ACCESS

Coordinate Descent Full Configuration Interaction for Excited States

Zhe Wang, Zhiyuan Zhang, Jianfeng Lu,* and Yingzhou Li*

)	Cite This: J.	Chem.	Theory	Comput.	2023,	19,	7731-	-7739
▼.									



Article Recommendations

ABSTRACT: An efficient excited state method, named xCDFCI, in the configuration interaction framework is proposed. xCDFCI extends the unconstrained nonconvex optimization problem in coordinate descent full configuration interaction (CDFCI) to a multicolumn version for low-lying excited states computation. The optimization problem is addressed via a tailored coordinate descent method. In each iteration, a determinant is selected based on an approximated gradient, and coefficients of all states associated with the selected determinant are updated. A deterministic compression is applied to limit memory usage. We test xCDFCI applied to H₂O and N₂ molecules under the cc-pVDZ basis set. For both systems, five low-lying excited states in the same symmetry sector are calculated, together with the ground state. xCDFCI also produces accurate binding curves of the carbon dimer in the cc-pVDZ basis with chemical accuracy, where the ground state and four excited states in the same symmetry sector are benchmarked.

III Metrics & More



1. INTRODUCTION

Excited state computations are of great importance in understanding and predicting many phenomena in photochemistry, spectroscopy, and others. Compared to the ground state computation, excited state computations are more challenging for wave function ansatz-based methods, including Hartree–Fock methods,^{1,2} configuration interaction methods,³ and coupled cluster methods^{4–6}. The excited states in general have multireference characters, and the wave function ansatzes in these methods limit the representation of dynamic correlations. Similarly, it is more challenging to calculate the excited states than the ground state using density functional theory (DFT) methods^{7–11} and time-dependent DFT methods.^{12,13}

Under the full configuration interaction (FCI) framework, it is also considered more challenging to calculate the excited states, but the difficulty is not as severe as the aforementioned methods. In general, there are two types of challenges for excited state computations under FCI. First, due to the natural multireference features of excited states, the discretization basis set should be of larger sizes than that in ground state computation and the corresponding FCI matrix size should be larger. Second, the energy gaps between excited states are, in general, smaller than those between the ground state and the first excited state, which would lead to more iterations in iterative eigensolvers before converging. In this article, we propose xCDFCI for excited state computation under the FCI framework. The method is closely related to the recently developed efficient FCI solver, by three of the authors, coordinate descent FCI (CDFCI).¹

Many modern FCI solvers have been developed for ground state computation in the past two decades, together with their extensions to excited state computations. Density matrix renormalization group (DMRG)¹⁵⁻¹⁸ uses matrix product state as the wave function ansatz and applies an iterative sweeping procedure as an eigensolver. Various strategies^{19,20} are proposed to address excited states one by one. FCI quantum Monte Carlo (FCIQMC) and its variants²¹⁻²³ use quantum Monte Carlo walker idea to reduce the computational cost. In its extension to excited state computations,²⁴ several groups of walkers are used to represent excited states, and an orthogonal projection is introduced between iterations to prevent groups from collapsing into the ground state. Selected-CI is a group of FCI solvers based on sequential configuration selections, including adaptive configuration interaction (ACI),²⁵ heat-bath configuration interaction (HCI),²⁶⁻²⁸ and adaptive sampling configuration interaction (ASCI) .²⁹ Extending selected-CI methods to excited state computations is straightforward. After a small modification of the selection criteria, $\frac{29-31}{29-31}$ the excited states are computed by solving the low-lying eigenstates of the reduced Hamiltonian matrix. FCI fast random iteration (FCI-FRI)³² adopts a biasfree sampling procedure to compress the wave function under the power method framework. In the excited state version of

Received:April 26, 2023Revised:October 7, 2023Accepted:October 9, 2023Published:October 23, 2023





FCI-FRI,³³ the iterative method is a multicolumn version power method, where the normalization is carried out every iteration and the orthogonalization is carried out every few iterations. Recent review papers^{34,35} summarize other FCIrelated methods as well. Coordinate descent FCI¹⁴ reformulates the eigenvalue problem as an unconstrained optimization problem, which is the single-column version of the optimization problem used in this paper (4). Then a coordinate-descent method is applied to address the optimization problem, where the coordinates are selected based on the magnitude of the gradient vector and the stepsize is calculated from an exact linesearch. Importantly, a tailored compression strategy is applied to limit the growth of nonzeros in the state vector and hence limit memory usage. The compression is not applied to the state vector c directly. Instead, it is applied to b = Hc for H to be the Hamiltonian matrix to truncate small updates that increase the memory cost. The compression in Coordinate descent FCI is carefully designed so that the Rayleigh quotient could be accurately evaluated and the second-order energy estimation becomes available.

Moreover, FCI problems have also attracted attention from the numerical linear algebra community in recent years. Many algorithms and analyses^{36–42} influence the developments above. Other works attempt to incorporate machine learning and reinforcement learning technique to accelerate the FCI calculation.^{43,44}

In this paper, we extend CDFCI to excited state computations and name the method as xCDFCI. The unconstrained optimization problem in CDFCI is extended to a multicolumn version to accommodate low-lying excited states. The coordinate-descent method used to optimize the objective function is replaced by a row-block descent scheme in xCDFCI and the compression is still carried out in an entrywise way. The multicolumn vector in xCDFCI does not converge to the ground state and low-lying excited states directly. Instead, it converges to a subspace formed by the ground and low-lying excited states. The eigenvectors can be recovered by a postprocessing procedure. Most importantly, all desired features of the original CDFCI are preserved. Symmetries, including time-reversal symmetry and angular momentum symmetry, are implemented to reduce both computational and memory costs when the computation is restricted to a symmetry sector. Finally, numerical results on H_2O , N_2 are included to demonstrate the efficiency of xCDFCI for excited state computations. We also report the binding curve of C₂ obtained using xCDFCI for singlet.

The rest of the article is organized as follows. Section 2 introduces xCDFCI for excited state computations and other related discussions. Section 3 provides numerical examples of xCDFCI. The paper is concluded in Section 4.

2. XCDFCI

We introduce xCDFCI in this section and discuss some of the implementation details. Notations are kept the same as that in Wang et al.¹⁴ as much as possible. In the following, we first propose the unconstrained optimization problem for excited state computations, then explain the xCDFCI algorithm step-by-step, and finally discuss its implementation details: initialization, stopping criteria, and symmetry.

2.1. Optimization Formula for Excited State Computations. Given a spin–orbital set $\{\chi_p\}$, we denote the creation and annihilation operator as \hat{a}_p^{\dagger} and \hat{a}_q respectively. The Hamiltonian operator, under the second quantization, is given by

$$\hat{H} = \sum_{p,q} t_{pq} \hat{a}_p^{\dagger} \hat{a}_q + \sum_{p,r,q,s} \nu_{prqs} \hat{a}_p^{\dagger} \hat{a}_r^{\dagger} \hat{a}_s \hat{a}_q$$
(1)

where t_{pq} and v_{prqs} are one-body and two-body integrals, respectively. The K low-lying states of the time-independent Schrödinger equation can be obtained by solving

$$\hat{H}|\Phi_k\rangle = E_k|\Phi_k\rangle \tag{2}$$

for k = 0, 1, ..., K - 1, where E_0 is the smallest eigenvalue associated with the ground state $|\Phi_0\rangle$, E_1 is the second smallest eigenvalue associated with the first excited state $|\Phi_1\rangle$, and so on, $\{|\Phi_k\rangle\}_{k=0}^{K-1}$ are orthogonal to each other.^{*a*} Throughout this paper, we assume that all E_0 , E_1 , ..., E_{K-1} are negative. This assumption can be made without loss of generality, as otherwise we can shift the Hamiltonian by a constant. We further denote the Slater determinants as $\{|D_i\rangle\}_{i=1}^N$ for $N = N_{\text{FCI}}$ being the size of the entire electron-preserving configuration space. Using $\{|D_i\rangle\}_{i=1}^N$ as the basis, the ground state and excited states are discretized as

$$|\Phi_k\rangle = \sum_i V_{i,k} |D_i\rangle \tag{3}$$

and coefficients $V_{i,k}$ forms a matrix V of size $N \times K$ satisfying the orthonormality constraint, $V^T V = I$ for I being an identity matrix of size $K \times K$. The Hamiltonian operator is discretized as the Hamiltonian matrix H with its (i, j)-th entry being $H_{ij} = \langle D_i | \hat{H} | D_j \rangle$. After the discretization, solving (2) is reduced to solving for the low-lying K eigenpairs of H, where the major computational difficulty comes from the factorial scaling of N_{FCI} with respect to the number of spin—orbitals and electrons.

Now we extend the unconstrained optimization problem in CDFCI^{14} to excited states. The optimization problem is extended as

$$\min_{C \in \mathbb{R}^{N \times K}} f(C) \tag{4}$$

for

$$f(C) = \|H + CC^{\mathrm{T}}\|_{\mathrm{F}}^{2}$$
(5)

where C is a matrix of size N by K. When K = 1, (4) is the same as the optimization problem in Wang et al.¹⁴ The gradient of f(C) admits

$$G = \nabla f = 4HC + 4C(C^{\mathrm{T}}C) \tag{6}$$

As has been analyzed in Gao et al.,⁴⁰ the unconstrained optimization problem (4) has many stationary points, but has no spurious local minima. All local minima are global minima of the form

$$V\sqrt{-\Lambda}Q\tag{7}$$

where $\Lambda \in \mathbb{R}^{K \times K}$ is a diagonal matrix with its diagonal entries being E_0 , E_1 , ..., E_{K-1} , $V \in \mathbb{R}^{N \times K}$ is the corresponding eigenvector matrix as defined in (3), and $Q \in \mathbb{R}^{K \times K}$ is an arbitrary orthogonal matrix such that $Q^T Q = QQ^T = I$.

Generally, gradient-based first-order methods, including the coordinate descent method, avoid saddle points and converge to a global minimum almost surely.⁴⁵ We remark that the minimizers of (4) give the eigenspace only due to the arbitrary

2.2. Algorithm. The algorithm we propose for excited state computations is a coordinate descent method applying to (4), where some specifics are designed to fully incorporate the properties of FCI problems. We introduce our algorithm step by step. Throughout the algorithm, two matrices *C* and *B* are kept: *C* is the iterator targeting (7) and *B* is used to track *HC*, i.e., $B \approx HC$. Further, we use superscript in the parentheses to denote iteration index, e.g., $C^{(I)}$ denotes the iterator at the *I*-th iteration. Colon notation is used to denote the entire row or column, e.g., $C_{i,i}^{(I)}$ denotes the *i*-th row of $C^{(I)}$.

The xCDFCI algorithm is composed of an iterative part with 5 steps and a postprocessing step. At each iteration, the first step selects a determinant with maximum absolute value in an approximated gradient (4). The second step then conducts a linesearch and updates C, where a fourth-order polynomial is minimized to determine the optimal stepsize. In the third and fourth steps, the corresponding update to B is calculated with compression and a row of B is recalculated to improve accuracy with minimal additional cost. In the last step of the iterative part, energies are estimated via a generalized Rayleigh quotient procedure. When the iteration converges according to some stopping criteria, a postprocessing step could be carried out to obtain the ground state vector and excited state vectors. In the following, we explain each step of xCDFCI in detail.

2.2.1. Step 1: Determinant Select. This step aims to select a determinant for the update, which potentially leads to the greatest decay in f(C). The determinant selection strategy is as follows

$$i^{(l+1)} = \arg \max_{j \in I_{H}(i^{(l)})} |4B_{j,k}^{(l)} + 4C_{j,:}^{(l)}[(C^{(l)})^{\mathrm{T}}C^{(l)}]_{:,k}|_{0 \le k < K}$$
(8)

where $i^{(l)}$ is the argument j achieving the maximum value. Here $I_H(i^{(l)})$ denotes the set of determinants connected to $i^{(l)}$ via H, i.e., for any $j \in I_H(i^{(l)})$, $H_{i^{(l)}j}$ is nonzero and for any $j \notin I_H(i^{(l)})$, $H_{i^{(l)}j}$ is zero. Due to the existence of zeros in one- and two-body integrals, $I_H(i^{(l)})$ is a subset of the single and double excitations from the $i^{(l)}$ determinant. The intuition behind (8) is related to the gradient of f(C) (6). Comparing (8) and (6), we notice that the determinant is selected to be the row containing the absolutely largest gradient entry so that it potentially leads to the greatest reduction of the objective function.

2.2.2. Step 2: Coefficient Update. Given a selected determinant $i^{(l+1)}$, we seek the best stepsize τ and move the $i^{(l+1)}$ -th row of the coefficient matrix $C^{(l)}$ along the gradient direction with the stepsize. The best stepsize τ is achieved via solving

$$\tau = \arg\min_{\tilde{\tau}} f(C^{(l)} + \tilde{\tau} e_{i^{(l+1)}} \tilde{G}_{i^{(l+1)},:})$$
(9)

where $e_{i^{(l+1)}}$ is a vector with $i^{(l+1)}$ -th entry being one and zero otherwise, and

$$\tilde{G}_{i^{(l+1)},:} = 4B_{i^{(l+1)},:}^{(l)} + 4C_{i^{(l+1)},:}^{(l)}(C^{(l)})^{\mathrm{T}}C^{(l)}$$
(10)

is the $i^{(l+1)}$ -th row of the approximated gradient (6). Solving (9) is actually minimizing a fourth-order polynomial of $\hat{\tau}$, and all polynomial coefficients can be evaluated in $O(K^2)$

operations (details can be found in Appendix A). Once the

stepsize
$$\tau$$
 is determined, we update $C^{(l)}$ as follows

$$C_{i,:}^{(l+1)} = \begin{cases} C_{i,:}^{(l)} + \tau \tilde{G}_{i,:} & \text{if } i = i^{(l+1)} \\ \\ C_{i,:}^{(l)} & \text{otherwise} \end{cases}$$
(11)

2.2.3. Step 3: Coefficient Compression. Throughout the algorithm, we keep all entries of *C*. While, for B = HC without compression, the number of nonzeros in *HC* is much larger than that in *C*. We cannot afford to store *HC* in memory. Hence, we compress the representation of *B*.

We use supp(B) to denote the set of determinants containing at least one nonzero coefficient, i.e., supp(B) = { $i: \max_k |B_{i,k}| > 0$ }. Then we update and compress $B^{(l)}$ as follows, for $i = i^{(l+1)}$

$$B_{j,:}^{(l+1)} = \begin{cases} B_{j,:}^{(l)} + \tau H_{j,i} \tilde{G}_{i,:} & \text{if } j \in \text{supp}(B^{(l)}) \\ & \text{if } j \notin \text{supp}(B^{(l)}) \text{ and} \\ \tau H_{j,i} \tilde{G}_{i,:} & & \\ & & \max_k |\tau H_{j,i} \tilde{G}_{i,k}| > \varepsilon \end{cases}$$
(12)

where ε is the predefined compression threshold. Equation 12 indicates that for all pre-existing determinants in *B*, the coefficients are updated accurately; while for new determinants, the coefficients are added only if they contain an important update. Obviously, the compression limits the growth of nonzeros in *B*, and thus the data storage cost.

Now we explain the indirect connection to compression of *C*. According to (8), when a determinant is not in supp $(B^{(l)})$, the corresponding gradient is zero, hence the determinant will not be selected, which in turn limits the growth of nonzeros in *C*. Therefore, all compressions are explicitly applied to *B* only, indirectly limiting the growth of nonzeros in *C*.

2.2.4. Step 4: Coefficient Recalculation. In (12), we already compute all nonzero entries in the $i^{(l+1)}$ -th column of *H*. Now, we reused these results to refine coefficients in *B*. The $i^{(l+1)}$ -th row in *B* is recalculated as follows

$$B_{i(l+1),:}^{(l+1)} = \sum_{j \in I_{H}(i^{(l+1)})} H_{i^{(l+1)},j} C_{j,:}^{(l+1)} = \sum_{j \in I_{H}(i^{(l+1)})} H_{j,i^{(l+1)}} C_{j,:}^{(l+1)}$$
(13)

where the second equality is due to the symmetry property of the Hamiltonian.^b This recalculation of $B_{i(l+1),:}^{(l+1)}$ is of essential importance when the $i^{(l+1)}$ -th determinant is added to $C^{(l)}$ for the first time. It removes potential errors made by compressions from earlier iterations and, together with (12), keeps $B_i^{(l+1)} \equiv H_i^{(l+1)}$, C for all later iterations. From a numerical analysis viewpoint, the recalculation also preserves the numerical accuracy. Since the number of iterations in xCDFCI could easily go beyond 10⁸ to 10¹⁰, the accumulation of the numerical error caused by the finite precision computations in the worst case grows linearly with respect to the number of operations and would destroy the accuracy of energies. Regularly recalculating $B_{i(l+1),:}^{(l+1)}$ keeps $B_i^{(l+1)} \equiv H_i^{(l+1)}$. C at a low level of numerical error. 2.2.5. Step 5: Energy Estimation. Given a coefficient matrix $C^{(l+1)}$, the energy estimation is conducted through a generalized Rayleigh quotient of second-order accuracy, which solves a generalized eigenvalue problem of matrix pair $((C^{(l+1)})^T H C^{(l+1)}, (C^{(l+1)})^T C^{(l+1)})$, i.e.,

$$((C^{(l+1)})^{\mathrm{T}}HC^{(l+1)})U = ((C^{(l+1)})^{\mathrm{T}}C^{(l+1)})U\Gamma$$
(14)

for *U* being eigenvectors and Γ being the eigenvalue matrix.^c A detailed discussion on the accuracy of the Rayleigh quotient refers to Appendix B. Since only the coefficients of a determinant are updated, both matrices can be updated accordingly

$$(C^{(l+1)})^{\mathrm{T}} C^{(l+1)} = (C^{(l)})^{\mathrm{T}} C^{(l)} + \tau ((C^{(l)}_{i(l+1),:})^{\mathrm{T}} \tilde{G}_{i^{(l+1)},:}^{(l+1),:} + \tilde{G}^{\mathrm{T}}_{i(l+1),:} C^{(l)}_{i(l+1),:}) + \tau^{2} \tilde{G}^{\mathrm{T}}_{i(l+1),:} \tilde{G}_{i^{(l+1)},:}^{(l+1),:}$$

$$(15)$$

and

$$(C^{(l+1)})^{\mathrm{T}}HC^{(l+1)} = (C^{(l)})^{\mathrm{T}}HC^{(l)} + \tau((B^{(l+1)}_{i(l+1),:})^{\mathrm{T}}\tilde{G}_{i^{(l+1)},:} + \tilde{G}^{\mathrm{T}}_{i(l+1),:}B^{(l+1)}_{i^{(l+1)},:}) - \tau^{2}H_{i^{(l+1)}i^{(l+1)}}\tilde{G}^{\mathrm{T}}_{i^{(l+1)},:}\tilde{G}_{i^{(l+1)},:}$$
(16)

Since $B_{i(l+1),:}^{(l+1)}$ was recalculated in the previous step, both matrices are numerically accurate and not affected by our compression. The updated matrix $(C^{(l+1)})^T C^{(l+1)}$ is also involved and reused in the gradient computation of the next iteration. After the energy estimation, we check the stopping criteria. If the criteria are satisfied, we move on to postprocessing; otherwise, we go back to the first step.

2.2.6. Postprocessing. When the algorithm converges, the energies of low-lying excited states are already available in Γ . If excited states are needed for the down stream tasks, e.g., reduced density matrix computations, the coefficient matrix C needs to be transformed back to eigenvectors V and the transformation is as simple as

 $V \approx CU$ (17)

where U is the eigenvector matrix in (14).

2.3. Implementation. We now discuss some implementation details, including the data structures of C and B, the stopping criteria, and the symmetry of molecular systems in the following.

2.3.1. Data Structure. In Wang et al.,¹⁴ several data structures have been implemented and discussed, including the hash table, black-red tree, and so forth. Among these data structures, the hash table is the one achieving the best computational performance for CDFCI. Thus, for xCDFCI, we also adopt hash tables as our overall data structure. For the single-threaded version of our implementation, a Robin Hood hash table is adopted,⁴⁶ whereas for the multithreaded version, a Cuckoo hash table is adopted.^{47,48} In both hash tables, the keys are the binary representations of the determinants. Given a key corresponding to a determinant with index *i*, the bucket of the hash table is composed of two vectors, B_{i} , and C_{i} . Based on our tests of CDFCI, the hash table access costs nearly half of the runtime. Hence, in designing the algorithm and data structure of xCDFCI, we balance the number of hash table accesses and the number of entry updates. For each iteration in xCDFCI, where the number of hash table accesses is the

number of nonzeros in the column of H, we update the entire row of B and C, i.e., update both ground state and excited states of the selected determinant. In xCDFCI, the hash table access costs less than half of the runtime, and the per-iteration cost of xCDFCI is less than K times that of CDFCI. The drawback of our data structure implementation is that it ignores the sparsity across states. For example, consider the scenario that for a given determinant, the value of an excited state is noncompressible while values of other states are all compressible. Our implementation would treat the values of all states as noncompressible and allocate memory for them. In

the trade-off of hash table access cost and memory efficiency,

we lean against the former in the implementation of xCDFCI. 2.3.2. Stopping Criteria. The stopping criteria for coordinate descent methods are usually more complicated than those for general gradient descent methods. In gradient descent methods, the norm of the gradient is often used as a stopping criterion. For nonstiff problems, when the norm is sufficiently small, we are confident that the iteration is close to a first-order stationery point. However, for coordinate descent methods, we often cannot afford to check through the entire gradient vector, as in xCDFCI. It is also risky to stop when the entry update $\tau \tilde{G}_{i,:}$ is small. Hence, in our implementation, we adopt accumulated entry updates as the stopping criterion, i.e.,

$$\operatorname{tol}_{l} = \sum_{l=1}^{n} \beta^{n-l} \left\| \tau^{(l)} \tilde{G}_{i^{(l)},:} \right\|$$
(18)

where *n* is the current iteration index, β is a discounting factor strictly smaller than one, and $\tau^{(l)}$ is the best stepsize at *l*-th iteration. The accumulated entry updates could be evaluated iteratively

$$\operatorname{tol}_{l} = \left\| \tau^{(n)} \tilde{G}_{i^{(n)}, \cdot} \right\| + \beta \cdot \operatorname{tol}_{l-1}$$
(19)

and only a single tol needs to be kept in memory. Throughout, discounting factor β is left as a hyperparameter. Given a β , we could calculate all discounting coefficients in (18) and estimate the number of entry updates whose coefficient is greater than 0.1. Specifically, there are about $-\frac{1}{\log_{10}\beta}$ entry updates with coefficients greater than 0.1. The suggested value for β would be in the range of [0.99, 0.999] such that about a few hundred to a few thousand entry updates are accumulated with coefficients of the same ordering.

3. NUMERICAL RESULTS

In this section, we perform a sequence of numerical experiments for H_2O , C_2 , and N_2 under the cc-pVDZ basis set. In all experiments, the one-body and two-body integrals are calculated by Psi4.⁴⁹ The FCI excited states are calculated by our homebrewed package CDFCI.⁵⁰ All energies are reported in Hartree (Ha).

3.1. H_2O Excited States. This section calculates the excited states of H_2O at the equilibrium geometry. The OH bonds are of length 0.9751 Å, and the HOH bond angle is 110.565°. The maximum memory for the CDFCI calculation is 480 GB and the compression tolerance is 0 (no compression). With the cc-pVDZ basis set, there are 10 electrons and 24 orbitals involved in the calculation. Throughout, the reference energy of the ground state is -76.2418601 Ha, and reference energies of excited states are numerical results at one hundred

number of Iterations					
10 ⁴	107	2×10^{7}	5×10^{7}		
-76.2312241	-76.2418569	-76.2418594	-76.2418600		
-75.8803222	-75.8943336	-75.8943364	-75.8943371		
-75.8452281	-75.8604822	-75.8604851	-75.8604858		
-75.6550559	-75.6731155	-75.6731187	-75.6731195		
-75.5669476	-75.5846740	-75.5846775	-75.5846783		
-75.3466894	-75.4844768	-75.4844824	-75.4844836		
67.14	19414.28	38477.38	90759.04		
	10 ⁴ -76.2312241 -75.8803222 -75.8452281 -75.6550559 -75.5669476 -75.3466894 67.14	104 107 -76.2312241 -76.2418569 -75.8803222 -75.8943336 -75.8452281 -75.8604822 -75.6550559 -75.6731155 -75.5669476 -75.846740 -75.3466894 -75.4844768 67.14 19414.28	10 ⁴ 10 ⁷ 2 × 10 ⁷ -76.2312241 -76.2418569 -76.2418594 -75.8803222 -75.8943336 -75.8943364 -75.8452281 -75.604822 -75.8604851 -75.6550559 -75.6731155 -75.6731187 -75.5669476 -75.5846740 -75.5846775 -75.3466894 -75.4844768 -75.4844824 67.14 19414.28 38477.38		

Table 1. Convergence of Energy of H_2O^a

million iterations of xCDFCI. Reference values are attached in Appendix C.

From Table 1 and Figure 1, we shall see that the energy error drops quickly to the level of 10^{-4} mHa accuracy at the



Figure 1. Convergence of energies of six low-lying excited states of H_2O vs the number of iterations.

beginning. It then has a slower but steady decay. According to Figure 1, in general, energies associated with lower excited states are of better accuracy. The only exception for H_2O is the energy associated with the third excited state, which achieves better accuracy than the first and second excited state energies. From Table 1, we find that each state can quickly converge to chemical accuracy. After a burn-in stage (first few thousand iterations), the runtime is linear with respect to the number of iterations. Hence, if Figure 1 is redone for energy errors against the runtime, the curves would behave similarly, and the decays remain linear against the runtime after the burn-in stage. **3.2.** N₂ Excited States. This section calculates the excited states of N₂ at the equilibrium geometry. Nitrogen dimer N₂ is more challenging than H₂O because the FCI problem size is much larger, so we use thresholds of 10^{-4} and 10^{-5} for compression. The N₂ molecule has a bond length of 1.12079. The maximum memory in this section is limited to 960 GB. With the cc-pVDZ basis set, there are 14 electrons and 28 orbitals. The results of N₂ are reported in Tables 2, 3, and Figure 2. Throughout, the reference energy of the ground state is -109.28210 Ha, and reference energies of excited states are numerical results of xCDFCI at one hundred million iterations. Reference values are attached in Appendix C.

The convergence trend of N₂ is similar to that of H₂O except that the convergence rate in N₂ is slower. Similarly, after the first million iterations, xCDFCI converges linearly, and the convergence rates are quite stable for both the ground state and excited states. Therefore, we conclude that xCDFCI is stable and efficient for various chemistry systems with different correlation strengths. For N2, xCDFCI takes about ten thousand seconds to achieve chemical accuracy. Convergence rates for all states are approximately the same. Unlike H₂O, where the runtime scales linearly with respect to the number of iterations, for N_2 , the runtime scales sublinearly. This is mainly due to the compression. When the compression criterion is activated, the computational cost for compressed determinants is far less than that of uncompressed ones. Comparing Tables 1 and 3, we notice that the runtime of N_2 is smaller than that of H₂O. Although the computational system of N₂ is larger, the compression with tolerance 10⁻⁵ reduces a lot of computations and the runtime is also reduced. Comparing Tables 2 and 3, we find that the accuracies for both ground and excited state energies are at the same level of the truncation threshold. When a smaller truncation threshold is used, the runtime is longer whereas the accuracies are consistently improved. Therefore, the compression technique is efficient and reliable.

3.3. Carbon Dimer Binding Curves. In this section, we test C_2 with bond lengths from 1 to 2.6. We computed five

Tał	ole	2.	Convergence	of Energy	of N ₂	with a	Thresh	old	1 10 ⁻⁴	4
-----	-----	----	-------------	-----------	-------------------	--------	--------	-----	--------------------	---

energy (Ha)	number of iterations					
	105	10 ⁶	10 ⁷	5×10^{7}		
ground state	-109.26880	-109.28079	-109.28202	-109.28205		
first excited state	-108.71630	-108.73197	-108.73388	-108.73393		
second excited state	-108.64476	-108.66052	-108.66288	-108.66294		
third excited state	-108.63841	-108.65936	-108.66081	-108.66085		
fourth excited state	-108.60955	-108.62886	-108.63105	-108.63110		
fifth excited state	-108.58148	-108.60142	-108.60365	-108.60372		
wall time (sec)	227.8	1012.46	6671.83	31974.3		

Table 3. Convergence of Energy of N_2 with Threshold 10^{-5}

number of iterations					
105	10 ⁶	10 ⁷	5×10^{7}		
-109.26836	-109.28077	-109.28204	-109.28215		
-108.71546	-108.73196	-108.73390	-108.73407		
-108.64376	-108.66050	-108.66291	-108.66311		
-108.63613	-108.65935	-108.66083	-108.66096		
-108.60848	-108.62885	-108.63110	-108.63130		
-108.58040	-108.60141	-108.60370	-108.60392		
315.75	2140.65	15129.69	57403.06		
	10^{5} -109.26836 -108.71546 -108.64376 -108.63613 -108.60848 -108.58040 315.75	number of 10 ⁵ 10 ⁶ -109.26836 -109.28077 -108.71546 -108.73196 -108.64376 -108.66050 -108.63613 -108.65935 -108.60848 -108.62885 -108.58040 -108.60141 315.75 2140.65	number of iterations 10 ⁵ 10 ⁶ 10 ⁷ -109.26836 -109.28077 -109.28204 -108.71546 -108.73196 -108.73390 -108.64376 -108.66050 -108.66291 -108.63613 -108.65935 -108.66083 -108.60848 -108.62885 -108.63110 -108.58040 -108.60141 -108.60370 315.75 2140.65 15129.69		



Figure 2. Convergence of energies of six low-lying excited states of N_2 against the number of iterations. The threshold is 10^{-5} .

low-lying energies of the singlet of C_2 . The symmetry in the basis set is implemented via the Hartree–Fock calculation, i.e., in the Psi4 calculation. More precisely, the singlet calculation is realized by setting the molecule as a singlet and its irreducible representations. The maximum memory in this section is 120 GB and the tolerance is 0. With the cc-pVDZ basis set, there are 12 electrons and 56 orbitals. We perform 1 million iterations for xCDFCI. In all configurations, the accuracies for all states are at the level of chemical accuracy.

The energies of five low-lying states of C_2 in the singlet $({}^{1}\Sigma_g)$ are shown in Table 4. Binding curves are depicted in Figure 3. In general, we observe that the binding curves for lower energy



Figure 3. Low-lying potential energy surfaces of the carbon dimer in the singlet the cc-pVDZ basis.

states are smoother, as shown in Figure 3. We find a lot of crossover points. Each crossover point corresponds to a configuration whose energies are degenerate. Lower energy binding curves have fewer crossover points. The binding curve for the fourth excited state has many crossover points with binding curves of higher excited states though they are not calculated.

4. CONCLUSIONS AND DISCUSSION

We proposed xCDFCI in this paper as an efficient low-lying excited-state solver under the FCI framework. xCDFCI adopts an extension of the objective function in the CDFCI method.

I able	4.	Energy	of Five	Low-Lying	States of	C_2 in	Singlet	

		ene	ergy of five low-lying states (1	Ha)	
R (Å)	zeroth	first	second	third	fourth
1.0	-75.55231	-75.37074	-75.34005	-75.25824	-75.24635
1.1	-75.67528	-75.52584	-75.52314	-75.42099	-75.40454
1.2	-75.7246	-75.6188	-75.61144	-75.51174	-75.46344
1.3	-75.73152	-75.66195	-75.65091	-75.55151	-75.4995
1.4	-75.71569	-75.67459	-75.66213	-75.56135	-75.5052
1.5	-75.68951	-75.67034	-75.65703	-75.55471	-75.49432
1.6	-75.66102	-75.65712	-75.64203	-75.53995	-75.47353
1.7	-75.64014	-75.63676	-75.62022	-75.52375	-75.4532
1.8	-75.62201	-75.6169	-75.59619	-75.5117	-75.45362
1.9	-75.60453	-75.59944	-75.57506	-75.50693	-75.44587
2.2	-75.56245	-75.55941	-75.53746	-75.51137	-75.46051
2.5	-75.53929	-75.53814	-75.52593	-75.51654	-75.49545

Article

More precisely, xCDFCI extends the single-column version (ground state) to a multicolumn version (low-lying excited states) and leads to (4). Then a tailored coordinate descent method is applied to address (4). xCDFCI first selected a determinant with the largest entry in magnitude in the approximated gradient, and then the selected row of the iteration variable C is updated, i.e., the coefficients of a determinant for all states are updated. To avoid memory overflow, a hard-thresholding type compression is applied to B \approx *HC* for *H* being the Hamiltonian matrix, which in turn limits the growth of nonzeros in C. Finally, we carefully maintain the double precision accuracy of $C^{T}C$ and $C^{T}HC = C^{T}B$, and we estimate the eigenvalues through a generalized Rayleigh quotient procedure. Based on results from the theory of numerical analysis,⁵¹ the ground state and low-lying excited states are of the first order accuracy, whereas the ground state energy and excited state energies are of the second order accuracy. In summary, xCDFCI extends CDFCI to calculate low-lying excited states and inherits almost all desired properties of CDFCI. Numerical results on various chemistry systems demonstrate the efficiency of xCDFCI.

Extending the CDFCI to higher-lying excited states is feasible but more challenging. Memory cost is a major concern. When more excited states are computed using CDFCI, the number of columns in C and B is increased. At the same time, the number of nonzero rows in C and B also needs to be increased to incorporate the sparsity of higher-lying excited states. Hence, the memory increases faster than linear scaling with respect to the number of excited states. Besides the memory cost, the degeneracy in higher-lying excited states would also cause trouble if K is not properly chosen.

There are a few promising future directions. First, xCDFCI has not fully exploited the sparsity of the low-lying excited states. Due to the nature of (4), the objective function is rotation invariant, i.e., the objective function remains the same for C and CQ with Q being an orthogonal matrix. Hence, xCDFCI can converge to the eigenspace formed by the desired ground state and low-lying excited states. While it is not guaranteed to converge to the sparse eigenvectors directly. Some recent works⁴⁰⁻⁴² provide promising paths to address the sparsity issue. Second, the basis sets remain the Hartree-Fock molecular orbitals. Applying orbital optimization methods like CASSCF⁵² or OptOrbFCI⁵³ with state-averaged idea together with xCDFCI would be a direct extension. Exploring various orbital rotations for different excited states coupled with xCDFCI would be an interesting future direction. Lastly, we did not fully incorporate the compressed evaluation of the Hamiltonian matrix and other perturbative approximations as in other FCI excited state work, 30,33 which could be combined with xCDFCI to further accelerate the proposed method.

APPENDIX A

Optimal Stepsize via Linesearch

The optimal stepsize τ could be obtained by solving (9). The function $f(C^{(l)} + \tau \tilde{e}_{i^{(l+1)}}G_{i^{(l+1)},:})$ could be rewritten as a fourth-order polynomial of $\hat{\tau}$. For the sake of notation, we omit all superscripts of the iteration index and obtain

$$f(C + \tilde{\tau}e_{i}\tilde{G}_{i,:}) = c_{0} + c_{1}\tilde{\tau} + c_{2}\tilde{\tau}^{2} + c_{3}\tilde{\tau}^{3} + c_{4}\tilde{\tau}^{4}$$
(20)

where the polynomial coefficients are

Article

$$c_0 = f(C) \tag{21}$$

$$c_1 = \|\tilde{G}_{i,:}\|^2$$
 (22)

pubs.acs.org/JCTC

$$c_{2} = 2H_{i,i} \|\tilde{G}_{i,:}\|^{2} + 2\tilde{G}_{i,:}(C^{T}C)\tilde{G}_{i,:}^{T} + 2(C_{i,:}\tilde{G}_{i,:}^{T})^{2} + 2 \\ \|C_{i,:}\|^{2} \|\tilde{G}_{i,:}\|^{2}$$
(23)

$$c_{3} = 4(C_{i,:}\tilde{G}_{i,:}^{\mathrm{T}}) \|\tilde{G}_{i,:}\|^{2}$$
(24)

$$c_4 = \|\tilde{G}_{i,:}\|^4 \tag{25}$$

Notice that coefficient c_2 could be evaluated in $O(K^2)$ operations, and coefficient c_1 , c_3 and c_4 could be evaluated in O(K) operations, where K is the number of states and length of all row vectors.

Finding the minimum of the fourth-order polynomial could be addressed via solving a third-order polynomial

$$c_1 + 2c_2\tilde{\tau} + 3c_3\tilde{\tau}^2 + 4c_4\tilde{\tau}^3 = 0$$
(26)

There are three scenarios in solving (26): (1) one root; (2) two roots; and (3) three roots. When there is only one root, it achieves the minimum of (20). When there are two roots, one of which is of multiplicity, it achieves the minimum. When there are three roots, the one farther from the middle one achieves the minimum. Through the above procedure, the linesearch problem (9) could be addressed efficiently in $O(K^2)$ operations.

APPENDIX B

Rayleigh Quotient

Let *H* be a symmetric matrix of the size *N*. The eigenvalues of *H* are denoted as $E_0 < E_1 < \cdots < E_{N-1}$. And the associated eigenvectors are V_0 , V_1 , ..., V_{N-1} . For simplicity, we assume that *H* is a gapped matrix. Given a vector $x \in \mathbb{R}^N$, the Rayleigh quotient is defined as

$$r(x) = \frac{x^{\mathrm{T}} H x}{x^{\mathrm{T}} x}$$
(27)

Obviously, the Rayleigh quotient is x scale-invariant, i.e., $r(x) = r(\alpha x)$ for any nonzero scaler α . Hence, we could focus on a normalized vector x such that ||x|| = 1.

An interesting and useful property of the Rayleigh quotient is that r(x) is a quadratically accurate estimate of an eigenvalue. More precisely, let V_j be one of the eigenvectors of H. We consider the case that x is sufficiently close to V_j , i.e., $||x - V_j|| = O(\epsilon)$ for ϵ small. Then an important consequence of the Rayleigh quotient is that⁵⁴

$$r(x) = r(V_j) + O(\epsilon^2) = E_j + O(\epsilon^2)$$
(28)

This is the second-order accuracy to which we are referring in the main paper.

In this paper, instead of the Rayleigh quotient of a single vector, we adopt a generalized Rayleigh quotient (or block Rayleigh quotient), as in (14). The eigenvalue estimation is performed via solving a generalized eigenvalue problem. We could view the generalized eigenvalue problem step as a normalization step so that each column in CU as in (17) is a normalized and aligned estimation of an eigenvector of H. Then the quadratically accurate property of the Rayleigh quotient remains valid in the generalized Rayleigh quotient case.

Reference Energies

The reference energies for $\rm H_2O$ and $\rm N_2$ under the cc-pVDZ basis are reported in Table 5.

Table 5. Reference Energies (Ha) for H_2O and N_2 under the cc-pVDZ Basis

	H ₂ O	N_2
0th ES	-76.241860063	-109.282165
first ES	-75.894337144	-108.734087
second ES	-75.860485864	-108.663124
third ES	-75.673119564	-108.660977
fourth ES	-75.584678392	-108.631318
fifth ES	-75.484483689	-108.603936

AUTHOR INFORMATION

Corresponding Authors

- Jianfeng Lu Department of Mathematics and Department of Chemistry and Department of Physics, Duke University, Durham, North Carolina 27708-0187, United States; Email: jianfeng@math.duke.edu
- Yingzhou Li School of Mathematical Sciences, Fudan University, Shanghai 200433, China; orcid.org/0000-0003-1852-3750; Email: yingzhouli@fudan.edu.cn

Authors

- Zhe Wang Department of Mathematics, Duke University, Durham, North Carolina 27708-0187, United States
- **Zhiyuan Zhang** School of Future Technology, University of Science and Technology of China, Hefei, Anhui 230026, China

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.jctc.3c00452

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

The work of Z.W. is supported by the US National Science Foundation under awards DMS-1454939 and OAC-1450280. This work is part of Z.W.'s PhD thesis at Duke University. Y.L. is supported in part by National Natural Science Foundation of China (12271109) and Shanghai Pilot Program for Basic Research—FuDan University 21TQ1400100 (22TQ017).

ADDITIONAL NOTES

^{*a*}With some abuse of terminology, we will also refer the ground state as the 0-th excited state when it is convenient to do so. ^{*b*}If the Hamiltonian matrix is complex Hermitian, then a complex conjugate is needed in the equation.

^cWe assume U is a $((C^{(l+1)})^{T}C^{(l+1)})$ orthonormalized eigenvector matrix, i.e., $U^{T}((C^{(l+1)})^{T}C^{(l+1)})U = I$.

REFERENCES

(1) Morokuma, K.; Iwata, S. Extended Hartree-Fock theory for excited states. *Chem. Phys. Lett.* **1972**, *16*, 192–197.

(2) Barca, G. M. J.; Gilbert, A. T. B.; Gill, P. M. W. Communication: Hartree-Fock description of excited states of H2. *J. Chem. Phys.* 2014, *141*, 111104. (3) Sherrill, C. D. An Introduction to Configuration Interaction Theory; Citeseer, 1995.

(4) Geertsen, J.; Rittby, M.; Bartlett, R. J. The equation-of-motion coupled-cluster method: Excitation energies of Be and CO. *Chem. Phys. Lett.* **1989**, *164*, 57–62.

(5) Stanton, J. F.; Bartlett, R. J. The equation of motion coupledcluster method. A systematic biorthogonal approach to molecular excitation energies, transition probabilities, and excited state properties. J. Chem. Phys. **1993**, 98, 7029–7039.

(6) Nooijen, M.; Bartlett, R. J. A new method for excited states: Similarity transformed equation-of-motion coupled-cluster theory. J. Chem. Phys. **1997**, 106, 6441–6448.

(7) Onida, G.; Reining, L.; Rubio, A. Electronic excitations: Densityfunctional versus many-body Green's-function approaches. *Rev. Mod. Phys.* **2002**, *74*, 601–659.

(8) Rocca, D.; Bai, Z.; Li, R. C.; Galli, G. A block variational procedure for the iterative diagonalization of non-Hermitian random-phase approximation matrices. *J. Chem. Phys.* **2012**, *136*, 034111.

(9) Yang, Y.; Peng, D.; Lu, J.; Yang, W. Excitation energies from particle-particle random phase approximation: Davidson algorithm and benchmark studies. *J. Chem. Phys.* **2014**, *141*, 124104.

(10) Lu, J.; Thicke, K. Cubic scaling algorithms for RPA correlation using interpolative separable density fitting. *J. Comput. Phys.* **2017**, 351, 187–202.

(11) Hu, W.; Liu, J.; Li, Y.; Ding, Z.; Yang, C.; Yang, J. Accelerating excitation energy computation in molecules and solids within linear-response time-dependent density functional theory via interpolative separable density fitting decomposition. *J. Chem. Theory Comput.* **2020**, *16*, 964–973.

(12) Runge, E.; Gross, E. K. U. Density-functional theory for timedependent systems. *Phys. Rev. Lett.* **1984**, *52*, 997–1000.

(13) Burke, K.; Werschnik, J.; Gross, E. K. U. Time-dependent density functional theory: Past, present, and future. *J. Chem. Phys.* **2005**, *123*, 062206.

(14) Wang, Z.; Li, Y.; Lu, J. Coordinate descent full configuration interaction. J. Chem. Theory Comput. 2019, 15, 3558–3569.

(15) White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **1992**, *69*, 2863–2866.

(16) Chan, G. K.-L.; Sharma, S. The density matrix renormalization group in quantum chemistry. *Annu. Rev. Phys. Chem.* **2011**, *62*, 465–481.

(17) Schollwöck, U. The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.* **2011**, 326, 96–192.

(18) Baiardi, A.; Reiher, M. The density matrix renormalization group in chemistry and molecular physics: Recent developments and new challenges. J. Chem. Phys. **2020**, 152, 040903.

(19) Chandross, M.; Hicks, J. C. Density-matrix renormalizationgroup method for excited states. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1999**, 59, 9699–9702.

(20) Sharma, S. A general non-Abelian density matrix renormalization group algorithm with application to the C2 dimer. *J. Chem. Phys.* **2015**, 142, 024107.

(21) Booth, G. H.; Thom, A. J. W.; Alavi, A. Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space. *J. Chem. Phys.* **2009**, *131*, 054106.

(22) Cleland, D.; Booth, G. H.; Alavi, A. Communications: Survival of the fittest: Accelerating convergence in full configurationinteraction quantum Monte Carlo. J. Chem. Phys. **2010**, *132*, 041103.

(23) Petruzielo, F. R.; Holmes, A. A.; Changlani, H. J.; Nightingale, M. P.; Umrigar, C. J. Semistochastic projector Monte Carlo method. *Phys. Rev. Lett.* **2012**, *109*, 230201.

(24) Blunt, N. S.; Smart, S. D.; Booth, G. H.; Alavi, A. An excitedstate approach within full configuration interaction quantum Monte Carlo. *J. Chem. Phys.* **2015**, *143*, 134117.

(25) Schriber, J. B.; Evangelista, F. A. Communication: An adaptive configuration interaction approach for strongly correlated electrons with tunable accuracy. *J. Chem. Phys.* **2016**, *144*, 161106.

(26) Holmes, A. A.; Tubman, N. M.; Umrigar, C. J. Heat-bath configuration interaction: An efficient selected configuration inter-

action algorithm inspired by heat-bath sampling. J. Chem. Theory Comput. 2016, 12, 3674–3680.

(27) Li, J.; Otten, M.; Holmes, A. A.; Sharma, S.; Umrigar, C. J. Fast semistochastic heat-bath configuration interaction. *J. Chem. Phys.* **2018**, *149*, 214110.

(28) Sharma, S.; Holmes, A. A.; Jeanmairet, G.; Alavi, A.; Umrigar, C. J. Semistochastic heat-bath configuration interaction method: Selected configuration interaction with semistochastic perturbation theory. J. Chem. Theory Comput. **2017**, *13*, 1595–1604.

(29) Tubman, N. M.; Lee, J.; Takeshita, T. Y.; Head-Gordon, M.; Whaley, K. B. A deterministic alternative to the full configuration interaction quantum Monte Carlo method. *J. Chem. Phys.* **2016**, *145*, 044112.

(30) Holmes, A. A.; Umrigar, C. J.; Sharma, S. Excited states using semistochastic heat-bath configuration interaction. *J. Chem. Phys.* **2017**, *147*, 164111.

(31) Schriber, J. B.; Evangelista, F. A. Adaptive configuration interaction for computing challenging electronic excited states with tunable accuracy. *J. Chem. Theory Comput.* **2017**, *13*, 5354–5366.

(32) Greene, S. M.; Webber, R. J.; Weare, J.; Berkelbach, T. C. Beyond walkers in stochastic quantum chemistry: Reducing error using fast randomized iteration. *J. Chem. Theory Comput.* **2019**, *15*, 4834–4850.

(33) Greene, S. M.; Webber, R. J.; Smith, J. E. T.; Weare, J.; Berkelbach, T. C. Full configuration interaction excited-state energies in large active spaces from subspace iteration with repeated random sparsification. *J. Chem. Theory Comput.* **2022**, *18*, 7218–7232.

(34) Eriksen, J. J. The shape of full configuration interaction to come. J. Phys. Chem. Lett. 2021, 12, 418–432.

(35) Loos, P.-F.; Scemama, A.; Jacquemin, D. The quest for highly accurate excitation energies: A computational perspective. *J. Phys. Chem. Lett.* **2020**, *11*, 2374–2383.

(36) Lim, L.-H.; Weare, J. Fast randomized iteration: Diffusion Monte Carlo through the lens of numerical linear algebra. *SIAM Rev.* **2017**, *59*, 547–587.

(37) Li, Y.; Lu, J.; Wang, Z. CoordinateWise descent methods for leading eigenvalue problem. *SIAM J. Sci. Comput.* **2019**, *41*, A2681–A2716.

(38) Lu, J.; Wang, Z. The Full Configuration Interaction Quantum Monte Carlo Method through the Lens of Inexact Power Iteration. *SIAM J. Sci. Comput.* **2020**, *42*, B1–B29.

(39) Hernandez, T. M.; Van Beeumen, R.; Caprio, M. A.; Yang, C. A greedy algorithm for computing eigenvalues of a symmetric matrix with localized eigenvectors. *Numer Linear Algebra Appl.* **2021**, *28*, e2341. DOI: 10.1002/nla.2341.

(40) Gao, W.; Li, Y.; Lu, B. Triangularized orthogonalization-free method for solving extreme eigenvalue problems. *Journal of Scientific Computing* **2022**, *93*, 63.

(41) Gao, W.; Li, Y.; Lu, B. Global convergence of triangularized orthogonalization-free method. *Commun. Math. Sci.* **2023**, *21*, 195–218.

(42) Gao, W.; Li, Y.; Shen, H. Weighted trace-penalty minimization for full configuration interaction. **2023**, arXiv:2301.07270. arXiv preprint.

(43) Coe, J. P. Machine learning configuration interaction for ab initio potential energy curves. J. Chem. Theory Comput. 2019, 15, 6179–6189.

(44) Goings, J. J.; Hu, H.; Yang, C.; Li, X. Reinforcement learning configuration interaction. *J. Chem. Theory Comput.* **2021**, *17*, 5482–5491.

(45) Chen, Z.; Li, Y.; Lu, J. On the global convergence of randomized coordinate gradient descent for non-convex optimization. *SIAM J. Optim.* **2023**, *33*, 713.

(46) Leitner-Ankerl, M. *Robin-Hood-Hashing*, 2022. https://github. com/martinus/robin-hood-hashing.

(47) Fan, B.; Andersen, D. G.; Kaminsky, M. MemC3: compact and concurrent MemCache with dumber caching and smarter hashing. 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13); USENIX Association, 2013.

(48) Li, X.; Andersen, D. G.; Kaminsky, M.; Freedman, M. J. Algorithmic improvements for fast concurrent Cuckoo hashing. *EuroSys '14: Proceedings of the Ninth European Conference on Computer Systems*; Association for Computing Machinery, 2014.

(49) Smith, D. G.; Burns, L. A.; Simmonett, A. C.; Parrish, R. M.; Schieber, M. C.; Galvelis, R.; Kraus, P.; Kruse, H.; Di Remigio, R.; Alenaizan, A.; et al. PSI4 1.4: Open-source software for highthroughput quantum chemistry. *J. Chem. Phys.* **2020**, *152*, 184108.

(50) Wang, Z.; Li, Y.; Lu, J. CDFCI, 2019. https://github.com/quantum/CDFCI.

(51) Golub, G. H.; Van Loan, C. F. *Matrix Computations*, 4th ed.; The Johns Hopkins University Press, 2013; pp 756-757.

(52) Olsen, J. The CASSCF method: A perspective and commentary. Int. J. Quantum Chem. 2011, 111, 3267-3272.

(53) Li, Y.; Lu, J. Optimal Orbital Selection for Full Configuration Interaction (OptOrbFCI): Pursuing the Basis Set Limit under a Budget. J. Chem. Theory Comput. **2020**, *16*, 6207–6221.

(54) Trefethen, L. N.; Bau, D. Numerical Linear Algebra; Society for Industrial and Applied Mathematics, 1997.