network output by projection onto $m = 1, \dots, N_{\text{det}}$ distinct sets of correlated orbitals for each determinant,

$$\phi_j^m(\mathbf{r}_i; \{\mathbf{r}_{/i}\}) = \mathbf{w}_{2j}^m \cdot \mathbf{h}_i^L + i \mathbf{w}_{2j+1}^m \cdot \mathbf{h}_i^L, \quad (22)$$

where the projection vectors $\mathbf{w}_{2j}^m, \mathbf{w}_{2j+1}^m$ are individually learnable.

In summary, the full wave-function *Ansatz* is of the form

$$\Psi(\mathbf{R}) = \sum_{m=1}^{N_{\text{det}}} \det_{ij} (\phi_j^m(\mathbf{r}_i; \{\mathbf{r}_{/i}\})), \quad (23)$$

where the determinant is taken over the argument matrix with indices i, j , correlated orbitals $\phi_j^m(\mathbf{r}_i; \{\mathbf{r}_{/i}\})$ are constructed from Eq. (22). In practice, we find using a few generalized Slater determinants usually performs better than using a single generalized Slater determinant.

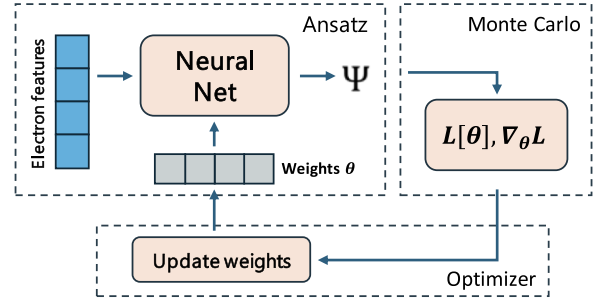


FIG. 2. Building blocks of variational Monte Carlo. In the Monte Carlo algorithm, the wave-function *Ansatz* Ψ_θ is constructed and sampled to efficiently evaluate the optimization goal $L[\theta] := \langle \Psi_\theta | \hat{H} | \Psi_\theta \rangle$ of minimizing the energy. Accordingly updating the weights of the variational *Ansatz* by an optimizer during each step of the iterative procedure ensures convergence to the ground state of \hat{H} .

IV. VARIATIONAL MONTE CARLO

Techniques based on a variational wave-function *Ansatz* require efficient numerical evaluation of their energy to achieve the optimization goal of approximating the ground state of the system Hamiltonian. This is achieved by Monte Carlo sampling of the particle configuration and corresponding energy according to the distribution determined by the variational wave function. This section summarizes the principal concepts of the variational Monte Carlo technique applicable to any variational wave-function *Ansatz* and refer to Ref. [3] for a detailed introduction. A schematic overview of the Monte Carlo method is presented in Fig. 2.

A. Sampling expectation values

The optimization goal to approximate the ground state of a correlated electron Hamiltonian with a variational wave-function *Ansatz* $\Psi_\theta(\mathbf{R})$ can be chosen to minimize the energy expectation value

$$E_\theta = \frac{\int d\mathbf{R} \Psi_\theta^*(\mathbf{R}) \hat{H}(\mathbf{R}) \Psi_\theta(\mathbf{R})}{\int d\mathbf{R} |\Psi_\theta(\mathbf{R})|^2}. \quad (24)$$

The numerically expensive integrals over configuration space $\mathbf{R} \in \mathbb{R}^{d \cdot N_{\text{el}}}$, where $\mathbf{R} = (\mathbf{r}_1, \dots, \mathbf{r}_N)$ are particle positions, can be efficiently evaluated by noticing the identity

$$E_\theta = \frac{\int d\mathbf{R} |\Psi_\theta(\mathbf{R})|^2 \Psi_\theta^{-1}(\mathbf{R}) \hat{H}(\mathbf{R}) \Psi_\theta(\mathbf{R})}{\int d\mathbf{R} |\Psi_\theta(\mathbf{R})|^2}. \quad (25)$$

This integral can be efficiently approximated by sampling configurations $\mathbf{R} = (\mathbf{r}_1, \dots, \mathbf{r}_N)$ according to the distribution $|\Psi_\theta(\mathbf{R})|^2$,

$$E_\theta \approx \mathbb{E}_{\mathbf{R} \sim |\Psi_\theta(\mathbf{R})|^2} [E_{\text{loc}, \theta}(\mathbf{R})], \quad (26)$$

where the *local energy* was introduced,

$$E_{\text{loc}, \theta} = \Psi_\theta^{-1}(\mathbf{R}) \hat{H}(\mathbf{R}) \Psi_\theta(\mathbf{R}), \quad (27)$$

and $\mathbb{E}_{\mathbf{R} \sim |\Psi_\theta(\mathbf{R})|^2} [E_{\text{loc}, \theta}] = \frac{1}{M} \sum_M E_{\text{loc}, \theta}$ denotes a summation of the local energy evaluated for M sampled configurations \mathbf{R} from the distribution $|\Psi_\theta(\mathbf{R})|^2$.

The variance of the local energy equals the energy uncertainty ΔE of the system Hamiltonian,

$$(\Delta E)^2 = \langle \hat{H}^2 \rangle - \langle \hat{H} \rangle^2 \approx \mathbb{E}_{\mathbf{R} \sim |\Psi_\theta(\mathbf{R})|^2} [E_{\text{loc}, \theta}^2(\mathbf{R})] - E_\theta^2, \quad (28)$$

where the equality is established by inserting $1 = \Psi_\theta(\mathbf{R})\Psi_\theta^{-1}(\mathbf{R})$ before and in-between the Hamiltonian operators when computing $\langle \hat{H}^2 \rangle$ using similar manipulations as the ones leading to Eq. (26). The energy uncertainty is proportional to the distance $|\langle \Psi_\theta | \Psi_0 \rangle|^2$ of the variational state Ψ_θ from an eigenstate Ψ_0 of the system Hamiltonian [45]. In an eigenstate, the energy uncertainty is zero.

The variational wave-function *Ansatz* Ψ_θ generally is unnormalized because computing the norm itself is numerically demanding. Unnormalized probability distributions can be sampled using the Metropolis-Hastings algorithm [46,47]; this technique is applied in the numerical results presented in this paper.

B. Optimizing variational parameters

To optimize the parameters of the wave function, the steepest descent direction is identified as the direction that minimizes the cost function $L[\theta + d\theta]$ for a step $d\theta$, where the distance of the step is measured on the variational wave-function space $\|d\theta\|_\Psi$ [48,49]:

$$\|d\theta\|_\Psi^2 = 1 - |\langle \Psi_{\theta+d\theta} | \Psi_\theta \rangle|^2 \approx \sum_{n,m} g_{nm}(\theta) d\theta_n d\theta_m \quad (29)$$

up to higher-order corrections, where the summation includes all variational parameters θ_n and $g_{nm}(\theta)$ is the quantum geometric tensor

$$g_{nm}(\theta) = \langle \partial_{\theta_n} \Psi_\theta | (1 - |\Psi_\theta\rangle\langle\Psi_\theta|) | \partial_{\theta_m} \Psi_\theta \rangle. \quad (30)$$

In this way, the information of how each parameter affects the wave function is included in the distance measure. The resulting optimal step then includes the inverse of the quantum geometric tensor [48]

$$d\theta = -\eta g^{-1}(\theta) \nabla_\theta L[\theta], \quad (31)$$

where η controls the length of the optimizer step $d\theta$ and is commonly known as the *learning rate*. This procedure is known as natural gradient descent.

In practice, for variational models with large number of parameters, computing the inverse of the quantum geometric tensor is numerically impractical. For neural network wave-function *Ansätze*, the inverse of the geometric tensor can be efficiently approximated by the Kronecker-factorized approximate curvature (KFAC) method [10,50], where correlations between different layers as well as input and output of individual layers are neglected. Following Ref. [10], this method is applied in our numerical calculations. The current implementation of KFAC operates only on the absolute magnitude of the wave function, while the phase is neglected. In this case, the quantum geometric tensor reduces to the Fisher information matrix.

We remark that around a second-order quantum phase transition, the Fisher information diverges in the parameter that

TABLE I. Comparison of ground-state energies of the 27-site system at $\frac{2}{3}$ filling and $\epsilon = 10$ and 5 obtained using different methods. For 27 sites, BP-ED is limited to a single band due to the prohibitively large Hilbert space dimension. The bold font highlights the lowest energy result.

N_e	ϵ	Self-attention NN	SlaterNet	BP-ED
18	10	-32.070(7)	-31.35(2)	-31.32443
18	5	-59.127(9)	-58.01(3)	-57.80848

drives the quantum phase transition [51,52]. These divergencies may pose difficulties in the numerical approximation of the Fisher information matrix so that optimization around a quantum critical point may be more challenging than deep within a phase, even if the variational *Ansatz* remains sufficiently expressive to describe the critical state.

V. RESULTS

This section presents the numerical results obtained from our variational Monte Carlo calculations of Hartree-Fock using SlaterNet and the minimal self-attention NN and the benchmark with band-projected exact diagonalization.

Throughout this article, we consider WSe₂/WS₂ as model system, with the following model parameters determined by first-principles calculations [53]: effective mass $m^* = 0.35m_e$, moiré potential strength $V_0 = 15$ meV, and moiré shape parameter $\varphi = \pi/4$. We consider a moiré lattice period of $a_M = 8.031$ nm which results from the lattice mismatch between WSe₂ and WS₂. In this case, moiré filling of $\nu = 1$ particles per moiré unit cell corresponds to a density of $n = 1.785 \times 10^{12}$ cm⁻². The relative dielectric constant for a surrounding dielectric hBN is $\epsilon \approx 5$. Devices with tunable dielectric constants are possible using tunable dielectrics such as SrTiO₃ [54]. All our calculations are performed for the spin-polarized system.

These units are converted to the dimensionless units of Eq. (1) as follows. In Eq. (1), distances are measured in effective Bohr radii $a_B^* = 4\pi\epsilon\epsilon_0 \frac{\hbar^2}{m^*e^2} = \epsilon \frac{m_e}{m^*} a_0$ and energies in effective Hartree $\text{Ha}^* = \frac{\hbar^2}{m^*a_0^2} = \frac{1}{\epsilon^2} \frac{m^*}{m_e} \text{Ha}$, where a_0 and Ha are the standard Bohr radius and Hartree energy defined in terms of the free-electron mass m_e , m^* is the effective mass for charge carriers, and ϵ is the relative dielectric constant of the surrounding medium.

We summarize the hyperparameters used for training the models in Table II in Appendix C. While architectural parameters must be adjusted for each specific system (as detailed below), we found that the remaining hyperparameters were largely independent of the problem's specifics and consistently yielded stable and optimal results. We particularly highlight the learning rate $\eta_0 = 10$, which, although unusually high, emerged as the optimal value for a wide variety of 2D systems we experimented with, including both free Fermi liquid and moiré systems.

A. Convergence and scaling with system size

In any numerical method, it is crucial to ensure convergence and address how the numerical complexity scales with

TABLE II. Table of default hyperparameters used in our numerical calculations with the self-attention neural network.

	Parameter	Value
Architecture	Network layers	3
	Attention heads per layer	6
	Attention dimension	16
	Perceptron dimension	64
	No. perceptrons per layer	1
	Determinants	4
Training	Training iterations	15×10^4
	Learning rate at time t	$\eta_0(1 + \frac{t}{t_0})^{-1}$
	Initial learning rate η_0	10
	Learning rate delay t_0	1×10^5
	Local energy clipping ρ	5.0
MCMC	Batch size	4096
KFAC	Norm constraint	1×10^{-3}
	Damping	1×10^{-3}

the system size. Here, we first demonstrate convergence of the energy estimate as a function of Monte Carlo steps. Next, we investigate the saturation of the converged energy with the network dimensions. This enables us to numerically estimate the numerical complexity in terms of the scaling of the required number of variational parameters with the number of electrons.

To ensure the convergence of our results, we verified that the learning curves (i.e., the local energy averaged of the batch size of each step) exhibited stable behavior and, on average, decreased monotonically. This trend remained consistent across different system sizes, as illustrated by the raw learning curves for the 9- and 27-site systems at $\nu = \frac{2}{3}$ filling in Fig. 3(a), shown without smoothing. We highlight the mean and standard deviation of the obtained mean-energy-per-batch using a dashed line and a stripe. The latter arises from the finite batch size used for sampling at each iteration and sets an important scale, as independent learning curves typically remain within one standard deviation of the mean.

To further validate our approach, we examine the convergence of the ground-state energy per particle as a function of the number of variational parameters for the 9-, 12-, 27-, and 36-site systems at $\nu = \frac{2}{3}$ filling. Figure 3(b) explores networks spanning three orders of magnitude in variational parameters and reveals a clear convergence pattern in all cases, with the ground-state energy approaching a lower bound as the network reaches 1×10^6 variational parameters.

This distinct converging pattern indicates *convergence saturation*, meaning that further increasing the number of parameters does not improve the ground-state energy. We define the saturation point N_{par}^* for a given system as the threshold beyond which the converged energies consistently fall within one standard deviation of the lowest observed energy. In Fig. 3(b), these values are marked with stars.

We extract these saturation points as a function of the number of electrons N and find that they follow a well-defined scaling law:

$$N_{\text{par}}^* \approx (400 \pm 49) \times N^{2.01 \pm 0.05}. \quad (32)$$

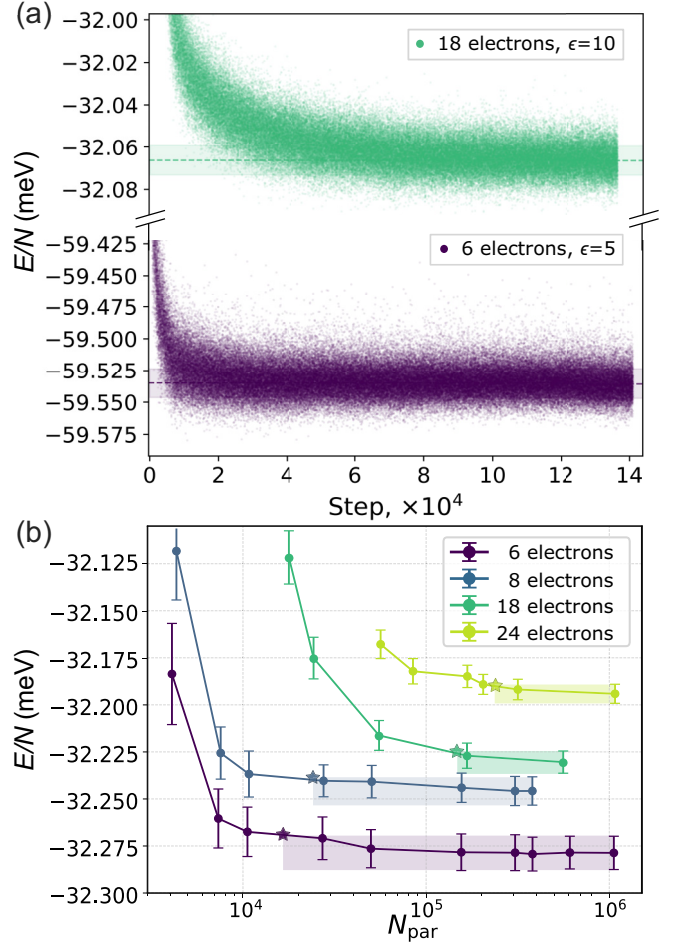


FIG. 3. (a) Raw (unsmoothed) learning curves for the energy per electron at $\nu = \frac{2}{3}$ filling for 6 electrons with $\epsilon = 5$, and 18 electrons with $\epsilon = 10$. (b) Converged energy of the moiré system at $\nu = \frac{2}{3}$ filling and $\epsilon = 10$ as a function of the number of variational parameters in the self-attention NN. Error bars indicate the standard deviation of the mean local energy, averaged over the batch size at the end of optimization. Stars mark the estimated number of parameters, N_{par}^* , required to reach convergence saturation. Further increasing in the number of parameters keeps the energies within one standard deviation (highlighted by the striped region). For sake of presentation, the 8e data were shifted up by 0.39 meV, and 18e and 24e data shifted down by 0.16 and 0.25 meV, respectively. Default simulation parameters are contained in Table II.

The fit is shown in Fig. 4. While this scaling behavior is derived for the present system and network architecture, we believe it may serve as a useful approximation for estimating parameter growth in other settings as well, though it may not directly generalize.

This scaling law provides a helpful estimate for the number of parameters required to achieve a given accuracy; however, the parameter count alone is not the only relevant factor. Since attention heads play a crucial role in capturing correlations, a minimum number of attention heads and layers (approximately above 3 in our tests) is necessary to reach the ground state. Simply increasing the total number of parameters without ensuring sufficient attention heads and layers would not improve performance. Our scaling law serves as a guideline




FIG. 4. Scaling law for the number of variational parameters at convergence saturation as a function of the number of electrons for $\epsilon = 10$. The extracted saturation points are well described by a power-law relation $N^* = CN^\alpha$, with best-fit parameters $\alpha = 2.01 \pm 0.05$ and $C = (4.00 \pm 0.49) \times 10^2$. This relation quantifies how the required parameter count grows with system size and provides a practical guideline for selecting an optimal network size while ensuring convergence.

within the regime where these minimal requirements are met. We leave a more detailed analysis of scaling laws as a function of individual network dimensions for future work.

In comparison, tensor network methods for two-dimensional lattices under area-law assumptions require a number of parameters scaling as $\propto e^{\sqrt{N}}$ [55]. For tensor networks, parameter scaling depends on the entanglement structure of the problem and tensor network *Ansatz*. Our empirical result of quadratic scaling of the parameters in the self-attention NN *Ansatz* suggests significantly better scaling compared to tensor network methods.

B. Benchmark for small system size

After ensuring convergence and saturation in the number of variational parameters, we benchmark our self-attention NN results with SlaterNet and BP-ED to assess the quantitative accuracy of predictions from the self-attention NN.

We study a variety of different system sizes and regimes at different interaction strengths. The comprehensive energy comparisons given in Fig. 5 graphically evince the difference in ground-state energy levels between all three methods for nine-unit cell system again at $\nu = \frac{2}{3}$. The self-attention NN energy is considerably lower than SlaterNet, with the difference precisely representing the correlation energy. We see that the correlation energy is about 2% in our moiré system, considerably higher than the correlation energy in ordinary molecular systems [3]. This enhanced correlation energy underscores the relevance of moiré systems for studying strongly correlated phenomena. Additionally, we observe that as the number of bands increases in BP-ED, the ground-state energy approaches but remains higher than that of the self-attention NN.

The advantage of self-attention NN becomes even more pronounced in larger systems, such as those with 18 electrons (Table I), where BP-ED is limited to a single band due to the enormous Hilbert space dimension. In this case, the self-attention NN energy is nearly 2.5% lower than that of BP-ED. In fact, the band mixing is so strong that even SlaterNet achieves a lower energy than BP-ED.

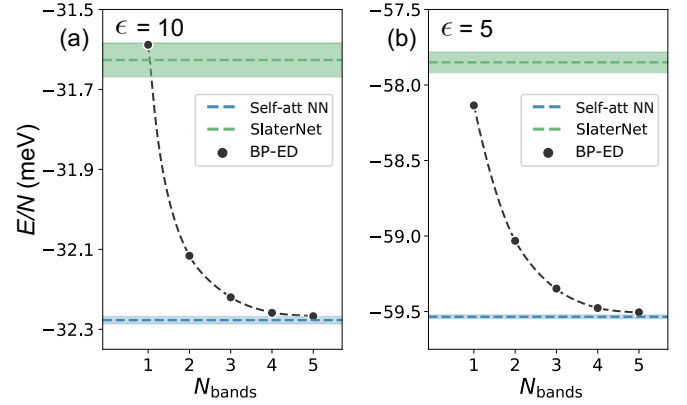


FIG. 5. Comparison of ground-state energies obtained from the self-attention NN (blue), Hartree-Fock from SlaterNet (green), and BP-ED (dots) in a 3×3 supercell with $\nu = \frac{2}{3}$ filling. Dashed lines for the self-attention NN and SlaterNet indicate mean of local energy obtained from the variational Monte Carlo algorithm. The shaded region indicates the standard deviation of the mean of the local energy averaged over the batch size per optimization step. Dashed gray line is a guide to the eyes.

These results demonstrate that self-attention NN consistently achieves lower energies than both BP-ED and SlaterNet, underscoring its capacity to express ground states of the moiré systems to higher accuracy, and revealing the variational superiority of our self-attention NN method.

C. Fermi liquids and generalized Wigner crystals

With the variational superiority of the self-attention NN method established, we now demonstrate the capability of the self-attention NN to not only reveal the essential physics of the strongly correlated system, but predict expected phenomenon. In real homobilayer systems, experimentalists can adjust displacement field to tune the level of electron-electron interactions, driving the system from a Fermi liquid state to generalized Wigner crystal state as interactions increase. While heterobilayer systems do not offer the same experimental freedom, a metal-insulator transition has been theoretically predicted as electron interactions grow stronger by previous ED and QMC studies [56,57].

Numerically, we can simulate this phenomenon by changing the dielectric constant ϵ which is inversely proportional to interaction strength. When interactions are small compared to kinetic energy, charge carriers are de-localized and the system is in the free-electron regime. In the presence of weak interactions, electrons will interact with each other at the boundary of the Brillouin zone, giving rise to a Fermi liquid. In real space, the delocalization of electrons results occupation of all potential minimum sites of the superlattice, producing the triangular lattice density pattern shown in Fig. 6(a).

As interactions increase, the Coulomb repulsion between the charge carriers makes it more energetically favorable to spread as far apart as possible, driving the system to transition into an electron crystalline state. The crystal breaks the translational symmetry of the underlying potential. At $\nu = \frac{2}{3}$ electrons per triangular potential unit cell, the electronic crystal has the shape of a honeycomb lattice as shown by the density pattern in Fig. 6(b).

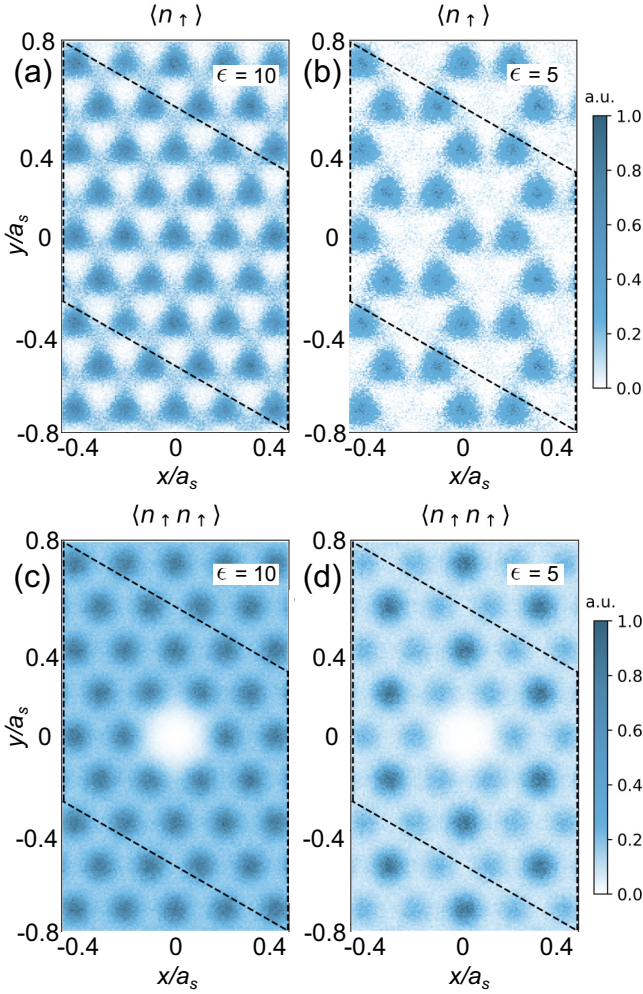


FIG. 6. Converged ground-state charge density [(a), (b)] and density-density correlation function [(c), (d)] the moiré system at $\nu = \frac{2}{3}$ filling of 27 sites per supercell with $\epsilon = 10$ (left) and $\epsilon = 5$ (right column). The dashed line indicates the boundary of the periodic simulation supercell.

In addition to the density profile, Figs. 6(c) and 6(d) show the pair correlation function for weak ($\epsilon = 10$) and strong interaction ($\epsilon = 5$), respectively. Both the density and pair correlation profile show clear signatures of a Fermi liquid state for $\epsilon = 10$ and a generalized Wigner crystal state for $\epsilon = 5$. This demonstrates the self-attention NN ability to capture the essential physics and strongly correlated behavior beyond simply lowering the ground-state energy.

Finally, the difference in energies from the self-attention NN and SlaterNet allows to estimate correlation energy equation (3). We include the extracted correlation energy as a function of dielectric constant in Appendix D. Furthermore, in Appendix E, we provide additional exact-diagonalization results on the metal-insulator phase transition.

VI. COMPARISON TO OTHER NEURAL NETWORK ANSÄTZE

In contrast to previous works studying periodic solids with neural quantum states [18,43], we do not multiply an envelope

function when constructing single-particle orbitals. Instead, the full functional dependence of the periodic orbitals is represented by the neural network. Reference [18] motivated the use of envelope functions by a better convergence for their system; however, our numerical experiments for our system did not show any improvements when envelope functions are included. Reference [43] found faster convergence using plane-wave envelopes when the system was close to a Hartree-Fock state. Envelope functions introduce human bias towards a certain functional dependence, therefore, we remove them in our network architecture.

Further, in our wave-function *Ansatz*, we do not multiply a Jastrow factor to the wave function constructed from generalized Slater determinants, thereby deviating from the PsiFormer architecture for molecular systems [11]. A Jastrow factor is usually motivated to capture correctly the wave-function cusps that occur when two particles approach each other, which is model specific and can be derived analytically (see Appendix B and Ref. [3] for details). We performed numerical experiments (Appendix B) including a simple Jastrow factor with a single learnable parameter that only enforces the correct wave-function shape as two particles approach each other. We did not find significant improvement by including the Jastrow factor, instead we observed that including the Jastrow increased the GPU time per step by around 10% to 20%. In contrast to molecular systems, we suspect that the Jastrow factor is less relevant for low-density electrons in a periodic potential as studied in this paper, because electrons are mostly bound to individual wells while overlap of two electrons is small.

We highlight that our results have been computed from a randomly initialized wave function. The only physical input is the determinant operation guaranteeing the antisymmetric structure of the fermionic wave function. In contrast, previous works [10,11] applied pretraining with Hartree-Fock to initialize the wave function before optimizing energy. Also, recent message-passing architectures [23] describe correlations starting from a chosen Slater determinant of single-particle orbitals.

In the future, our work can be extended to include computation of excited states [58] and zero-temperature Green's functions and spectral functions [59]. Finite-temperature effects can be studied by an extension of the neural network architecture to represent thermal density matrices [60,61]. Dynamical and open systems can be studied by applying the time-dependent variational principle [62] to a variational density matrix [63–68].

VII. CONCLUSION

In this work, we studied the performance of a wave-function *Ansatz* based on a many-parameter, self-attention neural network with minimal human input at a case study of a periodic moiré solid. Our numerical experiments revealed the following:

- (i) The self-attention neural network accurately predicts ground-state energy, outperforming our benchmark with band-projected exact diagonalization.
- (ii) The minimal number of parameters required for energy convergence scales roughly as N^2 with the particle

number, suggesting an efficient description of the ground state.

(iii) Charge density and correlation functions identify a Fermi liquid to Wigner crystal phase transition, confirming convergence on both sides of the phase transition.

Furthermore, Ref. [27] recently demonstrated that a similar self-attention neural network is able to accurately describe fractional quantum Hall ground-state wave functions, demonstrating that the ansatz is able to describe ground states whose correlations are fundamentally distinct from a state of uncorrelated orbitals.

The conceptual strength of the self-attention neural network *Ansatz* lies in its two pillars of construction: (i) it employs a large-scale parametrization without human bias, and (ii) the self-attention mechanism efficiently learns relations between electrons that encode their correlations. Altogether, these results put faith in the hope that a unified wave-function *Ansatz* may be able to accurately describe a wide range of, if not all, quantum phases in strongly interacting electron systems.

ACKNOWLEDGMENTS

We are grateful to A. Reddy for assistance with exact diagonalization study. We thank T. Teng, K. Lei, X. Li, and Y. Chen for helpful discussions. This work was primarily supported by a Simons Investigator Award from the Simons Foundation. M.G. acknowledges support from the German Research Foundation under the Walter Benjamin program (Grant Agreement No. 526129603) and by the Air Force Office of Scientific Research under Award No. FA2386-24-1-4043. T.Z. was supported by the MIT Dean of Science Graduate Student Fellowship. K.N. was supported by Air Force Office of Scientific Research under Award No. FA2386-21-1-4058. L.F. also acknowledges support from the NSF through Award No. PHY-2425180. This work made use of resources provided by subMIT at MIT Physics. The authors acknowledge the MIT SuperCloud and Lincoln Laboratory Supercomputing Center for providing resources [69].

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

Our software implementation is based on Google DeepMind's VMC package FERMINET [70] and makes use of the Forward Laplacian [13] implementation from [71].

APPENDIX A: COULOMB SYSTEM WITH PERIODIC BOUNDARY CONDITION

In the following, we employ and develop numerical techniques for solving the interacting Hamiltonian (1) on a finite-size cluster with periodic boundary condition, such that the N -electron wave function satisfies the condition

$$\Psi(\mathbf{r}_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_N) = \Psi(\mathbf{r}_1, \dots, \mathbf{r}_i + \mathbf{L}, \dots, \mathbf{r}_N) \quad (\text{A1})$$

for any particle i , where the two vectors $\mathbf{L} = n\mathbf{L}_1 + m\mathbf{L}_2$ with $n, m \in \mathbb{Z}$ specify the cluster size and geometry. For systems with an underlying periodic potential as in Eq. (1), the cluster must accommodate an integer number of unit cells of the periodic potential. Our finite-size system can also be viewed as one supercell within an infinitely large two-dimensional (2D) plane in which only particle configurations satisfying supercell periodicity are allowed.

For systems with periodic boundary conditions, the Coulomb interaction must be adapted to account for the interaction energy with periodic images of the charge configuration in the other “supercells.” The Coulomb interaction decomposes into two terms

$$H_{ee} = \frac{1}{2} \sum_i^N \sum_{j \neq i}^N \sum_{\mathbf{L}} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j + \mathbf{L}|} + \frac{1}{2} \sum_i^N \sum_{\mathbf{L} \neq 0} \frac{1}{|\mathbf{L}|} \quad (\text{A2})$$

where the first term describes interaction of the N distinct electrons in the supercell and their periodic images, while the second term describes the interaction of each individual electron in the supercell with its own periodic images. The second term is determined by the *Madelung constant* $\xi_M = \sum_{\mathbf{L} \neq 0} \frac{1}{|\mathbf{L}|}$. Because Coulomb interactions are long ranged, the summation contains a divergent, homogeneous contribution. This divergence is canceled when a neutralizing charge background is introduced, as we show later.

The remainder of this Appendix describes how the Coulomb interaction and Madelung constant are efficiently evaluated using *Ewald summation*. Our derivation closely follows Ref. [72]. Readers not interested in this derivation can refer to Eqs. (A11) and (A12) for the result of the electron-electron interaction and Madelung constant, respectively, that defines the Coulomb Hamiltonian for our numerical study.

The numerical procedure for the efficient computation of Eq. (A2) proceeds by evaluating the Coulomb potential induced by an individual electron at position \mathbf{r}_a in the supercell and its periodic images

$$\varphi_a(\mathbf{r}) = \sum_{\mathbf{L}} \frac{1}{|\mathbf{r} - \mathbf{r}_a - \mathbf{L}|}. \quad (\text{A3})$$

The principal idea of Ewald summation is to decompose the sum into quickly converging sums for short- and long-ranged contributions. This is achieved by Fourier transformation of each individual term, where the integration domain is taken as the infinite plane

$$\varphi_{a,L}(\mathbf{q}) = \int_{\mathbb{R}^2} d\mathbf{r} \frac{e^{i\mathbf{q}\mathbf{r}}}{|\mathbf{r} - \mathbf{r}_a - \mathbf{L}|} = 2\pi \frac{e^{-i\mathbf{q}(\mathbf{r}_a + \mathbf{L})}}{q}, \quad (\text{A4})$$

where $q = |\mathbf{q}|$. By applying the identity

$$\frac{1}{q} = \frac{2}{\sqrt{\pi}} \int_0^\infty dt e^{-q^2 t^2} = \frac{2}{\sqrt{\pi}} \left[\int_0^\eta + \int_\eta^\infty \right] dt e^{-q^2 t^2} \quad (\text{A5})$$

the integration is split into short- and long-ranged contributions. The first term describes the short-ranged contributions, as is seen by Fourier transformation back to

real space:

$$\begin{aligned}\varphi_{a,L}^S(\mathbf{r}) &= \frac{1}{\pi^{3/2}} \int_{\mathbb{R}^2} d^2\mathbf{q} e^{i\mathbf{q}(\mathbf{r}-\mathbf{r}_a-\mathbf{L})} \int_0^\eta dt e^{-q^2 t^2} \\ &= \frac{1}{|\mathbf{r}-\mathbf{r}_a-\mathbf{L}|} \left[\text{Erfc}\left(\frac{|\mathbf{r}-\mathbf{r}_a-\mathbf{L}|}{2\eta}\right) \right].\end{aligned}\quad (\text{A6})$$

The integral describing the long-ranged contributions can be directly evaluated in momentum space,

$$\varphi_{a,L}^L(\mathbf{q}) = 2\pi e^{-i\mathbf{q}\mathbf{r}_a} \frac{\text{Erfc}(\eta\mathbf{q})}{q}. \quad (\text{A7})$$

The complementary error function $\text{Erfc}(x)$ ensures fast convergence of the summation due to an exponential suppression for $x > 1$.

The long-range part contains a divergent contribution in the limit $q \rightarrow 0$. By requiring that the mean electrostatic potential is zero, the $\mathbf{q} = 0$ contributions of both short- and long-range potentials are canceled by a homogeneous charge background. This cancels the divergence of the long-range part, and requires to subtract the $\mathbf{q} \rightarrow 0$ contribution of the short-range part:

$$\begin{aligned}\varphi_{a,\mathbf{q} \rightarrow 0}^S &= \int_{\mathbb{R}^2} \frac{d^2\mathbf{q}}{2\pi} \sum_L e^{i\mathbf{q}(\mathbf{r}-\mathbf{r}_a-\mathbf{L})} \frac{\text{Erf}(\eta\mathbf{q})}{q} \\ &= \int_{\mathbb{R}^2} \frac{d^2\mathbf{q}}{2\pi} \sum_G \delta(\mathbf{q}-\mathbf{G}) e^{i\mathbf{q}(\mathbf{r}-\mathbf{r}_a)} \frac{\text{Erf}(\eta\mathbf{q})}{q} \\ &= \frac{2\pi}{A_{\text{u.c.}}} \sum_G e^{i\mathbf{G}(\mathbf{r}-\mathbf{r}_a)} \frac{\text{Erf}(\eta\mathbf{G})}{G} = \frac{2\pi}{A_{\text{u.c.}}} \frac{2\eta}{\sqrt{\pi}}.\end{aligned}\quad (\text{A8})$$

The total electrostatic potential of a charge at \mathbf{r}_a and its periodic images is given by the summation of short- and long-range contributions from all unit cells shifted by \mathbf{L} :

$$\begin{aligned}\varphi_a(\mathbf{r}) &= \sum_L \varphi_{a,L}^S(\mathbf{r}) - \varphi_{a,\mathbf{q} \rightarrow 0}^S \\ &\quad + \sum_L \int_{\mathbb{R}_{\neq 0}^2} \frac{d^2\mathbf{q}}{4\pi^2} e^{i\mathbf{q}\mathbf{r}} \varphi_{a,L}^L(\mathbf{q}).\end{aligned}\quad (\text{A9})$$

The long-range contribution can further be simplified

$$\begin{aligned}\sum_L \int_{\mathbb{R}_{\neq 0}^2} \frac{d^2\mathbf{q}}{4\pi^2} e^{i\mathbf{q}\mathbf{r}} \varphi_{a,L}^L(\mathbf{q}) \\ &= \frac{1}{2\pi} \int_{\mathbb{R}_{\neq 0}^2} d^2\mathbf{q} \sum_G \delta(\mathbf{q}-\mathbf{G}) e^{i\mathbf{q}(\mathbf{r}-\mathbf{r}_a)} \frac{\text{Erfc}(\eta\mathbf{q})}{q} \\ &= \frac{2\pi}{A_{\text{u.c.}}} \sum_{\mathbf{G} \neq 0} e^{i\mathbf{G}(\mathbf{r}-\mathbf{r}_a)} \frac{\text{Erfc}(\eta\mathbf{G})}{G},\end{aligned}\quad (\text{A10})$$

where \mathbf{G} are reciprocal supercell vectors $\mathbf{G} \cdot \mathbf{L} = 2\pi n$, $n \in \mathbb{N}$.

The total electrostatic energy of a configuration of N particles in a unit cell is then written as

$$H_{ee} = \frac{1}{2} \sum_b \sum_{a \neq b} \varphi_a(\mathbf{r}_b) + \frac{1}{2} \sum_b \xi_M, \quad (\text{A11})$$

where $a, b = 1, \dots, N$ enumerates the particles in the unit cell, the first term describes the energy of all particles b in the total potential of all other particles $a \neq b$, while the second

term describes the energy of all particles in the potential created by their images in the other unit cells, determined by the Madelung constant

$$\begin{aligned}\xi_M &= \sum_{L \neq 0} \varphi_{b,L}^S(\mathbf{r}_b) + \sum_{L \neq 0} \int_{\mathbb{R}_{\neq 0}^2} \frac{d^2\mathbf{q}}{4\pi^2} e^{i\mathbf{q}\mathbf{r}_b} \varphi_{b,L}^L(\mathbf{q}) \\ &= \sum_{L \neq 0} \frac{1}{|L|} \left[\text{Erfc}\left(\frac{|L|}{2\eta}\right) \right] - \varphi_{a,\mathbf{q} \rightarrow 0}^S \\ &\quad + \frac{2\pi}{A_{\text{u.c.}}} \sum_{\mathbf{G} \neq 0} \frac{\text{Erfc}(\eta\mathbf{G})}{G} - \xi_0^L,\end{aligned}\quad (\text{A12})$$

where the $\mathbf{q} \rightarrow 0$ part of the short-range potential is subtracted, and the term ξ_0^L arises from subtracting the $\mathbf{L} = 0$ term from the evaluation of the long-range potential

$$\xi_0^L = \int \frac{d^2\mathbf{q}}{2\pi} \frac{\text{Erfc}(\eta\mathbf{q})}{q} = \frac{1}{\eta\sqrt{\pi}}.$$

For exact diagonalization in momentum space as presented in Sec. II, we note that the reciprocal space Coulomb interaction [last term in Eq. (4)] arises from Fourier transform of the first term of Eq. (A2) describing the real-space Coulomb interaction in systems with periodic boundary condition. In the reciprocal space formulation of Eq. (4), the divergent $\mathbf{q} = 0$ term of the Coulomb interaction is set to zero by requiring cancellation with a homogeneous charge background; the same requirement was imposed when deriving the real-space Coulomb interaction [Eq. (A11)] using Ewald summation. Notice that the Madelung energy $N\xi_M/2$ needs to be explicitly included in the reciprocal space formulation [Eq. (4)] in order to capture the interaction of each electron in the supercell with its own periodic images.

APPENDIX B: JASTROW FACTORS

A conventional approach to describe electron-electron correlations beyond Hartree-Fock is to multiply a *Jastrow* factor to a Slater determinant [73]. No universal approximation theorem exists for this wave-function *Ansatz* to the knowledge of the authors at the time of publication. Jastrow factors are typically be model specific and physically motivated, allowing to capture a large fraction of the correlation energy [3].

A Jastrow factor is a bosonic (i.e., permutation equivariant) function that is multiplied to the many-body wave function

$$\Psi(\mathbf{R}) \rightarrow J(\mathbf{R})\Psi(\mathbf{R}).$$

For systems with Coulomb interactions, a typical choice for the Jastrow factor is a functional of the form

$$J(\mathbf{R}) = e^{-\sum_i^{N_{\text{elec}}} \sum_{i \neq j}^{N_{\text{elec}}} u(|\mathbf{r}_i - \mathbf{r}_j|)}, \quad (\text{B1})$$

where $u(r)$ depends on the relative electron-electron distances and is tailored to capture the wave-function cusps as two electrons approach each other: For Coulomb interaction, the energy diverges as two particles approach each other. This divergence is canceled by a divergence of kinetic energy contribution for a specific functional form of the wave function. Because eigenstates of the Coulomb gas have finite energy, this cancellation occurs for all eigenstates, and the functional form as two particles approach each other can be analytically




FIG. 7. Comparison of training curves for 27 sites with $\frac{2}{3}$ filling and (a) $\epsilon = 10$ and (b) $\epsilon = 5$ with simple Jastrow factor $J(\mathbf{R})$ and $u(r)$ defined by Eq. (B3) to fix cusp conditions (blue) and no Jastrow factor (green). Dashed and dotted lines indicate mean and standard deviation of the local energy averaged over batch size. These calculations were performed with batch size 1024. Each data point represents the local energy averaged over batch size at each step.

derived. This dependence can be explicitly enforced by the Jastrow factor.

Concretely, adapting the derivation of Ref. [3] to two dimensions, requiring that the local energy E_{loc} is finite for configurations with two particles approaching each other for the Coulomb gas Hamiltonian in dimensionless units [Eq. (1)] leads to the conditions [27]

$$\beta := \partial_r u(r) = \begin{cases} -1/3 & \text{for parallel spins,} \\ -1 & \text{for opposite spins,} \end{cases} \quad (\text{B2})$$

where ∂_r is the partial derivative in the radial direction. The conditions for parallel and opposite spins are different because for the former, the Slater determinant part of the wave function vanishes while for the latter it remains finite as two particles approach each other.

In Fig. 7 we present numerical calculations including a simple Jastrow factor

$$u(r) = -\beta \frac{\alpha^2}{\alpha + r} \quad (\text{B3})$$

with a single learnable parameter α and β determined by Eq. (B2) to enforce the analytically known behavior of $\Psi(\mathbf{R})$ as two particles approach each other. In comparison to our calculations without Jastrow factor, the inclusion of the




FIG. 8. Converged energy of the moiré system at $\nu = \frac{2}{3}$ filling and $\epsilon = 10$ as a function of number of variational parameters. The labels indicate the parameters (n_{layers} , n_{heads} , dim_{att} , dim_{perc}) used for each calculation. The total number of parameters are counted explicitly from the checkpoints saved during the training.

Jastrow factor did not significantly improve the results. Instead, we found faster training without the Jastrow factor.

In systems with periodic boundary condition, the electron-electron distance r that enters the Jastrow factor is calculated with respect to a periodic and smooth norm that reproduces the Euclidean norm in the limit $r \rightarrow 0$. The smooth and periodic norm $\|r\|_M$ with respect to the supercell vectors $M = (L_1, L_2)$ is defined as [18]

$$\|r\|_M = \frac{1}{2\pi} \sqrt{a^T M^T M a + b^T M^T M b}, \quad a = 1 - \cos(Gr), \quad b = \sin(Gr), \quad (\text{B4})$$

where $G = 2\pi M^{-1}$ are the reciprocal supercell vectors.

APPENDIX C: HYPERPARAMETERS

Table II presents the default hyperparameters used in our calculations, ensuring stable and monotonous convergence during training. As discussed in the main text, architectural parameters must be adjusted according to system size to achieve convergence. For the convergence and scaling law analysis, we explored various hyperparameter distributions. The table lists the set that successfully ensured convergence for the 18-electron system. These parameters are applied to generate the figures in the main text.

Additionally, we note that batch size is a crucial factor in controlling the standard deviation of the energy curves, with error bars decreasing as $1/\sqrt{\text{batch size}}$. In some cases, reducing the KFAC damping parameter resulted in slightly lower energies than those reported. However, we observed significant instabilities when training with reduced KFAC damping, often leading to corrupted training runs.

Figure 8 contains the precise network dimensions used to generate the parameter scaling in Fig. 3 in the main text. For the self-attention neural network as described in Sec. III C, the total number of variational parameters depends on the network




FIG. 9. Correlation energy per electron $E_{\text{corr}} = E_{\text{SELFATTN}} - E_{\text{SlaterNet}}$, at $\nu = \frac{2}{3}$ filling with 27 sites as a function of inverse dielectric constant ϵ^{-1} .

dimensions as

$$\begin{aligned} N_{\text{par}} = & 2d_{\text{dim}}d_L + L(2N_{\text{heads}}d_{\text{Attn}}d_L \\ & + 2N_{\text{heads}}d_{\text{AttnValues}}d_L \\ & + d_L(d_L + 1)) + 2d_LN_{\text{el}}N_{\text{det}}, \end{aligned}$$

where the first line is the embedding W_0 of the periodic features in the layer dimension d_L , the second line are key and query matrices W_k^{lh} , W_q^{lh} for each layer $l = 1, \dots, L$ and head $h = 1, \dots, N_{\text{heads}}$, the third line are the value matrices W_v^{lh} and projection matrices W_o^l onto the layer dimension, the fourth line are the perceptron weights and biases W^l , b_l , and the last line are the projection w_{2j}^m , w_{2j+1}^m on $j = 1, \dots, N_{\text{el}}$ complex orbitals for $m = 1, \dots, N_{\text{det}}$ determinants.

APPENDIX D: CORRELATION ENERGY

Here, we investigate the correlation energy per electron, defined as the ground-state energy difference between the Hartree-Fock (SlaterNet) and self-attention NN approaches. We analyze its dependence on the interaction strength, controlled by the dielectric constant ϵ (see Fig. 9 for the result). As ϵ decreases (corresponding to stronger interactions), the correlation energy gradually increases, reflecting the growing importance of many-body effects beyond mean-field theory. This trend highlights the limitations of Hartree-Fock in capturing strong correlations and underscores the role of self-attention NN in accurately describing the correlated electronic states.

APPENDIX E: BP-ED BENCHMARK

In addition to benchmarking our self-attention NN energies with BP-ED, we leverage BP-ED to reveal deeper details on our system, namely, identifying the Fermi-liquid (FL) and generalized Wigner crystal (GWC) phases, along with the order and location of the metal-insulator (MIT) transition. Figure 10 shows two distinct phase signatures at different interaction strengths $\epsilon = 10$ [Figs. 10(a) and 10(c)], and $\epsilon = 5$ [Figs. 10(b) and 10(d)]. By assigning each crystal momenta an integer index, the many-body spectra (MBS) in Figs. 10(a) and 10(b) show the lowest energies per momentum sector, with the corresponding momentum space positions of each sector graphically displayed in the moiré Brillouin

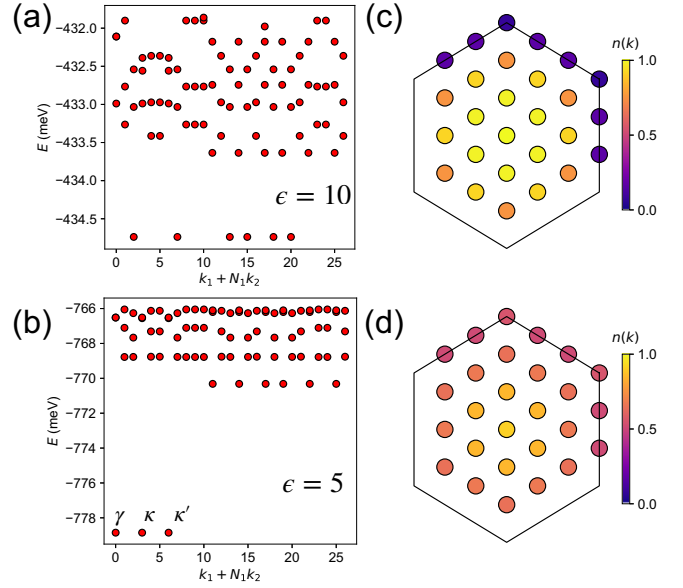


FIG. 10. Many-body spectrum (MBS) (a), (b) and ground-state occupation (c), (d) for 27 sites at $\frac{2}{3}$ filling for (above) $\epsilon = 10$ and (below) $\epsilon = 5$. The MBS of the left figure demonstrates the sixfold GS degeneracy, while the MBS on the right demonstrates the threefold GS degeneracy at the high-symmetry points, indicating a generalized Wigner crystal state. A clear Fermi surface is also visible in (c) as the system is in a Fermi-liquid state.

zone (MBZ) as a function of ground-state (GS) occupation in Figs. 10(c) and 10(d).

The sixfold rotational symmetry of the 27-site cluster allows the GS to accommodate a sixfold degeneracy, emblematic of a high-mobility, Fermi-liquid-like state. On the other hand, the 27-site cluster also accommodates a tripled unit cell where the corners of the MBZ κ , κ' fold back to γ in the BZ, indicating an expected threefold GS degeneracy at γ , κ , κ' for a GWC signature. Not only does the MBS in Figs. 10(a) and 10(b) show clear signatures of a FL and GWC state at




FIG. 11. Many-body energy gap ($E_{\text{gap}} = E_4 - E_3$, where E_i labels the state with the i th lowest energy) as a function of interaction strength ϵ^{-1} . The gap opens just after $\epsilon^{-1} = 0.10$, signaling an MIT at about $\epsilon^{-1} = 0.11$.




FIG. 12. $\langle H_0 \rangle$ [(a), (c)] and $S(\mathbf{q} = \kappa)$ [(b), (d)] as a function of ϵ^{-1} for the 21- [(a), (b)] and 27-site [(c), (d)] systems. Here, $\langle \dots \rangle_{\text{GS}}$ denotes the expectation value over a single ground state. The discontinuity gap increases with system size evincing a first-order metal-insulator transition.

respective interactions, but also Fig. 10(c) shows the presence of a Fermi surface while its $\epsilon^{-1} = 5$ counterpart does not.

The exact MIT point predicted by BP-ED is shown in Fig. 11, in which the many-body energy gap (difference between fourth and third lowest-energy values) are plotted versus interaction strength ϵ^{-1} . Our expectation of an MIT as interaction strength increases is validated by the sixfold GS degeneracy eventually giving way to a threefold GS degeneracy at about $\epsilon^{-1} = 0.11$, signaling the crossover from FL to GWC. Our BP-ED results are well consistent with our self-attention NN results in Fig. 6, as evidence for FL and GWC states are demonstrated at the appropriate interaction strengths, showcasing the ability of our self-attention NN to capture the physics of this moiré system.

BP-ED techniques can also be used to gain insight on the nature of this MIT. We can characterize the transition as first or second order by probing the discontinuity of observables across the transition point. First-order transitions are characterized by discontinuities across transition points while second-order transitions are continuous across transition points for large enough system size. Since we are dealing with finite system sizes in BP-ED, there will always be a discontinuity; however, if this discontinuity grows with system size, then we can deduce that correlation length will not tend to 0 with increasing system size, thereby showing evidence for

a first-order transition. Inversely, a decrease in the observable discontinuity across the transition point as system size grows is evidence for a second-order phase transition.

By probing the kinetic plus periodic potential one-body energy (H_0 in Eq. (1),

$$\langle H_0 \rangle = \sum_{\mathbf{k}} \epsilon_{\mathbf{k}} \langle c_{\mathbf{k}}^{\dagger} c_{\mathbf{k}} \rangle_{\text{GS}}, \quad (\text{E1})$$

and the band-projected structure factor

$$\bar{S}(\mathbf{q}) = \frac{1}{N_e} \langle \bar{n}(\mathbf{q}) \bar{n}(-\mathbf{q}) \rangle_{\text{GS}}, \quad (\text{E2})$$

across the MIT point ($\epsilon^{-1} = 0.11$) for 21- and 27-site clusters (21, 27 k -point mesh detailed in [74]), we can measure the discontinuity gap to shed light on the order of the phase transition. Since the projected structure factor has peaks along $\mathbf{q} = \kappa, \kappa'$, we plot $\bar{S}(\mathbf{q} = \kappa)$ as a function of interaction strength to isolate the effect of the transition.

Figure 12 not only demonstrates a discontinuity exactly at the critical point for each cluster, but also shows that the discontinuity gap grows as system size grows for both kinetic energy and the band-projected structure factor (for $\mathbf{q} = \kappa$). The discontinuity gap for $\langle H_0 \rangle$ increases from 0.61 to 0.87 meV and from 0.88 to 1.35 for $\bar{S}(\mathbf{q} = \kappa)$, demonstrating evidence for a first-order MIT.

- [1] D. R. Hartree, The wave mechanics of an atom with a non-coulomb central field. Part II. Some results and discussion, *Math. Proc. Cambridge Philos. Soc.* **24**, 111 (1928).
- [2] V. Fock, Näherungsmethode zur Lösung des quantenmechanischen Mehrkörperproblems, *Z. Phys.* **61**, 126 (1930).
- [3] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Quantum Monte Carlo simulations of solids, *Rev. Mod. Phys.* **73**, 33 (2001).
- [4] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [5] K. T. Schütt, P.-J. Kindermans, H. E. Sauceda, S. Chmiela, A. Tkatchenko, and K.-R. Müller, SchNet: a continuous-filter convolutional neural network for modeling quantum interactions, in *Guide Proceedings* (Curran Associates, Red Hook, New York, 2017), pp. 992–1002.
- [6] G. Carleo, K. Choo, D. Hofmann, J. E. T. Smith, T. Westerhout, F. Alet, E. J. Davis, S. Efthymiou, I. Glasser, S.-H. Lin, M. Mauri, G. Mazzola, C. B. Mendl, E. van Nieuwenburg, O. O'Reilly, H. Théveniaut, G. Torlai, F. Vicentini, and A. Wietek, NetKet: A machine learning toolkit for many-body quantum systems, *SoftwareX* **10**, 100311 (2019).
- [7] D. Luo and B. K. Clark, Backflow transformations via neural networks for quantum many-body wave functions, *Phys. Rev. Lett.* **122**, 226401 (2019).
- [8] K. Choo, A. Mezzacapo, and G. Carleo, Fermionic neural-network states for *ab-initio* electronic structure, *Nat. Commun.* **11**, 2368 (2020).
- [9] J. Hermann, Z. Schätzle, and F. Noé, Deep-neural-network solution of the electronic Schrödinger equation, *Nat. Chem.* **12**, 891 (2020).
- [10] D. Pfau, J. S. Spencer, A. G. D. G. Matthews, and W. M. C. Foulkes, *Ab initio* solution of the many-electron Schrödinger equation with deep neural networks, *Phys. Rev. Res.* **2**, 033429 (2020).
- [11] I. von Glehn, J. S. Spencer, and D. Pfau, A self-attention Ansatz for *ab-initio* quantum chemistry, *The Eleventh International Conference on Learning Representations* (2023).
- [12] L. L. Viteritti, R. Rende, and F. Becca, Transformer variational wave functions for frustrated quantum spin systems, *Phys. Rev. Lett.* **130**, 236401 (2023).
- [13] R. Li, H. Ye, D. Jiang, X. Wen, C. Wang, Z. Li, X. Li, D. He, J. Chen, W. Ren, and L. Wang, Forward Laplacian: A new computational framework for neural network-based variational Monte Carlo, [arXiv:2307.08214](https://arxiv.org/abs/2307.08214).
- [14] N. Gao and S. Günnemann, Generalizing neural wave functions, in *Proceedings of the 40th International Conference on Machine Learning, ICML'23* (JMLR.org, 2023), pp. 10708–10726.
- [15] J. R. Moreno, G. Carleo, A. Georges, and J. Stokes, Fermionic wave functions from neural-network constrained hidden states, *Proc. Natl. Acad. Sci. USA* **119**, e2122059119 (2022).
- [16] D. Luo, Z. Chen, K. Hu, Z. Zhao, V. M. Hur, and B. K. Clark, Gauge-invariant and anyonic-symmetric autoregressive neural network for quantum lattice models, *Phys. Rev. Res.* **5**, 013216 (2023).
- [17] A. Chen and M. Heyl, Empowering deep neural quantum states through efficient optimization, *Nat. Phys.* **20**, 1476 (2024).
- [18] G. Cassella, H. Sutterud, S. Azadi, N. D. Drummond, D. Pfau, J. S. Spencer, and W. M. C. Foulkes, Discovering quantum phase transitions with fermionic neural networks, *Phys. Rev. Lett.* **130**, 036401 (2023).
- [19] M. Wilson, S. Moroni, M. Holzmann, N. Gao, F. Wudarski, T. Vegge, and A. Bhowmik, Neural network Ansatz for periodic wave functions and the homogeneous electron gas, *Phys. Rev. B* **107**, 235139 (2023).
- [20] D. Luo, D. D. Dai, and L. Fu, Pairing-based graph neural network for simulating quantum materials, [arXiv:2311.02143](https://arxiv.org/abs/2311.02143).
- [21] J. Kim, G. Pescia, B. Fore, J. Nys, G. Carleo, S. Gandolfi, M. Hjorth-Jensen, and A. Lovato, Neural-network quantum states for ultra-cold Fermi gases, *Commun. Phys.* **7**, 148 (2024).
- [22] C. Smith, Y. Chen, R. Levy, Y. Yang, M. A. Morales, and S. Zhang, Unified variational approach description of ground-state phases of the two-dimensional electron gas, *Phys. Rev. Lett.* **133**, 266504 (2024).
- [23] G. Pescia, J. Nys, J. Kim, A. Lovato, and G. Carleo, Message-passing neural quantum states for the homogeneous electron gas, *Phys. Rev. B* **110**, 035108 (2024).
- [24] J. A. Sobral, M. Perle, and M. S. Scheurer, Physics-informed transformers for electronic quantum states, [arXiv:2412.12248](https://arxiv.org/abs/2412.12248).
- [25] X. Li, Y. Qian, W. Ren, Y. Xu, and J. Chen, Emergent Wigner phases in moiré superlattice from deep learning, [arXiv:2406.11134](https://arxiv.org/abs/2406.11134).
- [26] D. Luo, D. D. Dai, and L. Fu, Simulating moiré quantum matter with neural network, [arXiv:2406.17645](https://arxiv.org/abs/2406.17645).
- [27] Y. Teng, D. D. Dai, and L. Fu, Solving and visualizing fractional quantum Hall wavefunctions with neural network, *Phys. Rev. B* **111**, 205117 (2025).
- [28] Y. Qian, T. Zhao, J. Zhang, T. Xiang, X. Li, and J. Chen, Describing Landau level mixing in fractional quantum Hall states with deep learning, *Phys. Rev. Lett.* **134**, 176503 (2025).
- [29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, Attention is all you need, in *Advances in Neural Information Processing Systems*, Vol. 30, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., 2017).
- [30] E. C. Regan, D. Wang, C. Jin, M. I. Bakti Utama, B. Gao, X. Wei, S. Zhao, W. Zhao, Z. Zhang, K. Yumigeta *et al.*, Mott and generalized Wigner crystal states in WSe₂/WS₂ moiré superlattices, *Nature (London)* **579**, 359 (2020).
- [31] Y. Tang, L. Li, T. Li, Y. Xu, S. Liu, K. Barmak, K. Watanabe, T. Taniguchi, A. H. MacDonald, J. Shan, and K. F. Mak, Simulation of Hubbard model physics in WSe₂/WS₂ moiré superlattices, *Nature (London)* **579**, 353 (2020).
- [32] E. A. Arsennault, Y. Li, B. Yang, X. Wang, H. Park, E. Mosconi, E. Ronca, T. Taniguchi, K. Watanabe, D. Gamelin, A. Millis, C. R. Dean, F. de Angelis, X. Xu, and X. Y. Zhu, Two-dimensional moiré polaronic electron crystals, *Phys. Rev. Lett.* **132**, 126501 (2024).
- [33] H. Li, S. Li, E. C. Regan, D. Wang, W. Zhao, S. Kahn, K. Yumigeta, M. Blei, T. Taniguchi, K. Watanabe *et al.*, Imaging two-dimensional generalized Wigner crystals, *Nature (London)* **597**, 650 (2021).
- [34] W. Zhao, B. Shen, Z. Tao, Z. Han, K. Kang, K. Watanabe, T. Taniguchi, K. F. Mak, and J. Shan, Gate-tunable heavy fermions in a moiré Kondo lattice, *Nature (London)* **616**, 61 (2023).
- [35] F. Wu, T. Lovorn, E. Tutuc, and A. H. MacDonald, Hubbard model physics in transition metal dichalcogenide moiré bands, *Phys. Rev. Lett.* **121**, 026402 (2018).

- [36] Y. Zhang, H. Isobe, and L. Fu, Density functional approach to correlated moiré states: Itinerant magnetism, [arXiv:2005.04238](#).
- [37] A. P. Reddy, T. Devakul, and L. Fu, Artificial atoms, wigner molecules, and an emergent kagome lattice in semiconductor moiré superlattices, *Phys. Rev. Lett.* **131**, 246501 (2023).
- [38] K. Hornik, M. Stinchcombe, and H. White, Multilayer feed-forward networks are universal approximators, *Neural Netw.* **2**, 359 (1989).
- [39] H. Kim, Y. Zhou, Y. Xu, K. Varma, A. H. Karamlou, I. T. Rosen, J. C. Hoke, C. Wan, J. P. Zhou, W. D. Oliver, Y. D. Lensky, K. Q. Weinberger, and E.-A. Kim, Attention to quantum complexity, [arXiv:2405.11632](#).
- [40] G. Pescia, J. Han, A. Lovato, J. Lu, and G. Carleo, Neural-network quantum states for periodic systems in continuous space, *Phys. Rev. Res.* **4**, 023138 (2022).
- [41] R. P. Feynman and M. Cohen, Energy spectrum of the excitations in liquid helium, *Phys. Rev.* **102**, 1189 (1956).
- [42] Y. Kwon, D. M. Ceperley, and R. M. Martin, Effects of three-body and backflow correlations in the two-dimensional electron gas, *Phys. Rev. B* **48**, 12037 (1993).
- [43] X. Li, Z. Li, and J. Chen, *Ab initio* calculation of real solids via neural network ansatz, *Nat. Commun.* **13**, 7895 (2022).
- [44] M. J. Buehler, Graph-aware isomorphic attention for adaptive dynamics in transformers, *APL Mach. Learn.* **3**, 026108 (2025).
- [45] J. Anandan and Y. Aharonov, Geometry of quantum evolution, *Phys. Rev. Lett.* **65**, 1697 (1990).
- [46] N. Metropolis and S. Ulam, The Monte Carlo method, *J. Am. Stat. Assoc.* **44**, 335 (1949).
- [47] W. K. Hastings, Monte Carlo sampling methods using Markov chains and their applications, *Biometrika* **57**, 97 (1970).
- [48] S.-i. Amari, Natural gradient works efficiently in learning, *Neural Comput.* **10**, 251 (1998).
- [49] S. Sorella, Green function Monte Carlo with stochastic reconfiguration, *Phys. Rev. Lett.* **80**, 4558 (1998).
- [50] J. Martens and R. Grosse, Optimizing neural networks with Kronecker-factored approximate curvature, in *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML'15* (JMLR.org, 2015), pp. 2408–2417.
- [51] L. C. Venuti and P. Zanardi, Quantum critical scaling of the geometric tensors, *Phys. Rev. Lett.* **99**, 095701 (2007).
- [52] S.-J. Gu, Fidelity approach to quantum phase transitions, *Int. J. Mod. Phys. B* **24**, 4371 (2010).
- [53] Y. Zhang, N. F. Q. Yuan, and L. Fu, Moiré quantum chemistry: Charge transfer in transition metal dichalcogenide superlattices, *Phys. Rev. B* **102**, 201115(R) (2020).
- [54] D. Fuchs, C. W. Schneider, R. Schneider, and H. Rietschel, High dielectric constant and tunability of epitaxial SrTiO₃ thin film capacitors, *J. Appl. Phys.* **85**, 7362 (1999).
- [55] M. Ganahl, J. Beall, M. Hauru, A. G. M. Lewis, T. Wojno, J. H. Yoo, Y. Zou, and G. Vidal, Density matrix renormalization group with tensor processing units, *PRX Quantum* **4**, 010317 (2023).
- [56] N. Morales-Durán, A. H. MacDonald, and P. Potasz, Metal-insulator transition in transition metal dichalcogenide heterobilayer moiré superlattices, *Phys. Rev. B* **103**, L241110 (2021).
- [57] Y. Yang, M. A. Morales, and S. Zhang, Metal-insulator transition in a semiconductor heterobilayer model, *Phys. Rev. Lett.* **132**, 076503 (2024).
- [58] D. Pfau, S. Axelrod, H. Sutterud, I. von Glehn, and J. S. Spencer, Accurate computation of quantum excited states with neural networks, *Science* **385**, 6711 (2024).
- [59] D. Hendry, H. Chen, P. Weinberg, and A. E. Feiguin, Chebyshev expansion of spectral functions using restricted Boltzmann machines, *Phys. Rev. B* **104**, 205130 (2021).
- [60] N. Irikura and H. Saito, Neural-network quantum states at finite temperature, *Phys. Rev. Res.* **2**, 013284 (2020).
- [61] S. Lu, G. Giudice, and J. I. Cirac, Variational neural and tensor network approximations of thermal states, *Phys. Rev. B* **111**, 075102 (2025).
- [62] R. Jackiw and A. Kerman, Time-dependent variational principle and the effective action, *Phys. Lett. A* **71**, 158 (1979).
- [63] M. J. Hartmann and G. Carleo, Neural-network approach to dissipative quantum many-body dynamics, *Phys. Rev. Lett.* **122**, 250502 (2019).
- [64] A. Nagy and V. Savona, Variational quantum Monte Carlo method with a neural-network Ansatz for open quantum systems, *Phys. Rev. Lett.* **122**, 250501 (2019).
- [65] F. Vicentini, A. Biella, N. Regnault, and C. Ciuti, Variational neural-network ansatz for steady states in open quantum systems, *Phys. Rev. Lett.* **122**, 250503 (2019).
- [66] N. Yoshioka and R. Hamazaki, Constructing neural stationary states for open quantum many-body systems, *Phys. Rev. B* **99**, 214306 (2019).
- [67] M. Reh, M. Schmitt, and M. Gärtner, Time-dependent variational principle for open quantum systems with artificial neural networks, *Phys. Rev. Lett.* **127**, 230501 (2021).
- [68] J. Mellak, E. Arrigoni, and W. von der Linden, Deep neural networks as variational solutions for correlated open quantum systems, *Commun. Phys.* **7**, 268 (2024).
- [69] A. Reuther, J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron, V. Gadepally, M. Houle, M. Hubbell, M. Jones, A. Klein, L. Milechin, J. Mullen, A. Prout, A. Rosa, C. Yee, and P. Michaleas, Interactive supercomputing on 40,000 cores for machine learning and data analysis, in 2018 IEEE High Performance Extreme Computing Conference (HPEC) (IEEE, Piscataway, NJ, 2018), pp. 1–6.
- [70] D. P. James S. Spencer and F. Contributors, FermiNet (2020), <https://github.com/google-deepmind/ferminet>.
- [71] Y. Chen, *Poor Man's Forward Laplacian* (2024), <https://github.com/y1xiaoc/fwdlap>.
- [72] L. M. Fraser, W. M. C. Foulkes, G. Rajagopal, R. J. Needs, S. D. Kenny, and A. J. Williamson, Finite-size effects and Coulomb interactions in quantum Monte Carlo calculations for homogeneous systems with periodic boundary conditions, *Phys. Rev. B* **53**, 1814 (1996).
- [73] R. Jastrow, Many-body problem with strong forces, *Phys. Rev.* **98**, 1479 (1955).
- [74] P. Wilhelm, T. C. Lang, and A. M. Läuchli, Interplay of fractional chern insulator and charge density wave phases in twisted bilayer graphene, *Phys. Rev. B* **103**, 125406 (2021).