

ARTICLE OPEN



Tensor-structured algorithm for reduced-order scaling large-scale Kohn–Sham density functional theory calculations

Chih-Chuen Lin¹, Phani Motamarri^{1,2} and Vikram Gavini^{1,3}✉

We present a tensor-structured algorithm for efficient large-scale density functional theory (DFT) calculations by constructing a Tucker tensor basis that is adapted to the Kohn–Sham Hamiltonian and localized in real-space. The proposed approach uses an additive separable approximation to the Kohn–Sham Hamiltonian and an L_1 localization technique to generate the 1-D localized functions that constitute the Tucker tensor basis. Numerical results show that the resulting Tucker tensor basis exhibits exponential convergence in the ground-state energy with increasing Tucker rank. Further, the proposed tensor-structured algorithm demonstrated sub-quadratic scaling with system-size for both systems with and without a gap, and involving many thousands of atoms. This reduced-order scaling has also resulted in the proposed approach outperforming plane-wave DFT implementation for systems beyond 2000 electrons.

npj Computational Materials (2021)7:50; <https://doi.org/10.1038/s41524-021-00517-5>

INTRODUCTION

Density functional theory (DFT) has been the workhorse of ab-initio materials simulations for over three decades, providing many key insights into materials properties and materials behavior. In order to study ground-state properties, based on the Hohenberg–Kohn theorem¹ and the Kohn–Sham formulation², DFT reduces the Schrödinger equation in $3N_e$ spatial coordinates (N_e denoting the number of electrons) to an equivalent problem in the electron-density that only depends on three spatial coordinates. This reduces the exponential computational complexity (with system-size) of solving the Schrödinger equation to the cubic computational complexity of DFT. While DFT has enabled wide-ranging ab-initio calculations, with $\sim 1/4$ th of the computational resources on some public supercomputers utilized for DFT calculations, the cubic computational complexity has limited routine DFT calculations to typical system-sizes involving a few hundred atoms. In an attempt to enable DFT calculations on large-scale systems that are critical to understanding many aspects of complex materials phenomena, many efforts over the past three decades have focused on developing reduced-order scaling algorithms for electronic structure calculations^{3–13}. These approaches have either relied on a localized representation of the single-electron wavefunctions (such as Wannier functions⁵) or the exponential decay of the density-matrix in real-space, and have been demonstrated to provide close to linear-scaling complexity for materials with a gap. However, they have not been widely successful for metallic systems (without a gap) either due to the errors resulting from realizing locality in the wavefunctions in real-space or due to the higher prefactors that make these approaches computationally more expensive than the traditional cubic-scaling algorithms for system-sizes of interest. In this work, we present an alternative direction by using tensor-structured ideas to achieve systematically convergent and efficient DFT calculations that exhibit sub-quadratic scaling for systems with and without a gap over system-sizes spanning many thousands of atoms.

This line of work is motivated from a study that revealed a low-rank representation for the electronic structure using Tucker and canonical tensor decomposition^{14–16}. Another study based on a posteriori analysis showed that the rank required to approximate the electron-density is only weakly dependent on the system-size¹⁷. These studies have thereby prompted the development of a tensor-structured approach for DFT calculations¹⁸, where a Tucker tensor basis adapted to the Kohn–Sham Hamiltonian was employed to solve the Kohn–Sham equations. Importantly, the rank of the Tucker tensor basis was only weakly dependent with system-size for materials systems with and without a gap, which revealed the potential for realizing a reduced-order scaling approach for DFT calculations. However, since the constructed Tucker tensor basis is global, the representation of the Kohn–Sham Hamiltonian in this basis was dense. An L_2 localization scheme that was explored to localize the Tucker tensor basis did not provide sufficient locality to exploit sparsity in the Kohn–Sham Hamiltonian. With increasing system-size, albeit the slow growth in the size of the Hamiltonian matrix due to the weak rank dependence, the approach became computationally prohibitive for large systems and limited the system-sizes to a few hundred atoms.

In this work, L_1 localization is used to overcome the aforementioned drawbacks, and we demonstrate systematically convergent, efficient and reduced-order scaling large-scale DFT calculations using tensor-structured techniques (cf. Fig. 1 for an overview of the approach). The L_1 localization utilizes the idea of auto-encoder commonly recognized in machine learning to construct a series of 1-D functions that are localized yet closely approximate the function space of interest. The 1-D localized functions that are a close approximation to the eigensubspace of a suitably constructed additive separable approximation of the Kohn–Sham Hamiltonian are used to generate the localized Tucker tensor basis for the DFT problem. The locality of the Tucker tensor basis results in a sparse discrete Kohn–Sham Hamiltonian matrix, which is exploited in the solution of the Kohn–Sham equations using the Chebyshev filtering subspace iteration scheme.

¹Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI, USA. ²Department of Computational and Data Sciences, Indian Institute of Science, Bangalore, India. ³Department of Materials Science & Engineering, University of Michigan, Ann Arbor, MI, USA. ✉email: vikramg@umich.edu

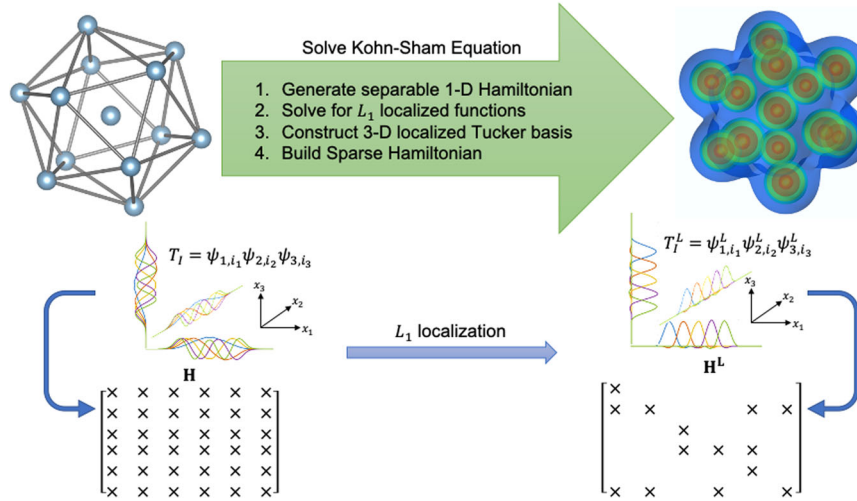


Fig. 1 Overview of the tensor-structured algorithm for Kohn-Sham DFT using L_1 localized functions. The tensor-structured algorithm seeks to construct a systematically convergent reduced-order tensor-structured basis for efficiently solving the Kohn-Sham equations. To this end, an additive separable approximation to the Kohn-Sham Hamiltonian is constructed, whose eigenbasis presents a suitable reduced-order basis, given by $T_l = \psi_{1,i_1} \psi_{2,i_2} \psi_{3,i_3}$, where l is a composite index $l = (i_1, i_2, i_3)$. However, the discrete Kohn-Sham Hamiltonian in this basis is dense due to the global nature of the 1-D functions ψ_{k,i_k} , $k = 1, 2, 3$ (lower-left). The L_1 localization is applied to alleviate this bottleneck, where localized tensor-structured basis functions are constructed such that the subspace spanned by this localized basis is a close approximation to the eigensubspace of the separable Hamiltonian. Denoting the localized 1-D functions by ψ_{k,i_k}^L , the 3-D localized basis functions are given by $T_l^L = \psi_{1,i_1}^L \psi_{2,i_2}^L \psi_{3,i_3}^L$. The 3-D localized basis functions, being compactly supported, yields a sparse Hamiltonian (lower-right). The resulting sparse Hamiltonian, in conjunction with the slow growth of the Tucker rank with system-size to accurately represent the electronic structure, has provided sub-quadratic scaling with system-size for both insulating and metallic systems spanning over many thousands of atoms.

The sparsity of the Kohn-Sham Hamiltonian matrix represented in the localized Tucker tensor basis improves both the computational efficiency and the memory footprint. Further, as will be demonstrated, the proposed approach has enabled sub-quadratic scaling DFT calculations on large-scale systems involving many thousands of atoms. The approach is generic and treats both systems with and without a gap on an equal footing. Importantly, this translates to substantial speed-ups over Quantum Espresso (QE), a widely used state-of-the-art plane-wave DFT code^{19,20}, with speed-ups of ~ 8 -fold for metallic nano-particles containing ~ 2000 atoms.

RESULTS

The Kohn-Sham DFT formulation

The ground state energy in Kohn-Sham DFT (spin-independent formulation) of a materials system with N_a atoms and N_e electrons is computed by solving a non-interacting single-particle Schrödinger equation in a mean-field determined by the effective potential $V_{\text{eff}}(\mathbf{x})$:

$$\left(-\frac{1}{2}\nabla^2 + V_{\text{eff}}(\mathbf{x})\right)\Psi_i = \epsilon_i\Psi_i, \quad i \in \{1, \dots, N\}. \quad (1)$$

Equation (1) represents a non-linear eigenvalue problem with $\mathcal{H} := -\frac{1}{2}\nabla^2 + V_{\text{eff}}$ being the Kohn-Sham Hamiltonian, and ϵ_i denoting the i th Kohn-Sham eigenvalue and Ψ_i denoting the corresponding Kohn-Sham orbital (eigenvector). The effective potential V_{eff} and the subsequent Hamiltonian matrix elements can be evaluated efficiently using a low-rank Tucker tensor approximation and tensor-structured techniques (cf. Supplementary Method: Kohn-Sham effective potential computation using tensor-structured techniques). The electron-density $\rho(\mathbf{x})$ is computed in terms of the Kohn-Sham eigenstates as $\rho(\mathbf{x}) = 2 \sum_{i=1}^N f(\epsilon_i; \mu) |\Psi_i(\mathbf{x})|^2$, where $f(\epsilon; \mu)$ denotes the orbital occupancy factor given by the Fermi-Dirac distribution $f(\epsilon; \mu) = 1 / (1 + \exp(\frac{\epsilon - \mu}{k_B T}))$ with the Boltzmann constant k_B , the Fermi

energy μ , and the smearing temperature T . We note that Eq. (1) represents a non-linear eigenvalue problem, as the Kohn-Sham Hamiltonian depends on ρ , which in turn depends on the eigenstates. Thus, Eq. (1) is solved self-consistently via a self-consistent field (SCF) iteration²¹.

Tensor-structured algorithm using L_1 localized 1-D functions

In our previous work¹⁸, it was suggested that an additive separable approximation to the Kohn-Sham Hamiltonian can be used to construct a Tucker tensor basis that is systematically convergent. In particular, using a tensor-structured cuboidal domain Ω spanned by the tensor product of 1-D domains $\omega_{k=1,2,3}$, an additive separable approximation to the Kohn-Sham Hamiltonian ($\mathcal{H}_1(x_1) + \mathcal{H}_2(x_2) + \mathcal{H}_3(x_3) \approx \mathcal{H}(\mathbf{x})$) retains some features of the Hamiltonian, and thus presents a useful operator to generate reduced-order basis functions. To this end, the eigenfunctions of the additive separable approximation to the Hamiltonian, which constitute a Tucker tensor basis formed from the 1-D eigenfunctions of the separable parts of the Hamiltonian (\mathcal{H}_k , $k = 1, 2, 3$), are used to solve the Kohn-Sham equations. While an efficient basis, the global nature of the ensuing Tucker tensor basis limits the computational efficiency of the algorithms to solve the Kohn-Sham equations. In the proposed work, in place of the 1-D eigenfunctions of \mathcal{H}_k , we instead construct compressed modes preserving the subspace spanned by the 1-D eigenfunctions using the L_1 localization technique²². The obtained 1-D localized functions are then used to generate the 3-D Tucker tensor basis, which is localized in real-space and allows us to exploit the sparsity of the Kohn-Sham Hamiltonian represented in this basis for both computational efficiency and realizing reduced-order scaling in solving the Kohn-Sham equations. The various aspects of our tensor-structured algorithm are now presented, which include the generation of the additive separable approximation of the Kohn-Sham Hamiltonian, the evaluation of the L_1 localized 1-D functions, the construction of the localized Tucker tensor basis, the projection of the Kohn-Sham problem onto the

localized Tucker tensor basis, and the solution of the Kohn–Sham equations.

Construction of separable Hamiltonian

We seek to construct a separable approximation to the Kohn–Sham Hamiltonian $\mathcal{H}_1(x_1) + \mathcal{H}_2(x_2) + \mathcal{H}_3(x_3) \approx \mathcal{H}(\mathbf{x})$ based on a rank-1 approximation of the eigenfunction corresponding to the lowest eigenvalue. To this end, we consider the rank-1 representation for the eigenfunction as $\Psi'(\mathbf{x}) = \psi_1(x_1)\psi_2(x_2)\psi_3(x_3)$. Thus, the problem of computing the smallest eigenvalue of the Kohn–Sham Hamiltonian using the rank-1 approximation is given by the variational problem

$$\min_{\psi_k} L(\Psi') \quad \text{subject to: } \langle \Psi' | \Psi' \rangle = 1 \quad (2)$$

with the Lagrangian $L(\Psi') = \langle \Psi' | -\frac{1}{2}\nabla^2 + V_{\text{eff}}(\mathbf{x}) | \Psi' \rangle$. Upon taking the variations of the functional with respect to ψ_1 , ψ_2 and ψ_3 , we obtain three simultaneous 1-D problems

$$\mathcal{H}_k \psi_k = \alpha_k \psi_k, \quad k = 1, 2, 3. \quad (3)$$

As \mathcal{H}_k and α_k are parametrized by $\psi_{j \neq k}$, the three simultaneous 1-D problems represent a non-linear problem that can be solved self-consistently via SCF iteration (cf. Supplementary Method: Minimization problem for computing separable Hamiltonian approximation). Upon achieving self-consistency, the 1-D Hamiltonians (\mathcal{H}_k) we obtain represent the additive separable approximation of the Kohn–Sham Hamiltonian that we seek. The eigenfunctions of this

additive separable approximation to the Hamiltonian, which can be obtained as the tensor product of the 1-D eigenfunctions of \mathcal{H}_k ($k = 1, 2, 3$), constitute a complete basis, thus providing systematic convergence as will be demonstrated subsequently.

We note that the proposed approach represents one possibility of systematically constructing an additive separable approximation to the Kohn–Sham Hamiltonian, and other possibilities may exist. We also note that the resulting tensor-structured basis—the eigenbasis of $\mathcal{H}_1 + \mathcal{H}_2 + \mathcal{H}_3$ —is expected to be better than the plane-wave basis. To elaborate, the plane-wave basis is the eigenbasis of the Laplace operator (which is additive separable), whereas the additive separable approximation obtained via the proposed approach includes both the Laplace operator and an additive separable approximation of the Kohn–Sham potential V_{eff} , thus retaining some additional features of the Kohn–Sham Hamiltonian and providing a better basis than the plane-wave basis. The superior approximation properties of the proposed tensor-structured basis over the plane-wave basis will be demonstrated subsequently via numerical benchmark studies (cf. Tables 1 and 2).

We further note that the Kohn–Sham Hamiltonian changes during the course of the SCF iteration. In principle, the separable approximation to the Kohn–Sham Hamiltonian can be computed for each SCF iteration adapting the tensor-structured basis to the Hamiltonian in a given SCF iteration. However, in this work, we choose to keep the basis fixed after computing the tensor-structured basis in the first SCF iteration using the Kohn–Sham

Table 1. Comparison of the computational performance for Aluminum nano-particles.

Nano particle	Tucker				QE			Reference energy/atom
	Rank	# basis/atom	E/atom	Time	#basis/atom	E/atom	Time	
Al ₁₃	40	4923	−55.9965	0.00067	12403	−55.9993	0.00022	−56.0034*
Al ₁₄₇	70	2333	−56.6179	0.069	5078	−56.6198	0.028	−56.6274*
Al ₅₆₁	85	1095	−56.8119	0.69	4490	−56.8122	1.24	−56.8191*
Al ₂₀₅₇	120	840	−56.9192	7.96	5015	N/A	66.68	−56.9284 [†]
Al ₆₅₂₅	150	517	−57.0013	55.08	N/A	N/A	N/A	−57.0090 [†]

Comparison of the computational performance of the tensor-structured approach with Quantum Espresso (QE) for the benchmark aluminum nano-particles systems. All energies are reported in eV, and the computational times are reported in node-hrs per SCF iteration. The plane-wave cut-off employed for QE calculations to target the desired accuracy is 25 Ha. Full ground-state calculations were performed using the tensor-structured approach for all systems. In the case of QE, full-ground-state calculations were performed for the systems where the ground-state energies are provided, whereas for Al₂₀₅₇ only a few SCF iterations were performed to compute the stable SCF time due to significantly increased computational cost. Al₆₅₂₅ system was beyond reach using QE. The reference energies are computed using QE (*) with higher plane-wave cut-off 55 Ha for smaller systems. The reference energies for the larger systems are obtained using DFT-FE (†).

Table 2. Comparison of the computational performance for silicon quantum dots.

Si quantum dots	Tucker				QE			Reference energy/atom
	Rank	# basis/atom	E/atom	Time	# basis/atom	E/atom	Time	
Si ₁₀ H ₁₆	45	3505	−51.0271	0.0065	7048	−51.0279	0.00014	−51.0339*
Si ₂₂₀ H ₁₄₄	80	1407	−71.3841	0.094	2534	−71.3839	0.096	−71.3930*
Si ₅₂₅ H ₂₇₆	90	910	−76.1182	0.96	2251	−76.1194	1.12	−76.1279*
Si ₁₂₁₄ H ₅₀₄	100	582	−80.8627	3.85	2132	N/A	20.01	−80.8717 [†]
Si ₆₀₄₇ H ₁₃₀₈	140	373	−91.5659	67.49	N/A	N/A	N/A	−91.5741 [†]

Comparison of the computational performance of the tensor-structured approach with Quantum Espresso (QE) for the benchmark Silicon quantum dots systems. All energies are reported in eV, and the computational times are reported in node-hrs per SCF iteration. The plane-wave cut-off employed for QE calculations to target the desired accuracy is 20 Ha. Full ground-state calculations were performed using the tensor-structured approach for all systems. In the case of QE, full-ground-state calculations were performed for the systems where the ground-state energies are provided, whereas for Si₁₂₁₄H₅₀₄ only a few SCF iterations were performed to compute the stable SCF time due to significantly increased computational cost. Si₆₀₄₇H₁₃₀₈ system was beyond reach using QE. The reference energies are computed using QE (*) with higher plane-wave cut-off 50 Ha for smaller systems. The reference energies for the larger systems are obtained using DFT-FE (†).

potential obtained from the superposition of atomic densities. This is motivated from our numerical studies on the benchmark examples involving both metallic and insulating systems, which demonstrate that the approximation properties of the resulting tensor-structured basis are not substantially altered. In particular, the difference in the ground-state energies obtained by either using a fixed basis constructed in the first SCF or adapting the basis in every SCF is substantially smaller than the basis discretization error. We refer to Supplementary Note 1 for data supporting this observation (cf. Supplementary Table 1).

Computation of L_1 localized 1-D functions

The tensor-structured basis computed using the 1-D eigenfunctions of \mathcal{H}_k represents an efficient basis. However, the global nature of the basis limits the computational efficiency and scaling (with system-size) of the solution to the Kohn–Sham equations. To this end, we use the L_1 localization approach²² to construct a spatially localized tensor-structured basis that is a close approximation to the original tensor-structured basis. The localized basis is obtained by solving the following variational problem (for $k = 1, 2, 3$):

$$\min_{\Psi'_k \in \mathbb{R}^{n \times N_k}} \frac{1}{\mu} |\Psi'_k| + \text{Tr}(\Psi'_k{}^T \mathbf{H}_k \Psi'_k) \quad \text{s.t.} \quad \Psi'_k{}^T \Psi'_k = \mathbf{I}, \quad (4)$$

where \mathbf{H}_k is the matrix representation of \mathcal{H}_k in a suitable orthogonal basis with dimension n , Ψ'_k denotes the representation of N_k trial localized functions in the chosen basis, and μ is a parameter controlling the trade-off between the representability of the original eigensubspace and the locality of the 1-D functions, with $|\cdot|$ denoting the L_1 norm of the matrix. The minimizer of this variational problem, henceforth denoted as Ψ'_k , provides localized functions whose span closely approximates the eigensubspace of the lowest N_k eigenfunctions of \mathcal{H}_k , as will be demonstrated subsequently. We refer to the “Methods” section for the solution procedure employed to solve the aforementioned variational problem.

Construction of the localized 3-D Tucker tensor basis \mathbb{T}^L

The 1-D localized functions whose span is a close approximation to the subspace spanned by the 1-D eigenfunctions of \mathcal{H}_k are subsequently used to construct the 3-D Tucker tensor basis. Denoting the 1-D localized functions as $\psi_{1,i_1}^L(x_1)$, $\psi_{2,i_2}^L(x_2)$, $\psi_{3,i_3}^L(x_3)$, the 3-D localized tensor-structured basis functions T_l^L are given by

$$T_l^L(\mathbf{x}) = \psi_{1,i_1}^L(x_1) \psi_{2,i_2}^L(x_2) \psi_{3,i_3}^L(x_3), \quad (5)$$

where $1 \leq i_d \leq R_d$ and l is the composite index $l = (i_1, i_2, i_3)_{1 \leq i_d \leq R_d}$. The rank of the Tucker tensor basis is given by (R_1, R_2, R_3) which denotes the number of localized 1-D functions in each direction. The space spanned by the 3-D localized tensor-structured basis functions is denoted as \mathbb{T}^L .

Projection of the Kohn–Sham Hamiltonian onto \mathbb{T}^L

The Kohn–Sham Hamiltonian is projected onto \mathbb{T}^L spanned by the 3-D tensor-structured localized basis functions. We note that the Kohn–Sham effective potential V_{eff} is a functional of the electron-density ρ , and is comprised of a local-part $V_{\text{eff}}^{\text{loc}}$ (local in real-space) and a non-local part $V_{\text{ext}}^{\text{nl}}$. $V_{\text{eff}}^{\text{loc}}$ includes the Hartree potential (V_H), the exchange–correlation potential and the local-part of the pseudopotential, whereas $V_{\text{ext}}^{\text{nl}}$ comprises of the non-local projectors of the pseudopotential (cf. Supplementary Method: Kohn–Sham effective potential computation using tensor-structured techniques). The convolution integral involved in the evaluation of V_H can be efficiently computed using a low-rank Tucker tensor decomposition of the electron-density (R_ρ denoting the rank of the decomposition) and approximating the Coulomb integral by a

series of Gaussian functions^{14,23} (cf. Supplementary Method: Kohn–Sham effective potential computation using tensor-structured techniques). Subsequently, a low-rank Tucker approximation of $V_{\text{eff}}^{\text{loc}}$ and $V_{\text{ext}}^{\text{nl}}$ is utilized, with R_V and R_V^{nl} denoting the corresponding ranks, respectively. Denoting the low-rank Tucker approximation of the effective potential V_{eff} by \tilde{V}_{eff} , whose approximation error decays exponentially with the Tucker rank^{14,15}, the projection of the Kohn–Sham Hamiltonian onto \mathbb{T}^L is given by

$$\tilde{H}_{l,j}^L = \langle T_l^L | -\frac{1}{2} \nabla^2 + \tilde{V}_{\text{eff}}(\rho; \mathbf{R}) | T_j^L \rangle. \quad (6)$$

We note that the low-rank representation \tilde{V}_{eff} reduces the integrals involved in Eq. (6) to tensor products of one-dimensional integrals, thereby facilitating the efficient evaluation of Hamiltonian matrix elements (cf. Supplementary Method: Computation of projected Hamiltonian).

The projected Kohn–Sham Hamiltonian matrix elements are computed, and a truncation tolerance is introduced to zero out the Hamiltonian matrix elements below the tolerance. This truncation is performed in every SCF iteration, which improves the sparsity of the Hamiltonian matrix and thereby reducing the memory footprint of the calculation. Furthermore, the sparsity of the Hamiltonian matrix also reduces the computational complexity of the algorithm employed to solve the Kohn–Sham equations, which is discussed subsequently. We note that the error in the ground-state energy, corresponding to the truncation of the Hamiltonian matrix elements, systematically decreases with tighter truncation tolerance. We refer to Supplementary Note 2 (Supplementary Table 2) which provides data to support this observation. Furthermore, we note that a truncation tolerance of 10^{-4} Ha provides excellent sparsity in the Hamiltonian, with the error resulting from the truncation being significantly smaller than the basis discretization error and the desired chemical accuracy in ground-state energy (cf. Supplementary Note 2).

Computation of the occupied eigenstates

The discretized Kohn–Sham problem, corresponding to Eq. (1), in the localized orthonormal tensor-structured basis is given by the standard eigenvalue problem

$$\mathbf{H}^L \Psi_i = \epsilon_i \Psi_i, \quad i \in \{1, \dots, N\} \quad (7)$$

where \mathbf{H}^L denotes the truncated sparse Kohn–Sham Hamiltonian matrix. We use the Chebyshev filtering-based subspace iteration (ChFSI)²⁴ to efficiently solve the Kohn–Sham equations. The ChFSI method has been demonstrated to be efficient with good parallel scalability for real-space implementations of DFT^{25–27}. In the ChFSI method, in each SCF iteration, a suitably constructed Chebyshev filter using \mathbf{H}^L is employed to construct a close approximation to the relevant eigensubspace of the occupied states. The action of the Chebyshev filter on a given subspace can be cast as a recursive iteration involving matrix-vector multiplications between \mathbf{H}^L and vectors obtained during the course of recursive iteration. Since \mathbf{H}^L is sparse, the computational complexity of the Chebyshev filtering operation scales as $\mathcal{O}(R^3 N)$, where $R = \max\{R_1, R_2, R_3\}$. In ChFSI, the Chebyshev filtered vectors are orthogonalized using a Gram–Schmidt orthogonalization procedure, and subsequently the Kohn–Sham eigenstates are computed by projecting \mathbf{H}^L onto the Chebyshev-filtered subspace and diagonalizing this projected Hamiltonian. The computational complexity of the orthogonalization procedure and the subspace projection scales as $\mathcal{O}(R^3 N^2)$ while the diagonalization cost scales as $\mathcal{O}(N^3)$. As demonstrated in Supplementary Table 3 (cf. Supplementary Note 3), Chebyshev filtering, which scales linearly with N , remains the dominant cost even at 25,000 electrons for the various benchmark examples considered in this work. We note that, at even larger system sizes, other costs that exhibit quadratic-scaling (orthogonalization and subspace projection) and cubic-scaling (diagonalization) with N can




Fig. 2 1-D functions from separable Hamiltonian. 1-D functions in x_1 direction constructed from the additive separable approximation of the Kohn–Sham Hamiltonian for Al_{147} nano-particle. Top: Lowest five eigenfunctions of \mathcal{H}_1 . Bottom: The corresponding L_1 localized 1-D functions.

start to compete. However, at such a point, explicit diagonalization can be avoided, and already developed ideas¹³ of localizing the Chebyshev-filtered vectors in conjunction with Fermi-operator expansion can be adopted to retain the reduced-order scaling for systems with or without a gap. We refer to the methods section for details on the various numerical parameters employed in conducting the DFT calculations using the SCF iteration approach via the ChFSI technique.

Eigensubspace representability of L_1 localized 1-D functions

We now demonstrate the ability of the L_1 localized functions to closely approximate the eigensubspace of \mathcal{H}_k using Al_{147} nano-particle with icosahedral symmetry. We compute the additive separable approximation of the Kohn–Sham Hamiltonian for this nano-particle, and, then compute the lowest 70 eigenstates of \mathcal{H}_k . We subsequently use the L_1 localization approach to compute the localized functions that are a close approximation to the eigensubspace. Figure 2 shows the lowest five eigenfunctions of \mathcal{H}_1 (one of the 1-D separable Hamiltonian) (top) and the corresponding 1-D localized functions (bottom). We refer to Supplementary Note 4 (Supplementary Fig. 1) for an illustration of all 70 eigenstates and the corresponding 1-D localized functions. It is evident that, while the eigenfunctions are global in nature, the functions obtained from the L_1 localization approach are localized in real-space. This locality is key to the sparsity of the Kohn–Sham Hamiltonian matrix in the Tucker tensor basis, and the resulting computational efficiency.

In order to demonstrate the accuracy of the L_1 localization approach in closely approximating the eigensubspace of the separable Hamiltonian, we consider the first 70 eigenstates of \mathcal{H}_1 and the eigenvalues of the matrix $K_{ij} = \langle \psi_{1,i}^L | \mathcal{H}_1 | \psi_{1,j}^L \rangle$, $1 \leq i, j \leq 70$. Figure 3 shows the eigenvalues of \mathcal{H}_1 and the eigenvalues of K_{ij} . It is interesting to note that the eigenvalues of the first 65 states are almost identical, with only slight deviations for the higher states. This demonstrates that the space spanned by the localization functions obtained using the L_1 localization approach is a close approximation to the eigensubspace of \mathcal{H}_k . We also note here that better accuracy can be achieved, when necessary, by simply increasing the size of N_k to be solved for in Eq. (4). In order to assess the accuracy afforded by the localization procedure in the ground-state energy, we computed the ground-state energy of Al_{147} using the 3-D localized basis with rank 70, and compared that with the energy obtained using the eigenbasis of $\mathcal{H}_1 + \mathcal{H}_2 + \mathcal{H}_3$ (i.e., without localization) of the same rank. The energy obtained without localization is -56.61882 eV per atom in comparison to -56.61893 eV per atom obtained using localization. Thus, the error introduced due to localization is ~ 0.1 meV per




Fig. 3 Representability of the 1-D localized functions. Comparison of the eigenvalues of the 1-D separable Hamiltonian in x_1 direction of Al_{147} nano-particle (marked with blue circle) with the eigenvalues of $K_{ij} = \langle \psi_{1,i}^L | \mathcal{H}_k | \psi_{1,j}^L \rangle$ (marked with red cross and labeled as SOC).

atom, which is substantially smaller than the basis discretization error of ~ 8.5 meV per atom corresponding to rank 70 (reference energy is -56.6274 eV per atom; cf. Table 1).

Convergence of the tensor-structured basis

We next investigate the convergence properties of the 3-D Tucker tensor basis constructed from the 1-D localized functions. For the convergence study we consider two benchmark problems: (i) C_{60} (fullerene) molecule; (ii) tris (bipyridine) ruthenium, a transition metal complex. We note that these systems have no tensor structure symmetry and serve as stringent benchmarks to assess the convergence and accuracy afforded by the proposed Tucker tensor basis. The ground-state energy for these molecules is computed for various Tucker tensor ranks R ($R_1 = R_2 = R_3 = R$), with the 3-D Tucker tensor basis getting systematically refined with increasing R . In this study, R_ρ , R_V and R_V^{nl} are chosen stringently such that the resulting errors are significantly smaller than the basis discretization errors, and they are held fixed for increasing R . In particular, we used $R_\rho = 45$, $R_V = 65$ and $R_V^{\text{nl}} = 20$ for fullerene, and $R_\rho = 80$, $R_V = 55$ and $R_V^{\text{nl}} = 20$ for tris (bipyridine) ruthenium. The basis discretization error (convergence with respect to R) is measured with respect to a well-converged QE result. The converged QE ground-state energies for fullerene molecule is taken to be -155.1248 eV per atom ($E_{\text{cut}} = 60$ Ha) and that of tris (bipyridine) ruthenium is taken to be -118.2128 eV per atom ($E_{\text{cut}} = 65$ Ha).

Figure 4a, b shows the relative error in the ground-state energy for the various ranks of the Tucker tensor basis. It is evident from these results that the Tucker tensor basis constructed using our approach provides an exponential convergence in the ground-state energy with increasing Tucker rank. The convergence study of these molecules suggests that the proposed tensor-structured technique provides systematic convergence with high accuracy and is capable of handling generic materials systems, including those involving transition metals.

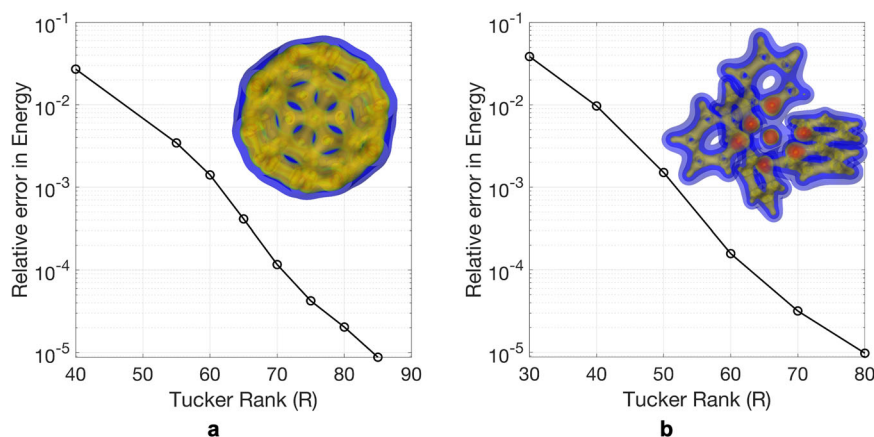


Fig. 4 Convergence with respect to the Tucker rank. **a** Fullerenes and **b** Tris (bipyridine) ruthenium. The electron-density isocontours are provided in the inset.

Performance and scaling analysis

To study the performance and scaling with system-size of the proposed tensor-structured approach for DFT calculations, we consider two classes of benchmark systems: (i) Aluminum nano-particles of various sizes ranging from 13 atoms to 6525 atoms; (ii) Silicon quantum dots with system-sizes ranging from 26 atoms to 7355 atoms. These benchmark systems constitute materials systems with and without a gap, thus allowing us to assess the system-size scaling for both classes of materials. In order to compare the efficiency of the proposed tensor-structured approach with the widely used plane-wave DFT calculations, we also conducted the DFT calculations using QE wherever possible. For the sake of estimating the computational efficiency, the energy cut-off for QE and the Tucker rank are chosen such that the ground-state energy is converged to within 10 meV per atom measured with respect to a highly converged reference calculation. The reference ground-state energies are obtained from QE (using a high-energy cut-off) for smaller systems and the DFT-FE code²⁶—a massively parallel real-space code for large-scale DFT calculations—for larger system-sizes. The cell size for plane-wave calculations is chosen such that each atom is at least 10 Bohr away from the boundary, which was needed to obtain the desired accuracy.

In these benchmark calculations, the additive separable approximation to the Kohn–Sham Hamiltonian is computed only in the first SCF iteration and the resulting 3-D Tucker basis is held fixed for subsequent SCF iterations. We note that the approximation properties of an adaptive Tucker basis (where the basis is regenerated for every SCF iteration) and the fixed Tucker basis are similar, with the differences in the accuracy being substantially smaller than the basis discretization error for a given Tucker rank. We refer to Supplementary Note 1 (Supplementary Table 1) which provides data supporting this observation. In the tensor-structured calculations reported subsequently, all the numerical parameters—ranks for approximating electron-density and effective Kohn–Sham potential (R_ρ, R_V, R_V^{nl}) in \mathbf{H}^L , and the truncation tolerance adopted in computing \mathbf{H}^L —are chosen such that the resulting errors are substantially smaller than the basis discretization error in the ground-state energy associated with the Tucker rank R of the localized 3-D Tucker tensor basis and the desired chemical accuracy. In particular, for the Aluminum nano-particles we used $R_\rho = 40$, $R_V = 50$ and $R_V^{nl} = 25$. In the case of Silicon quantum dots, we used $R_\rho = 55$, $R_V = 55$ and $R_V^{nl} = 25$. The truncation tolerance in computing \mathbf{H}^L was chosen to be 10^{-4} Ha for all the calculations, which provides excellent sparsity for \mathbf{H}^L , and, importantly, the sparsity is either steady or improves with

increasing system-size (cf. Supplementary Tables 2 and 3). The error in ground-state energy associated with this choice of truncation tolerance is ~ 1 meV per atom (cf. Supplementary Table 2), as opposed to the targeted accuracy in this study of being within 10 meV per atom of reference energies.

Aluminum nano-particles

The computational efficiency afforded by the proposed tensor-structured approach in comparison to QE for the various aluminum nano-particles with icosahedral symmetry considered in this work is provided in Table 1. We note that the Tucker rank required to achieve the desired accuracy only grows slowly with increasing system-size. Importantly, we note that the number of basis functions needed to achieve the desired accuracy using the localized Tucker basis is smaller than the plane-wave basis. As previously discussed, this is a consequence of the superior approximation properties of the Tucker tensor basis generated as the eigenbasis of an additive separable approximation of the Kohn–Sham Hamiltonian that in addition to the Laplace operator retains some characteristics of the Kohn–Sham potential, as opposed to the plane-wave basis which corresponds to the eigenbasis of the Laplace operator. In terms of the computational time, while QE is more efficient for the smaller system-sizes, the tensor-structured approach starts to substantially outperform for larger system-sizes. Notably, for Al_{2057} , the tensor-structured approach is 8-fold more efficient. Furthermore, using the computational times, the scaling of the proposed tensor-structured approach is estimated to be around $\mathcal{O}(N_e^{1.78})$ with N_e denoting the number of electrons (cf. Fig. 5a). Notably, the scaling with system-size is sub-quadratic for this metallic system over system-sizes spanning many thousands of atoms, as opposed to the cubic-scaling complexity for plane-wave DFT calculations. We note that this is a consequence of the slow growth of the Tucker rank with system-size that results in a sub-linear growth of the total number of basis functions with system-size. The breakdown of the computational costs for the various steps of the calculation is provided in Supplementary Note 3 (cf. Supplementary Table 3).

Silicon quantum dots

Table 2 compares the computational performance of the proposed tensor-structured approach with QE for a wide range of silicon quantum dots passivated with hydrogen. As in the case of aluminum nano-particles, the Tucker tensor basis is more efficient than the plane-wave basis in terms of the number of basis functions to attain the desired accuracy. In terms of computational

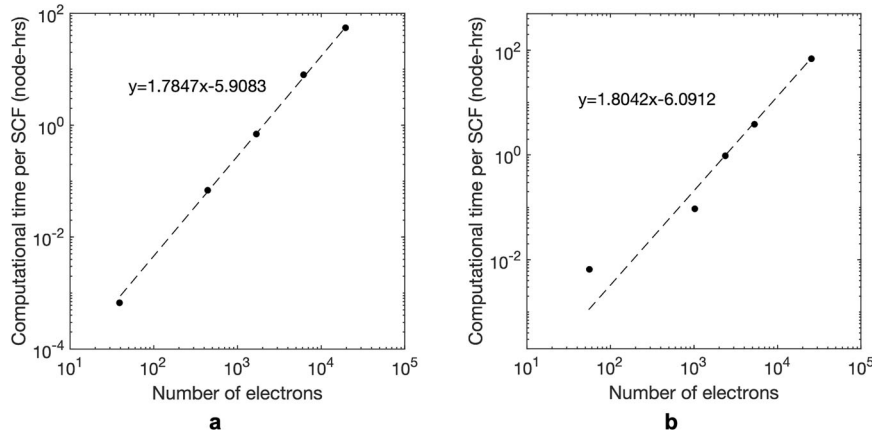


Fig. 5 Computational complexity of the tensor-structured approach. **a** Aluminum nano-particles $\mathcal{O}(N_e^{1.8})$ and **b** Silicon quantum dots $\mathcal{O}(N_e^{1.8})$.

time, the proposed tensor-structured approach starts competing with QE beyond a few hundred atoms, and significantly outperforms for larger systems. Moreover, the scaling with system-size for the tensor-structured algorithm, for a range of system-sizes with the largest containing 7355 atoms, is estimated to be $\mathcal{O}(N_e^{1.8})$ (cf. Fig. 5b). Notably, this scaling is similar to that obtained for aluminum nano-particles as the algorithm treats systems with and without a gap on a similar footing.

DISCUSSION

We have presented a tensor-structured algorithm, where the Tucker tensor basis is constructed as a tensor product of localized 1-D functions whose span closely approximates the eigensubspace of a suitably constructed additive separable approximation to the Kohn–Sham Hamiltonian. The resulting localized Tucker tensor basis, that is adapted to the Kohn–Sham Hamiltonian, provides a systematically convergent basis as evidenced by the exponential convergence of the ground-state energy with increasing Tucker rank. Our numerical studies on the computational performance suggest that the proposed approach exhibits sub-quadratic scaling (with system-size) over a wide range of system-sizes with the largest involving many thousands of atoms. Importantly, sub-quadratic scaling is realized for both systems with and without a gap, as the algorithm treats both metallic and insulating systems on an equal footing. Further, comparing the computational efficiency of the proposed approach with QE, we observe significant outperformance for system-sizes beyond 5000 electrons.

We note that the sub-quadratic scaling is a consequence of the slow growth of the Tucker rank with system-size, with the resulting number of basis functions growing sub-linearly with system-size even for systems containing many thousands of atoms. By combining the proposed approach with reduced-order scaling techniques that exploit the locality of the wavefunctions in real-space, there is further room to reduce the scaling with system-size and is a useful future direction to pursue. Further, the proposed tensor-structured approach is amenable to GPU acceleration that can further substantially enhance the computational efficiency of the approach, and is currently being pursued. We note that the benchmark systems presented here were restricted to non-periodic calculations as the proposed tensor-structured approach was implemented in a non-periodic setting as a first step of an ongoing effort. However, we remark that the ideas presented here are generic and can be extended to periodic calculations.

METHODS

Tucker tensor representation

Tucker tensor representation is a higher-order generalization of principal component analysis for a tensor^{28–30}. An N-way tensor is approximated by a Tucker tensor through Tucker decomposition with a smaller N-way core tensor and N factor matrices whose columns are the rank-1 components from the decomposition^{28–32}. In the scope of this work, the discussion is restricted to three-way tensor. Let $A \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ be a real-valued three-way tensor of size $I_1 \times I_2 \times I_3$ indexed by a set of integers (i_1, i_2, i_3)

$$A_{(i_1, i_2, i_3)} = a_{i_1 i_2 i_3}, \quad (8)$$

where $i_d \in \{1, 2, \dots, I_d\}$, $I_d \in \mathbb{N}$ and $d \in \{1, 2, 3\}$ denotes the dimensions. A Tucker tensor representation of the tensor A with decomposition rank $\mathbf{R} = (R_1, R_2, R_3)$ for each direction has the form

$$A \approx A^{(\mathbf{R})} = \sum_{i_1=1}^{R_1} \sum_{i_2=1}^{R_2} \sum_{i_3=1}^{R_3} \sigma \mathbf{u}_1^{i_1} \circ \mathbf{u}_2^{i_2} \circ \mathbf{u}_3^{i_3}, \quad (9)$$

where $\sigma \in \mathbb{R}^{R_1 \times R_2 \times R_3}$ is the core tensor, $\mathbf{u}_d^{i_d} \in \mathbb{R}^{I_d}$ forms the factor matrix $\mathbf{U}_d \in \mathbb{R}^{I_d \times R_d}$, and "o" denotes the vector outer product $(\mathbf{u}_1^{i_1} \circ \mathbf{u}_2^{i_2} \circ \mathbf{u}_3^{i_3})_{i_1, i_2, i_3} := u_{1, i_1}^{i_1} u_{2, i_2}^{i_2} u_{3, i_3}^{i_3}$. The core tensor stores the coefficients σ_{i_1, i_2, i_3} for each rank-1 tensor $\mathbf{u}_1^{i_1} \circ \mathbf{u}_2^{i_2} \circ \mathbf{u}_3^{i_3}$. The core tensor and the factor matrices can be viewed as the higher-order correspondence of the singular values and unitary matrices. The tensor representation of the Tucker form can be obtained with the higher-order singular value decomposition (HOSVD). The HOSVD flattens the given tensor in three directions and employs singular-value decomposition to obtain the factor matrices. The factor matrices are then used to contract with the given tensor to obtain the core tensor. We refer to³² and³³ for details of HOSVD and further review on tensor decomposition and tensor analysis. In this work, an MPI implementation in C++ for Tucker decomposition is used^{34,35}.

Computation of L_1 localized 1-D functions

The variational problem in Eq. (4) is solved by the splitting orthogonality constraint algorithm (SOC). The SOC algorithm introduces two auxiliary variables controlling the orthonormality and the locality constraints and translates the constrained minimization problem into unconstrained problems. The unconstrained problems are then optimized iteratively. The SOC algorithm is capable of providing a set of compressed modes with good locality, yet preserving the orthogonality. We refer to Ozoliński et al.²² and Lai and Osher³⁶ for a detailed discussion, and present the algorithm in the context of this work in the Supplementary Method: SOC Algorithm for the sake of completeness.

Ground-state DFT calculations

All calculations are performed using the norm-conserving Troullier–Martin pseudopotentials in Kleinmann–Bylander form^{37,38}, and a local density approximation (LDA) for the exchange–correlation functional^{39–41}. In the ChFSI method employed to solve the Kohn–Sham equations, we use a Chebyshev filter constructed using polynomial degree of 10–20 for the

various materials systems reported in this study. The n -stage Anderson mixing scheme⁴² is employed in the SCF iteration. We used Fermi–Dirac smearing with $T = 500$ K for computing the orbital occupancies. The performance benchmarks are obtained on compute nodes comprising of 68-core Intel Xeon Phi Processor 7250 and 96 GB memory per node. All calculations were performed in the good parallel-scaling regime to ensure that the obtained computational times (node-hours) are representative of the computational efficiency of the approach.

DATA AVAILABILITY

The authors declare that all the data supporting the results of this study are available upon reasonable request to the corresponding author.

CODE AVAILABILITY

The code used to perform Tucker DFT calculations is available upon reasonable request to the corresponding author.

Received: 12 October 2020; Accepted: 8 March 2021;

Published online: 12 April 2021

REFERENCES

- Hohenberg, P. & Kohn, W. Inhomogeneous electron gas. *Phys. Rev.* **136**, B864–B871 (1964).
- Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133–A1138 (1965).
- Goedecker, S. Linear scaling electronic structure methods. *Rev. Mod. Phys.* **71**, 1085–1123 (1999).
- Bowler, D. R. & Miyazaki, T. $\mathcal{O}(N)$ methods in electronic structure calculations. *Rep. Prog. Phys.* **75**, 036503 (2012).
- Li, X.-P., Nunes, R. W. & Vanderbilt, D. Density-matrix electronic-structure method with linear system-size scaling. *Phys. Rev. B* **47**, 10891–10894 (1993).
- Mauri, F., Galli, G. & Car, R. Orbital formulation for electronic-structure calculations with linear system-size scaling. *Phys. Rev. B* **47**, 9973–9976 (1993).
- Goedecker, S. & Colombo, L. Efficient linear scaling algorithm for tight-binding molecular dynamics. *Phys. Rev. Lett.* **73**, 122–125 (1994).
- Kim, J., Mauri, F. & Galli, G. Total-energy global optimizations using non-orthogonal localized orbitals. *Phys. Rev. B* **52**, 1640–1648 (1995).
- Stephan, U. & Drabold, D. A. Order- N projection method for first-principles computations of electronic quantities and wannier functions. *Phys. Rev. B* **57**, 6391–6407 (1998).
- Haynes, P. D., Skylaris, C. K., Mostofi, A. A. & Payne, M. C. ONETEP: linear scaling density functional theory with local orbitals and plane waves. *Phys. Stat. Sol. (b)* **243**, 2489–2499 (2006).
- Garcia-Cervera, C. J., Lu, J., Xuan, Y. & E, W. Linear-scaling subspace-iteration algorithm with optimally localized nonorthogonal wave functions for Kohn–Sham density functional theory. *Phys. Rev. B* **79**, 115110 (2009).
- Lin, L., Chen, M., Yang, C. & He, L. Accelerating atomic orbital-based electronic structure calculation via pole expansion and selected inversion. *J. Phys.: Condens. Matter* **25**, 295501 (2013).
- Motamarri, P. & Gavini, V. Subquadratic-scaling subspace projection method for large-scale Kohn–Sham density functional theory calculations using spectral finite-element discretization. *Phys. Rev. B* **90**, 115127 (2014).
- Khoromskij, B., Khoromskaia, V., Chinnamsetty, S. & Flad, H.-J. Tensor decomposition in electronic structure calculations on 3d cartesian grids. *J. Comput. Phys.* **228**, 5749–5762 (2009).
- Hackbusch, W. & Khoromskij, B. N. Tensor-product approximation to operators and functions in high dimensions. *J. Complexity* **23**, 697–714 (2007).
- Khoromskij, B. N. & Khoromskaia, V. Multigrid accelerated tensor approximation of function related multidimensional arrays. *SIAM J. Sci. Comput.* **31**, 3002–3026 (2009).
- Blesgen, T., Gavini, V. & Khoromskaia, V. Approximation of the electron density of aluminum clusters in tensor-product format. *J. Comput. Phys.* **231**, 2551–2564 (2012).
- Motamarri, P., Gavini, V. & Blesgen, T. Tucker-tensor algorithm for large-scale Kohn–Sham density functional theory calculations. *Phys. Rev. B* **93**, 125104 (2016).
- Giannozzi, P. et al. QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials. *J. Phys.: Condens. Matter* **21**, 395502 (2009).
- Giannozzi, P. et al. Advanced capabilities for materials modelling with QUANTUM ESPRESSO. *J. Phys.: Condens. Matter* **29**, 465901 (2017).
- Martin, R. M. *Electronic Structure: Basic Theory and Practical Methods* (Cambridge University Press, 2004).
- Ozoliņš, V., Lai, R., Caffisch, R. & Osher, S. Compressed modes for variational problems in mathematics and physics. *Proc. Natl Acad. Sci. USA* **110**, 18368–18373 (2013).
- Braess, D. & Hackbusch, W. On the efficient computation of high-dimensional integrals and the approximation by exponential sums. In *Multiscale, Nonlinear and Adaptive Approximation* (eds DeVore, R. A. & Kunoth, A.) 39–74 (Springer, 2009).
- Zhou, Y., Saad, Y., Tiago, M. L. & Chelikowsky, J. R. Parallel self-consistent-field calculations via Chebyshev-filtered subspace acceleration. *Phys. Rev. E* **74**, 066704 (2006).
- Motamarri, P., Nowak, M., Leiter, K., Knap, J. & Gavini, V. Higher-order adaptive finite-element methods for Kohn–Sham density functional theory. *J. Comput. Phys.* **253**, 308–343 (2013).
- Motamarri, P. et al. DFT-FE—A massively parallel adaptive finite-element code for large-scale density functional theory calculations. *Comput. Phys. Commun.* **246**, 106853 (2020).
- Das, S. et al. Fast, scalable and accurate finite-element based ab initio calculations using mixed precision computing: 46 pflops simulation of a metallic dislocation system. In *Proc. International Conference for High Performance Computing, Networking, Storage and Analysis* (eds Taufer, M., Balaji, P. & Peña, A. J.) SC '19 (Association for Computing Machinery, 2019).
- Tucker, L. R. Implications of factor analysis of three-way matrices for measurement of change. In *Problems in Measuring Change* (ed. Harris, C. W.) 122–137 (University of Wisconsin Press, 1963).
- Tucker, L. R. The extension of factor analysis to three-dimensional matrices. In *Contributions to Mathematical Psychology* (eds Frederiksen, N. & Gulliksen, H.) 110–127 (Holt, Rinehart and Winston, 1964).
- Tucker, L. R. Some mathematical notes on three-mode factor analysis. *Psychometrika* **31**, 279–311 (1966).
- De Lathauwer, L., De Moor, B. & Vandewalle, J. A multilinear singular value decomposition. *SIAM J. Matrix Anal. Appl.* **21**, 1253–1278 (2000).
- Kolda, T. G. & Bader, B. W. Tensor decompositions and applications. *SIAM Rev.* **51**, 455–500 (2009).
- Hackbusch, W. *Tensor Spaces and Numerical Tensor Calculus*, Vol. 42 of *Springer Series in Computational Mathematics* (Springer, 2012).
- Austin, W., Ballard, G. & Kolda, T. G. Parallel tensor compression for large-scale scientific data. In *IPDPS'16: Proc. 30th IEEE International Parallel and Distributed Processing Symposium* (ed. Faraj, D. A.) 912–922 (IEEE, 2016).
- Ballard, G., Klinvex, A. & Kolda, T. G. TuckerMPI: a parallel C++/MPI software package for large-scale data compression via the tucker tensor decomposition. *ACM Trans. Math. Softw.* **46**, 13:1–13:31 (2020).
- Lai, R. & Osher, S. A splitting method for orthogonality constrained problems. *J. Sci. Comput.* **58**, 431–449 (2014).
- Troullier, N. & Martins, J. L. Efficient pseudopotentials for plane-wave calculations. *Phys. Rev. B* **43**, 1993–2006 (1991).
- Kleinman, L. & Bylander, D. M. Efficacious form for model pseudopotentials. *Phys. Rev. Lett.* **48**, 1425–1428 (1982).
- Ceperley, D. M. & Alder, B. J. Ground state of the electron gas by a stochastic method. *Phys. Rev. Lett.* **45**, 566–569 (1980).
- Perdew, J. P. & Zunger, A. Self-interaction correction to density-functional approximations for many-electron systems. *Phys. Rev. B* **23**, 5048–5079 (1981).
- Perdew, J. P. & Wang, Y. Accurate and simple analytic representation of the electron-gas correlation energy. *Phys. Rev. B* **45**, 13244–13249 (1992).
- Anderson, D. G. Iterative procedures for nonlinear integral equations. *J. ACM* **12**, 547–560 (1965).

ACKNOWLEDGEMENTS

We gratefully acknowledge the support of the Air Force Office of Scientific Research through grant number FA-9550-17-0172 under the auspices of which this work was conducted. V.G. also gratefully acknowledges the support of the Army Research office through the DURIP grant W911NF1810242, which provided computational resources for this work.

AUTHOR CONTRIBUTIONS

C.-C.L. conducted the research work under the guidance of P.M. and V.G. All authors contributed and reviewed the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41524-021-00517-5>.

Correspondence and requests for materials should be addressed to V.G.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021