

Simple Fermionic backflow states via a systematically improvable tensor decomposition

Massimo Bortone^{1*}, Yannic Rath^{1,2} and George H. Booth^{1*}

¹Department of Physics and Thomas Young Centre, King’s College London, Strand, London, WC2R 2LS, United Kingdom.

²National Physical Laboratory, Teddington, TW11 0LW, United Kingdom.

*Corresponding author(s). E-mail(s): massimo.bortone@kcl.ac.uk; george.booth@kcl.ac.uk;
Contributing authors: yannic.rath@npl.co.uk;

Abstract

Strongly correlated electrons give rise to an array of electronic properties increasingly exploited in many emerging materials and molecular processes. However, the reliable numerical simulation of this quantum many-body problem still poses an outstanding challenge, in particular when accounting for the fermionic statistics of electrons. In this work, we introduce a compact and systematically improvable fermionic wave function based on a CANDECOMP/PARAFAC (CP) tensor decomposition of backflow correlations in second quantization. This ansatz naturally encodes many-electron correlations without the ordering dependence of other tensor decompositions. We benchmark its performance against standard models, demonstrating improved accuracy over comparable methods in Fermi-Hubbard and molecular systems and competitive results with state-of-the-art density matrix renormalization group (DMRG) in ab initio 2D hydrogenic lattices. By considering controllable truncations in the rank and range of the backflow correlations, as well as screening the local energy contributions for realistic Coulomb interactions, we obtain a scalable and interpretable approach to strongly correlated electronic structure problems that bridges tensor factorizations and machine learning-based representations.

1 Introduction

Computational methods that can solve the physics of strongly correlated electrons play an important role in the study of molecular and condensed matter systems, where common perturbative or empirical density functional approaches fail. In these systems, the interactions between electrons in some or all of their degrees of freedom contend with the kinetic energy of the electrons, leading to competition between localization and delocalization of the electronic structure, and the emergence of many remarkable properties and

low-energy phases. Electronic structure poses a particular challenge amongst the broader umbrella of quantum many-body problems [1] due to both charge and spin degrees of freedom, as well as the requirement for antisymmetry, which significantly complicates the form of the solution. However, these features are crucial in the understanding of the emergent physical behaviour in many technologically relevant advanced materials, from high-temperature superconductors to catalytic transition metal complexes.

A recent trend has emerged in the use of systematically improvable parameterized ansätze for

quantum states, which hold the promise of an exact limit, providing confidence and the ability to internally validate results. Naturally, the universal approximators devised in the field of machine learning (ML) have inspired ansätze for this purpose, leading to the development of Neural Quantum States (NQS) [2] with a wide range of network architectures and of differing depths and widths [3–7]. Concurrently, approaches based on kernel methods have also been considered [8, 9]. All these parameterized ansätze can in principle approximate the complex functional dependencies between the probability amplitudes of the electronic configurations, and have proven capable of obtaining accurate results with minimal user intervention over a variety of systems relying on an optimization based on the techniques of variational Monte Carlo (VMC).

It should be stressed that the use of systematically improvable ansätze in electronic structure is certainly not a new phenomenon with the emergence of NQS. One of the most successful variational methods relies on tensor network states, which provide an improvable tensor factorization of the many-body amplitudes. For a one-dimensional network, the efficient contraction of these amplitudes has led to the prominence of Matrix Product State (MPS) descriptions of correlated systems. The single (hyper)parameter controlling the expressivity of the model is the ‘bond dimension’, which can be quasi-continuously enlarged to describe higher levels of entanglement towards a complete model. Importantly, the simple structure of these states also allows for additional probes and insights into the emergent many-body physics of the model, and is able to characterize entanglement measures and structures [10, 11]. Insights into the nature of the correlations and the entanglement can be harder to quantify for NQS, where the diversity of different architectures and model parameters can also cloud a clear path towards practical improvability for the states, while it can also be unclear how to precisely design optimal parameterizations.

In this work we return to tensor factorizations to develop an alternative wave function parameterization, inspired by the developments in the class of NQS descriptions of Fermionic quantum matter. There has been much research to indicate that the considerable flexibility of

complex NQS architectures is not being fully exploited for many correlated problems, due to the challenges in their optimization and initialization within the VMC framework [12, 13]. Many simple parameterizations have performed as accurately as more complex forms, and a premium is placed on the compactness of the ansatz for ease of practical optimization alongside the overall flexibility. The simpler form for these models can also potentially provide tools for easier interpretability and improvability of the many-body physics, and open avenues to alternative optimization strategies. An example of this is the Gaussian process state (GPS), which was originally motivated via Bayesian regression as a systematically improvable kernel model with a single model parameter controlling the expressibility. It was shown to achieve similar quality results to NQS, while often being more compact and open to novel insights [9, 14, 15]. Further development exposed a duality of the GPS wave function model to an exponential of a CANDECOMP/PARAFAC (CP) tensor-rank decomposition of the wave function amplitudes in second quantization [8]. This represented an interesting simplification of the model, and suggests further developments in the use of tensor decompositions for systematically improvable descriptions of correlated states. The potential for synergies between the two domains of variational state parameterizations in NQS and tensor decompositions has in fact already motivated the development of NQS architectures that incorporate MPS parameterizations [16, 17].

This is the topic of this work where we consider a simple and systematically improvable variational quantum state based on tensor factorization, with application to general Fermionic systems, which have proven a particular challenge for NQS methods to date. In this work we consider a fixed basis and second quantization, in contrast to the real-space formulations of other Fermionic ansatz [18–20]. While this introduces a (necessarily incomplete) basis set approximation, it also allows for more flexibility in the choice of model (permutational invariance and antisymmetry are automatically enforced). This also allows for problems to be defined in a finite and discrete space for the stochastic sampling where additional approximations can be devised and chemical insights from atomic orbital correlators are easily accessible. Furthermore, this

formulation allows for a straightforward treatment of core electrons and direct comparison to established quantum chemical methods, as well as natural application to multi-resolution and quantum embedding methodologies [21–23]. The basis set approximation, in common with traditional quantum chemical methods, is also much studied with a number of approaches available which can substantially ameliorate it [24–27].

Direct application of NQS-like ansätze in second quantization has often struggled to clearly extend beyond state-of-the-art quantum chemistry, such as coupled-cluster methods (CCSD) or exact diagonalization (FCI), with results often restricted to small molecules and/or minimal basis sets [4, 5, 28, 29]. The commutation relations of second quantized operators enforced by a necessarily unphysical choice of ordering of the degrees of freedom can induce highly non-local and high-rank parity flips to the probability amplitudes. In principle, these long-range structures can be described by NQS, but in practice are very difficult to model and to appropriately optimize within VMC frameworks, which have mainly been developed for quantum spin systems [4]. As such, finding better Fermion to spin (qubit) mappings to reduce the rank or range of these non-local parity changes is an active area of research [30–33]. Alternatively, NQS-like states can be multiplied by an explicitly antisymmetric state (e.g. Slater determinant, Pfaffian or antisymmetrized geminal power) that will subsume much of the impact of these parity flips, at the cost of potentially limiting the rigorous systematic improvability of the resulting state. Nevertheless, this approach in combination with symmetry-breaking and restoration has achieved impressive results in Fermionic models [3, 8, 15].

Parallel to these developments, backflow transformations have been parameterized via neural networks as an alternative approach to describe Fermionic correlations in strongly interacting systems [34, 35]. This approach modifies the single-electron functions of a Slater determinant (or other antisymmetric function) to depend parametrically on many (potentially all N) electron coordinates in a configuration-dependent way. A closely related approach modifies the Slater determinant by coupling the physical degrees of freedom to a set of configuration-dependent auxiliary or ‘hidden’ fermions, which can similarly be

parameterized as a neural network [36, 37]. The parameterization of these backflow-type states has undergone a similar development to other classes of variational wave functions, starting initially from physically-motivated few-body parameterizations [38–41], to a more general ML architecture which allows (in principle) for systematic improvability to exactness where each orbital and electron can arbitrarily change based on all other electronic positions. These configuration-dependent orbitals in backflow states have been defined by a number of different ML architectures, both in real-space and discrete Fock space models, with and without an additional Jastrow factor in the parameterization [18, 37, 42, 43]. They were shown to be effective in describing the ground states of Fermi-Hubbard models [36, 44, 45], homogeneous electron gases [42], ultra-cold Fermi gases [46] and (primarily in first quantization) ab initio molecular systems [18–20, 47], achieving energies comparable or surpassing those from Diffusion Monte Carlo, as well as high accuracy coupled-cluster quantum chemical methods.

Here, we consider a particularly simple CP tensor rank decomposition for these configuration-dependent backflow orbitals, which allows for a straightforward yet systematically improvable form for the introduction of explicit many-body correlations into the overall state [48, 49]. We also develop a practical approach for ab initio systems to truncate the length scale of the backflow correlations, providing a further compression of the model with minimal loss of accuracy. We apply this variational ansatz to find the ground state of (doped) Fermi-Hubbard models and the water molecule, outperforming comparable neural network backflow parameterizations, as discussed in the ‘Fermi-Hubbard model’ and ‘Water molecule’ subsections of the Results and Discussion. In common with other studies, we find increasing the sampling of the VMC optimization important to improve results, indicating that despite the simple form of the state it is still challenging to optimize to the expressibility limit of the ansatz [4, 47]. In ‘Towards hydrogen materials’, we consider a 6×6 2D lattice of hydrogen atoms as a step towards extended systems, with the CPD backflow state comparing favourably to state-of-the-art density matrix renormalization group (DMRG) calculations and significantly beyond the scope

of exact approaches. Finally, in ‘Scaling’ we discuss the computational scaling of the method and approaches to reduce this as an outlook towards larger systems and widespread application.

2 Methods

2.1 Backflow determinants via CP tensor-rank decomposition

The wave function for a system of N interacting electrons can be defined in first quantization by assigning a unique label to each electron, and introducing real-space $\mathbf{r}_\alpha \in \mathbb{R}^3$ and spin $\sigma_\alpha \in \{\uparrow, \downarrow\}$ coordinates, so that $\Psi(\mathbf{x}) = \Psi(\mathbf{x}_1, \dots, \mathbf{x}_N)$, with $\mathbf{x}_\alpha = (\mathbf{r}_\alpha, \sigma_\alpha)$. The simplest wave function that satisfies the required antisymmetry is a Slater determinant of N single-particle spin-orbitals, $\phi_i(\mathbf{x}_\alpha)$:

$$\begin{aligned} \Phi_0(\mathbf{x}_1, \dots, \mathbf{x}_N) &= \frac{1}{\sqrt{N!}} \begin{vmatrix} \phi_i(\mathbf{x}_1) & \phi_j(\mathbf{x}_1) & \cdots & \phi_k(\mathbf{x}_1) \\ \phi_i(\mathbf{x}_2) & \phi_j(\mathbf{x}_2) & \cdots & \phi_k(\mathbf{x}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_i(\mathbf{x}_N) & \phi_j(\mathbf{x}_N) & \cdots & \phi_k(\mathbf{x}_N) \end{vmatrix}, \\ &\quad (1) \\ &= \mathcal{A}[\phi_i(\mathbf{x}_1)\phi_j(\mathbf{x}_2)\dots\phi_k(\mathbf{x}_N)], \end{aligned} \quad (2)$$

where $\mathcal{A}$ antisymmetrizes and normalizes the subsequent product of orbitals with respect to exchange of their arguments. We can consider these single-particle (molecular) orbitals as linear combinations of an underlying basis (e.g. atomic orbitals, AOs) $\chi_\mu(\mathbf{r})$, as

$$\phi_i(\mathbf{r}) = \sum_{\mu=1}^L \varphi_{\mu i} \chi_\mu(\mathbf{r}), \quad (3)$$

where L is the size of this basis and $\varphi_{\mu i}$ are the coefficients of the linear combination.

The key-idea of backflow ansätze is to extend the Slater determinant by generalizing the single-particle orbitals to functions with non-linear parametric dependencies on all electron coordinates. Historically this meant transforming the electron coordinates $\mathbf{r}_\alpha$ with a new set of coordinates $\mathbf{r}_\alpha^{bf} = \mathbf{r}_\alpha + \sum_{\beta \neq \alpha} \eta(|\mathbf{r}_\beta - \mathbf{r}_\alpha|)(\mathbf{r}_\beta - \mathbf{r}_\alpha)$, where the function $\eta(\mathbf{r})$ describes the effective displacement of the α electron due to the instantaneous position of the other electrons [50]. This

configurational-dependence on all other electron positions can also be directly encoded into the variational parameters of a linear expansion of single-particle orbitals, as first introduced by Ref. [34] for lattice models, yielding a new set of backflow orbitals $\phi_i^{bf}(\mathbf{r}_\alpha; \{\mathbf{r}_{/\alpha}\})$. In an effort to improve the systematic description of these configuration-dependent backflow orbitals, recent work has proposed to model $\phi_i^{bf}(\mathbf{r}_\alpha; \{\mathbf{r}_{/\alpha}\})$ using neural networks [18, 20, 44, 51], ensuring that these functions are invariant under permutation of the electron labels in $\{\mathbf{r}_{/\alpha}\}$ to retain overall antisymmetry of the state.

Within a second quantization representation, the permutational invariance of electrons and antisymmetry of the state is automatically ensured by the action and commutation relations of the second quantized operators, independent of the ansatz chosen. A Slater determinant can thus be obtained from the vacuum state $|0\rangle$ by creating N electrons in the corresponding single-particle orbitals as

$$|\Phi_0\rangle = \prod_{i=1}^N \hat{c}_i^\dagger |0\rangle = \prod_{i=1}^N \left(\sum_{\mu=1}^L \varphi_{\mu i} \hat{c}_\mu^\dagger \right) |0\rangle, \quad (4)$$

where $\hat{c}_\mu^\dagger$ ($\hat{c}_\mu$) is now the operator that creates (annihilates) an electron in the μ -th basis state. To model electron correlation, Eq. 4 can now be straightforwardly extended via analogy to the backflow transformations by including in each orbital a parametric dependence on the full instantaneous orbital occupation vector, $\mathbf{n} = (n_1, \dots, n_L)$, where n_μ indexes instantaneous occupancy of the four Fock states of spin- $\frac{1}{2}$ fermions in the chosen orthonormal representation of degree of freedom μ . This modifies the creation operator of orbital i to be

$$\hat{c}_i^\dagger(\mathbf{n}) = \sum_{\mu=1}^L \varphi_{\mu i; \mathbf{n}} \hat{c}_\mu^\dagger, \quad (5)$$

resulting in an exact model, as each orbital can vary independently according to the instantaneous occupation over the full state. However, it is of limited use as it is an over-parameterization of the full state, with an exponential number of variables. We therefore consider a specific tensor-rank decomposition, the Canonical Decomposition (CANDECOMP) or Parallel Factor (PARAFAC)

decomposition (CPD)[48, 49]. This allows for a systematic and improvable decomposition of this tensor for each orbital into a polynomial and low-rank form that is independent of the choice of ordering of the degrees of freedom defining the occupation vector, $\mathbf{n}$. The CP decomposition factorizes the occupation number vector over all states of Eq. 5 into a sum of M tensor products, with each term in the product depending on each degree of freedom in the full occupation number vector, as

$$\varphi_{\mu i; \mathbf{n}}^{\text{CPD}} = \sum_{m=1}^M \prod_{\nu=1}^L \epsilon_{\mu i; n_{\nu} \nu m}. \quad (6)$$

We now have a polynomially complex tensor of variational parameters for each orbital, $\epsilon_{\mu i; n_{\nu} \nu m}$, which encodes the correlation-driven modifications to orbital i for the specific occupied degree of freedom μ , based on the fact that state ν has a local occupation of n_{ν} . M represents an improvable parameter describing the systematic coupling of the occupations across all possible occupation strings, providing an increasingly flexible description of higher-rank correlations in the state towards exactness. We denote this single parameter controlling the flexibility of the model as its ‘support dimension’, by analogy with the CP decomposition within Gaussian process states and kernel model definitions of quantum states [8, 9, 14]. This CP decomposition splits the L -dimensional indices indicating the $\mathbf{n}$ -dependence of the orbital into a sum of products of rank-3 tensors, depending on each orbital and its occupation. Since this is a simple product rather than matrix product, there is no change in the flexibility of these backflow orbitals with the ordering of the degrees of freedom, ensuring that there should be no explicit dependence on this choice (as found in tensor network states) or dimensionality of the system.

The proposed ‘CPD’ backflow wave function is obtained by replacing the orbitals of the Slater determinant in Eq. 4 by those of Eq. 6, giving an explicitly antisymmetric state where all orbitals depend on the instantaneous occupation of all degrees of freedom,

$$|\Psi^{\text{CPD}}\rangle = \sum_{\mathbf{n}} \Psi^{\text{CPD}}(\mathbf{n}) |\mathbf{n}\rangle, \quad (7)$$

with

$$\Psi^{\text{CPD}}(\mathbf{n}) = \mathcal{A}[\varphi_{\mu_1 1; \mathbf{n}}^{\text{CPD}} \varphi_{\mu_2 2; \mathbf{n}}^{\text{CPD}} \cdots \varphi_{\mu_N N; \mathbf{n}}^{\text{CPD}}], \quad (8)$$

where the antisymmetrizer acts with respect to the N occupied orbitals of the configuration $\mathbf{n}$, given by $\mu_1, \mu_2 \dots \mu_N$. This model can be evaluated naively via building a matrix and computing a determinant in $\mathcal{O}[N^2 ML + N^3]$ cost, with each orbital evaluated according to Eq. 6. However, for low-rank changes to $\mathbf{n}$ where only $\mathcal{O}[1]$ orbital occupations change, a fast updating scheme can be devised to reduce the scaling in the matrix build by a factor of L . The update for each orbital in Eq. 6 can be found in $\mathcal{O}[M]$ time by dividing out contributions from the previous occupations and multiplying by the new occupations, analogous to the approach in Ref. [15]. Since all configurational updates in VMC can be formulated in this way, the evaluation of configurational amplitudes of this CPD state can be reduced to $\mathcal{O}[N^2 M + N^3]$.

The total number of variational parameters in this state (which in this work are all real) is therefore $\mathcal{O}[4L^2 NM]$, where L is the size of the underlying basis, N is the number of electrons, and M the ‘support dimension’ controlling the flexibility of the model. This scaling in terms of the evaluation of the model and number of parameters allows the standard techniques of VMC to be used for its sampling, optimization and extraction of observables. We note that extending this CPD form to an explicit antisymmetrization of geminal two-particle states within a Pfaffian or antisymmetrized geminal power rather than single-particle orbitals of a determinant would also be possible in this framework and will be explored in the future [46, 51].

Unless otherwise indicated, in this work, we conserve a definite spin-polarization quantum number for each of the N orbitals labelled $1, 2, 3, \dots, N$ in the product in Eq. 8, in which case the overall state must conserve $\hat{S}_z$ symmetry. This is ensured by only allowing spin-orbital degrees of freedom with the same spin-polarization to be included in the expansion coefficients of the orbital (i.e. the μ labels in Eq. 6). This allows the state for each $\mathbf{n}$ to factorize into a product of spin-up and spin-down determinants (which are allowed to independently optimize, analogous to an ‘unrestricted’ single determinant). An alternative (which is considered in the ‘Fermi-Hubbard

model’ subsection of the Results and Discussion) is to form a ‘generalized’ determinant by allowing spin-polarization to mix in each orbital definition, formally breaking $\hat{S}_z$ symmetry in the state. This symmetry is nevertheless restored via the sampling of configurations with definite $\hat{S}_z$ in the Markov chain during the VMC procedure. This $\hat{S}_z$ symmetry-breaking and projective restoration can improve results by allowing further flexibility in the state, but increases the cost in the evaluation of the determinant defining the amplitude by a factor of eight, and doubles the number of parameters (as μ labels spin-orbitals, not spatial orbitals). Importantly, regardless of whether this spin symmetry is broken or not in the orbital definition, the backflow correlations act both for same-spin and opposite-spin correlations, with the orbital dependence in Eq. 6 running over the spin-full occupations of all other degrees of freedom, n_ν , ensuring that spin-dependent correlated physics is captured.

Finally, we note that, although the functional form of the configuration-dependent orbitals of Eq. 6 is linear in M , it does not reduce to an uncorrelated determinant in the limit of $M = 1$. Correlated physics such as that captured via Gutzwiller or Jastrow correlators are included even in this limit, since the dependence between the instantaneous occupation of sites μ and ν can be independently addressed in a product form. Indeed, full non-trivial N -body correlations are included even at $M = 1$, as the exponentially large sum of products of these orbitals formed from the determinant in Eq. 8 builds in an exponential sum of these N -fold products of variational parameters for each orbital. This results in an expressive state even for very low M , which is systematically improvable to exactness as M is increased.

Two previous approaches to combine tensor decompositions with backflow parameterizations have been considered in the literature previously; a backflow-inspired extension of the MPS ansatz [52] to build non-local entanglement beyond the native MPS tensor ordering constraints for spin systems, as well as a fixed tensor representation of two-body Fermionic backflow form [45]. The latter study did not factorize these backflow correlations and was constrained to a two-body form for these correlations. In contrast, the CP decomposition of the backflow parameterization introduced in this work overcomes these

issues, providing a simple and improvable form for arbitrary rank correlations which is invariant to orbital ordering. We consider the expressibility of these states from an formal perspective further below, separating the practical challenges associated with the faithful optimization of these states in correlated systems.

2.2 Universality of the CPD backflow ansatz

In this section, we consider the formal universality of the proposed CPD backflow ansatz, where ‘universality’ in this context refers to the ability to describe any antisymmetric state within the defined Hilbert space of the problem. The universality is simple to prove, and directly stems from the universality of the CP decomposition. This means that in the large- M limit, the CP decomposition employed to model the orbitals $\varphi_{\mu i; \mathbf{n}}^{\text{CPD}}$ according to Eq. (6), can be chosen such that they are allowed to vary independently for each many-electron configuration $\mathbf{n}$. This limit also implies a mapping between basis states (Slater determinants) and wavefunction amplitudes after anti-symmetrization which can represent any anti-symmetric state within the Hilbert space defined for the problem without approximation error, according to the definition of the CPD backflow ansatz,

$$\Psi^{\text{CPD}}(\mathbf{n}) = \mathcal{A}[\varphi_{\mu_1 1; \mathbf{n}}^{\text{CPD}} \varphi_{\mu_2 2; \mathbf{n}}^{\text{CPD}} \cdots \varphi_{\mu_N N; \mathbf{n}}^{\text{CPD}}]. \quad (9)$$

This however formally requires M to scale with the number of many-body configurations in the space (as expected for any exact parameterization), as we expand on below.

To show this and make contact with other forms of parameterized states, we consider the subset of CPD states in which the backflow (many-electron) orbitals $\varphi_{\mu i; \mathbf{n}}^{\text{CPD}}$ can factorize into a term which is independent of the specific configuration (an $\mathbf{n}$ -independent ‘static’ molecular orbital), and a term which is independent of the site index, but yet can depend on a CP decomposition of the specific many-electron configuration. This can be written as

$$\varphi_{\mu i; \mathbf{n}}^{\text{factored-CPD}} = \varphi_{\mu i} \times \left(\sum_{m=1}^M \prod_{\nu=1}^L \epsilon_{i; n_\nu \nu m} \right) \quad (10)$$

With this construction, a product of CP decompositions can be factored out of the determinant, bringing the backflow ansatz into the form of a Slater-Jastrow wavefunction,

$$\Psi(\mathbf{n}) = \left(\sum_{m=1}^M \prod_{\nu=1}^L \epsilon_{1;n_\nu \nu m} \right) \left(\sum_{m=1}^M \prod_{\nu=1}^L \epsilon_{2;n_\nu \nu m} \right) \cdots \left(\sum_{m=1}^M \prod_{\nu=1}^L \epsilon_{N;n_\nu \nu m} \right) \times \mathcal{A}[\varphi_{\mu_1 1} \varphi_{\mu_2 2} \cdots \varphi_{\mu_N N}]. \quad (11)$$

In this, a Slater determinant common to all configurations (defined by the N orbitals labelled $1, 2, \dots, N$) is multiplied by a product of N CP decompositions, each of which depends on the local occupations of each site (n_ν). This product of CP decompositions takes the place of the Jastrow factor. Assuming that no two rows of the configuration-independent Slater determinant orbital matrix, $\varphi_{\mu i}$, are linearly dependent such that the determinant always evaluates to a non-zero value [36], the universal approximator property of the CP decompositions in the prefactor allows this ansatz to define a one-to-one mapping from each many-electron configuration to an arbitrary wavefunction amplitude.

2.3 The CPD backflow ansatz in Slater-Jastrow form

This approach to factoring out a CPD decomposition from the Slater determinant shows that we formally require a support dimension M which scales as the size of the Hilbert space for a universal approximator, and is therefore of little practical use. Nonetheless, the representation according to Eq. (11) still provides insights into the ability of the ansatz to represent electronic quantum states of interest at smaller support dimensions, M . From the consideration of the restricted version of the CPD state in this Slater-Jastrow form as shown in Eq. 11, we find that $M = 1$ is sufficient to represent any single Slater determinant within the given basis, as well as a site-dependent penalty function depending on its local occupation. This encapsulates physically-relevant electronic correlation beyond the mean-field picture. As a specific example, we can consider a parameterization of a

Gutzwiller factor of

$$\epsilon_{i;n_\nu \nu m} = \begin{cases} e^{g_\nu} & \text{if } i = 1 \text{ and } n_\nu \equiv \uparrow\downarrow \\ 1 & \text{otherwise} \end{cases}, \quad (12)$$

where $n_\nu \equiv \uparrow\downarrow$ indicates a double occupancy of the ν^{th} site [53, 54]. This modulates the Slater determinant with a factor depending on the double occupancy of the sites in each configuration, $\Psi(\mathbf{n}) \sim e^{\sum_\nu g_\nu n_{\nu, \uparrow} n_{\nu, \downarrow}} \times \mathcal{A}[\varphi_{\mu_1 1} \varphi_{\mu_2 2} \cdots \varphi_{\mu_N N}]$, with parameters g_ν . General forms for the $\epsilon_{i;n_\nu \nu m}$ parameters even at $M = 1$ will however also admit factorized non-local dependence on the site occupations beyond Gutzwiller form.

In general, this simple factorization of the CPD state only represents a small subset of the parametrizations possible, which has a significantly larger variational flexibility even for $M = 1$. This is because the factorization into a site-dependent term and Slater determinant does not need to be imposed, allowing a non-trivial coupling between these ‘orbital’ and ‘site’ effects at the level of the ansatz. This enlarges the span of states accessible within this decomposition at small M , and can therefore outperform ‘Slater-Jastrow’ type factorizations where the Jastrow is taken to have a flexible form, such as those previously considered within the GPS family of states [15].

2.4 Initialization

A practical bottleneck in working with parameterized quantum states with many variational parameters can often be their reliable stochastic optimization. This can be particularly sensitive to the initialization of the state, since random initialization of the parameters does not always guarantee a good overlap with the ground state, which can slow down or even prevent the optimization from converging to the true ground state. The simple functional form of the CPD backflow orbitals in Eq. 6 allows for a straightforward and effective initialization of the variational parameters, without the requirement for pre-training [18, 47]. Specifically, the tensor $\epsilon_{\mu i; n_\nu \nu m}$ can be initialized to ensure that the CPD wave function exactly spans a given single determinant such as that found from a prior mean-field solution. For most practical cases, this provides a good starting point for the VMC optimization.

In this work, we initialize from a restricted Hartree–Fock state, extracting the molecular orbital coefficients $\varphi_{\mu i}^{\text{HF}}$ in the basis in which the state is to be sampled. The CPD variational parameter tensor can then be initialized as follows:

$$\epsilon_{\mu i; n_\nu \nu m} = \mathcal{N}(0, \sigma) + \begin{cases} \varphi_{\mu i}^{\text{HF}} & \text{if } m = 1 \text{ and } \nu = 1, \\ 1 & \text{if } m = 1 \text{ and } \nu > 1 \\ 0 & \text{if } m > 1 \text{ and } \nu \geq 1 \end{cases} \quad (13)$$

where $\mathcal{N}(0, \sigma)$ is a random number drawn from a normal distribution with standard deviation σ . This small amount of random noise is optional, but is added to the initialization in case the Hartree–Fock solution is too close to a local minimum of the optimization surface.

2.5 Backflow truncation via exchange cutoff

While the CPD backflow state only has a polynomial number of parameters, the $\mathcal{O}[4L^2NM]$ scaling is still significantly higher than the native non-backflow (e.g. GPS) state, and there are significant benefits in attempting to reduce this further with a controllable compromise on the flexibility of the state. Largely redundant parameters in VMC add to statistical noise without improving accuracy and can be particularly deleterious in the optimization of the state [55, 56]. In particular, the scaling with respect to the underlying basis size (L) is quadratic in the CPD state, and in this section we motivate a physical and black-box truncation of this scaling to further improve the overall performance of the state and enable access to larger systems.

We do this by restricting the number of degrees of freedom that the backflow parameterization considers for each μ -indexed site, reducing it from L to a new parameter K . This can be motivated as a range-truncation of the backflow correlations, as has also been considered in other truncated expansions [35]. If application of this methodology was purely to local lattice models, then strictly truncating by a distance criteria would likely be sufficient to capture the dominant correlations. However, we intend the methodology to be applied equally across lattice models and ab initio systems and therefore seek an alternative proxy to define the choice of entangled orbital subspace in which these backflow correlations are defined for each

degree of freedom. This is because an ab initio basis will necessarily be extended in space and perhaps not even able to be uniquely associated with an atomic centre. Additionally, the inclusion of the long-range Coulomb interaction in these systems does not necessarily favour purely distance-based criteria. We therefore take inspiration from ab initio formulations of the Density Matrix Renormalization Group (DMRG), where heuristics for the entanglement between two orbitals are necessary in order to find an approximately optimal ordering of the extended orbitals for an effective MPS ansatz. While there are a number of options in the literature, it has been found that the importance of one orbital in describing the dominant correlations with another can be reasonably quantified by the magnitude of the exchange integral between them [57], as

$$\mathcal{K}_{\mu\nu} = \int \int d\mathbf{r}_1 d\mathbf{r}_2 \chi_\mu^*(\mathbf{r}_1) \chi_\nu^*(\mathbf{r}_1) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \chi_\mu(\mathbf{r}_2) \chi_\nu(\mathbf{r}_2). \quad (14)$$

This exchange-based metric should decay exponentially between localized orbitals, tending towards a flexible locality based truncation in the limit of fully local orbitals, while including the full range of the Coulomb interaction in the kernel. More rigorous definitions of entanglement between orbitals such as their mutual information (pair entanglement entropy) [58] could also be used, but require an initial correlated level of theory on which to build these metrics. Since we initialize the CPD backflow molecular orbitals from a Hartree–Fock calculation, the exchange matrix $\mathcal{K}_{\mu\nu}$ is readily available for no additional cost. The set of K most entangled orbitals for each orbital $\chi_\mu(\mathbf{r})$ according to this metric are selected, defining an $L \times K$ lookup table which maps to the relevant orbital indices $x_{\mu\nu} \in \{1, \dots, L\}$. The choice of orbitals in the CPD decomposition of Eq. 6 are therefore restricted as

$$\varphi_{\mu i; \mathbf{n}} = \sum_{m=1}^M \prod_{\nu=1}^K \epsilon_{\mu i; n_{x_{\mu\nu}} \nu m}, \quad (15)$$

thus reducing the number of variational parameters to $\mathcal{O}(LKMN)$ and formally linear with the size of the system, assuming that K is sufficiently large to capture the range of correlations around each degree of freedom. As K tends to L , the state




Fig. 1: Electron density of orbital subspaces selected via exchange truncation. The $K = 5$ orbital subspace (red) is shown for a central atom in a 6×6 lattice of hydrogen atoms in a STO-6G basis at both compressed (1.0 Å) and extended (2.0 Å) geometries.

returns to the original definition (albeit with an inconsequential reordering of sites in the backflow) giving the full flexibility of backflow correlations.

To illustrate the action of this exchange cutoff heuristic, in Figure 1 we consider the electron density of the $K = 5$ most entangled orbitals about a specific atom for a 6×6 square grid of ab initio hydrogen atoms in a Boys localized basis [59] at two different interatomic distances, d . This truncation is used later for numerical results in the ‘Towards hydrogen materials’ subsection of the Results and Discussion to assess the accuracy of the truncation scheme. A choice of $K = 5$ respects the local symmetries of each atom, as it enables each atom to be explicitly correlated via the backflow transformations with its four nearest neighbour atoms. As hoped, we find that the exchange cutoff protocol described automatically performs this selection of the nearest-neighbor atomic-localized orbitals around the chosen hydrogen atom in both geometries considered, providing a black-box metric to select the backflow subspace of correlations for each orbital via exploitation of locality of these correlations. We note again that the product structure of the CPD ansatz will build longer-ranged and higher-rank correlations outside the chosen subset implicitly, albeit no longer explicitly for each orbital independently.

3 Results and Discussion

In all results below we initialize the CPD backflow state from the restricted Hartree–Fock solution as outlined in Eq. 13, with a noise scale value $\sigma = 0.01$. We optimize parameters using the Stochastic Reconfiguration (SR) method [60], and when the number of parameters is larger than

that of the samples, we take advantage of the recently introduced kernel formulation from Ref. [61] to improve the computational cost of the optimization, as outlined further in the ‘Scaling’ subsection.. On the Fermi-Hubbard model and the water molecule, we found that a SR optimizer with RMSProp momentum regularization, as introduced in Ref. [62], outperforms standard SR, and we therefore use this optimizer for the results presented in the ‘Fermi-Hubbard model’ and ‘Water molecule’ subsections. The final energies presented in the results are computed as averages over 50 independent energy evaluations with the final optimized parameters and a large sample size (2^{16} for the Fermi-Hubbard model and the water molecule, and 2^{14} for the 6×6 hydrogen lattice). Error bars are computed as the standard error of these independent energy evaluations.

The VMC calculations are implemented in the NetKet package [63, 64], which we interface with our own plugin module, GPSKet for the required custom functionality. For ab initio systems, Hartree–Fock orbital coefficients and Hamiltonians are supplied from PySCF [65, 66].

3.1 Fermi-Hubbard model

While the main ambition of this work is to apply the newly developed CPD backflow ansatz to ab initio systems, we first consider a small Fermi-Hubbard model on a 2D square lattice as a prototypical system for strongly correlated electrons, where comparison to exact results and neural-network parameterized backflow states from the literature are both available. The Hamiltonian for this system is defined as

$$\hat{H} = -t \sum_{\langle i,j \rangle, \sigma} \hat{c}_{i,\sigma}^\dagger \hat{c}_{j,\sigma} + U \sum_i \hat{n}_{i,\uparrow} \hat{n}_{i,\downarrow}, \quad (16)$$

where $\hat{c}_{i,\sigma}^\dagger$ ($\hat{c}_{i,\sigma}$) is the operator that creates (annihilates) a fermion with spin σ on site i , $\hat{n}_{i,\sigma} = \hat{c}_{i,\sigma}^\dagger \hat{c}_{i,\sigma}$ is the number operator, t is the hopping amplitude, and U is the on-site interaction strength. We apply the CPD backflow state, allowing for spin-polarization breaking and restoration of the orbitals as described in Methods, in the strong interaction regime at $U/t = 8$ on a 4×4 lattice with periodic boundary conditions, at half-filling ($n = N/L = 1.0$) and in the hole doped case ($n = 0.875$). This hole-doped case is of particular




Fig. 2: Performance of the CPD and neural network backflow on the Fermi-Hubbard model. Percentage relative energy error for the ground state of the 4×4 square Fermi-Hubbard model at $U/t = 8$ compared to exact diagonalization results [69] at (a) a hole-doped filling of $n = 0.875$ and (b) half-filling. CPD backflow results (Ψ^{CPD}) are shown as a function of the configurational sample size in the optimization of the parameters, for two different model complexities of $M = 1$ (blue circles) and $M = 2$ (green circles). Neural network backflow results (NNB, red dashed lines) are taken from Ref. [44]. CPD backflow energies are obtained as averages over 50 independent evaluations using the optimized parameters and a sample size of 2^{16} , with error bars represented by the standard error across these evaluations.

interest as the point at which superconductivity and striped orders strongly compete and is much debated in the literature to date [67, 68]. We compare our results with those obtained by backflow ansätze based on neural networks (NNB) with similar numbers of parameters ($\sim 35,000$) taken from Ref. [44] as well as exact diagonalization (ED) [69].

In Fig. 2 we show the percentage relative energy error compared to exact diagonalization for the CPD state of this system, plotted against the number of samples used in the Markov chain for each update of the parameters in the SR steps. Our results significantly improve upon the comparable published neural-network backflow results for this system, even when these are extrapolated with respect to the complexity of the network

architecture in the NNB ansatz. We find percentage relative errors as low as 0.5% for the doped case and 0.1% for the half-filled case, which is competitive and within the scatter of other state-of-the-art techniques in the literature for this correlation regime [70], albeit with this system too small to be compared in the thermodynamic limit.

We also show the variational improvability as the support dimension M of the CPD decomposition is increased from $M = 1$ to $M = 2$, with a systematic lowering of all energies found, leading to a maximum of 65,536 parameters. Nevertheless, we unfortunately find that it is still generally more advantageous to increase the number of configurational samples in the Markov chain than to formally increase the flexibility of the state by increasing M . This is due to noise in the estimates of the expectation values required for the optimization of the CPD parameters, which amongst other things affects the inversion of the sampled quantum geometric tensor. This indicates that we cannot be confident of a complete optimization to the global minimum of this state, despite the simple parameterization of the CPD form, with the optimization still limited more by noise in the samples than flexibility in the model, as found in many other studies of comparable states. We will consider this behaviour more in the following section, but note that emerging optimization approaches, such as the SPRING algorithm [71], will be able to be transferred to this setting and hold promise to boost the resulting performance of the CPD backflow state. However, despite these current limitations we do find a reliable a systematic improvement in the optimized state as the number of samples is increased, and a high level of accuracy overall for this correlated state.

3.2 Water molecule

We now consider ab initio molecular systems, which are described in second quantization by an electronic Hamiltonian of the form

$$\hat{H} = \sum_{ij,\sigma} h_{ij}^{(1)} \hat{c}_{i,\sigma}^\dagger \hat{c}_{j,\sigma} + \frac{1}{2} \sum_{ijkl,\sigma\tau} h_{ijkl}^{(2)} \hat{c}_{i,\sigma}^\dagger \hat{c}_{j,\tau}^\dagger \hat{c}_{l,\tau} \hat{c}_{k,\sigma}, \quad (17)$$

where the sums run over the degrees of freedom in the system and σ, τ are binary spin variables. The $h_{ij}^{(1)}$ matrix elements describe the kinetic energy

operator and interaction with the external potential in these degrees of freedom, while the $h_{ijkl}^{(2)}$ terms model the Coulomb interaction between particles. Compared to Fermi-Hubbard models, the computational complexity of these Hamiltonians is significantly increased by the $N^2(2L - N)^2$ scaling of the connected configurations required in evaluating the local energy (compared to $\mathcal{O}[N]$ terms in Hubbard and other lattice models). Since the evaluation of the CPD wave function model at each configuration is $\mathcal{O}[N^2M + N^3]$ with the fast update (see Methods) this constrains the number of configurational samples that can be afforded.

For our initial benchmark system, we consider the water molecule in the 6-31G basis set at the equilibrium geometry used in Ref. [4]. While this seems an unassuming system from an electron correlation perspective, it has emerged as somewhat of a benchmark system in the Quantum Monte Carlo (QMC) community, where it has been studied extensively using a variety of ansätze [55, 72, 73]. Recently developed NQS architectures have struggled to reach state-of-the-art accuracy for this molecule, despite it still being of a size where exact diagonalization is possible. Part of the issue with this comes from the fact that the weakly correlated physics and compact nature of the molecule mean that it is hard to define an appropriate representation for the basis which can enable efficient, faithful and representative sampling of the state with few configurational samples.

Minimizing the number of samples relies on finding a representation which can be faithfully approximated by a small stochastic selection of configurations, necessitating an orbital representation of the basis in which the wave function amplitudes are as flat as possible throughout the Hilbert space. This maximizes the acceptance rates of the Metropolis-Hastings Markov chain growth, and ensures that as small a sample as possible can represent the wave function distribution. Canonical bases of mean-field (e.g. Hartree-Fock) theories are therefore particularly poorly suited, as they are (away from very strong correlation) dominated by the configuration of a single Slater determinant. These bases have been found for NQS with restricted Boltzmann machine architectures to obtain relatively large correlation energy errors ($\approx 5 - 10\%$) despite scaling up to

10^6 configurational samples [4]. The development of autoregressive NQS models has been able to improve upon this by allowing a direct sampling algorithm of unique configurations that is not constrained by the limitations of the Metropolis-Hastings algorithm [5]. However, these models still require large sample sizes and have only been benchmarked in a STO-3G minimal basis set for this system [5, 28, 29]. Rather than changing the sampling algorithm, in a previous work, we considered the effect of different orbital representations for the configurational basis [15]. Following this, we consider orthogonal Foster-Boys orbitals for the configurations, localized over all degrees of freedom to minimize the physical spread of the resulting orbitals [59].

The results in Fig. 3(a) show that the CPD backflow ansatz formulated in this local basis exhibits a clear systematic improvability, with the error decreasing inversely with the support dimension of the model, M . We can use this empirical scaling to extrapolate the results to the infinite support dimension limit, which results in a relative correlation energy error of below 2%, for the ansatz optimized with $\mathcal{O}[10^4]$ configurational samples. At infinite M the model is complete, and the error therefore must arise from the incomplete optimization of the finite- M models. We therefore also consider the improvability in M for two different numbers of configurational samples, N_S in the Markov chains used for each optimization step. The $1/M$ decay of the error is clearly seen in both of these sample sizes, with the extrapolated model result decreasing towards exactness for increasing N_S .

We analyse this trend more systematically in Fig. 3(b), where we show the convergence in N_S for two different model complexity parameters M , showing a relatively robust $N_S^{-\frac{1}{2}}$ scaling in the error. This indicates that doubling the support dimension has a similar effect on reducing the error as quadrupling N_S . This robustness and reliability in the error reduction is to be expected with increasing M , but is more surprising with increasing N_S . It indicates that the noise introduced into the sampling at finite N_S values is not simply changing the variance in the resulting energy, or indeed resulting in different optimized states due to convergence to different local minima in the landscape (where we would expect a wider scatter




Fig. 3: Systematic improvability of the CPD backflow ansatz for the water molecule. Relative ground state correlation energy error (compared to exact diagonalization) for the CPD backflow ansatz (Ψ^{CPD}) on the water molecule (6-31G basis, equilibrium geometry as specified in Ref. [4]) as a function of (a) support dimension M and (b) number of configurational samples N_S . We also extrapolate these results to infinite M or N_S to provide the values in their infinite limit as shown by diamonds. CPD backflow energies are obtained as averages over 50 independent evaluations using the optimized parameters and a sample size of 2^{16} , with error bars represented by the standard error across these evaluations.

of optimized energies). Instead, the robustness and systematic trend in the results indicates that N_S is controlling the intrinsic error of the optimization of the state in a more systematic fashion. This could potentially arise from non-linear steps in the optimization protocol, and is something which requires further scrutiny going forwards.

Comparing the accuracies obtained to previous state-of-the-art results in Figure 4, we find that

the CPD backflow state with $M = 1$ (6.7k parameters) already outperforms both the Gaussian process state augmented by a symmetry-broken Pfaffian (585 parameters) [15] and the restricted Boltzmann machine NQS state (728 parameters) [4] when optimized with $\mathcal{O}[10^3]$ configurational samples. The accuracy is further improved when larger support dimensions and sample sizes are considered, with the CPD model at $M = 4$ (27k parameters) and $N_S \sim \mathcal{O}[10^4]$ outperforming the best NQS by 2% in the relative correlation energy error. While we still don’t quite reach the level of accuracy of coupled cluster methods with singles and doubles (and the significant ‘chemical accuracy’ hurdle – albeit defined with respect to the finite basis set energy), to the best of our knowledge, these results represent the state-of-the-art for an NQS-like variational ansatz for this system.

3.3 Towards hydrogen materials

Extending the CPD backflow ansatz beyond benchmark studies and comparison to exact results, we consider a two-dimensional ab initio lattice of hydrogen atoms as a step towards combining strong correlation, long-range interactions and extended systems. These hydrogenic systems have been studied by a variety of methods in the recent years given their simple specification and challenge of realistic interactions, whilst maintaining a close connection to the Fermi-Hubbard model [15, 74–79]. In particular, different correlation regimes can be probed by simply changing the interatomic distance of the lattice, similar to tuning the interaction strength in the Fermi-Hubbard model. However, crucially these hydrogen lattices require the accurate treatment of realistic long-range Coulomb interactions and their effects, which are not present in the Fermi-Hubbard model. Accurately capturing the ground state of these systems for different interatomic distance is thus a challenging task for most quantum chemistry methods, as it requires a flexible and expressive model with a treatment of long-range interactions and high-energy scattering physics that gives rise to states of significantly different character.

In Figure 5, we report the ground state energy per atom obtained from the CPD backflow ansätze (with and without the rank and range truncation introduced in the ‘Backflow truncation via




Fig. 4: Performance of the CPD backflow and other models on the water molecule. Relative ground state correlation energy error (compared to exact diagonalization) for CPD backflow ansätze (Ψ^{CPD}) on the water molecule (6-31G basis, equilibrium geometry as specified in Ref. [4]) as the number of configurational samples N_S in each Markov chain is increased in the parameter optimization steps. Two support dimensions corresponding to $M = 1$ (blue circles) and $M = 4$ (green circles) are shown. Comparison energies for the Gaussian process state augmented by a symmetry-broken Pfaffian ($\Psi^{GPS} \times |\text{Pf}\rangle$), red diamond) and restricted Boltzmann machine NQS (Ψ^{RBM} , orange squares) are taken from Refs. [15] and [4], respectively, while the CCSD energy (solid line) is calculated with PySCF [65, 66]. CPD backflow energies are obtained as averages over 50 independent evaluations using the optimized parameters and a sample size of 2^{16} , with error bars represented by the standard error across these evaluations.

exchange cutoff’ subsection of the Methods) for a 6×6 hydrogen lattice in a minimal basis (STO-6G) with open boundary conditions. We choose $K = 5$ for the range cutoff of the backflow to ensure that the local symmetries of quantum fluctuations about each atomic site are preserved. We consider both compact (lower effective U/t) and extended (higher effective U/t) lattice structures by varying the interatomic distances all the way to essentially dissociated non-interacting hydrogen atoms. We compare our results to energies obtained with restricted Hartree-Fock (RHF) and unrestricted coupled-cluster with single and double excitations (UCCSD), as well as an efficient ab initio




Fig. 5: Performance of the CPD backflow and other methods on the hydrogen lattice. **a** Ground state energy per atom of a 6×6 square hydrogen lattice in a STO-6G basis for increasing lattice constants. Shown are RHF (orange line) and UCCSD energies (violet line), while the horizontal grey line indicates the exact energy of the fully dissociated limit in this basis. The CPD backflow ansatz (Ψ^{CPD}) with support dimension $M = 1$ (solid black line) is compared to a GPS ansatz augmented with a Slater determinant ($\Psi^{GPS} \times |\Phi\rangle$) with support dimension $M = 72$ [15] (dotted black line), as well as energies obtained with DMRG (red dots and line) using a bond dimension of $M = 1024$. A further CPD backflow ansatz with support dimension $M = 1$ and a truncation in the backflow subspace to $K = 5$ orbitals is shown as a dashed black line. The backflow CPD (GPS augmented by Slater determinant) VMC ansätze were optimized with 4096 (10,000) samples, while DMRG results were obtained using the `block2` package [80, 81]. **b** Energy difference per atom relative to DMRG of the CPD backflow ansätze and the GPS times Slater determinant ansatz. To aid the comparison, we have removed the outlier DMRG energy at 2.5 Å. The energies of the CPD backflow and GPS times Slater determinant ansätze are obtained as averages over 50 independent evaluations using the optimized parameters and a sample size of 2^{14} , with error bars represented by the standard error across these evaluations.

implementation of density matrix renormalization group (DMRG) going up to bond dimension of 1024 in a fully spin-adapted basis implemented in the `block2` package [80, 81]. For an additional comparison between contrasting approaches to Fermionic variational wave functions, we also include the results obtained from a GPS ansatz with support dimension $M = 72$ acting as a Jastrow in front of a co-optimized Slater determinant [15], to compare the CPD backflow to this approach. We optimize the CPD backflow and GPS multiplied by Slater determinant ansätze in a Boys localized basis for the orbitals, whereas for the DMRG results we rely on a split-localized basis, in which occupied and virtual orbitals are localized separately. We found this choice to give the most consistent results for DMRG across the range of geometries studied. Each optimization of the CPD backflow wave functions took ≈ 250 GPU hours across 4 Nvidia A100 devices, whereas the DMRG runs took a total of ≈ 500 CPU hours on an Intel(R) Core(TM) i9 device.

The RHF and UCCSD description of this equation of state qualitatively breaks down quite early in this stretching coordinate, with UCCSD failing to converge beyond $1.5 \text{ \AA}$. Furthermore, the UCCSD exhibits quantitative error of $\sim 2 \text{ mE}_h$ per atom even around equilibrium geometries, confirming that substantial correlation effects are present even in this regime. As another point of reference, the fully dissociated limit can be computed via exact diagonalization, where the assumption of simple energy extensivity from a single atom can be applied. In this limit the energy is $\approx 0.03 \text{ E}_h$ above the analytic result for the hydrogen atom due to the basis set incompleteness error, which nevertheless will exhibit a large degree of cancellation for energy differences along this changing geometry. The DMRG provides the best variational comparison for this system, with (apart from $2.5 \text{ \AA}$) the CPD and GPS results being within 2 mE_h per atom of this value.

The CPD backflow ansatz manages to quantitatively capture the features of the expected potential energy surface, reaching the correct dissociation limit at large interatomic distances, and showing an overall smooth transition from weak to strong correlation regimes. We can directly compare different systematically improvable variational ansätze (DMRG, GPS multiplied by a

Slater determinant and the CPD backflow), all of which are competitive and variationally optimal at different points in the changing physics of this system. Around the equilibrium of this system, the CPD backflow and DMRG states are almost identical and variationally optimal amongst the comparison. In the intermediate regime ($1.5 \leq d < 2.0 \text{ \AA}$), the GPS ansatz augmented with a Slater determinant provides the best variational energies, despite (or perhaps because of) the smaller number of parameters ($\approx 12k$ vs. $\approx 187k$ for the CPD backflow ansatz without truncation and $\approx 26k$ for the one with). An outlier appears to be the nearly dissociated limit of $2.5 \text{ \AA}$ interatomic distance, where the DMRG energy appears erroneously high. This could be due to a particularly large impact on the one-dimensional MPS topology used, the choice of basis for the orbitals or the DMRG sweep getting stuck in a local minimum. Nevertheless, the other variational ansatz largely agree at this point.

Comparing the CPD backflow curves with and without the backflow truncation, we find (as expected) that the $K = 5$ results are all variationally higher than the parent CPD backflow. This truncation has a very small effect on the energies at larger interatomic distances, but becomes more significant around the equilibrium distance and mildly stretched geometries where it reaches a maximum error of 2 mE_h per atom. This is expected as the range of the correlations in the compressed lattice will extend further than the stretched limit. Nonetheless, even with this restriction the ansatz is able to reach the coupled-cluster level of accuracy around equilibrium, and to outperform it on stretched geometries, with a reduction in the number of parameters compared to the parent model by more than a factor of seven. This validates the exchange cutoff as a practical parameter reduction scheme for the CPD backflow ansatz, suggesting benefits in the study of larger systems, and potentially allowing for an increase in the support dimension of the model.

To further compare the physical properties of the potential energy surface of this lattice as described by the different levels of theory, we fit a simple Morse potential at different interatomic distances (r), given by

$$V(r) = D_e \left(1 - e^{-a(r-r_e)}\right)^2 + u, \quad (18)$$

Method	D_e (eV)	r_e (Å)	w_e (cm ⁻¹)	$w_e x_e$ (cm ⁻¹)
UCCSD	0.163	1.225	2177.988	9.007
DMRG	0.031	1.223	1900.038	36.409
Ψ^{CPD}	0.029	1.219	1928.083	39.672
$\Psi^{\text{CPD}}(K=5)$	0.028	1.229	1865.462	38.586
$\Psi^{\text{GPS}} \times \Phi\rangle$	0.029	1.237	1748.939	33.247

Table 1: Physical properties of the hydrogen lattice as obtained by different methods. Dissociation energy (D_e), equilibrium bond length (r_e), harmonic (w_e) and anharmonic ($w_e x_e$) frequencies in wavenumbers for the 6×6 hydrogen lattice obtained by fitting a Morse potential to data from different quantum chemistry and variational methods described in this work.

where D_e and a control the depth and width of the well, r_e is the equilibrium bond length, and u is the energy offset. Although the Morse potential is generally used for diatomic molecules, the symmetric stretching coordinate of this system is nevertheless well modelled by this form. The differences in the quantum chemistry and VMC methods used to obtain the potential energy data are reflected in the variations of dissociation energy (D_e), equilibrium bond length (r_e), harmonic vibrational frequency (w_e), and anharmonicity constant ($w_e x_e$) presented in Table 1.

UCCSD is expected to accurately describe the correlated physics near equilibrium geometries, however the rapid divergence after this point renders even the harmonic vibrational frequencies unreliable. In contrast, DMRG, which handles both the strong and weak correlations on a consistent level, presents a more accurate dissociation energy ($D_e = 0.031$ eV) and harmonic vibrational frequency ($w_e = 1900.038$ cm⁻¹), while yielding anharmonicity of the vibrational motion of the atomic lattice describing the beyond-parabolic nature of the binding as $w_e x_e = 36.409$ cm⁻¹. The values obtained from CPD backflow ansätze with and without truncation closely track those from DMRG. On the other hand, while agreeing on the dissociation energy, the GPS multiplied by Slater determinant ansatz stands out amongst the variational methods with a marginally softer bond, with a larger equilibrium lattice parameter ($r_e = 1.237$ Å) and the lowest harmonic vibrational frequency ($w_e = 1748.939$ cm⁻¹). Overall, the methods agree on an equilibrium bond length around 1.22 Å and a dissociation energy of 0.03 eV (except UCCSD), with variations in the harmonic and anharmonic wavenumbers.

These variations highlight the strengths and limitations of each method in modelling the potential energy surfaces and vibrational properties of hydrogen materials. Taking all these results into consideration, the CPD backflow ansatz emerges as a competitive method for the study of strong electron correlation, providing a variational description of the ground state of a two-dimensional lattice of hydrogen atoms that is in good agreement with other state-of-the-art methods, while being able to capture strong correlations and anharmonic effects in the system in a low-energy basis.

3.3.1 Spin-spin correlations

The local atomic basis framework of the CPD backflow ansatz allows for the straightforward computation of atom-resolved expectation values for further insights into the electronic structure. In the context of hydrogen materials, local spin-spin correlation functions are of particular interest, as they can provide insights into the nature of the ground state of the system, and the emergence of magnetic order. By analogy with Hubbard models, we would expect some anti-ferromagnetic order to emerge in the electronic structure of this system, with this order decaying algebraically in the thermodynamic limit. However, in the presence of long-range interactions this behaviour is far from confirmed in two-dimensions. While admittedly far from this thermodynamic limit, we consider the two-point spin-spin correlation function $C(r)$ between the centre of the 6×6 hydrogen lattice, and atoms at a distance r from the centre. We can define this function via instantaneous (equal-time) spin-spin correlators $\langle \hat{S}_{\vec{r}_a}^z \hat{S}_{\vec{r}_b}^z \rangle$ between two atoms as

$$C(r) = \frac{1}{N_{\text{bulk}}} \sum_{\vec{r}_a \in \text{bulk}} \sum_{|\vec{r}_a - \vec{r}_b| = r} \langle \hat{S}_{\vec{r}_a}^z \hat{S}_{\vec{r}_b}^z \rangle, \quad (19)$$

where $\vec{r}_a$ and $\vec{r}_b$ are the positions of atoms a and b , and N_{bulk} is the number of equivalent atoms in the centre that we average over (four). We use the atom-centered atomic orbitals themselves as natural projectors for the spin operators of each atom. More details about the calculation of the instantaneous spin-spin correlation function can be found in the Supplementary Note 1.




Fig. 6: Spin-spin correlations in the hydrogen lattice. Two-point instantaneous spin-spin correlation function for the ground state of the 6×6 hydrogen lattice between atoms at two different interatomic distances $d = 1.2 \text{ \AA}$ (blue) and $d = 3.0 \text{ \AA}$ (green). Correlation function is given as a function of normalized radial interatomic distance between atoms in the correlator, showing nearest, next-nearest and next-next-nearest magnetic correlations. The correlators were computed from the optimized CPD backflow ansatz with $M = 1$.

In Figure 6, we compare the radial spin-spin correlation function for the ground state approximation obtained with the CPD backflow ansatz at near-equilibrium interatomic distance $d = 1.2 \text{ \AA}$, and at a large stretching of $d = 3.0 \text{ \AA}$, normalized for the changing inter-atomic distances. We find the emergence of the short-range antiferromagnetic order in the material, as anticipated by analogy with Hubbard models. The magnitude of this antiferromagnetic order increases with increasing interatomic separation, again keeping with anticipated Hubbard behaviour of increasing U/t values. However, this order is very short ranged, with the spin in the extended lattice not directly affecting a lattice site beyond its nearest neighbours. At more compressed geometries, this order does extend beyond this to the outer atoms in the lattice (next-next-nearest-neighbours), due to the shorter distance in real space, but the overall magnitude of these magnetic correlations is reduced.

3.4 Scaling

As illustrated in the results above, the CPD backflow ansatz performs well on small Fermionic systems, but further developments for scaling to significantly larger systems are still required for this to become a clearly competitive method for the wider electronic structure community. Simplifying the scaling to assume a general growth of both the basis and electron number such that $N \sim L$ and assuming M and K are independent of system size, the parameters grow with system size as $N_P \sim \mathcal{O}[N^3]$ for the full ansatz, with the subspace truncation of the backflow reducing this asymptotically to $N_P \sim \mathcal{O}[N^2]$ (see the ‘Backflow truncation via exchange cutoff’ subsection of the Methods). The fast updating of backflow orbitals also enables the evaluation of the wave function log-amplitudes to be performed in $\mathcal{O}[MN^2 + N^3]$ (regardless of whether a backflow truncation is applied). However, we find in practice the determinant evaluation has a significantly smaller prefactor than the construction of the orbitals, so that the dominant scaling is rather $\mathcal{O}[N^2]$ for small to medium-sized systems. Given that the number of terms in general second quantized ab initio hamiltonians scales as $\mathcal{O}[N^4]$, the resulting scaling of the local energy evaluation is then $\mathcal{O}[N^7]$ for the CPD backflow state in the asymptotic limit, or $\mathcal{O}[N^6]$ for small to medium-sized systems. While this should be competitive with accurate quantum chemical methods such as coupled-cluster, it is clear that the prefactor is significantly larger.

Rather than just the local energy evaluation, we should also consider the computational scaling for the update of the parameters. For larger numbers of parameters (such as the $\sim 187,000$ of the hydrogen lattice above), their update used to be the main bottleneck for VMC large-scale ansätze when using the original SR algorithm [60]. For a model with N_P parameters, SR would scale as $\mathcal{O}(N_P^3)$, since it involves inverting the $N_P \times N_P$ quantum geometric tensor matrix. This is not the case for recently introduced alternative formulations of SR, such as minimum-step SR [82], the kernel formulation of SR [61] or SPRING [71]. In particular for minimum-step SR and the kernel formulation of SR simple linear algebra identities were used to reduce the dimension of the matrix that is inverted in the SR algorithm from

N_P to N_S , i.e. the number of samples used during the optimization. When $N_P \gg N_S$, as in large-scale models, the scaling of the parameter update becomes $\mathcal{O}(N_S^2 N_P + N_S^3)$, i.e. linear in the number of parameters. Thus, the evaluation of the local energy remains the computational bottleneck of the algorithm in the case of ab initio systems.

We show this scaling explicitly in Fig. 7, where we measure the mean runtime for a full VMC parameter update step for the CPD backflow ansatz, including the Markov chain sampling of a fixed number of configurations, evaluation of the local energy, and the subsequent parameter update. By increasing the number of hydrogen atoms in a chain with fixed equilibrium inter-atomic distances ($d = 1.68 \text{ \AA}$) in a STO-6G basis up to 90 atoms we can extract a realistic asymptotic scaling of the approach. We set the support dimension of the ansatz to $M = 1$ and choose a sample size of $N_S = 128$, in order to fit the data in memory even for the largest system sizes. For each system size, we let the VMC algorithm run for 50 iterations on a single Nvidia A100 GPU with 40GB of memory. Extracting the scaling from the large system limit gives a scaling of $\mathcal{O}[N^{6.5}]$, which is only evident for this system when we reach > 40 atoms.

As discussed in the context of the GPS model in Ref. [15] and building on other works in this area [74, 83–85], this scaling for ab initio systems can be reduced further by truncating the number of terms in the sum over connected configurations at each evaluation of the local energy. This truncation is performed on the magnitude of the Hamiltonian matrix element connecting the configurations. By presorting the electron repulsion integrals between the degrees of freedom, this truncation can be implemented without having to consider the entire set of Hamiltonian matrix elements for each evaluation of the local energy. Formally, the exponentially decreasing overlap between the orbitals in the sampled space should reduce the number of connected determinants which contribute to the local energy asymptotically to $\mathcal{O}[N^2]$ rather than $\mathcal{O}[N^4]$ – a scaling which then matches the scaling of the local energy evaluation in a first quantized perspective. This results in a practical scaling of the CPD ansatz of $\sim \mathcal{O}[N^{4-6}]$. However, since this method comes with a certain overhead in terms of data structures, the lower-bounds on this scaling only materialize after




Fig. 7: Scalability of the CPD backflow ansatz. Mean VMC step runtime for the CPD backflow ansatz as a function of the number of atoms in a hydrogen chain with fixed inter-atomic distances ($1.68 \text{ \AA}$) in a STO-6G basis. The support dimension is $M = 1$ and the sample size is $N_S = 128$. The blue points are obtained with the full Hamiltonian, while the green and orange points are obtained with the local energy evaluated with a Hamiltonian pruned with a threshold of 10^{-9} and $10^{-5} E_h$. The dashed lines show the observed scaling for the largest system sizes. Average runtimes are obtained over 50 iterations, with error bars represented by the standard error across these runtimes.

a certain crossover system size, which depends on the system and on the truncation threshold on the Hamiltonian matrix elements.

Figure 7 also shows the analogous results including this energetic threshold of 10^{-5} and $10^{-9} E_h$, with this tighter threshold expected to incur negligible change in the sampled energy for a given state. The results show that a practical crossover point, after which the pruning of Hamiltonian elements below a small threshold yields a speed-up, is reached already around system sizes as low as 20-30 electrons for this ansatz, noting that a one-dimensional chain is advantageous in terms of affecting an advantage from this approach. However, for large system sizes the speed-up is more than an order of magnitude and provides the expected asymptotic quadratic improvement, giving an overall scaling of $\mathcal{O}[N^{3-4}]$ up to ~ 100 electrons, a scaling competitive with hybrid Density Functional Theory (DFT) techniques. Overall, the combination of this ansatz

with the various scaling reductions outlined represent a real potential towards a second quantized, systematically improvable, VMC algorithm with a practical and competitive $\mathcal{O}(N^4)$ scaling for medium to large ab initio systems. Clearly, further developments for prefactor reductions are key to take advantage of this improved scaling and access these system sizes.

4 Conclusions

In this work, we introduce a general and simple ansatz suitable for ab initio fermions, based on a systematically improvable tensor rank decomposition of a general backflow form. This systematically builds configuration-dependent orbitals of a single antisymmetric Slater determinant in second quantization, directly encoding non-trivial N -body electron-electron correlations with a parameter scaling of $\mathcal{O}[N^{2-3}]$. We have shown that the ansatz can achieve competitive accuracy on small Fermionic systems, such as the Fermi-Hubbard model and the water molecule, and that it can be used to model larger strongly correlated lattices of ab initio hydrogen atoms with an accuracy comparable to state-of-the-art DMRG techniques. Finally, we have discussed the scalability of the ansatz and shown that we can affect various reductions in a practical fashion to demonstrate $\mathcal{O}(N^4)$ scaling on medium to large ab initio systems.

We are working on further improvements in the accuracy and efficiency of the ansatz, as well as taking advantage of the benefits of working in a second quantized formalism to integrate with multiscale methods and quantum embedding methodologies to provide a practical route in the modelling of truly extended systems within this CPD backflow framework [86, 86]. These techniques could also be integrated within the ‘hidden fermion’ model of correlated states as an alternative parameterization of the correlations [36, 37]. It is also natural to ask whether alternative tensor factorization techniques could be applied within the context of describing second-quantized backflow correlations. These could naturally be fitted into the framework described above, and will also be explored in the future.

Data availability

All the results presented in this work can be fully reproduced with the publicly available source code and input configurations. The data supporting this article is openly available from the King’s College London research data repository, KORDS, at <https://doi.org/10.18742/c.7699007>. Model parameters are available upon request to the corresponding author.

Code Availability

The code for this project is implemented as part of the publicly available [GPSKet](#) plugin for [NetKet](#) [63, 64]. It is made available, together with configurations files to reproduce the figures in the paper, at <https://github.com/BoothGroup/GPSKet/tree/master/scripts/cpd-backflow>.

Acknowledgments

The authors gratefully acknowledge support from the Air Force Office of Scientific Research under award number FA8655-22-1-7011. We are grateful to the UK Materials and Molecular Modelling Hub for computational resources, which is partially funded by EPSRC (EP/P020194/1 and EP/T022213/1). Furthermore, we acknowledge the use of the high performance computing environment CREATE at King’s College London [87]. YR also acknowledges the support of the Engineering and Physical Sciences Research Council (EP/Y005090/1).

Author Contributions

All authors jointly developed the methodology and wrote the manuscript. MB and YR implemented the approach and performed the numerical experiments. GHB supervised the project.

Competing interests

The authors declare no competing interests.

References

- [1] Choo, K., Neupert, T. & Carleo, G. Two-dimensional frustrated $\{J\}_1\text{-}\{J\}_2$

- model studied with neural network quantum states. *Physical Review B* **100**, 125124 (2019).
- [2] Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).
 - [3] Nomura, Y., Darmawan, A. S., Yamaji, Y. & Imada, M. Restricted Boltzmann machine learning for solving strongly correlated quantum systems. *Physical Review B* **96**, 205152 (2017).
 - [4] Choo, K., Mezzacapo, A. & Carleo, G. Fermionic neural-network states for ab-initio electronic structure. *Nature Communications* **11**, 2368 (2020).
 - [5] Barrett, T. D., Malyshev, A. & Lvovsky, A. I. Autoregressive neural-network wavefunctions for ab initio quantum chemistry. *Nature Machine Intelligence* **4**, 351–358 (2022).
 - [6] Hibat-Allah, M., Ganahl, M., Hayward, L. E., Melko, R. G. & Carrasquilla, J. Recurrent neural network wave functions. *Physical Review Research* **2**, 023358 (2020).
 - [7] Sprague, K. & Czischek, S. Variational Monte Carlo with large patched transformers. *Communications Physics* **7**, 1–11 (2024).
 - [8] Rath, Y. & Booth, G. H. Quantum Gaussian process state: A kernel-inspired state with quantum support data. *Physical Review Research* **4**, 023126 (2022).
 - [9] Giuliani, C., Vicentini, F., Rossi, R. & Carleo, G. Learning ground states of gapped quantum Hamiltonians with Kernel Methods. *Quantum* **7**, 1096 (2023).
 - [10] Orús, R. Tensor networks for complex quantum systems. *Nature Reviews Physics* **1**, 538–550 (2019).
 - [11] Eisert, J., Cramer, M. & Plenio, M. B. Area laws for the entanglement entropy. *Reviews of Modern Physics* **82**, 277–306 (2010).
 - [12] Bukov, M., Schmitt, M. & Dupont, M. Learning the ground state of a non-stoquastic quantum Hamiltonian in a rugged neural network landscape. *SciPost Physics* **10**, 147 (2021).
 - [13] Inack, E. M., Morawetz, S. & Melko, R. G. Neural Annealing and Visualization of Autoregressive Neural Networks in the Newman–Moore Model. *Condensed Matter* **7**, 38 (2022).
 - [14] Glielmo, A., Rath, Y., Csányi, G., De Vita, A. & Booth, G. H. Gaussian Process States: A Data-Driven Representation of Quantum Many-Body Physics. *Physical Review X* **10**, 041026 (2020).
 - [15] Rath, Y. & Booth, G. H. Framework for efficient ab initio electronic structure with Gaussian Process States. *Physical Review B* **107**, 205119 (2023).
 - [16] Chen, Z., Newhouse, L., Chen, E., Luo, D. & Soljacic, M. ANTN: Bridging Autoregressive Neural Networks and Tensor Networks for Quantum Many-Body Simulation. *Advances in Neural Information Processing Systems* **36**, 450–476 (2023).
 - [17] Wu, D., Rossi, R., Vicentini, F. & Carleo, G. From tensor-network quantum states to tensorial recurrent neural networks. *Physical Review Research* **5**, L032001 (2023).
 - [18] Pfau, D., Spencer, J. S., Matthews, A. G. D. G. & Foulkes, W. M. C. Ab initio solution of the many-electron Schrödinger equation with deep neural networks. *Physical Review Research* **2**, 033429 (2020).
 - [19] Hermann, J., Schätzle, Z. & Noé, F. Deep-neural-network solution of the electronic Schrödinger equation. *Nature Chemistry* **12**, 891–897 (2020).
 - [20] von Glehn, I., Spencer, J. S. & Pfau, D. A Self-Attention Ansatz for Ab-initio Quantum Chemistry (2022). [2211.13672](https://arxiv.org/abs/2211.13672).
 - [21] Nusspickel, M. & Booth, G. H. Systematic Improvability in Quantum Embedding

- for Real Materials. *Physical Review X* **12**, 011046 (2022).
- [22] Sun, Q. & Chan, G. K.-L. Quantum Embedding Theories. *Accounts of Chemical Research* **49**, 2705–2712 (2016).
- [23] Kotliar, G. *et al.* Electronic structure calculations with dynamical mean-field theory. *Reviews of Modern Physics* **78**, 865–951 (2006).
- [24] de Lara-Castells, M. P., Krems, R. V., Buchachenko, A. A., Delgado-Barrio, G. & Villarreal, P. Complete basis set extrapolation limit for electronic structure calculations: Energetic and nonenergetic properties of HeBr and HeBr₂ van der Waals dimers. *The Journal of Chemical Physics* **115**, 10438–10449 (2001).
- [25] Grüneis, A., Hirata, S., Ohnishi, Y.-y. & Ten-no, S. Perspective: Explicitly correlated electronic structure theory for complex systems. *The Journal of Chemical Physics* **146**, 080901 (2017).
- [26] Booth, G. H., Cleland, D., Alavi, A. & Tew, D. P. An explicitly correlated approach to basis set incompleteness in full configuration interaction quantum Monte Carlo. *The Journal of Chemical Physics* **137**, 164112 (2012).
- [27] Yanai, T. & Shiozaki, T. Canonical transcorrelated theory with projected Slater-type geminals. *The Journal of Chemical Physics* **136**, 084107 (2012).
- [28] Wu, Y., Guo, C., Fan, Y., Zhou, P. & Shang, H. *NNQS-Transformer: An Efficient and Scalable Neural Network Quantum States Approach for Ab initio Quantum Chemistry*, SC '23, 1–13 (Association for Computing Machinery, New York, NY, USA, 2023).
- [29] Zhao, T., Stokes, J. & Veerapaneni, S. Scalable neural quantum states architecture for quantum chemistry. *Machine Learning: Science and Technology* **4**, 025034 (2023).
- [30] Bravyi, S. B. & Kitaev, A. Y. Fermionic Quantum Computation. *Annals of Physics* **298**, 210–226 (2002).
- [31] Nys, J. & Carleo, G. Variational solutions to fermion-to-qubit mappings in two spatial dimensions. *Quantum* **6**, 833 (2022).
- [32] Nys, J. & Carleo, G. Quantum circuits for solving local fermion-to-qubit mappings. *Quantum* **7**, 930 (2023).
- [33] Derby, C., Klassen, J., Bausch, J. & Cubitt, T. Compact fermion to qubit mappings. *Physical Review B* **104**, 035118 (2021).
- [34] Tocchio, L. F., Becca, F., Parola, A. & Sorella, S. Role of backflow correlations for the nonmagnetic phase of the t-t' Hubbard model. *Physical Review B* **78**, 041101 (2008).
- [35] Tocchio, L. F., Becca, F. & Gros, C. Backflow correlations in the Hubbard model: An efficient tool for the study of the metal-insulator transition and the large- U limit. *Physical Review B* **83**, 195138 (2011).
- [36] Moreno, J. R., Carleo, G., Georges, A. & Stokes, J. Fermionic wave functions from neural-network constrained hidden states. *Proceedings of the National Academy of Sciences* **119**, e2122059119 (2022).
- [37] Liu, Z. & Clark, B. K. Unifying view of fermionic neural network quantum states: From neural network backflow to hidden fermion determinant states. *Physical Review B* **110**, 115124 (2024).
- [38] Feynman, R. P. & Cohen, M. Energy Spectrum of the Excitations in Liquid Helium. *Physical Review* **102**, 1189–1204 (1956).
- [39] Wigner, E. & Seitz, F. On the Constitution of Metallic Sodium. II. *Physical Review* **46**, 509–524 (1934).
- [40] Kwon, Y., Ceperley, D. M. & Martin, R. M. Effects of three-body and backflow correlations in the two-dimensional electron gas. *Physical Review B* **48**, 12037–12046 (1993).

- [41] Kwon, Y., Ceperley, D. M. & Martin, R. M. Effects of backflow correlation in the three-dimensional electron gas: Quantum Monte Carlo study. *Physical Review B* **58**, 6800–6806 (1998).
- [42] Pescia, G., Nys, J., Kim, J., Lovato, A. & Carleo, G. Message-passing neural quantum states for the homogeneous electron gas. *Physical Review B* **110**, 035108 (2024).
- [43] Romero, I., Nys, J. & Carleo, G. Spectroscopy of two-dimensional interacting lattice electrons using symmetry-aware neural backflow transformations (2024). [2406.09077](#).
- [44] Luo, D. & Clark, B. K. Backflow Transformations via Neural Networks for Quantum Many-Body Wave Functions. *Physical Review Letters* **122**, 226401 (2019).
- [45] Zhou, Y.-T., Zhou, Z.-W. & Liang, X. Solving Fermi-Hubbard-type Models by Tensor Representations of Backflow Corrections. *Physical Review B* **109**, 245107 (2024).
- [46] Kim, J. *et al.* Neural-network quantum states for ultra-cold Fermi gases. *Communications Physics* **7**, 1–12 (2024).
- [47] Liu, A.-J. & Clark, B. K. Neural network backflow for ab initio quantum chemistry. *Physical Review B* **110**, 115137 (2024).
- [48] Kiers, H. A. L. Towards a standardized notation and terminology in multiway analysis. *Journal of Chemometrics* **14**, 105–122 (2000).
- [49] Kolda, T. G. & Bader, B. W. Tensor Decompositions and Applications. *SIAM Review* **51**, 455–500 (2009).
- [50] López Ríos, P., Ma, A., Drummond, N. D., Towler, M. D. & Needs, R. J. Inhomogeneous backflow transformations in quantum Monte Carlo calculations. *Physical Review E* **74**, 066701 (2006).
- [51] Lou, W. T. *et al.* Neural Wave Functions for Superfluids. *Physical Review X* **14**, 021030 (2024).
- [52] Lami, G., Carleo, G. & Collura, M. Matrix product states with backflow correlations. *Physical Review B* **106**, L081111 (2022).
- [53] Gutzwiller, M. C. Effect of correlation on the ferromagnetism of transition metals. *Physical Review Letters* **10**, 159 (1963).
- [54] Misawa, T. *et al.* mvmc—open-source software for many-variable variational monte carlo method. *Computer Physics Communications* **235**, 447–462 (2019).
- [55] Casula, M., Attaccalite, C. & Sorella, S. Correlated geminal wave function for molecules: An efficient resonating valence bond approach. *The Journal of Chemical Physics* **121**, 7110–7126 (2004).
- [56] Park, C.-Y. & Kastoryano, M. J. Geometry of learning neural quantum states. *Physical Review Research* **2**, 023232 (2020).
- [57] Olivares-Amaya, R. *et al.* The ab-initio density matrix renormalization group in practice. *The Journal of Chemical Physics* **142**, 034102 (2015).
- [58] Rissler, J., Noack, R. M. & White, S. R. Measuring orbital interaction using quantum information theory. *Chemical Physics* **323**, 519–531 (2006).
- [59] Foster, J. M. & Boys, S. F. Canonical Configuration Interaction Procedure. *Reviews of Modern Physics* **32**, 300–302 (1960).
- [60] Sorella, S. Generalized Lanczos algorithm for variational quantum Monte Carlo. *Physical Review B* **64**, 024512 (2001).
- [61] Rende, R., Viteritti, L. L., Bardone, L., Becca, F. & Goldt, S. A simple linear algebra identity to optimize large-scale neural network quantum states. *Communications Physics* **7**, 1–8 (2024).
- [62] Lovato, A., Adams, C., Carleo, G. & Rocco, N. Hidden-nucleons neural-network quantum states for the nuclear many-body problem. *Physical Review Research* **4**, 043178 (2022).

- [63] Carleo, G. *et al.* NetKet: A machine learning toolkit for many-body quantum systems. *SoftwareX* **10**, 100311 (2019).
- [64] Vicentini, F. *et al.* NetKet 3: Machine Learning Toolbox for Many-Body Quantum Systems. *SciPost Physics Codebases* 007 (2022).
- [65] Sun, Q. *et al.* PySCF: The Python-based simulations of chemistry framework. *WIREs Computational Molecular Science* **8**, e1340 (2018).
- [66] Sun, Q. *et al.* Recent developments in the PySCF program package. *The Journal of Chemical Physics* **153**, 024109 (2020).
- [67] Zheng, B.-X. *et al.* Stripe order in the underdoped region of the two-dimensional Hubbard model. *Science* **358**, 1155–1160 (2017).
- [68] Sorella, S. Systematically improvable mean-field variational ansatz for strongly correlated systems: Application to the Hubbard model. *Physical Review B* **107**, 115133 (2023).
- [69] Dagotto, E., Moreo, A., Ortolani, F., Poilblanc, D. & Riera, J. Static and dynamical properties of doped Hubbard clusters. *Physical Review B* **45**, 10741–10760 (1992).
- [70] Simons Collaboration on the Many-Electron Problem *et al.* Solutions of the Two-Dimensional Hubbard Model: Benchmarks and Results from a Wide Range of Numerical Algorithms. *Physical Review X* **5**, 041041 (2015).
- [71] Goldshlager, G., Abrahamsen, N. & Lin, L. A Kaczmarz-inspired approach to accelerate the optimization of neural network wavefunctions. *Journal of Computational Physics* **516**, 113351 (2024).
- [72] Gurtubay, I. G. & Needs, R. J. Dissociation energy of the water dimer from quantum Monte Carlo calculations. *The Journal of Chemical Physics* **127**, 124306 (2007).
- [73] Clark, B. K., Morales, M. A., McMinis, J., Kim, J. & Scuseria, G. E. Computing the energy of a water molecule using multideterminants: A simple, efficient algorithm. *The Journal of Chemical Physics* **135**, 244105 (2011).
- [74] Hachmann, J., Cardoen, W. & Chan, G. K.-L. Multireference correlation in long molecules with the quadratic scaling density matrix renormalization group. *The Journal of Chemical Physics* **125**, 144101 (2006).
- [75] Simons Collaboration on the Many-Electron Problem *et al.* Towards the Solution of the Many-Electron Problem in Real Materials: Equation of State of the Hydrogen Chain with State-of-the-Art Many-Body Methods. *Physical Review X* **7**, 031059 (2017).
- [76] Stella, L., Attaccalite, C., Sorella, S. & Rubio, A. Strong electronic correlation in the hydrogen chain: A variational Monte Carlo study. *Physical Review B* **84**, 245117 (2011).
- [77] Sinitskiy, A. V., Greenman, L. & Mazziotti, D. A. Strong correlation in hydrogen chains and lattices using the variational two-electron reduced density matrix method. *The Journal of Chemical Physics* **133**, 014104 (2010).
- [78] Tsuchimochi, T. & Scuseria, G. E. Strong correlations via constrained-pairing mean-field theory. *The Journal of Chemical Physics* **131**, 121102 (2009).
- [79] Motta, M. *et al.* Ground-State Properties of the Hydrogen Chain: Dimerization, Insulator-to-Metal Transition, and Magnetic Phases. *Physical Review X* **10**, 031058 (2020).
- [80] Zhai, H. *et al.* Block2: A comprehensive open source framework to develop and apply state-of-the-art DMRG algorithms in electronic structure and beyond. *The Journal of Chemical Physics* **159**, 234801 (2023).
- [81] Sharma, S. & Chan, G. K.-L. Spin-adapted density matrix renormalization group algorithms for quantum chemistry. *The Journal of Chemical Physics* **136**, 124121 (2012).

- [82] Chen, A. & Heyl, M. Empowering deep neural quantum states through efficient optimization. *Nature Physics* 1–6 (2024).
- [83] Mahajan, A. & Sharma, S. Efficient local energy evaluation for multi-Slater wave functions in orbital space quantum Monte Carlo. *The Journal of Chemical Physics* **153**, 194108 (2020).
- [84] Wei, H. & Neuscamman, E. Reduced scaling Hilbert space variational Monte Carlo. *The Journal of Chemical Physics* **149**, 184106 (2018).
- [85] Sabzevari, I. & Sharma, S. Improved Speed and Scaling in Orbital Space Variational Monte Carlo. *Journal of Chemical Theory and Computation* **14**, 6276–6286 (2018).
- [86] Nusspickel, M., Ibrahim, B. & Booth, G. H. Effective Reconstruction of Expectation Values from Ab Initio Quantum Embedding. *Journal of Chemical Theory and Computation* **19**, 2769–2791 (2023).
- [87] King’s College London e-Research team. King’s Computational Research, Engineering and Technology Environment (CREATE). <https://doi.org/10.18742/rnvf-m076> (2022).