Variationally localized search direction method for constrained optimization of non-orthogonal, localized orbitals in electronic structure calculations

Álvaro Ruiz-Serrano and Chris-Kriton Skylaris

Citation: The Journal of Chemical Physics **139**, 164110 (2013); doi: 10.1063/1.4826164
View online: http://dx.doi.org/10.1063/1.4826164
View Table of Contents: http://scitation.aip.org/content/aip/journal/jcp/139/16?ver=pdfcov
Published by the AIP Publishing

Variationally localized search direction method for constrained optimization of non-orthogonal, localized orbitals in electronic structure calculations

Álvaro Ruiz-Serrano and Chris-Kriton Skylaris[a)]
School of Chemistry, University of Southampton, Highfield, Southampton SO17 1BJ, United Kingdom

(Received 2 July 2013; accepted 4 October 2013; published online 23 October 2013)

A new method for the constrained optimization of non-orthogonal, spatially localized orbitals using direct energy minimization techniques, in the context of electronic structure calculations, is presented. The variationally localized search direction (VLSD) method, as it was named, ensures that strict localization constraints are imposed upon the search direction vectors exactly, analytically and in a fully variational fashion. In contrast, the truncated search direction (TSD) method, of standard use in many electronic structure approaches with localization constraints, relies on the approximation that the truncated search direction vectors of the unconstrained problem resemble the exact search direction vectors of the constrained problem. With the TSD method, in order to maintain the localization constraints, a part of the pre-calculated information that is stored in the search direction vectors has to be deleted via an *ad hoc*, non-variational truncation step. The results on an extensive set of test molecules show that, in general, calculations with the VLSD method require less iterations to converge than with the TSD method for any size of the localization region. It was found that in calculations on certain systems where the TSD method is forced to delete a very large amount of information, the VLSD method is capable of achieving convergence in up to three times less iterations. Validation tests show that structural and electronic properties calculated with either method are accurate and in agreement with other electronic structure approaches. © *2013 AIP Publishing LLC*. [http://dx.doi.org/10.1063/1.4826164]

I. INTRODUCTION

Strict localization of atomic and molecular orbitals is utilized as a technique to achieve linear-scaling cost in electronic structure calculations.[1–7] This approximation is based on the perception that electronic interactions tend to occur within a finite distance from the nuclei. In methods based on the electronic density matrix, such as Kohn-Sham density functional theory,[8,9] the notion of locality emerges as a result of the exponential decay with the distance of the one-particle electronic density matrix in systems with a non-zero bandgap.[10] This observation is in the core of linear-scaling density functional theory approaches[11–13] in which the computational cost scales proportionally to the number of atoms in the system.

To fully exploit localization, the Kohn-Sham orbitals can be constructed as a linear combination of non-orthogonal, atom-centered, localized functions.[14–19] Non-orthogonality is not a necessary requirement, but, in some approaches, it can be more convenient to allow the localized functions to be non-orthogonal and to adopt the tensorial algebraic manipulation of the mathematical equations typical of curved manifolds.[20–22] One possibility is to keep the localized functions constant during the calculation, and consider only the Kohn-Sham orbitals expansion coefficients as the degrees of freedom of the calculation. While this approach can yield accurate results, it requires a non-minimal basis set of localized functions. Pseudoatomic orbitals (PAOs) can be used to systematically build multiple-ζ basis set that might include higher-order polarization orbitals.[23–26]

This work focuses on an alternative approach, in which chemical accuracy is obtained by describing the Kohn-Sham orbitals as a linear combination of a minimal number of localized functions that are self-consistently optimized in terms of a high-resolution basis set associated to a computational grid.[27–29] A standard technique to impose localization constraints with linear-scaling cost is to initially calculate the search direction vectors associated to the unconstrained problem (i.e., ignoring the localization constraints), and then perform an *ad hoc* truncation of the resulting delocalized search direction vectors to restore localization.[27,28] Hereafter, this method will be referred to as the *truncated search direction* (TSD) method. Experience dictates that the TSD method is capable of yielding successful convergence and of producing results which are in agreement with other electronic structure approaches.[30]

It is very important for what follows to establish the difference between localization and truncation. Localization is a condition or constraint imposed upon the model that must be maintained during the energy minimization process. In this study, localization of the Kohn-Sham states by means of an expansion in terms of a set of atom-centered spatially localized functions is a well-controlled approximation: as the localization region becomes larger, the description of the problem converges to that of the unconstrained problem with fully delocalized Kohn-Sham states.[31,32] On the other hand, truncation is the act of deleting part of the data that has been previously calculated (the delocalized part of the search direction

a)Electronic mail: C.Skylaris@soton.ac.uk

vectors, in the case of the TSD method). Truncation does not aim to reduce the energy of the system, instead it aims to reimpose a constraint that has been previously neglected (localization of the search direction vectors, in this case). The point is that it is not guaranteed that this is the best choice of localized search direction vector, even if it leads to a decrease of the total energy of the system. In this sense, the TSD method is only approximate: it implies the assumption that the truncated search direction vector of the unconstrained problem (where localization is ignored) resembles the exact search direction vector of the constrained problem (where localization is always considered). Furthermore, truncation introduces an error into the numerical description of the problem as a consequence of information deletion. If the error is very large, convergence can be slow and the results inaccurate. Therefore, the operation of truncation of the search direction vectors, and by extension the TSD method, cannot be considered as a strictly variational scheme.

There is no reason to believe that localization could not be imposed exactly and by construction. In fact, this is a much more desirable option which will completely avoid truncation and any form of numerical error introduced upon the model. What the present work shows is a method to do precisely that: localization is imposed without truncation, exactly and analytically, upon the search direction vectors. The idea is to follow a constrained minimization procedure were localization is considered throughout, and where all the steps taken during the optimization of the localized functions are based on the variational principle, i.e., they aim to minimize the total energy functional. This new method will be referred to hereafter as the *variationally localized search direction* (VLSD) method. In this work, the principles of the VLSD method are described and discussed. As it stands, it represents a novel realization of the energy minimization problem with strict localization constraints and a non-orthogonal representation, which gives new insights into how the localization constraints propagate throughout the covariant and contravariant spaces based on a formal tensorial algebra analysis. An algorithm for VLSD calculations is also presented. Currently, the scaling of VLSD calculations is $\mathcal{O}(N^4)$ with the number of atoms in the system, N, although it could potentially be reduced to $\mathcal{O}(N^2)$ by introducing linear-scaling techniques to deal with inverse matrices.[33,34] More research is required to devise a fully linear-scaling VLSD approach. The results show that the VLSD method requires less iterations to converge than the TSD method for any size of the localization region.

This work is organized as follows. In Sec. II, the self-consistent method for calculating the total energy with non-orthogonal, localized orbitals is outlined. It includes a summary of the tensorial algebra necessary for the manipulation of the vectors and matrices with a non-orthogonal representation, as well as an analysis of the constraints imposed upon the model (localization, orthonormality of the Kohn-Sham states, and conservation of the number of electrons), in the context of direct energy minimization approaches. In Sec. III, the TSD method is discussed and examined step by step. The new VLSD method is introduced in Sec. IV. First, the mathematical principle behind the VLSD method is enunciated and proven. Then, the numerical algorithm for the calculation of search direction vectors (which are directly localized by construction) is described in full detail. A serial version of the VLSD algorithm was implemented within the ONETEP[17] program for density functional theory calculations, which, by default, uses an implementation of the TSD method. The extension of the VLSD method to other electronic structure methods should also be viable. The results of calculations with the TSD and VLSD methods on many test systems are shown in Sec. V. First, the TSD method is analyzed quantitatively, in terms of the amount of information that is deleted during the truncation of the search direction vectors. The study revealed that, in many cases, as the iterative minimization progresses, a large portion of the total pre-calculated information stored in the delocalized search direction vectors has to be deleted in order to maintain the localization constraints. Also, a number of calculations that compare the convergence properties of the TSD and VLSD methods were completed, showing that, in general, VLSD outperforms TSD by converging towards the ground state in less iterations (up to three times less iterations, in some cases). Further validation calculations show that both the TSD and VLSD methods are capable of producing correct results for the electronic and structural properties of the test molecules, compared to other computational electronic structure approaches. A set of calculations on poly-ethylene fragments of increasing length was performed to test the scaling of the VLSD algorithm with the system size. To conclude, Sec. VI contains some remarks about how the VLSD method could be expanded and adapted to other relevant problems of constrained optimization using a non-orthogonal representation.

II. KOHN-SHAM DENSITY FUNCTIONAL THEORY WITH NON-ORTHOGONAL, LOCALIZED FUNCTIONS

A. Total energy

In the context of Kohn-Sham density functional theory, the total energy functional of a system of N interacting electrons is calculated as

$$E[\hat{\rho}] = E_K[\hat{\rho}] + E_{ext}[\hat{\rho}] + E_H[\hat{\rho}] + E_{XC}[\hat{\rho}], \quad (1)$$

where E_K, E_{ext}, E_H, and E_{XC} are the kinetic, external, Hartree, and exchange-correlation energies, respectively, and $\hat{\rho}$ is the one-particle density matrix operator

$$\hat{\rho} = \sum_{i=1} f_i \, |\psi_i\rangle \, \langle\psi_i| \, , \quad (2)$$

where $\{\psi_i\}$ are orthonormal Kohn-Sham orbitals with occupation numbers $\{f_i\}$. The total number of electrons in the system, N, is equal to the trace of $\hat{\rho}$:

$$N_e = \mathrm{tr}[\hat{\rho}] = \sum_{i=1} f_i . \quad (3)$$

The Kohn-Sham orbitals $\{\psi_i\}$, represented in the real space, can be expanded in terms of a set of N_ϕ non-orthogonal, spatially localized functions $\{\phi_\alpha\}$ as

$$\psi_i(\mathbf{r}) = \sum_{\alpha=1}^{N_\phi} \phi_\alpha(\mathbf{r}) M^\alpha{}_i , \quad (4)$$

where $M^{\alpha}_{\ i}$ is the transformation matrix between the $\{\psi_i\}$ and $\{\phi_\alpha\}$ representations. The one-particle density matrix, $\rho(\mathbf{r}, \mathbf{r}')$, can be written as[35]

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_{\alpha,\beta=1}^{N_\phi} \phi_\alpha(\mathbf{r}) K^{\alpha\beta} \phi^*_\beta(\mathbf{r}'), \tag{5}$$

where the matrix $K^{\alpha\beta}$, referred to as the density kernel, is equal to

$$K^{\alpha\beta} = \sum_{i=1} M^{\alpha}_{\ i} f_i M^{\dagger\beta}_{\ i}. \tag{6}$$

The ground state of the system can be found by minimizing the total energy self-consistently with respect to $\hat{\rho}$, under the constraints of orthonormality of $\{\psi_i\}$ and conservation of N. In insulators and semi-conductors, as a consequence of the distinct bandgap at the Fermi level, the occupation numbers $\{f_i\}$ take integer values, resulting in idempotency of $K^{\alpha\beta}$. Also, in these systems, locality can be exploited by setting to zero the matrix elements of $K^{\alpha\beta}$ corresponding to pairs $(\phi_\alpha, \phi_\beta)$ associated to atoms further than a cut-off distance.[16,36] It is often more efficient to separate the self-consistent minimization into two nested iterative minimization processes. First, an inner loop optimizes $K^{\alpha\beta}$ variationally, under the constraints of idempotency and locality, with fixed $\{\phi_\alpha\}$. Common approaches to undertake this optimization are the Li-Nunes-Vanderbilt (LNV) algorithm,[37,38] and density mixing.[39] Once the inner loop has completed, an outer loop optimizes $\{\phi_\alpha\}$ variationally, subject to strict localization constraints.[28,40] In order to do that, $\{\phi_\alpha\}$ can be expanded in terms of a linear combination of systematically improvable, high-resolution basis set of N_k functions, $\{D_k\}$, centered on the points of a uniform grid with coordinate vectors $\{\mathbf{r}_k\}$:

$$\phi_\alpha(\mathbf{r}) = \sum_{k=1}^{N_k} D_k(\mathbf{r}) c_{k\alpha}, \tag{7}$$

where $\{c_{k\alpha}\}$ are the expansion coefficients. The outer loop is twice-constrained by orthonormality of $\{\psi_i\}$ and spatial localization of $\{\phi_\alpha\}$. Convergence of the inner and outer loops results in the ground state solution of the Kohn-Sham equations.

B. Tensorial properties of the non-orthogonal representation

When working with non-orthogonal representations, the vectors involved in the calculation acquire a non-trivial tensorial character that must be taken into account for consistency in the equations.[20–22] The set of Kohn-Sham states $\{\psi_i\}$ are the eigenvectors of the N-electron Hamiltonian $\hat{H}$, and form a complete basis of the complex Hilbert space $\mathcal{H}$. On the other hand, the localized functions $\{\phi_\alpha\}$, defined after the transformation $M^{\alpha}_{\ i}$ in Eq. (4), also form a complete basis set of $\mathcal{H}$, albeit non-orthogonal. A common choice of inner product in $\mathcal{H}$ is

$$\langle f|g\rangle = \int d\mathbf{r}\, f^*(\mathbf{r}) g(\mathbf{r}), \tag{8}$$

where $f, g \in \mathcal{H}$. The overlap matrix associated to $\{\phi_\alpha\}$ is defined as

$$S_{\alpha\beta} = \langle \phi_\alpha | \phi_\beta \rangle . \tag{9}$$

A set of dual functions $\{\phi^\alpha\}$ that satisfy the biorthonormality condition,

$$\langle \phi_\alpha | \phi^\beta \rangle = \delta_\alpha^{\ \beta}, \tag{10}$$

can be calculated from the duality relations

$$\phi^\alpha(\mathbf{r}) = \sum_{\beta=1}^{N_\phi} \phi_\beta(\mathbf{r}) S^{\beta\alpha}, \tag{11}$$

$$\phi_\alpha(\mathbf{r}) = \sum_{\beta=1}^{N_\phi} \phi^\beta(\mathbf{r}) S_{\beta\alpha}. \tag{12}$$

where

$$S^{\alpha\beta} = \langle \phi^\alpha | \phi^\beta \rangle , \tag{13}$$

which implies that $\sum_{\beta=1}^{N_\phi} S^{\alpha\beta} S_{\beta\gamma} = \delta^\alpha_{\ \gamma}$. The dual functions $\{\phi^\alpha\}$ also form a complete basis set of $\mathcal{H}$. The functions $\{\phi_\alpha\}$ are rank-one covariant tensors (identified with Greek subscript), while $\{\phi^\alpha\}$ are rank-one contravariant tensors (identified with a Greek superscript). Any operator represented in terms of $\{\phi_\alpha\}$ or $\{\phi^\alpha\}$ will inherit the covariant or contravariant tensorial character, respectively. Contravariant indices can be made covariant by multiplication with $S_{\alpha\beta}$, whereas covariant indices can be made contravariant by multiplication with $S^{\alpha\beta}$. Multiplication of tensors can only involve a pair of covariant and contravariant indices, and tensors of the same kind can be multiplied after mediation of the appropriate metric tensor. The energy functional, ultimately calculated as the trace of the product of covariant and contravariant tensors, is tensor-invariant (or a rank-zero tensor). Derivatives with respect to covariant functions $\{\phi_\alpha\}$ produce rank-one contravariant tensors:

$$\frac{\delta E}{\delta \langle \phi_\alpha |} = |g^\alpha\rangle , \tag{14}$$

while derivatives with respect to contravariant functions $\{\phi^\alpha\}$ produce rank-one covariant tensors:

$$\frac{\delta E}{\delta \langle \phi^\alpha |} = |g_\alpha\rangle . \tag{15}$$

The Kohn-Sham orbitals $\{\psi_i\}$ can be written in terms of $\{\phi^\alpha\}$ as

$$\psi_i(\mathbf{r}) = \sum_{\alpha=1}^{N_\phi} \phi^\alpha(\mathbf{r}) M_{\alpha i}, \tag{16}$$

where

$$M_{\alpha i} = \sum_{\beta=1}^{N_\phi} S_{\alpha\beta} M^{\beta}_{\ i}. \tag{17}$$

The overlap matrices $S_{\alpha\beta}$ and $S^{\alpha\beta}$, and the transformations $M^{\alpha}_{\ i}$ and $M_{\alpha i}$, are related as

$$S_{\alpha\beta} = \sum_{i=1} M_{\alpha i} M^{\dagger}_{i\beta}, \tag{18}$$

$$S^{\alpha\beta} = \sum_{i=1} M^{\alpha}{}_{i} M_{i}^{\dagger\beta}, \tag{19}$$

which implies that $(M^{-1})_{\alpha i} = M_i^{\dagger\alpha}$. The contravariant functions $\{\phi^{\alpha}\}$ can be represented in the $\{D_k\}$ basis set as

$$\phi^{\alpha}(\mathbf{r}) = \sum_{k=1}^{N_k} D_k(\mathbf{r}) c_k^{\alpha}, \tag{20}$$

where $\{c_k^{\alpha}\}$ are the expansion coefficients. Duality applies to the expansion coefficients point-wise:

$$c_{k\alpha} = \sum_{\beta=1}^{N_\phi} c_k^{\beta} S_{\beta\alpha}, \tag{21}$$

$$c_k^{\alpha} = \sum_{\beta=1}^{N_\phi} c_{k\beta} S^{\beta\alpha}. \tag{22}$$

These relations form the basic algebra that must be applied when working with a non-orthogonal representation, and they shall be referred to at different stages of this work.

C. Orthonormality of the Kohn-Sham states

The set of Kohn-Sham states $\{\psi_i\}$ forms an orthonormal basis set of the Hilbert space $\mathcal{H}$. At convergence, both the Hamiltonian and the one-particle density operators are diagonal in this representation, and hence their commutator $[\hat{H}, \hat{\rho}]$ is zero. Orthonormality of $\{\psi_i\}$ is a necessary physical constraint on the mathematical model that must be strictly maintained during the self-consistent optimization. Using Eq. (4), orthonormality of $\{\psi_i\}$ can be written as

$$\langle\psi_i|\psi_j\rangle = \sum_{\alpha,\beta=1}^{N_\phi} M_i^{\dagger\alpha} \langle\phi_\alpha|\phi_\beta\rangle M^{\beta}{}_{j} = \delta_{ij}. \tag{23}$$

Left-multiplying by $M_{\gamma i}$ and right-multiplying by $M_j^{\dagger\epsilon}$, and summing over i and j, the following result is obtained

$$\sum_{i,j=1} \sum_{\alpha,\beta=1}^{N_\phi} M_{\gamma i} M_i^{\dagger\alpha} \langle\phi_\alpha|\phi_\beta\rangle M^{\beta}{}_{j} M_j^{\dagger\epsilon} = \sum_{i,j=1} M_{\gamma i} \delta_{ij} M_j^{\dagger\epsilon}. \tag{24}$$

Using Eq. (19) and the property $(M^{-1})_{\alpha i} = M_i^{\dagger\alpha}$, Eq. (24) yields $\langle\phi_\alpha|\phi^{\beta}\rangle = \delta_{\alpha}{}^{\beta}$, which is the biorthonormality condition as defined in Eq. (10). The interpretation of this result is that the transformation $M^{\alpha}{}_{i}$ converts the constraint of orthonormality of the Kohn-Sham states $\{\psi_i\}$ into a constraint of biorthonormality between the dual sets $\{\phi_\alpha\}$ and $\{\phi^{\alpha}\}$. Biorthogonality is possible if and only if Eqs. (11) and (12) hold exactly and are equivalent for the current sets $\{\phi_\alpha\}$ and $\{\phi^{\alpha}\}$. This condition implies that $S^{\alpha\beta}$, calculated by direct inversion of $S_{\alpha\beta}$, must match the matrix calculated as the result of the inner products $\langle\phi^{\alpha}|\phi^{\beta}\rangle$. During the steepest-descent optimization of $\{\phi_\alpha\}$, biorthonormality holds only approximately (to first order). This point is explained in further detail in Sec. II E.

D. Localization of the non-orthogonal functions

In principle, the Kohn-Sham states are defined and single-valued at every point of the real space. Many electronic structure methods represent $\{\psi_i\}$ as a linear combination of delocalized basis functions (plane-waves,[41] for instance), which typically makes the cost of building the electronic density scale with the cube of the number of atoms. The constraint of strict localization of $\{\phi_\alpha\}$ is an extra requirement that must be imposed in order to achieve linear-scaling cost with the number of atoms (another option is to use numerical thresholding and a Gaussian basis set[42,43]). In such cases, the search towards the minimum in the potential energy surface (PES) is transformed into a search along the metasurface of the PES compatible with exact localization constraints. As a consequence, at convergence, $E_{loc} \geq E_{deloc}$, where E_{loc} and E_{deloc} are the ground state energies of the problem with and without localization constraints, respectively.[31,32] Thus, approaches that use localized functions are variational with respect to the size of the localization regions, and are equivalent to approaches with delocalized functions in the limit of infinitely large localization regions. Demanding strict localization of $\{\phi_\alpha\}$ should not result in a loss of orthonormality of $\{\psi_i\}$, or else unphysical results would be obtained. To avoid this scenario, as discussed in Sec. II C, Eqs. (11) and (12) must be verifiable and equivalent at the ground state.

A mathematically concise way to define the strict localizaton constraints, based on the expansion of $\{\phi_\alpha\}$ in the $\{D_k\}$ basis set as in Eq. (7), is

$$c_{k\alpha} = 0, \ \forall k \in [1, \ldots, N_k] \mid \mathbf{r}_k \notin \mathrm{LR}(\phi_\alpha(\mathbf{r})), \quad \forall\alpha \in [1, \ldots, N_\phi], \tag{25}$$

where $\mathrm{LR}(\phi_\alpha(\mathbf{r}))$ is the localization region of $\phi_\alpha(\mathbf{r})$. Without loss of generality, $\mathrm{LR}(\phi_\alpha(\mathbf{r}))$ can be considered as a well-delimited, finite region of the real-space simulation cell. It is important to realize that the contravariant functions $\{\phi^{\alpha}\}$ are also spatially localized in a well-delimited region of the real-space simulation cell. According to Eq. (11), the contravariant functions are constructed as a linear combination of strictly localized covariant functions centered on the atomic nuclei (often spherical regions), with expansion coefficients $S^{\beta\alpha}$. Therefore, the localization region of $\phi^{\alpha}(\mathbf{r})$ is equal to the volume occupied by the union of all the localization regions of $\{\phi_\beta\}$ for which $S^{\beta\alpha} \neq 0$:

$$\mathrm{LR}(\phi^{\alpha}(\mathbf{r})) = \bigcup_{\substack{\beta=1 \\ S^{\beta\alpha}\neq 0}}^{N_\phi} \mathrm{LR}(\phi_\beta(\mathbf{r})). \tag{26}$$

One can conclude from this result that the union of all the localization regions of $\{\phi_\alpha\}$ and the union of all the localization regions of $\{\phi^{\alpha}\}$ occupy exactly the same volume in real space. The individual localization regions of $\{\phi^{\alpha}\}$ will generally be much larger than the individual localization regions of $\{\phi_\alpha\}$ and will cover all the atom centers of the system under study. Thus, the contravariant functions do not (in general) extend to every point in the real space, but still are delocalized enough to cover the entire molecule. Figure 1 illustrates

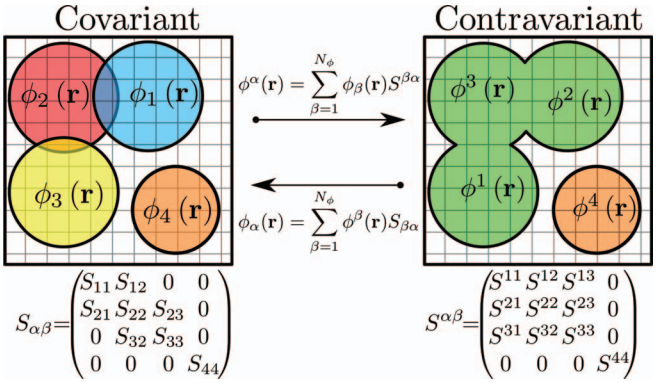

FIG. 1. Two-dimensional depiction of the real-space localization regions in a calculation involving four non-orthogonal localized functions. The covariant functions $\phi_1(\mathbf{r})$, $\phi_2(\mathbf{r})$, $\phi_3(\mathbf{r})$, and $\phi_4(\mathbf{r})$ are localized within different circumferences, producing a distinctive sparsity pattern in $S_{\alpha\beta}$. In contravariant space, $\phi^1(\mathbf{r})$, $\phi^2(\mathbf{r})$, and $\phi^3(\mathbf{r})$ are localized within the same localization region, made up from the union of the localization regions of $\phi_1(\mathbf{r})$, $\phi_2(\mathbf{r})$, $\phi_3(\mathbf{r})$, while $\phi^4(\mathbf{r})$ is localized in a circumference, as a result of $\phi_4(\mathbf{r})$ not overlapping any other function. The volume covered by all the localization regions is the same both in the covariant and contravariant representations. All the expansion coefficients associated to grid points outside of this volume are zero by definition of the localization constraint.

the localization properties of the covariant and contravariant functions with an example. The inverse overlap matrix, $S^{\beta\alpha}$, is typically dense,[33,44] unless one or more of the covariant functions $\{\phi_\alpha\}$ do not overlap any of the other covariant functions (which is not normally the case). Therefore, in the $\{D_k\}$ representation (Eq. (20)), the localization constraint on $\{\phi^\alpha\}$ can be defined as

$$c_k^{\ \alpha} = 0,\ \forall k \in [1, \ldots, N_k] \mid \mathbf{r}_k \notin \mathrm{LR}(\phi^\alpha(\mathbf{r})),\ \forall \alpha \in [1, \ldots, N_\phi]. \tag{27}$$

Imposing orthonormality of $\{\psi_i\}$ and localization of $\{\phi_\alpha\}$ simultaneously does not overconstrain the Kohn-Sham equations. To prove this assertion, the number of irreducible degrees of freedom (DOF), defined as the minimum number of linearly independent variational parameters capable of describing the system under these constraints, was counted. The number of DOF can be calculated as the sum of DOF of the unconstrained problem minus the number of constraints. In the unconstrained problem, each covariant function represented in real space, $\phi_\alpha(\mathbf{r})$, is expanded in terms of N_k coefficients $\{c_{k\alpha}\}$, and each contravariant function $\phi^\alpha(\mathbf{r})$ in terms of N_k coefficients $\{c_k^{\ \alpha}\}$. These add up to a total of $2 \times N_\phi \times N_k$ coefficients. Orthonormality of $\{\psi_i\}$ is imposed by ensuring that Eqs. (11) and (12) are equivalent, so only one of them needs to be taken into account when counting the constraints. Also, they must hold point-wise, for each different grid-point $\mathbf{r}_k$, which amounts for a total of $N_\phi \times N_k$ constraints. The localization constraint is imposed via Eq. (25). Assuming that there are N_α grid points inside $\mathrm{LR}(\phi_\alpha(\mathbf{r}))$, the number of coefficients that are constrained to be equal to zero is $\sum_{\alpha=1}^{N_\phi}(N_k - N_\alpha) = N_\phi \times N_k - \sum_{\alpha=1}^{N_\phi} N_\alpha$. Therefore, the number of irreducible degrees of freedom, N_{DOF}, is

$$N_{DOF} = 2 \times N_\phi \times N_k - N_\phi \times N_k - \left(N_\phi \times N_k - \sum_{\alpha=1}^{N_\phi} N_\alpha\right) = \sum_{\alpha=1}^{N_\phi} N_\alpha. \tag{28}$$

Perhaps the most intuitive and natural choice of N_{DOF} variational parameters to describe the system are the set of covariant coefficients associated to the points inside the localization region of each of the covariant functions $\{\phi_\alpha\}$:

$$\Omega = \{c_{k\alpha},\ \forall k \in [1, \ldots, N_k] \mid \mathbf{r}_k \in \mathrm{LR}(\phi_\alpha(\mathbf{r})),\ \forall \alpha \in [1, \ldots, N_\phi]\}. \tag{29}$$

Technically speaking, however, there are an infinite number of choices of N_{DOF} linearly independent variational parameters that can also provide an equivalent description of the mathematical problem.

E. Steepest-descent optimization of the non-orthogonal functions

From this point onwards, it is assumed that the density kernel $K^{\alpha\beta}$ is an idempotent matrix that has been fully optimized during an inner loop. The focus will be put on the optimization of $\{\phi_\alpha\}$ in an outer loop, based on the steepest-descent algorithm for direct energy minimization.[28,40,45–47] The extension to the more efficient conjugate gradients method is straightforward. At every iteration n of the outer loop, the covariant functions $\{\phi_\alpha^{(n)}\}$ are updated as

$$\left|\phi_\alpha^{(n+1)}\right\rangle = \left|\phi_\alpha^{(n)}\right\rangle + \lambda^{(n)} \left|\Delta_\alpha^{(n)}\right\rangle, \tag{30}$$

where $\{\Delta_\alpha^{(n)}\}$ are the search direction vectors compatible with the constraints applied onto $\{\phi_\alpha^{(n)}\}$, and $\lambda^{(n)}$ is the optimum step length, calculated after a line-search process. The property of tensor-invariance of the energy functional implies that a steepest-descent update of the contravariant functions $\{\phi^{\alpha(n)}\}$ of the type:

$$|\phi^{\alpha(n+1)}\rangle = |\phi^{\alpha(n)}\rangle + \lambda^{(n)} |\Delta^{\alpha(n)}\rangle, \tag{31}$$

must produce an equivalent energy minimization path. The covariant and contravariant search directions are mutually dual

$$\Delta_\alpha^{(n)}(\mathbf{r}) = \sum_{\beta=1}^{N_\phi} \Delta^{\beta(n)}(\mathbf{r}) S_{\beta\alpha}^{(n)}, \tag{32}$$

$$\Delta^{\alpha(n)}(\mathbf{r}) = \sum_{\beta=1}^{N_\phi} \Delta_\beta^{(n)}(\mathbf{r}) S^{\beta\alpha(n)}. \tag{33}$$

In practice, the steepest-descent update takes place in the space of coefficients in terms of $\{D_k\}$:

$$c_{k\alpha}^{(n+1)} = c_{k\alpha}^{(n)} + \lambda^{(n)} d_{k\alpha}^{(n)}, \tag{34}$$

$$c_k^{\alpha(n+1)} = c_k^{\alpha(n)} + \lambda^{(n)} d_k^{\alpha(n)}, \tag{35}$$

where $\{d_{k\alpha}^{(n)}\}$ and $\{d_k^{\alpha(n)}\}$ are the expansion coefficients of $\{\Delta_\alpha^{(n)}\}$ and $\{\Delta^{\alpha(n)}\}$:

$$\Delta_\alpha^{(n)}(\mathbf{r}) = \sum_{k=1}^{N_k} D_k(\mathbf{r}) d_{k\alpha}^{(n)}, \tag{36}$$

$$\Delta^{\alpha(n)}(\mathbf{r}) = \sum_{k=1}^{N_k} D_k(\mathbf{r}) d_k^{\alpha(n)}, \tag{37}$$

The coefficients $\{d_{k\alpha}^{(n)}\}$ and $\{d_k^{\alpha(n)}\}$ are mutually dual

$$d_{k\alpha}^{(n)} = \sum_{\beta=1}^{N_\phi} d_k^{\beta(n)} S_{\beta\alpha}^{(n)}, \tag{38}$$

$$d_k^{\alpha(n)} = \sum_{\beta=1}^{N_\phi} d_{k\beta}^{(n)} S^{\beta\alpha(n)}. \tag{39}$$

Orthonormality of $\{\psi_i^{(n)}\}$ must be maintained during the steepest-descent update. As discussed in Sec. II C, imposing orthonormality of $\{\psi_i\}$ is equivalent to imposing biorthonormality between the dual sets $\{\phi_\alpha^{(n)}\}$ and $\{\phi^{\alpha(n)}\}$. After updating the non-orthogonal functions as in Eqs. (30) and (31), The biorthonormality condition, Eq. (10), becomes

$$\begin{aligned}\left\langle \phi_\alpha^{(n+1)} \middle| \phi^{\beta(n+1)} \right\rangle = {} & \left\langle \phi_\alpha^{(n)} \middle| \phi^{\beta(n)} \right\rangle \\ & + \lambda^{(n)} \left(\langle \phi_\alpha^{(n)} | \Delta^{\beta(n)} \rangle + \left\langle \Delta_\alpha^{(n)} \middle| \phi^{\beta(n)} \right\rangle \right) \\ & + (\lambda^{(n)})^2 \left\langle \Delta_\alpha^{(n)} \middle| \Delta^{\beta(n)} \right\rangle. \end{aligned} \tag{40}$$

Exact biorthonormality can be imposed only if the following two equations are simultaneously true

$$\left\langle \phi_\alpha^{(n)} \middle| \Delta^{\beta(n)} \right\rangle = \delta_\alpha{}^\beta, \tag{41}$$

$$\left\langle \Delta_\alpha^{(n)} \middle| \Delta^{\beta(n)} \right\rangle = \delta_\alpha{}^\beta. \tag{42}$$

The task of imposing these two conditions simultaneously is highly non-trivial. It is often more practical to impose only Eq. (41) and assume a small error of order $\mathcal{O}(\lambda^2)$. Thus, after the steepest-descent update, biorthonormality is maintained only to first order, and orthonormality of $\{\psi_i^{(n+1)}\}$ is not exact. An immediate consequence of this is that the steepest-descent updates Eqs. (30) and (31) are not equivalent. Another manifestation of the lack of exact orthonormality of $\{\psi_i^{(n+1)}\}$ is that the total number of electrons in the system is not maintained

$$\tilde{N}_e = \mathrm{tr}[\hat{\rho}] = \sum_{\alpha,\beta=1}^{N_\phi} K^{\alpha\beta} S_{\beta\alpha}^{(n+1)} \neq N_e, \tag{43}$$

where $S_{\beta\alpha}^{(n+1)}$ is the up-to-date overlap matrix. To fix the problem of electron conservation, a rescale factor[48] can be applied to the density kernel as

$$\tilde{K}^{\alpha\beta} = \frac{Ne}{\tilde{N}_e} K^{\alpha\beta}, \tag{44}$$

which guarantees that $\sum_{\alpha,\beta=1}^{N_\phi} \tilde{K}^{\alpha\beta} S_{\beta\alpha}^{(n+1)} = N_e$. This technique reduces the negative effect of the lack of orthonormality of $\{\psi_i^{(n)}\}$, which can be completely restored after re-entering the inner loop by allowing changes to the transformation matrix $M^\alpha{}_i$.

Regarding the localization constraint on $\{\phi_\alpha^{(n+1)}\}$, inspection of Eq. (30) reveals that the search directions $\{\Delta_\alpha^{(n)}\}$ must also be spatially localized within the same region as $\{\phi_\alpha^{(n)}\}$. This observation also applies to the contravariant functions. Therefore,

$$\mathrm{LR}(\Delta_\alpha(\mathbf{r})) = \mathrm{LR}(\phi_\alpha(\mathbf{r})), \tag{45}$$

$$\mathrm{LR}(\Delta^\alpha(\mathbf{r})) = \mathrm{LR}(\phi^\alpha(\mathbf{r})). \tag{46}$$

In the representation in terms of $\{D_k\}$, the localization constraints on $\{\Delta_\alpha^{(n)}\}$ and $\{\Delta^{\alpha(n)}\}$ can be written as

$$\begin{aligned} & d_{k\alpha}^{(n)} = 0, \ \forall k \in [1, \ldots, N_k] \mid \mathbf{r}_k \notin \mathrm{LR}(\Delta_\alpha(\mathbf{r})), \\ & \forall \alpha \in [1, \ldots, N_\phi], \end{aligned} \tag{47}$$

$$\begin{aligned} & d_k^{\alpha(n)} = 0, \ \forall k \in [1, \ldots, N_k] \mid \mathbf{r}_k \notin \mathrm{LR}(\Delta^\alpha(\mathbf{r})), \\ & \forall \alpha \in [1, \ldots, N_\phi]. \end{aligned} \tag{48}$$

As demonstrated in Sec. II D, imposing orthonormality of $\{\psi_i\}$ and localization of $\{\phi_\alpha\}$ does not overconstrain the Kohn-Sham equations. However, orthonormality of $\{\psi_i\}$ can only be imposed to first order during the optimization of $\{\phi_\alpha\}$. Methods such as kernel rescaling can stabilize the algorithm, but, ultimately, it is the task of the inner loop to completely fix and maintain orthonormality of $\{\psi_i\}$. Therefore, one can choose to introduce a larger error in the orthonormality of $\{\psi_i\}$ if in doing so, imposing localization of $\{\phi_\alpha\}$ becomes simpler or faster. The risk of such an approach is that the inner loop might converge at a slower rate or even diverge. Thus, it is interesting to explore methods for which imposing localization of $\{\phi_\alpha\}$ implies the lowest alteration to the biorthonormality constraint. When the self-consistent solution is reached, the search direction is perpendicular to the energy gradient, and the slope along the search direction compatible with the localization constraints is zero:

$$\left\langle \Delta_\alpha^{(n)} \middle| \frac{\delta E}{\delta \left| \phi_\alpha^{(n)} \right\rangle} \right. = 0, \tag{49}$$

When this occurs, a minimum on the energy metasurface that respects the biorthonormality and localization constraints has been reached.

In Secs. III and IV, the TSD and VLSD methods will be analyzed. For clarity in the equations, in what follows, the iteration index n will be dropped, but it must be clear that all the steps presented hereafter have to be repeated at every new iteration of the outer loop.

III. TRUNCATED SEARCH DIRECTION (TSD) METHOD

The derivative of the energy functional with respect to $\{c_{k\alpha}\}$ is

$$g_k^\alpha = \frac{\partial E}{\partial c_{k\alpha}} = \left[\sum_{\beta=1}^{N_\phi} \hat{H} \phi_\beta(\mathbf{r}) K^{\beta\alpha} \right]_{\mathbf{r}=\mathbf{r}_k}, \tag{50}$$

where $\{g_k^{\alpha}\}$ are the expansion coefficients associated to the gradients $\{\Gamma^{\alpha}\}$:

$$\Gamma^{\alpha}(\mathbf{r}) = \sum_{k=1}^{N_k} D_k(\mathbf{r}) g_k^{\alpha}. \tag{51}$$

The coefficients $\{g_k^{\alpha}\}$ are contravariant tensors that cannot be used to update $\{c_{k\alpha}\}$ directly. Furthermore, they cannot yet be used to update $\{c_k^{\alpha}\}$ either, since they do not account for the constraints of orthonormality of $\{\psi_i\}$ and localization of $\{\phi_{\alpha}\}$. Orthonormality of $\{\psi_i\}$ can be imposed to first order by projecting out the components of $\{\Gamma^{\alpha}\}$ parallel to $\{\phi_{\beta}\}$,[45,46] so that Eq. (41) holds

$$\tilde{g}_k^{\alpha} = \left[\sum_{\beta=1}^{N_{\phi}} \hat{H}\phi_{\beta}(\mathbf{r}) K^{\beta\alpha} - \sum_{\beta,\gamma,\delta=1}^{N_{\phi}} \phi_{\beta}(\mathbf{r}) K^{\beta\gamma} H_{\gamma\delta} S^{\delta\alpha} \right]_{\mathbf{r}=\mathbf{r}_k}, \tag{52}$$

where $\{\tilde{g}_k^{\alpha}\}$ are the expansion coefficients of $\{\tilde{\Gamma}^{\alpha}\}$:

$$\tilde{\Gamma}^{\alpha}(\mathbf{r}) = \sum_{k=1}^{N_k} D_k(\mathbf{r}) \tilde{g}_k^{\alpha}. \tag{53}$$

LNV-type methods can avoid this projection, as accounting for idempotency of $K^{\alpha\beta}$ is equivalent to maintaining orthonormality of $\{\psi_i\}$ to first order.[49] At convergence of LNV, the derivative of the energy functional with respect to $\{c_{k\alpha}\}$ returns $\{\tilde{g}_k^{\alpha}\}$ directly. The covariant representation of $\{\tilde{g}_k^{\alpha}\}$ can be calculated by multiplying with $S_{\alpha\beta}$ as in Eq. (38). Occupancy preconditioning[50] followed by kinetic energy preconditioning[45,51] can be applied, resulting in

$$\tilde{d}_{k\alpha} = -\hat{P} \left[\sum_{\beta=1}^{N_{\phi}} \hat{H}\phi_{\alpha}(\mathbf{r}) - \sum_{\beta,\gamma=1}^{N_{\phi}} \phi_{\beta}(\mathbf{r}) S^{\beta\gamma} H_{\gamma\alpha} \right]_{\mathbf{r}=\mathbf{r}_k}, \tag{54}$$

where $\hat{P}$ is a generic kinetic energy preconditioner. The search directions $\{\tilde{\Delta}_{\alpha}\}$ expanded by $\{\tilde{d}_{k\alpha}\}$ as

$$\tilde{\Delta}_{\alpha}(\mathbf{r}) = \sum_{k=1}^{N_k} D_k(\mathbf{r}) \tilde{d}_{k\alpha}, \tag{55}$$

are delocalized in real space, and therefore not yet suitable for updating $\{\phi_{\alpha}\}$ with the steepest-descent method. To solve the issue of delocalization, an *ad hoc* truncation of $\{\tilde{\Delta}_{\alpha}\}$ can be utilized,[28,40] setting to zero all the expansion coefficients corresponding to grid points outside $\mathrm{LR}(\Delta_{\alpha}(\mathbf{r}))$:

$$d_{k\alpha} = \begin{cases} \tilde{d}_{k\alpha}, & \text{if } \mathbf{r}_k \in \mathrm{LR}(\Delta_{\alpha}(\mathbf{r})) \\ 0, & \text{if } \mathbf{r}_k \notin \mathrm{LR}(\Delta_{\alpha}(\mathbf{r})) \end{cases}. \tag{56}$$

A steepest-descent update of $\{\phi_{\alpha}\}$ that is consistent with the constraint of localization is now possible using the search directions $\{\Delta_{\alpha}\}$, which are represented in the $\{D_k\}$ basis set as

$$\Delta_{\alpha}(\mathbf{r}) = \sum_{k=1}^{N_k} D_k(\mathbf{r}) d_{k\alpha}. \tag{57}$$

After truncation, the constraint of orthonormality of $\{\psi_i\}$ is not maintained, not even to first order. This can be seen by the fact that the inner product $\langle \Delta_{\alpha} | \phi^{\beta} \rangle$ is not zero for any pair $\alpha \neq \beta$. A new projection of the kind of Eq. (52) to re-impose biorthonormality to first order would not result in a suitable search direction, since it would delocalize $\{\Delta_{\alpha}\}$ again. The algorithm then continues with the updating $\{\phi_{\alpha}\}$ and the overlap matrices, $S_{\alpha\beta}$ and $S^{\alpha\beta}$, followed by the rescaling of the density kernel as in Eq. (44). The TSD algorithm relies upon the inner loop to restore orthonormality of $\{\psi_i\}$ during the next cycle.

In the TSD method, delocalization originates when the gradient coefficients $\{g_k^{\alpha}\}$ are calculated in Eq. (50). At this step, the derivatives with respect to all the covariant coefficients $\{c_{k\alpha}\}$ are calculated, without taking into consideration that not all $c_{k\alpha}$ are variational degrees of freedom listed in Ω (Eq. (29)). After imposing orthonormality to first order, multiplying with the overlap matrix, and applying occupancy and kinetic energy preconditioning, the covariant search directions $\{\tilde{\Delta}_{\alpha}\}$ are still delocalized in real space. If they were used to update $\{\phi_{\alpha}\}$ in a steepest-descent method, the algorithm would fall out of the metasurface of the PES compatible with the localization constraints. The only solution at this stage is to truncate $\{\tilde{\Delta}_{\alpha}\}$, which is equivalent to eliminating all the information stored in the form of $\{\tilde{d}_{k\alpha}\}$ coefficients associated to $\mathbf{r}_k$ outside the localization regions.

There are two major caveats with the TSD method. First, the search directions vectors $\{\tilde{\Delta}_{\alpha}\}$ do not correspond to the search direction vectors of the problem with explicit localization constraints. Instead, they are the search direction vectors corresponding to the unconstrained problem, where delocalization is allowed. The second caveat is that the TSD method does not guarantee that $\{\Delta_{\alpha}\}$, after truncation, resemble the correct search direction vectors of the problem with strict localization constraints and, therefore, the existence of a minimum along $\{\Delta_{\alpha}\}$ is not guaranteed either. Truncation of $\{\tilde{\Delta}_{\alpha}\}$ is a non-variational artifact that has to be introduced in order to obey the localization constraints, and the resulting search direction vector is only approximate. In other words, there is always an error associated to the truncation of $\{\tilde{\Delta}_{\alpha}\}$. Only in the cases where the amount of information that is deleted during truncation is negligible compared to the amount of information that remains in the search direction vectors $\{\Delta_{\alpha}\}$, the TSD method could be considered as a valid approximation.

IV. VARIATIONALLY LOCALIZED SEARCH DIRECTION (VLSD) METHOD

A. Mathematical principle of the VLSD method

The optimization of the localized functions $\{\phi_{\alpha}\}$ is twice-constrained by the requirements of biorthonormality of $\{\phi_{\alpha}\}$ and $\{\phi^{\alpha}\}$, and localization of $\{\phi_{\alpha}\}$. While biorthonormality is a strong requirement upon the physical model, because it is equivalent to imposing orthonormality of the Kohn-Sham orbitals $\{\psi_i\}$, the localization constraints are optional. It was shown in Sec. II D that requiring these two constraints simultaneously does not overconstrain the Kohn-Sham equations. The localization constraints, however, reduce the variational degrees of freedom of the optimization problem to $N_{DOF} \leq N_k$, as shown in Eq. (28).

The principle of the VLSD method can be enunciated as follows. Given the biorthonormality constraints, Eqs. (21) and (22), and the localization constraint, Eq. (25), all the non-zero covariant coefficients $\{c_{k\alpha} \in \Omega\}$, where Ω was defined in Eq. (29) as the set of N_{DOF} irreducible degrees of freedom of the constrained optimization problem, can be uniquely and unambiguously calculated from the subset of N_{DOF} contravariant coefficients

$$\{c_k^\alpha, \forall k \in [1, \ldots, N_k] \,|\, \mathbf{r}_k \in \mathrm{LR}(\phi_\alpha(\mathbf{r})), \quad \forall \alpha \in [1, \ldots, N_\phi]\}, \tag{58}$$

which correspond to grid points inside the volume of the individual localization regions of the covariant functions, provided that the overlap matrix, $S_{\alpha\beta}$, is known.

The proof of this statement begins by considering the group of all the covariant coefficients associated to a particular grid point, $\mathbf{r}_k$. It must be noted that there could be more than one $\{c_{k\alpha}\}$ coefficient associated to the same $\mathbf{r}_k$, originating from different $\{\phi_\alpha\}$. As a matter of fact, there is a finite number of $\{\phi_\alpha\}$ for which $\mathbf{r}_k$ is inside their localization regions, and, conversely, also a finite number of $\{\phi_\alpha\}$ for which $\mathbf{r}_k$ is outside their localization regions. Based on this criteria, for every single grid point $\mathbf{r}_k$, the following two subsets of tensorial indices can be formed

$$\xi_k = \{\alpha \in [1, \ldots, N_\phi] \mid \mathbf{r}_k \in \mathrm{LR}(\phi_\alpha(\mathbf{r}))\}, \tag{59}$$

$$\chi_k = \{\alpha \in [1, \ldots, N_\phi] \mid \mathbf{r}_k \notin \mathrm{LR}(\phi_\alpha(\mathbf{r}))\}. \tag{60}$$

This separation is unique and complete, in the sense that for every $\mathbf{r}_k$, a given index α always belongs to one, and only one, of these subsets. Hence, $N_\phi = N_\xi + N_\chi$, where N_ξ and N_χ are the number of indices included in each subset. The notation used in this work is as follows. Expressions such as $\{\phi_\eta \mid \forall \eta \in \xi_k\}$, refer to the subset of covariant functions ϕ_η for which $\mathbf{r}_k \in \mathrm{LR}(\phi_\eta(\mathbf{r}))$. The notation extends to the contravariant functions. For example, $\{\phi^\tau \mid \forall \tau \in \chi_k\}$ refers to the subset of contravariant functions ϕ^τ for which $\mathbf{r}_k \notin \mathrm{LR}(\phi_\tau(\mathbf{r}))$, with ϕ_τ being the dual of ϕ^τ, calculated as $\phi_\tau = \sum_{\alpha=1}^{N_\phi} \phi^\alpha S_{\alpha\tau}$. In the same fashion, the notation also extends to the search direction vectors and, in general, to any tensor. For example, $\{d_k^\iota \mid \forall \iota \in \xi_k\}$ refers to the subset of contravariant coefficients $\{d_k^\iota\}$ for which $\mathbf{r}_k \in \mathrm{LR}(\phi_\iota(\mathbf{r}))$, with $d_{k\iota} = \sum_{\alpha=1}^{N_\phi} d_k^\alpha S_{\alpha\iota}$. Figure 2 further clarifies this notation with an example.

Based on this point-wise subset separation, three distinct overlap matrices can be defined for every single grid point $\mathbf{r}_k$:

$$A_{k,\eta\iota} = \langle \phi_\eta | \phi_\iota \rangle, \quad \forall \eta, \iota \in \xi_k, \tag{61}$$

$$B_{k,\iota\sigma} = \langle \phi_\iota | \phi_\sigma \rangle. \quad \forall \iota \in \xi_k, \ \forall \sigma \in \chi_k, \tag{62}$$

$$C_{k,\sigma\tau} = \langle \phi_\sigma | \phi_\tau \rangle, \quad \forall \sigma, \tau \in \chi_k. \tag{63}$$

The matrices $A_{k,\eta\iota}$, $B_{k,\iota\sigma}$, and $C_{k,\sigma\tau}$ are unique different mappings of the full overlap matrix $S_{\alpha\beta}$. $A_{k,\eta\iota}$ is a square matrix of dimensions $N_\xi \times N_\xi$; $B_{k,\iota\sigma}$ is a rectangular matrix of dimensions $N_\xi \times N_\chi$; and $C_{k,\sigma\tau}$ is a square matrix of dimensions $N_\chi \times N_\chi$. If a grid point $\mathbf{r}_k$ is inside all the localization regions, then χ_k is an empty subset, and $B_{k,\iota\sigma}$ and $C_{k,\sigma\tau}$ are matrices of rank zero. If $\mathbf{r}_k$ is outside all the localization regions, then ξ_k is an empty subset, and $A_{k,\eta\iota}$ and $B_{k,\iota\sigma}$ are matrices of rank zero. Equation (21) can now be split into two expressions

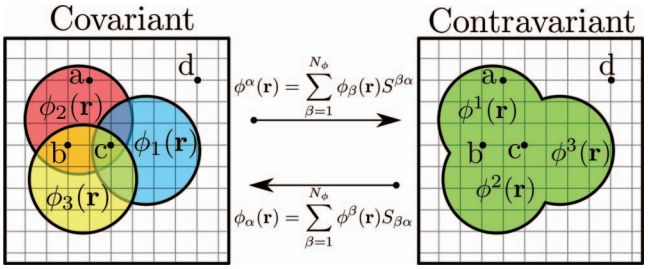

FIG. 2. Example of point-wise subset separation in a hypothetical calculation with a two-dimensional real space simulation cell, containing three covariant functions localized within circumferences. The grid point $\mathbf{r}_a$ is inside $\mathrm{LR}(\phi_2(\mathbf{r}))$, and outside $\mathrm{LR}(\phi_1(\mathbf{r}))$, and $\mathrm{LR}(\phi_3(\mathbf{r}))$. Therefore, $\xi_a = \{2\}$ and $\chi_a = \{1, 3\}$. The grid point $\mathbf{r}_b$ is inside $\mathrm{LR}(\phi_2(\mathbf{r}))$ and $\mathrm{LR}(\phi_3(\mathbf{r}))$, and outside $\mathrm{LR}(\phi_1(\mathbf{r}))$. Therefore, $\xi_b = \{2, 3\}$ and $\chi_b = \{1\}$. The grid point $\mathbf{r}_c$ is inside all the localization regions. Therefore, $\xi_c = \{1, 2, 3\}$ and χ_c is an empty subset. The grid point $\mathbf{r}_d$ is outside all the localization regions. Therefore, ξ_d is an empty subset and $\chi_d = \{1, 2, 3\}$.

$$c_{k\iota} = \sum_{\eta=1}^{N_\xi} c_k^\eta A_{k,\eta\iota} + \sum_{\sigma=1}^{N_\chi} c_k^\sigma B_{k,\sigma\iota}^\dagger, \quad \forall \eta, \iota \in \xi_k, \quad \forall \sigma \in \chi_k, \tag{64}$$

$$c_{k\tau} = \sum_{\eta=1}^{N_\xi} c_k^\eta B_{k,\eta\tau} + \sum_{\sigma=1}^{N_\chi} c_k^\sigma C_{k,\sigma\tau}, \quad \forall \eta \in \xi_k, \quad \forall \sigma, \tau \in \chi_k. \tag{65}$$

No approximation was made in the above separation, and duality between $\{\phi_\alpha\}$ and $\{\phi^\alpha\}$ is rigorously maintained. In this notation, the localization constraint on $\{\phi_\alpha\}$, as given in Eq. (25), can be written exactly as

$$c_{k\tau} = 0, \quad \forall \tau \in \chi_k. \tag{66}$$

Substituting Eq. (66) in the left-hand side of Eq. (65) leads to

$$c_k^\sigma = -\sum_{\eta=1}^{N_\xi} \sum_{\tau=1}^{N_\chi} c_k^\eta B_{k,\eta\tau} \left(C_k^{-1}\right)^{\tau\sigma}, \quad \forall \eta \in \xi_k, \forall \sigma, \tau \in \chi_k. \tag{67}$$

Equation (67) can now be introduced into (64) to obtain

$$c_{k\iota} = \sum_{\eta=1}^{N_\xi} c_k^\eta \left[A_{k,\eta\iota} - \sum_{\tau,\sigma=1}^{N_\chi} B_{k,\eta\tau} \left(C_k^{-1}\right)^{\tau\sigma} B_{k,\sigma\iota}^\dagger \right], \quad \forall \eta, \iota \in \xi_k, \quad \forall \sigma, \tau \in \chi_k. \tag{68}$$

Equations (66)–(68) show that, given the overlap matrix, $S_{\alpha\beta}$, the values of $\{c_k^\sigma, \ \forall \sigma \in \chi_k\}$ and $\{c_{k\iota}, \forall \iota \in \xi_k\}$ are uniquely determined by the values of the coefficients $\{c_k^\eta, \ \forall \eta \in \xi_k\}$. Therefore, the optimization problem with biorthonormality and localization constraints is completely and unambiguously determined by the contravariant coefficients:

$$\Omega' = \left\{c_k^\eta, \ \forall \eta \in \xi_k, \ \forall k \in [1, \ldots, N_k]\right\}. \tag{69}$$

A more detailed inspection of Eq. (69) reveals that the $\{c_k^\eta\}$ coefficients listed in Ω' correspond exactly to all the contravariant coefficients inside the localization regions of the covariant functions. The set Ω' can be redefined in formal terms as

$$\Omega' = \{c_k^\alpha,\ \forall k \in [1,\ldots,N_k] \mid \mathbf{r}_k \in \mathrm{LR}(\phi_\alpha(\mathbf{r})),\ \forall \alpha \in [1,\ldots,N_\phi]\}. \tag{70}$$

Hence, the initial hypothesis is demonstrated. The question of whether Ω' is a valid set of N_{DOF} irreducible degrees of freedom is left unresolved in this work. To prove such condition, it would be enough to demonstrate that the overlap matrix $S_{\alpha\beta}$ can be fully built from the coefficients in Ω'.

B. The VLSD algorithm

The equations introduced in Sec. IV A can be used to design an algorithm for the constrained optimization of $\{\phi_\alpha\}$ with the steepest-descent iterative method. The search direction vectors $\{\Delta_\alpha\}$ are calculated in the current $\{\phi_\alpha\}$ representation, of which the overlap matrix $S_{\alpha\beta}$ is fully known. Also, the search direction vectors are subject to the biorthonormality constraints, Eqs. (38) and (39), and the localization constraint, Eq. (47). The set of covariant search direction coefficients $\{d_{k\alpha}\}$ that can be used to update the variational degrees of freedom $\{c_{k\alpha} \in \Omega\}$ is denoted as

$$\Upsilon = \{d_{k\alpha},\ \forall k \in [1,\ldots,N_k] \mid \mathbf{r}_k \in \mathrm{LR}(\Delta_\alpha(\mathbf{r})),\ \forall \alpha \in [1,\ldots,N_\phi]\}, \tag{71}$$

while the set of contravariant search direction coefficients $\{d_k^\alpha\}$ that can be used to update $\{c_k^\alpha \in \Omega'\}$ is denoted as

$$\Upsilon' = \{d_k^\alpha,\ \forall k \in [1,\ldots,N_k] \mid \mathbf{r}_k \in \mathrm{LR}(\Delta_\alpha(\mathbf{r})),\ \forall \alpha \in [1,\ldots,N_\phi]\}. \tag{72}$$

Under these conditions, the equations in Sec. IV A demonstrate that the values of $\{d_{k\alpha} \in \Upsilon\}$ are fully and univocally determined by the values of $\{d_k^\alpha \in \Upsilon'\}$, via the recursive application of Eqs. (66)–(68) for every grid point $\mathbf{r}_k$.

The VLSD algorithm begins by calculating the values of the contravariant search direction coefficients in Υ'. First, the energy gradients $\{g_k^\alpha\}$ are calculated as

$$g_k^\alpha = \frac{\partial E}{\partial c_{k\alpha}} = \left[\sum_{\beta=1}^{N_\phi} \hat{H}\phi_\beta(\mathbf{r})\, K^{\beta\alpha}\right]_{\mathbf{r}=\mathbf{r}_k},\ \forall c_{k\alpha} \in \Omega. \tag{73}$$

A central difference between of the VLSD algorithm with respect to the TSD algorithm is that the derivative of the energy is calculated exclusively with respect to $\{c_{k\alpha} \in \Omega\}$, and not with respect to all the $\{c_{k\alpha}\}$ coefficients, which is consistent with the choice of variational degrees of freedom. Orthonormality of $\{\psi_i\}$ is maintained to first order with the projection:

$$\tilde{g}_k^\alpha = \left[\sum_{\beta=1}^{N_\phi} \hat{H}\phi_\beta(\mathbf{r}) K^{\beta\alpha} - \sum_{\beta,\gamma,\delta=1}^{N_\phi} \phi_\beta(\mathbf{r}) K^{\beta\gamma} H_{\gamma\delta} S^{\delta\alpha}\right]_{\mathbf{r}=\mathbf{r}_k},\ \forall c_{k\alpha} \in \Omega. \tag{74}$$

Occupancy[50] and kinetic energy preconditioning[45,51] are applied in the contravariant representation point-wise:

$$d_k^\alpha = -\hat{P}\left[\sum_{\beta=1}^{N_\phi} \hat{H}\phi_\beta(\mathbf{r}) S^{\beta\alpha} - \sum_{\beta,\gamma,\delta=1}^{N_\phi} \phi_\beta(\mathbf{r}) S^{\beta\gamma} H_{\gamma\delta} S^{\delta\alpha}\right]_{\mathbf{r}=\mathbf{r}_k},\ \forall c_{k\alpha} \in \Omega. \tag{75}$$

The resulting contravariant search direction coefficients correspond to the entire set $\{d_{k\alpha} \in \Upsilon\}$. At this stage, a loop runs over all the grid points $\mathbf{r}_k$. At each point, the subset separation (59) and (60) is established, followed by the determination of the overlap matrices $A_{k,\eta\iota}$, $B_{k,\iota\sigma}$, and $C_{k,\sigma\tau}$ with Eqs. (61)–(63). Then, Eqs. (66)–(68) are applied recursively to obtain the remaining non-zero covariant and contravariant search direction coefficients.

The localization constraint is introduced using Eq. (66):

$$d_{k\sigma} = 0,\ \forall \sigma \in \chi_k. \tag{76}$$

The contravariant search direction coefficients $\{d_k^\sigma, \forall \sigma \in \chi_k\}$ are calculated using Eq. (67):

$$d_k^\sigma = -\sum_{\eta=1}^{N_\xi}\sum_{\tau=1}^{N_\chi} d_k^\eta B_{k,\eta\tau}\left(C_k^{-1}\right)^{\tau\sigma},\quad \forall \eta \in \xi_k, \forall \sigma,\tau \in \chi_k. \tag{77}$$

And the non-zero covariant search direction coefficients $\{d_{k\iota}, \forall \iota \in \xi_k\}$ are calculated using Eq. (68):

$$d_{k\iota} = \sum_{\eta=1}^{N_\xi} d_k^\eta \left[A_{k,\eta\iota} - \sum_{\tau,\sigma=1}^{N_\chi} B_{k,\eta\tau}\left(C_k^{-1}\right)^{\tau\sigma} B_{k,\sigma\iota}^\dagger\right],\ \forall \eta,\iota \in \xi_k, \forall \sigma,\tau \in \chi_k. \tag{78}$$

At the conclusion of the loop over $\mathbf{r}_k$, the full covariant and contravariant search direction vectors, $\{\Delta_\alpha\}$ and $\{\Delta^\alpha\}$, can be built on the grid as

$$\Delta_\alpha(\mathbf{r}) = \sum_{k=1}^{N_k} D_k(\mathbf{r}) d_{k\alpha}, \tag{79}$$

$$\Delta^\alpha(\mathbf{r}) = \sum_{k=1}^{N_k} D_k(\mathbf{r}) d_k^\alpha, \tag{80}$$

and used to update $\{\phi_\alpha\}$ and $\{\phi^\alpha\}$, respectively, using the steepest-descent method. The formalism of the VLSD method ensures that Eqs. (38) and (39) hold, and hence $\{\Delta_\alpha\}$ and $\{\Delta^\alpha\}$ are mutually dual. More importantly, $\{\Delta_\alpha\}$ and $\{\Delta^\alpha\}$ are fully localized by construction. The VLSD algorithm is depicted in Figure 3.

Practical methods based on strictly localized covariant functions avoid operations that involve the contravariant representation. The reason is that $\{\phi^\alpha\}$ are non-zero in a much larger region of space than $\{\phi_\alpha\}$, and hence they need many more grid coefficients $\{c_k^\alpha\}$ to be correctly discretized. A complete energy minimization scheme can be designed so that only $\{\phi_\alpha\}$ are updated in a steepest-descent fashion at every iteration, which only requires the computation of the expansion coefficients $\{d_{k\alpha}\}$ corresponding to $\{\Delta_\alpha\}$. With the VLSD algorithm, $\{d_{k\alpha}\}$ can be calculated directly using

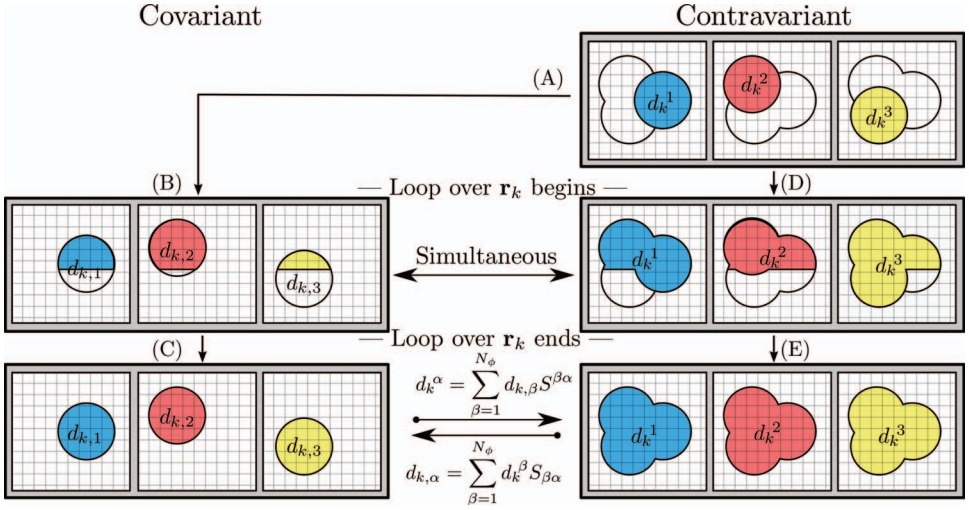

FIG. 3. Schematic representation of the VLSD method. In this example, there are three mutually overlapping functions $\{\phi_\alpha\}$ localized within circumferences in a two-dimensional simulation cell. The algorithm begins in (A), calculating the derivative of the energy with respect to the variational parameters in Ω, imposing orthogonality of $\{\psi_i\}$ and applying occupancy and kinetic energy preconditioning Eq. (75). The result are the values of $\{d_k^\alpha \in \Upsilon'\}$, which correspond to the grid points confined within the circumferences. Then, a loop over all the grid points $\mathbf{r}_k$ starts. The path (A)→(B)→(C) calculates $\{d_{k1}\}$, $\{d_{k2}\}$ and $\{d_{k3}\}$ directly using Eq. (76) and Eq. (78). Alternatively, if requested, the path (A)→(D)→(E) calculates the remaining contravariant search direction coefficients using Eq. (77). The VLSD algorithm is capable of calculating the covariant and contravariant search direction coefficients simultaneously by running steps (B) and (D) as independent processes. At the conclusion of the loop over $\mathbf{r}_k$, the resulting search directions are directly localized within the established localization regions and duality between $\{d_{k\alpha}\}$ and $\{d_k^\alpha\}$ holds exactly.

Eq. (78), following the path (A)→(B)→(C) shown in Fig. 3. After updating $\{\phi_\alpha\}$, $S_{\alpha\beta}$, and $S^{\alpha\beta}$, the density kernel is rescaled to assure electron conservation and to reduce the impact of second-order error terms, and the total energy functional is evaluated. The inner loop is then in charge of restoring orthonormality of $\{\psi_i\}$ beyond second-order terms during the next cycle.

V. RESULTS AND DISCUSSION

A. Description of the calculations

The VLSD algorithm was implemented in the ONETEP program[17] for density functional theory calculations, which, by default, uses an implementation of the TSD method. The current implementation of the VLSD algorithm is not optimized for computational performance. Furthermore, at the moment, calculations using the VLSD algorithm can only be run in serial (one processor), which limits the size of the molecules that can be simulated to a maximum of a few tens of atoms. Nevertheless, this version can be use to test the VLSD method, including its convergence properties and, to some extent, its scaling with the number of atoms. For comparison, TSD calculations were also run in serial. In the current implementation, the covariant search direction coefficients $\{d_{k\alpha}\}$ are calculated using Eq. (78) directly. This is equivalent to follow the path (A)→(B)→(C) in Fig. 3. In the ONETEP approach, the localized functions $\{\phi_\alpha\}$ are non-orthogonal generalized Wannier functions (NGWFs),[40] which are spatially localized within spheres of radii R_α in real space. The NGWFs are represented in a basis set $\{D_k\}$ of orthonormal periodic-sinc (psinc) functions,[27] centered on the points of a uniform Cartesian grid. The psinc functions are related to plane waves via a unitary transformation, which allows systematic control of the basis set resolution by varying the plane-wave kinetic-energy cut-off.

The calculations with the ONETEP program used a kinetic energy cut-off of 800 eV, the PBE[52] exchange-correlation functional and norm-conserving pseudopotentials.[45] An inner loop performed 3 LNV iterations to optimize the density kernel. The size of the localization regions of $\{\phi_\alpha\}$ is taken into account by performing calculations with different R_α in the range of 3.0 Å to 7.0 Å. Periodic boundary conditions were required. Interactions between periodic images are largely avoided by confining the systems within a cubic simulation cell of 52.91 Å width, ensuring that a vacuum space of at least 45.46 Å exists between periodic images. The possibility of long-range ionic interactions affecting the convergence of the calculations was discarded after performing calculations with spherical Coulomb cut-off[53] and obtaining identical results. Fast Fourier transforms were calculated using cubic FFTboxes[27,54] of size 44.45 Å and 189 grid points on each lattice vector direction.

A set of sixteen molecular systems was chosen and organized in four groups, each of which contains molecules that share a common property. These groups will be referred to hereafter as G1, G2, G3, and G4, respectively, and are constructed as follows. G1 contains molecules with at least one chlorine atom: Cl_2 (chlorine diatomic), CH_3Cl (methyl chloride), HCl (hydrogen chloride), and NaCl (sodium chloride). The first three of these chlorinated molecules contain covalent bonds, while NaCl is an ion pair with a long ionic bond. CH_3Cl, HCl, and NaCl are expected to show non-zero electric dipole. Calculations on these systems with Gaussian basis set approaches require the addition of diffuse functions

in order to achieve chemical accuracy in the prediction of their electronic properties.[55,56] Therefore, a certain degree of natural delocalization of their molecular orbitals is expected. G2 is formed by the phosphoric acid molecule H_3PO_4 and the subsequent anions $H_2PO_4^-$, HPO_4^{2-}, and PO_4^{3-}, of increasing total negative charge. Also, for such molecules, calculations with Gaussian basis sets need the addition of diffuse functions to correctly describe electronic properties.[55,56] Delocalization of their molecular orbitals is expected to increase with the total charge. G3 encompasses a set of common small molecules that form short covalent and hydrogen bonds. These are H_2O (water), NH_3NH_3 (ammonia dimer, hydrogen bonded), C_6H_6 (benzene, aromatic), and NH_4^+ (ammonium cation). The small size and reduced electric dipole of these test systems is expected to result in a high natural localization of their molecular orbitals. Finally, G4 contains four amino acids that are often found in proteins and other organic complexes: lysine (lys^+), glutamic acid (glu^-), glycine (gly), and serine (ser). Lysine and glutamic acid are positively and negatively charged, respectively, glycine is non-polar and serine is polar. Natural localization of the molecular orbitals is expected to be observed in these systems. The number of NGWFs per atom is as follows: one for H, four for C, N, O, and Na, and nine for Cl and P.

B. Analysis of the TSD method

The effect of truncation of the search direction vectors in calculations with the TSD method was analyzed. The coefficient μ, defined as

$$\mu = \sum_{\alpha=1}^{N_\phi} \frac{\sum_{\mathbf{r}_k \in \mathrm{LR}(\Delta_\alpha(\mathbf{r}))} d_{k\alpha}^2}{\sum_{k=1}^{N_k} \tilde{d}_{k\alpha}^2}, \tag{81}$$

compares the sum of the squares of $\{d_{k\alpha}\}$ inside $\mathrm{LR}(\Delta_\alpha(\mathbf{r}))$ and the sum of the squares of $\{\tilde{d}_{k\alpha}\}$ in the entire simulation cell prior to truncation. By definition, μ is confined to be in the range $0 \leq \mu \leq 1$, where $\mu = 0$ means that the truncation step deletes all the information available, and $\mu = 1$ means that no truncation happens at all. It must be emphasized that calculations with the VLSD method always result in $\mu = 1$, by construction. The value of μ was plotted against the iteration counter n for all the test systems and all the different values of R_α. Figure 4 shows the results. A first observation is that, as the number of iterations increases, μ saturates to a constant value, denoted by μ_0, indicating that the ratio of the amount of information that is deleted to the amount of information that is kept does no longer change. This trend was observed in calculations on all the four groups of test systems with any value of R_α. The pace at which μ saturates to μ_0 is

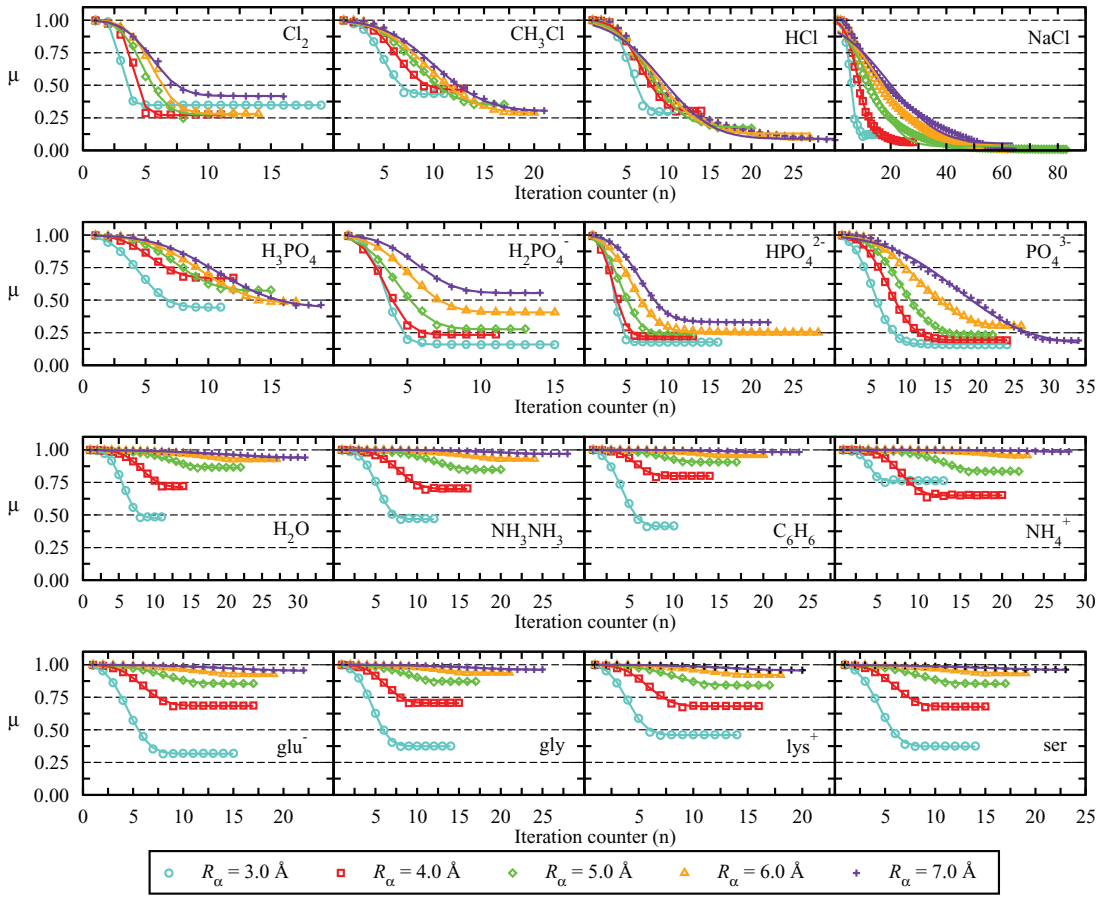

FIG. 4. Evolution of the parameter μ with the outer loop iteration counter, n, for different values of R_α. The solid lines are the Gaussian broadening curves used to fit the results.

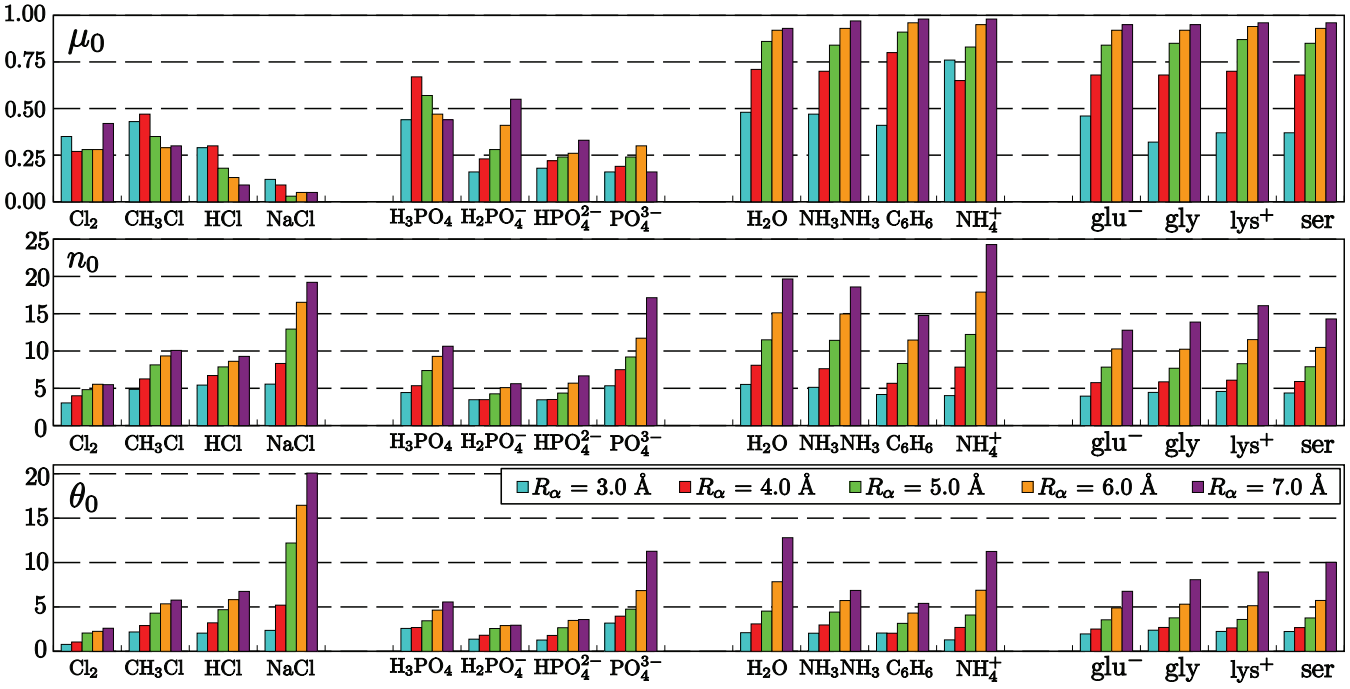

FIG. 5. Progression of μ_0, n_0, and θ_0, calculated after fitting a Gaussian broadening function through the results, for different values of R_α.

determined by the slope of the curve. To quantify and analyze this behavior, a Gaussian broadening curve with equation

$$\mu(n) = \mu_0 + \frac{(1-\mu_0)}{2}\left(1 - \mathrm{erf}\left[\frac{n-n_0}{\theta_0}\right]\right), \qquad (82)$$

was fitted to the results, where erf is the error function, n_0 is the iteration number for which $\mu = \frac{\mu_0}{2}$, and θ_0 is the broadening factor (inversely proportional to the slope). The choice of the Gaussian broadening function was motivated after trialing other alternatives, such as an exponential function, hyperbolic tangent, cumulative Gamma distribution, Gaussian distribution, and Fermi-Dirac distribution. It was found that the Gaussian broadening curve resulted in the smallest fitting error, with a correlation coefficient superior to 0.99 in all cases.

The values of μ_0, n_0, and θ_0 for each of the test cases in every group were plotted for the different values of R_α. Figure 5 shows the results. The relation between μ_0 and R_α, and between n_0, θ_0, and R_α, must be analyzed separately. Calculations on molecules in G1 and G2 resulted in a very low μ_0, typically below 0.50, and in some cases like HCl, NaCl, and PO_4^{3-}, below 0.25 for most of the values of R_α. The values of μ_0 for NaCl are remarkably low, reaching a minimum of 0.03 with $R_\alpha = 5.0$ Å. Such low values indicate that, in calculations on molecules in G1 and G2, during the last iterations, comparatively, there is more information contained outside the localization spheres than inside. The value of μ_0 does not increase with R_α for molecules in G1 and G2, however, it consistently increases with R_α for molecules in G3 and G4. Regarding n_0 and θ_0, it seems plausible to think that these two variables are in some way inter-related, for a smoother decay is expected to result in a larger value of n_0 and θ_0, and *vice versa*. It is clear from Fig. 4 that n_0 and θ_0 increase with R_α in every system in every group. Therefore, the non-variational truncation of the search direction results in the deletion of less information per iteration as R_α increases, measured in relative terms.

With the TSD method, the search direction is calculated as unconstrained and allowed to be delocalized until the final stage, where the truncation step imposes the localization constraints. According to the results, during the first few iterations, a small fraction of the total information is deleted. Then μ decays to the saturation value μ_0 at speed that is inversely proportional to R_α. A low value of μ_0 can be interpreted as a failure of the implicit assumption of TSD: that the truncated search direction of the unconstrained problem resembles the search direction of the problem with exact localization constraints. Truncation has a stronger effect in calculations on molecules in G1 and G2 than in calculations on molecules in G3 and G4. This is in consonance with the expectations of more diffuse molecular orbitals on systems in G1 and G2 than on G3 and G4.

C. Convergence properties

The convergence rate of the TSD and VLSD methods, defined as the number of iterations that is necessary to achieve self-consistent convergence, was compared. Convergence is considered to be achieved when the root-mean-square (RMS) of the gradient along the search direction falls below a tolerance threshold of 3.8×10^{-5} eV, and the change in the total energy below a tolerance threshold of 1×10^{-7} eV/atom. All the TSD and VLSD calculations converged according to these criteria, with the only exception of $H_2PO_4^-$ with $R_\alpha = 3.0$ Å, in which the total energy converged to the established tolerance threshold with an RMS gradient of 6.8×10^{-4} eV, both in TSD and VLSD calculations. The total number of iterations

TABLE I. Number of outer loop iterations, n, required to achieve self-consistent convergence in calculations with the TSD and VLSD methods. The values refer to the ratio n_{TSD}/n_{VLSD}.

R_α (Å)	Cl_2	CH_3Cl	HCl	NaCl	H_3PO_4	$H_2PO_4^-$	HPO_4^{2-}	PO_4^{3-}	H_2O	NH_3NH_3	C_6H_6	NH_4^+	glu^-	gly	lys^+	ser
3.0	19/12	11/9	11/10	16/13	11/10	15/14	16/16	24/11	11/11	12/10	10/8	13/11	15/15	14/14	14/14	14/13
4.0	11/8	13/11	14/12	28/17	12/10	11/14	13/14	24/11	14/15	16/14	14/15	20/14	17/16	15/14	16/16	15/14
5.0	12/12	17/14	20/18	81/25	15/13	13/11	13/14	22/13	22/20	20/17	17/17	22/24	17/18	17/16	17/16	17/16
6.0	14/12	20/15	27/32	62/28	17/15	15/12	28/12	26/15	27/21	24/21	20/18	23/26	19/19	21/21	18/17	19/17
7.0	16/9	21/18	43/22	64/31	19/16	14/12	22/12	34/20	31/29	28/23	24/22	28/27	22/20	25/22	20/19	23/22

to achieve convergence with the TSD and VLSD methods, denoted as n_{TSD} and n_{VLSD}, respectively, is shown in Table I. The coefficient $\zeta = n_{TSD}/n_{VLSD}$ gives a comparison of the relative convergence properties of the two methods: $\zeta > 1$ indicates that VLSD needs less iterations than TSD to converge, whereas $\zeta < 1$ indicates otherwise. The value of ζ in every calculation is shown in Fig. 6.

The results show that, with some exceptions, calculations with the VLSD method require less iterations to converge than calculations with the TSD method. The analysis of the value of ζ indicates that calculations on molecules in G1 and G2 seem to benefit more of the VLSD method than calculations on molecules in G3 and G4. The largest values of ζ are obtained for NaCl, HPO_4^{2-}, and PO_4^{3-}, which also show some of the lowest values of μ_0 in TSD calculations. The maximum ζ is obtained for NaCl with $R_\alpha = 5.0$ Å, which has the lowest μ_0 coefficient of all in TSD calculations, and corresponds to a value of $\zeta = 3.24$, meaning that TSD needs more than three times the number of VLSD iterations to converge. On the other hand, calculations on molecules in G3 and G4 seem to require approximately the same number of iterations with TSD than with VLSD. These results seem to point out that the VLSD method is more robust compared to TSD, especially in molecules where a certain degree of delocalization of the unconstrained molecular orbitals is expected.

D. Validation tests

A number of tests were performed to confirm that both the TSD and VLSD methods are capable of predicting structural and electronic properties correctly. According to the results in Sec. V C, NaCl converges very differently with the TSD and VLSD methods, whereas H_2O shows a similar convergence pattern with both. These two molecules were chosen to perform validation tests. The structures were optimized using the implementation of the BFGS algorithm in ONETEP,[57] using either the TSD or the VLSD methods to calculate the total energy and the atomic forces. The equilibrium bond length, l_0, of the Na–Cl bond in NaCl and of the O–H bonds in H_2O were determined. Then, the magnitude of the total electric dipole moment at the equilibrium geometry, p_0, was calculated. The results were compared to the values of l_0 and p_0 obtained with all-electron calculations using the NWCHEM program[58] and the near-complete aug-cc-pVQZ Gaussian basis set, which contains a large number of polarization and diffuse functions.[59,60] The values of l_0 and p_0 are shown in Table II. The optimized geometries are converged within ± 0.0005 Å maximum deviation in all cases.

The equilibrium bond lengths l_0 of NaCl and H_2O calculated with the TSD and VLSD methods in ONETEP and with NWCHEM agree very well, within a <1% deviation margin. Convergence of l_0 to a single value with respect to to R_α is fast both in NaCl and H_2O calculations.

On the other hand, the value of the total electric dipole moment seems to be more sensitive to changes in the accuracy of the basis set and the numerical method. In ONETEP calculations, p_0 converges to a single value when R_α increases, both in NaCl and H_2O calculations. The convergence of p_0 is faster in H_2O calculations than in NaCl, which might be seen as a consequence of the greater degree of delocalization of the Kohn-Sham orbitals in NaCl. The converged value of p_0 in ONETEP calculations, with either TSD or VLSD methods, is around 8.40 D for NaCl and 1.86 D for H_2O, which differ from the values obtained with NWCHEM by approximately −3% and 4%, respectively. This disagreement can be attributed to several numerical factors, such as the difference in the basis sets (localized NGWFs and psinc functions in ONETEP, and delocalized Gaussian functions in NWCHEM), the fineness of the ONETEP computational grid (defined by the choice of kinetic energy cut-off) or the treatment of the core electrons (non-local pseudopotentials in ONETEP and all-electron in NWCHEM).

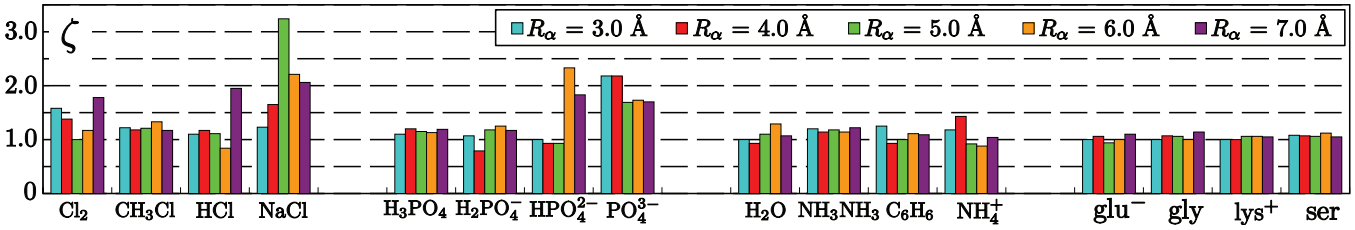

FIG. 6. Calculated value of the coefficient $\zeta = n_{TSD}/n_{VLSD}$ for each test system and increasing R_α.

TABLE II. Calculated equilibrium bond length, l_0, and total electric dipole moment at the equilibrium geometry, p_0, of NaCl and H_2O. The two O–H bonds in H_2O have the same calculated length.

		NaCl		H_2O	
	R_α (Å)	l_0 (Å)	p_0 (D)	l_0 (Å)	p_0 (D)
TSD	3.0	2.384	9.54	0.967	1.90
	4.0	2.375	8.95	0.967	1.90
	5.0	2.376	8.55	0.967	1.88
	6.0	2.375	8.42	0.968	1.86
	7.0	2.375	8.40	0.968	1.86
VLSD	3.0	2.381	9.51	0.967	1.90
	4.0	2.367	9.00	0.969	1.90
	5.0	2.370	8.68	0.968	1.88
	6.0	2.377	8.49	0.968	1.87
	7.0	2.376	8.42	0.968	1.86
NWChem		2.380	8.60	0.969	1.80

E. Scaling with the system size

Calculations on a set of poly-ethylene fragments of up to 92 atoms were performed to test the scaling of the VLSD algorithm as it was implemented. The calculations shown here used 800 eV kinetic energy cut-off and the PBE exchange-correlation functional. All the molecules under study were centered in a rectangular box of size 14.55 Å × 14.55 Å × 50.27 Å. Four NGWFs were placed on each carbon atom and one on each hydrogen atom, all with radii of 4.0 Å. It must be recalled that the current implementation of the VLSD algorithm is not optimized for performance and that it can only be executed on one core. The discussion of the scaling of the VLSD algorithm with the system size will first focus on the scaling of the unoptimized version of VLSD. Then, the optimization of the VLSD method for computational performance will be discussed. Figure 7 shows the timings registered for the unoptimized VLSD algorithm as well as the projection of the expected timings with an optimized VLSD code.

The unoptimized version of VLSD requires a loop over every single grid point that is inside any of the localization regions. The grid points that are outside all the localization regions are ignored since both $\{d_{k\alpha}\}$ and $\{d_k^\alpha\}$ are identically

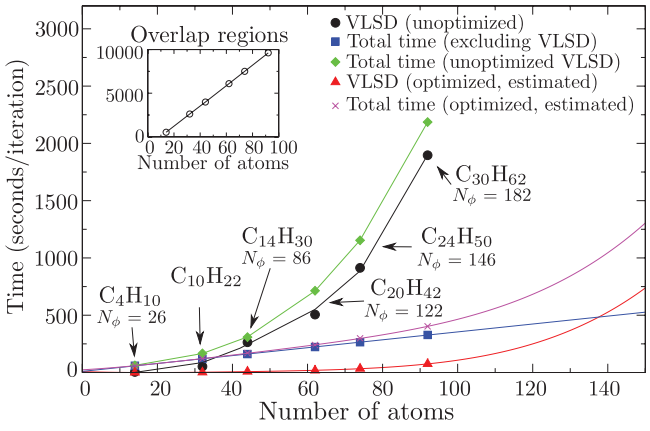

FIG. 7. Scaling of the VLSD algorithm with the system size. The number of overlap regions as a function of the number of atoms is shown in the inset.

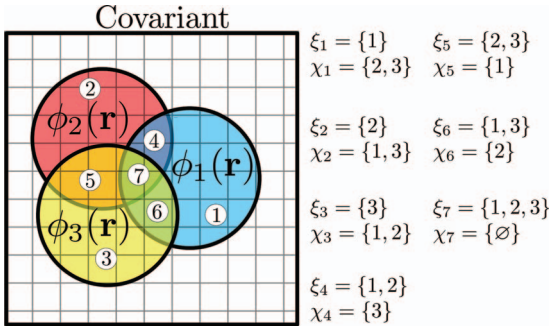

FIG. 8. Distinct overlap regions in a calculation with three localized covariant functions. The grid points within each of the seven overlap regions share the same subsets ξ_k and χ_k.

zero at these points. The total number of valid grid points is proportional to the system size, and scales as $\mathcal{O}(N)$. During the loop over the valid grid points, the major computational bottleneck is due to the matrix $C_{k,\sigma\tau}$ (63), which must be created, inverted and then used in Eq. (78). The size of $C_{k,\sigma\tau}$ depends on the grid point k. Its maximum possible size is $(N-1) \times (N-1)$ (which occurs when k is inside only one covariant sphere). Therefore, the cost of inverting $C_{k,\sigma\tau}$ is proportional to $\mathcal{O}(N^3)$. Overall, the cost of the VLSD loop over the valid grid points scales as $\mathcal{O}(N^4)$. This is what Fig. 7 shows, where the results of the series "VLSD (unoptimized)" were fitted to a fourth-order polynomial.

In a standard calculation, there are millions of grid points that are located inside at least one localization region. Looping over all of them is highly inefficient. Returning to Eqs. (59) and (60), which mark the subset separation that is in the core of the VLSD method, it must be noticed that the same subsets ξ_k and χ_k can (and will) be associated to a set of distinct grid points k which are located in the same overlap region. The number of distinct overlap regions increases linearly with the system size (as shown in the inset plot in Fig. 7). Figure 8 illustrates this idea with an example regarding a calculation with three covariant functions that gives rise to seven different overlap regions. All the grid points within the same overlap region share the same $A_{k,\eta\iota}$, $B_{k,\iota\sigma}$, and $C_{k,\sigma\tau}$ matrices, which have to be calculated (and in the case of $C_{k,\sigma\tau}$, inverted) only once per overlap region. Therefore, the VLSD algorithm can be largely optimized by transforming the loop over the grid points into a loop over the distinct overlap regions. In the calculations in poly-ethylenes, the ratio between the total number of overlap regions and the total number of valid grid points is of the order of 10^{-2}. The estimated timings shown in the series "VLSD (optimized)" shown in Fig. 7, correspond to the unoptimized timings multiplied with the above-mentioned ratio.

The scaling of the optimized VLSD method is still $\mathcal{O}(N^4)$, but the prefactor due to the execution of the VLSD part of the code has been greatly reduced. It must be noted that all the remaining parts of the calculation (excluding VLSD) are executed at $\mathcal{O}(N)$ cost (in this case, following the standard ONETEP methodology). Extrapolating the series "VLSD (optimized)" in Fig. 7 results in a crossover point near 140 atoms before the quartic scaling of VLSD becomes more expensive than the rest of the calculation. A further reduction to

the total time per iteration of the VLSD method could be obtained with parallel algebra libraries, such as SCALAPACK,[61] to invert $C_{k,\sigma\tau}$ using a large number of cores.

Although, as shown in Sec. V C, calculations with the VLSD method require less iterations to converge than with the TSD method, each VLSD iteration requires a longer time to complete than a TSD iteration. The TSD method is linear-scaling (the cost scales as the blue squares series in Fig. 7), and therefore it is suitable for large-scale calculations. On the other hand, it is clear that further steps must be taken before the VLSD method can be considered for large-scale calculations, due to its $\mathcal{O}(N^4)$ cost. For example, the cost of inversion and subsequent algebraic manipulations of $C_{k,\sigma\tau}$ could be reduced to $\mathcal{O}(N)$ using well-tested techniques common in standard linear-scaling methods, such as the Cholesky factorization method proposed by Challacombe,[33] the recursion method developed by Ozaki[34] or the Hotelling algorithm.[47] These methods are approximate and require the truncation of some of the elements of $(C_k^{-1})^{\tau\sigma}$ in order to sufficient sparsity.[44] The overall cost of the VLSD method could be reduced to $\mathcal{O}(N^2)$ with the use of these techniques (which may, however, carry a large prefactor). More research is required in order to devise a fully linear-scaling VLSD-based approach.

VI. CONCLUSIONS

A new method for imposing localization constraints in self-consistent energy minimization approaches $(C_k^{-1})^{\tau\sigma}$ for electronic structure calculations, called the VLSD, was described. With the VLSD method, the constraint of localization is included naturally in the mathematical formalism by means of a close analysis of the variational degrees of freedom and the use of formal tensorial algebra. As a result, the search direction vectors calculated with the VLSD method are fully localized by construction. The VLSD method is capable of achieving rapid convergence (measured as the total number of iterations) regardless of the size of the localization region. The conventional TSD method, in contrast, relies on the approximation that the truncated search direction vectors of the unconstrained problem resemble the exact search direction vectors of the constrained problem. The results of calculations on a number of molecules show that the TSD method can result in the deletion of a non-negligible portion of information, which was pre-calculated and stored in the delocalized search direction vectors, in order to keep the localization constraints. It was observed that the VLSD method often converges in less iterations than the TSD method. More particularly, in calculations on systems where the deletion of information by the TSD method is remarkably large, the VLSD method clearly outperforms TSD by converging in up to a third of the number of iterations. Validation tests confirm that structural and electronic properties can be accurately calculated with both methods. A direct implementation of the VLSD algorithm yields a method that scales as $\mathcal{O}(N^4)$ with the number of atoms. The cost could be reduced to $\mathcal{O}(N^2)$ by using indirect or approximate matrix inversion techniques. Further investigations are necessary to attain linear-scaling cost in VLSD calculations.

In this work, the VLSD algorithm has been introduced as a technique to optimize non-orthogonal, localized functions in real space. However, this is only one particular application of the method. VLSD could be generalized to any optimization problem in a non-orthogonal representation in which the search direction vector is subject to a set of analytical constraints. For example, VLSD can be used to impose that a certain subset of localized, non-orthogonal functions must remain constant while the rest are optimized, by ensuring that the relevant associated subset of search direction vectors is identically zero. Moreover, the VLSD method should not be limited to the case of localized functions. The optimization of any *n*th order tensor in a non-orthogonal representation can benefit from the VLSD algorithm. An example of a second-order tensor optimization problem of special relevance in linear-scaling density functional theory approaches is the optimization of the density matrix in a non-orthogonal representation,[35,36] which is expected to be localized, and in practice it is made localized by truncating its associated search direction vector.

ACKNOWLEDGMENTS

A.R.-S. acknowledges the support of the Engineering and Physical Sciences Research Council (EPSRC) (Grant No. EP/F038038/1) for a High End Computing Studentship through the UKCP consortium. C.-K.S. acknowledges support from the Royal Society in the form of a University Research Fellowship. The authors are grateful for the computing resources provided by Southampton University's iSolutions unit (Iridis3 supercomputer) which have enabled all the calculations presented here.

[1] W. Kohn, Phys. Rev. **115**, 809 (1959).
[2] W. Kohn, Phys. Rev. Lett. **76**, 3168 (1996).
[3] G. Galli and M. Parrinello, Phys. Rev. Lett. **69**, 3547 (1992).
[4] E. Tsuchida and M. Tsukada, J. Phys. Soc. Jpn. **67**, 3844 (1998).
[5] F. Aryasetiawan and O. Gunnarsson, Rep. Prog. Phys. **61**, 237 (1998).
[6] J. E. Subotnik, A. Sodt, and M. Head-Gordon, J. Chem. Phys. **125**, 074116 (2006).
[7] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, Comput. Phys. Commun. **180**, 2175 (2009).
[8] P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964).
[9] W. Kohn and L. J. Sham, Phys. Rev. **140**, A1133 (1965).
[10] E. Prodan and W. Kohn, Proc. Natl. Acad. Sci. U.S.A. **102**, 11635 (2005).
[11] W. Yang and T.-S. Lee, J. Chem. Phys. **103**, 5674 (1995).
[12] S. Goedecker, Rev. Mod. Phys. **71**, 1085 (1999).
[13] D. R. Bowler and T. Miyazaki, Rep. Prog. Phys. **75**, 036503 (2012).
[14] J.-L. Fattebert and J. Bernholc, Phys. Rev. B **62**, 1713 (2000).
[15] J. Soler, E. Artacho, J. Gale, A. Garcia, J. Junquera, P. Ordejon, and D. Sanchez-Portal, J. Phys. Condens. Matter **14**, 2745 (2002).
[16] M. J. Gillan, D. R. Bowler, A. S. Torralba, and T. Miyazaki, Comput. Phys. Commun. **177**, 14 (2007).
[17] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne, J. Chem. Phys. **122**, 084119 (2005).
[18] L. Genovese, A. Neelov, S. Goedecker, T. Deutsch, S. A. Ghasemi, A. Willand, D. Caliste, O. Zilberberg, M. Rayson, A. Bergman, and R. Schneider, J. Chem. Phys. **129**, 014109 (2008).
[19] T. Ozaki and H. Kino, Phys. Rev. B **72**, 045121 (2005).
[20] E. Artacho and L. Miláns del Bosch, Phys. Rev. A **43**, 5770 (1991).
[21] C. A. White, P. Maslen, M. S. Lee, and M. Head-Gordon, Chem. Phys. Lett. **276**, 133 (1997).
[22] D. D. O'Regan, "Optimised projections for the *ab initio* simulation of large and strongly correlated systems," Ph.D. thesis (University of Cambridge, 2011).
[23] O. F. Sankey and D. J. Niklewski, Phys. Rev. B **40**, 3979 (1989).
[24] J. Junquera, Ó. Paz, D. Sánchez-Portal, and E. Artacho, Phys. Rev. B **64**, 235111 (2001).

[25]A. S. Torralba, M. Todorovic, V. Brazdova, R. Choudhury, T. Miyazaki, M. J. Gillan, and D. R. Bowler, J. Phys. Condens. Matter **20**, 294206 (2008).
[26]A. Ruiz-Serrano, N. D. M. Hine, and C.-K. Skylaris, J. Chem. Phys. **136**, 234101 (2012).
[27]C.-K. Skylaris, A. A. Mostofi, P. D. Haynes, C. J. Pickard, and M. C. Payne, Comput. Phys. Commun. **140**, 315 (2001).
[28]E. Hernández, M. J. Gillan, and C. M. Goringe, Phys. Rev. B **55**, 13485 (1997).
[29]J.-L. Fattebert, J. Phys. Condens. Matter **20**, 294210 (2008).
[30]C.-K. Skylaris and P. D. Haynes, J. Chem. Phys. **127**, 164712 (2007).
[31]F. Mauri and G. Galli, Phys. Rev. B **50**, 4316 (1994).
[32]P. Ordejón, D. A. Drabold, R. M. Martin, and M. P. Grumbach, Phys. Rev. B **51**, 1456 (1995).
[33]M. Challacombe, J. Chem. Phys. **110**, 2332 (1999).
[34]T. Ozaki, Phys. Rev. B **64**, 195110 (2001).
[35]E. Hernández and M. J. Gillan, Phys. Rev. B **51**, 10157 (1995).
[36]P. D. Haynes, C.-K. Skylaris, A. A. Mostofi, and M. C. Payne, J. Phys.: Condens. Matter **20**, 294207 (2008).
[37]X. Li, R. Nunes, and D. Vanderbilt, Phys. Rev. B **47**, 10891 (1993).
[38]R. W. Nunes and D. Vanderbilt, Phys. Rev. B **50**, 17611 (1994).
[39]P. Pulay, Chem. Phys. Lett. **73**, 393 (1980).
[40]C.-K. Skylaris, A. A. Mostofi, P. D. Haynes, O. Dieguez, and M. C. Payne, Phys. Rev. B. **66**, 035119 (2002).
[41]S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. J. Probert, K. Refson, and M. C. Payne, Z. Kristallogr. Kristallgeom. Kristallphys. **220**, 567 (2005).
[42]C. J. Tymczak and M. Challacombe, J. Chem. Phys. **122**, 134102 (2005).
[43]E. Rudberg, E. H. Rubensson, and P. Sałek, J. Chem. Theory Comput. **7**, 340 (2011).
[44]E. H. Rubensson and E. Rudberg, J. Comput. Chem. **32**, 1411 (2011).
[45]M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, Rev. Mod. Phys. **64**, 1045 (1992).
[46]C. K. Gan, P. D. Haynes, and M. C. Payne, Comput. Phys. Commun. **134**, 33 (2001).
[47]W. H. Press, W. T. Vetterling, S. A. Teulosky, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge University Press, 1994).
[48]A. D. Rabuck and G. E. Scuseria, J. Chem. Phys. **110**, 695 (1999).
[49]E. Hernández, M. J. Gillan, and C. M. Goringe, Phys. Rev. B **53**, 7147 (1996).
[50]N. Marzari, D. Vanderbilt, and M. C. Payne, Phys. Rev. Lett. **79**, 1337 (1997).
[51]A. A. Mostofi, P. D. Haynes, C.-K. Skylaris, and M. C. Payne, J. Chem. Phys. **119**, 8842 (2003).
[52]J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**, 3865 (1996).
[53]N. D. M. Hine, J. Dziedzic, P. D. Haynes, and C.-K. Skylaris, J. Chem. Phys. **135**, 204103 (2011).
[54]A. A. Mostofi, C.-K. Skylaris, P. D. Haynes, and M. C. Payne, Comput. Phys. Commun. **147**, 788 (2002).
[55]E. R. Davidson and D. Feller, Chem. Rev. **86**, 681 (1986).
[56]J. G. Hill, Int. J. Quantum Chem. **113**, 21 (2013).
[57]B. G. Pfrommer, M. Cote, S. G. Louie, and M. L. Cohen, J. Comput. Phys. **131**, 233 (1997).
[58]M. Valiev, E. J. Bylaska, N. Govind, K. Kowalski, T. P. Straatsma, H. J. J. van Dam, D. Wang, J. Nieplocha, E. Apra, T. L. Windus, and W. A. de Jong, Comput. Phys. Commun. **181**, 1477 (2010).
[59]J. Thom H. Dunning, J. Chem. Phys. **90**, 1007 (1989).
[60]D. E. Woon and J. Thom H. Dunning, J. Chem. Phys. **98**, 1358 (1993).
[61]L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley, *ScaLAPACK Users' Guide* (Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997).