

Linear Response Equations Revisited: A Simple and Efficient Iterative Algorithm

Riccardo Alessandro, Ivan Gianni, Federica Pes, Tommaso Nottoli,* and Filippo Lipparini*



Cite This: *J. Chem. Theory Comput.* 2023, 19, 9025–9031



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

ABSTRACT: We present an algorithm to solve the linear response equations for Hartree–Fock, Density Functional Theory, and the Multiconfigurational Self-Consistent Field method that is both simple and efficient. The algorithm makes use of the well-established symmetric and antisymmetric combinations of trial vectors but further orthogonalizes them with respect to the scalar product induced by the response matrix. This leads to a standard, symmetric block eigenvalue problem in the expansion subspace that can be solved by diagonalizing a symmetric, positive definite matrix half the size of the expansion space. Numerical tests showed that the algorithm is robust and stable.



Linear response calculations¹ are commonly encountered in computational chemistry, because they are used to compute excitation energies and transition properties. Such calculations are available at virtually any level of theory, from Hartree–Fock (HF) and Kohn–Sham Density Functional Theory (KS-DFT),^{2–4} to coupled-cluster theory,^{5–8} to Multiconfigurational Self-Consistent Field (MCSCF),^{7,9–11} just to name a few. Given the importance and usefulness of such calculations, algorithms that allow for efficient implementation are paramount. In this short communication, we focus on self-consistent field (SCF) methods, i.e., HF, density functional theory (DFT), and MCSCF. In particular, we present the equations and detail the algorithm for MCSCF, as HF and DFT can be considered as special cases of the former.

The linear response equations for MCSCF have the general form, assuming that real basis functions are used:^{1,7,9–13}

$$\begin{pmatrix} A & B \\ B & A \end{pmatrix} \begin{pmatrix} \mathbf{y} \\ \mathbf{z} \end{pmatrix} = \omega \begin{pmatrix} \Sigma & \Delta \\ -\Delta & -\Sigma \end{pmatrix} \begin{pmatrix} \mathbf{y} \\ \mathbf{z} \end{pmatrix} \quad (1)$$

where A , B , and Σ are symmetric matrices in $\mathbb{R}^{n \times n}$, Δ is an antisymmetric matrix in $\mathbb{R}^{n \times n}$, and \mathbf{y} , \mathbf{z} are vectors in \mathbb{R}^n . Here, n represents the size of the problem, e.g., the number of occupied times number of virtual orbitals for HF and KS-DFT, or the number of orbital rotations and independent CI parameters for MCSCF. We further assume that Σ is positive definite, as are the combinations $A + B$ and $A - B$, conditions that are met if the ground-state solution is stable. The eigenvalue problem (eq 1) has been investigated in the standard form as well as in the generalized one: in refs 14–16, the authors developed a minimization principle to find a few smallest eigenvalues and the corresponding eigenvectors. For HF and KS-DFT response theory, eq 1 is simplified, as Σ

becomes the identity matrix and Δ becomes the zero matrix. For the sake of brevity, we also write the response equations as

$$\Lambda \mathbf{x} = \omega \Omega \mathbf{x} \quad (2)$$

eq 2 is a generalized eigenvalue problem with a nonpositive-definite metric Ω , which means that the standard procedure used in other quantum chemical calculations, where either a Cholesky decomposition of the metric or $\Omega^{-1/2}$ are computed in order to transform eq 2 into a standard eigenvalue problem, cannot be used. The various algorithms that have been proposed to tackle such a problem must, therefore, deal with this complication. A few of the main strategies proposed in the literature are discussed in the following.

The linear response equations have an important property. If $\{\omega, (\mathbf{y}, \mathbf{z})^T\}$ is a solution, then $\{-\omega, (\mathbf{z}, \mathbf{y})^T\}$ is also a solution to the same problem. In other words, the eigenvalues of eq 2 always appear as positive–negative pairs. It has been suggested in the literature^{17–20} that this symmetry should be encoded in the iterative algorithm used to solve the response equations, usually a modified version of Davidson’s method,²¹ by expanding each eigenvector \mathbf{x} as the linear combination of two sets of vectors:


$$\mathbf{x} = \sum_{i=1}^k (u_i^+ \mathbf{b}_i(+)) + u_i^- \mathbf{b}_i(-) \quad (3)$$

Received: September 9, 2023

Revised: October 30, 2023

Accepted: December 8, 2023

Published: December 11, 2023



where $u_i^\pm \in \mathbb{R}$ and the expansion vectors are defined as follows:

$$\mathbf{b}_i(+)=\begin{pmatrix} \mathbf{b}_i^+ \\ \mathbf{b}_i^+ \end{pmatrix} \in \mathbb{R}^{2n}, \quad \mathbf{b}_i(-)=\begin{pmatrix} \mathbf{b}_i^- \\ -\mathbf{b}_i^- \end{pmatrix} \in \mathbb{R}^{2n} \quad (4)$$

with $\mathbf{b}_i^+, \mathbf{b}_i^- \in \mathbb{R}^n$. In the following, we use the notation, for a generic vector $\mathbf{v} \in \mathbb{R}^{2n}$, $\mathbf{v}(+) = (\mathbf{v}^+, \mathbf{v}^+)^T$ and $\mathbf{v}(-) = (\mathbf{v}^-, -\mathbf{v}^-)^T$, with $\mathbf{v}^+, \mathbf{v}^- \in \mathbb{R}^n$ to denote symmetric and antisymmetric vectors, respectively. Furthermore, to keep the notation simple, we present the derivation of the equations for the case where only one eigenpair is sought: the generalization to many eigenpairs is straightforward. This choice of expansion spaces is particularly advantageous, as it gives rise to blocked reduced matrices for problem 1. In fact,

$$\begin{aligned} \Lambda \mathbf{b}(+) &= \begin{pmatrix} (A+B)\mathbf{b}^+ \\ (A+B)\mathbf{b}^+ \end{pmatrix} = \boldsymbol{\sigma}(+), \quad \boldsymbol{\sigma}^+ = (A+B)\mathbf{b}^+ \\ \Lambda \mathbf{b}(-) &= \begin{pmatrix} (A-B)\mathbf{b}^- \\ -(A-B)\mathbf{b}^- \end{pmatrix} = \boldsymbol{\sigma}(-), \quad \boldsymbol{\sigma}^- = (A-B)\mathbf{b}^- \end{aligned} \quad (5)$$

and

$$\begin{aligned} \Omega \mathbf{b}(+) &= \begin{pmatrix} (\Sigma+\Delta)\mathbf{b}^+ \\ -(\Sigma+\Delta)\mathbf{b}^+ \end{pmatrix} = \boldsymbol{\tau}(-), \quad \boldsymbol{\tau}^- = (\Sigma+\Delta)\mathbf{b}^+ \\ \Omega \mathbf{b}(-) &= \begin{pmatrix} (\Sigma-\Delta)\mathbf{b}^- \\ (\Sigma-\Delta)\mathbf{b}^- \end{pmatrix} = \boldsymbol{\tau}(+), \quad \boldsymbol{\tau}^+ = (\Sigma-\Delta)\mathbf{b}^- \end{aligned} \quad (6)$$

In the spirit of Davidson's method, the response equations can be therefore solved as follows.

Let \mathcal{V}_k^+ and \mathcal{V}_k^- be the symmetric and antisymmetric expansion subspaces, respectively, whose dimension increases with k . Here k represents the number of iterations already performed. Let

$$V_k(+)=\left(\mathbf{b}_1(+)\dots\mathbf{b}_k(+)\right), \quad V_k(-)=\left(\mathbf{b}_1(-)\dots\mathbf{b}_k(-)\right)$$

be matrices in $\mathbb{R}^{2n \times k}$ where the columns are the symmetric and antisymmetric expansion vectors, respectively. Let the columns of the matrices $V_k(\pm)$ form an orthonormal basis for the subspaces \mathcal{V}_k^\pm . Let us also introduce the spaces

$$\begin{aligned} \mathcal{L}\mathcal{V}_k(\pm) &= \text{span}\{\Lambda \mathbf{b}_1(\pm), \dots, \Lambda \mathbf{b}_k(\pm)\} \\ &= \text{span}\{\boldsymbol{\sigma}_1(\pm), \dots, \boldsymbol{\sigma}_k(\pm)\} \end{aligned}$$

and

$$\begin{aligned} \mathcal{B}\mathcal{V}_k(\pm) &= \text{span}\{\Omega \mathbf{b}_1(\mp), \dots, \Omega \mathbf{b}_k(\mp)\} \\ &= \text{span}\{\boldsymbol{\tau}_1(\pm), \dots, \boldsymbol{\tau}_k(\pm)\} \end{aligned}$$

that collect the applications of the matrices Λ and Ω to the expansion vectors. We remark that, in practice, only the vectors $\mathbf{b}^\pm \in \mathbb{R}^n$ are stored in memory. Specifically, we consider the $n \times k$ matrices $V_k^\pm = (\mathbf{b}_1^\pm \dots \mathbf{b}_k^\pm)$, the spaces $\mathcal{L}\mathcal{V}_k^\pm = \text{span}\{\boldsymbol{\sigma}_1^\pm, \dots, \boldsymbol{\sigma}_k^\pm\}$ and $\mathcal{B}\mathcal{V}_k^\pm = \text{span}\{\boldsymbol{\tau}_1^\pm, \dots, \boldsymbol{\tau}_k^\pm\}$, and the corresponding matrices $LV_k^\pm = (\boldsymbol{\sigma}_1^\pm \dots \boldsymbol{\sigma}_k^\pm)$ and $BV_k^\pm = (\boldsymbol{\tau}_1^\pm \dots \boldsymbol{\tau}_k^\pm)$.

The Rayleigh–Ritz variational procedure is then performed to compute the u^\pm coefficients from eq 3 by solving the projection of eq 1 into the subspace union of V_k^+ and V_k^- . We get the $2k \times 2k$ problem

$$\begin{pmatrix} E^+ & 0 \\ 0 & E^- \end{pmatrix} \begin{pmatrix} \mathbf{u}^+ \\ \mathbf{u}^- \end{pmatrix} = \omega_{k+1} \begin{pmatrix} 0 & S^T \\ S & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u}^+ \\ \mathbf{u}^- \end{pmatrix} \quad (7)$$

where

$$E^\pm = (V_k^\pm)^T L V_k^\pm, \quad S = (V_k^-)^T B V_k^+ \quad (8)$$

that is, for $i = 1, \dots, k$ and $j = 1, \dots, i$,

$$\begin{aligned} E_{ij}^\pm &= \langle \mathbf{b}_i^\pm, (A \pm B)\mathbf{b}_j^\pm \rangle \\ S_{ij} &= \langle \mathbf{b}_i^-, (\Sigma + \Delta)\mathbf{b}_j^+ \rangle \end{aligned}$$

We note here that eq 7 has the same structure of the original response equation, with a nonpositive definite metric. However, assuming that $\omega \neq 0$, it is possible to divide both sides of the reduced equation by ω obtaining

$$\begin{pmatrix} 0 & S^T \\ S & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u}^+ \\ \mathbf{u}^- \end{pmatrix} = \frac{1}{\omega_{k+1}} \begin{pmatrix} E^+ & 0 \\ 0 & E^- \end{pmatrix} \begin{pmatrix} \mathbf{u}^+ \\ \mathbf{u}^- \end{pmatrix} \quad (9)$$

where now the reduced response matrix, which is symmetric and positive definite, acts as a metric, making it possible to use standard linear algebra techniques.

Once the eigenvalue ω_{k+1} is computed in the reduced space and the current approximation to the eigenvector \mathbf{x}_{k+1} has been computed as in eq 3,

$$\mathbf{x}_{k+1} = V_k^+(\mathbf{u}^+) + V_k^-(\mathbf{u}^-), \quad \mathbf{u}^\pm \in \mathbb{R}^k \quad (10)$$

we can assemble the residual vector R_{k+1} in the full space \mathbb{R}^{2n}

$$R_{k+1} = (\Lambda - \omega_{k+1}\Omega)\mathbf{x}_{k+1} \quad (11)$$

which emerges naturally as the sum of a symmetric term $R_{k+1}(+)$ and an antisymmetric term $R_{k+1}(-)$, whose components are expressed in terms of elements of the spaces LV_k^\pm and BV_k^\pm , i.e.,

$$R_{k+1}^\pm = LV_k^\pm \mathbf{u}^\pm - \omega_{k+1} B V_k^\pm \mathbf{u}^\mp \quad (12)$$

The expansion subspaces are extended, in the next iteration, using the standard Davidson algorithm, i.e., by solving

$$(D_\Lambda - \omega_{k+1} D_\Omega) \mathbf{b}_{k+1} = -R_{k+1} \quad (13)$$

where D_Λ and D_Ω are diagonal matrices whose elements are the diagonal elements of Λ and Ω , respectively, and $\mathbf{b}_{k+1} = \mathbf{b}_{k+1}(+) + \mathbf{b}_{k+1}(-)$. Splitting eq 13 into a symmetric and an antisymmetric part we get, through simple algebra,

$$\tilde{\mathbf{b}}_{k+1}^\pm = -(D_A^2 - \omega_{k+1}^2 D_\Sigma^2)^{-1} (D_A R_{k+1}^\pm + \omega_{k+1} D_\Sigma R_{k+1}^\mp) \quad (14)$$

where D_A and D_Σ are the diagonals of A and Σ , respectively. The $\tilde{\mathbf{b}}_{k+1}^\pm$ vectors are then orthogonalized to V_k^\pm , respectively, and the procedure is iterated until convergence is reached. Each iteration requires four matrix-vector products (MVP), to compute $\boldsymbol{\sigma}_k^+$, $\boldsymbol{\sigma}_k^-$, $\boldsymbol{\tau}_k^+$, and $\boldsymbol{\tau}_k^-$, where the first two MVP are typically the cost-dominating step (see eqs 5 and 6). However, the cost of solving the generalized eigenvalue problem (eq 7) and the cost of orthogonalizing the new expansion vectors can also become computational bottlenecks, especially if a large number n_{eig} of eigenvalues and eigenvectors are sought. To

lower the cost of the former operation, Helmich-Paris²⁰ recently proposed a method to reduce the generalized eigenvalue problem into a half-dimensional problem, i.e., the size of one of the blocks E^\pm and of S . The proposed algorithm scales with $O(m^3)$, where $m = k \cdot n_{\text{eig}}$ is the size of the subspace, which is an important improvement with respect to the $O((2m)^3)$ solution to eq 7. However, it requires a sizable number of linear algebra operations, including two singular value decompositions (SVD), which can become rather expensive.

In this paper, we propose a different approach that results in a much simpler implementation and that performs, in a m -dimensional space, only one diagonalization and one matrix–matrix multiplication. Here, we only consider the case where a single eigenpair is searched, i.e., $m = k$.

We start by noting that, if the ground-state solution is stable, problem (2) has no vanishing eigenvalue, and, therefore, it is possible to rewrite it as²²

$$\Omega \mathbf{x} = \lambda \Lambda \mathbf{x}, \quad \lambda = \frac{1}{\omega} \quad (15)$$

Problem (15) is much easier to solve than problem (2), because it is a generalized eigenvalue problem with a symmetric, positive definite metric Λ . It is therefore possible to introduce a positive definite dot product

$$\langle \mathbf{x}, \mathbf{y} \rangle_\Lambda = \mathbf{x}^\top \Lambda \mathbf{y} \quad (16)$$

and the induced Λ -norm $\|\mathbf{x}\|_\Lambda = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle_\Lambda}$. We proceed similarly to what has already been presented and expand the eigenvector \mathbf{x} as the linear combination of symmetric and antisymmetric expansion vectors as in eq 3, where, in our approach, we choose the expansion vectors such that

$$\langle \mathbf{b}_i^+, \mathbf{b}_j^+ \rangle_\Lambda = \delta_{ij}, \quad \langle \mathbf{b}_i^-, \mathbf{b}_j^- \rangle_\Lambda = \delta_{ij} \quad (17)$$

that is, the expansion vectors are Λ -orthogonal.

With the definition of S given in eq 8, we obtain the following reduced problem:

$$\begin{pmatrix} 0 & S^\top \\ S & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u}^+ \\ \mathbf{u}^- \end{pmatrix} = \lambda_{k+1} \begin{pmatrix} \mathbf{u}^+ \\ \mathbf{u}^- \end{pmatrix} \quad (18)$$

which, due to the choice of Λ -orthogonal expansion vectors, is no longer a generalized eigenvalue problem. Furthermore, its block structure allows us to solve it by first computing \mathbf{u}^+ from

$$S^\top S \mathbf{u}^+ = \lambda^2 \mathbf{u}^+ \quad (19)$$

a symmetric and positive definite eigenvalue problem of size m , and then by recovering \mathbf{u}^- as

$$\mathbf{u}^- = \frac{1}{\lambda} S \mathbf{u}^+ \quad (20)$$

After assembling the current approximation to the eigenvector, we compute the residual

$$R_{k+1} = (\Omega - \lambda_{k+1} \Lambda) \mathbf{x}_{k+1} \quad (21)$$

which can be split into a symmetric and an antisymmetric term, whose components are

$$R_{k+1}^\pm = B V_k^\pm \mathbf{u}^\mp - \lambda_{k+1} L V_k^\pm \mathbf{u}^\pm \quad (22)$$

The new pair of expansion vectors are then obtained by solving

$$(D_\Omega - \lambda_{k+1} D_\Lambda) \mathbf{b}_{k+1} = -R_{k+1} \quad (23)$$

Splitting the new vector and the residual in the sum of symmetric and antisymmetric components, after some computation, we get

$$\tilde{\mathbf{b}}_{k+1}^\pm = -(\lambda_{k+1}^2 D_\Lambda^2 - D_\Sigma^2)^{-1} (\lambda_{k+1} D_A R_{k+1}^\pm + D_\Sigma R_{k+1}^\mp) \quad (24)$$

Once the $\tilde{\mathbf{b}}_{k+1}^\pm$ vectors have been computed, we enlarge the expansion subspaces \mathcal{V}_k^\pm by determining vectors that are orthogonal to V_k^\pm . To this end, the new vectors $\tilde{\mathbf{b}}_{k+1}^\pm$ are first orthogonalized to V_k^\pm and then orthonormalized. This is done as follows: first, we project out V_k^\pm from the vectors

$$\tilde{\mathbf{b}}_{k+1}^\pm - V_k^\pm (V_k^\pm)^\top \tilde{\mathbf{b}}_{k+1}^\pm \quad (25)$$

and then we orthonormalize the resulting vectors using the Cholesky decomposition of their overlap, as described in a recent paper by some of us.²³ To ensure the numerical stability of this procedure, these two steps are iterated until the norm of $(V_k^\pm)^\top \tilde{\mathbf{b}}_{k+1}^\pm$ is smaller than a (tight) threshold. In all of our numerical experiments, two iterations were sufficient to obtain vectors orthogonal to machine precision. The resulting vectors $\hat{\mathbf{b}}_{k+1}^\pm$ are then orthogonalized with respect to Λ by computing the Cholesky decomposition of

$$M^\pm = (\hat{\mathbf{b}}_{k+1}^\pm)^\top \sigma_{k+1}^\pm = L^\pm (L^\pm)^\top \quad (26)$$

and then solving the triangular linear system

$$\mathbf{b}_{k+1}^\pm (L^\pm)^\top = \hat{\mathbf{b}}_{k+1}^\pm \quad (27)$$

We remark that, if we are seeking one eigenpair, eq 26 is a trivial case, since M^\pm represents a scalar and L^\pm is the square root of M^\pm . This step becomes significant when we want to compute multiple eigenvalues–eigenvectors. Performing the Λ -orthogonalization after regular orthogonalization may look unnecessary, but it vastly improves the conditioning of M , thus ensuring the overall numerical stability and robustness of the procedure. Summarizing, first, we compute $\tilde{\mathbf{b}}_{k+1}^\pm$; second, we compute $\hat{\mathbf{b}}_{k+1}^\pm$, and finally the expansion vectors \mathbf{b}_{k+1}^\pm , which we store on the matrices $V_{k+1}^\pm = (V_k^\pm \mathbf{b}_{k+1}^\pm)$. Note that, to perform the latter passage, we must apply $A + B$ and $A - B$ to the new test vectors. In other words, the number of MVP in our algorithm is the same as in the traditional ones, but the linear algebra operations are both simpler and less expensive. In particular, the only $O(m^3)$ operations that we need to perform are the matrix–matrix multiplication $S^\top S$ and the diagonalization in eq 19. The price to pay is that we need to perform two more orthogonalizations (eqs 26 and 27); however, these have a cost of order $O(2n n_{\text{eig}}^2 + n_{\text{eig}}^3)$, n_{eig} being the number of sought eigenvectors, and should therefore be less critical.

The overall procedure, which we named Swapped-Metric-Orthogonal Generalized Davidson (SMO-GD), is summarized in Algorithm 1. The algorithms of the primitives used for the orthogonalizations are given in the Supporting Information. An open-source (LGPL3), Fortran implementation of the algorithm, including the orthogonalization primitives, can be found in the DIAGLIB library,²³ which is available on GITHUB at <https://github.com/Molecolab-Pisa/diaglib>. All of the code used to generate the numerical examples presented in the following can be accessed by downloading the code.

Algorithm 1 is general and can deal with any generalized eigenvalue problem as in eq 1. It exhibits monotonic convergence of the reduced eigenvalues, as the Hylleras–

Algorithm 1 SWAPPED-METRIC-ORTHOGONAL GENERALIZED DAVIDSON

Input: Initial guess $\mathbf{x}_0 = (y_0, \mathbf{z}_0)^T$, procedures to apply $A + B$, $A - B$, $\Sigma + \Delta$, $\Sigma - \Delta$ to a given vector, procedure to apply the preconditioner in eq. (24) to a vector, tolerance τ , maximum size of the subspaces m_{\max} , maximum number of iterations k_{\max} . Primitives $\mathbf{b}_{\text{ortho}}$ and $\mathbf{b}_{\text{ortho_vs_x}}$
Output: \mathbf{x} , λ , the lowest pencils of the generalized eigenvalue problem $\Omega\mathbf{x} = \lambda\Lambda\mathbf{x}$.

```

1: Assemble  $\hat{\mathbf{b}}_1^\pm = \mathbf{y}_0 \pm \mathbf{z}_0$ 
2:  $\hat{\boldsymbol{\sigma}}_1^\pm = (A \pm B)\hat{\mathbf{b}}_1^\pm$ 
3:  $\mathbf{b}_1^\pm, \boldsymbol{\sigma}_1^\pm = \mathbf{b}_{\text{ortho}}(\hat{\mathbf{b}}_1^\pm, \hat{\boldsymbol{\sigma}}_1^\pm)$ 
4:  $k = 0, m = 0$ 
5: Set  $V_1^\pm = \mathbf{b}_1^\pm, LV_1^\pm = \boldsymbol{\sigma}_1^\pm, BV_0^\pm = \emptyset$ 
6: while  $k < k_{\max}$  do
7:    $k = k + 1, m = m + 1$ 
8:    $\boldsymbol{\tau}_k^\pm = (\Sigma \pm \Delta)\mathbf{b}_k^\pm$ 
9:   Set  $BV_k^\pm = (BV_{k-1}^\pm, \boldsymbol{\tau}_k^\pm)$ 
10:   $S = (V_k^-)^T BV_k^-$ 
11:  Solve  $S^T S \mathbf{u}^\pm = \lambda^2 \mathbf{u}^\pm$  and get  $\lambda_{k+1} = \sqrt{\lambda^2}$ 
12:  Compute  $\mathbf{u}^\pm = \frac{1}{\lambda} S \mathbf{u}^\pm$ 
13:  Build the current approximation to the eigenvector
14:   $R_{k+1}^\pm = BV_k^\pm \mathbf{u}^\pm - \lambda_{k+1} LV_k^\pm \mathbf{u}^\pm$ 
15:  Lock converged eigenvectors, if done, return
16:  if  $m < m_{\max}$  then
17:     $\hat{\mathbf{b}}_{k+1}^\pm = \text{precd}(\lambda_{k+1}, R_{k+1}^\pm)$ 
18:     $\hat{\mathbf{b}}_{k+1}^\pm = \mathbf{b}_{\text{ortho\_vs\_x}}(\hat{\mathbf{b}}_{k+1}^\pm, V_k^\pm)$ 
19:     $\hat{\boldsymbol{\sigma}}_{k+1}^\pm = (A \pm B)\hat{\mathbf{b}}_{k+1}^\pm$ 
20:     $\mathbf{b}_{k+1}^\pm, \boldsymbol{\sigma}_{k+1}^\pm = \mathbf{b}_{\text{ortho}}(\hat{\mathbf{b}}_{k+1}^\pm, \hat{\boldsymbol{\sigma}}_{k+1}^\pm)$ 
21:    Set  $V_{k+1}^\pm = (V_k^\pm, \mathbf{b}_{k+1}^\pm), LV_{k+1}^\pm = (LV_k^\pm, \boldsymbol{\sigma}_{k+1}^\pm)$ 
22:  else
23:     $m = 1$ 
24:    Restart Davidson
25:  end if
26: end while

```

Undheim–McDonald theorem applies,^{24,25} and can be implemented efficiently using BLAS and LAPACK routines.

In the following, we report a few numerical tests of the proposed algorithm and compare it with the original algorithm proposed by Olsen et al.,¹⁷ where the $2m \times 2m$ reduced problem is solved, and the algorithm recently proposed by Helmich-Paris,²⁰ which uses a series of transformation to reduce the size of the problem to $m \times m$. To compare the three algorithms in a fair way, they all have been implemented in the DIAGLIB library (RESPONSE branch). As our focus is on the iterative algorithm and not on the implementation of MCSCF (or TD-DFT), we generate test matrices with the correct structure and use in-core matrix–vector multiplications to compute the required matrix–vector products. This allows us not only to validate our implementation against dense LAPACK routines but also to create common grounds to compare the three algorithms in a fair way. We focus our analysis on the time required to solve the reduced problem and to orthogonalize the new test vectors against the previous ones, as all of the other steps are identical. We generate symmetric and positive definite $A + B$ and $A - B$ matrices by putting

$$(A + B)_{ij} = \begin{cases} S + i & \text{if } i = j \\ \frac{1}{i + j} & \text{otherwise} \end{cases}$$

and

$$(A - B)_{ij} = \begin{cases} 2 + i & \text{if } i = j \\ \frac{0.2}{i + j} & \text{otherwise} \end{cases}$$

A and B are then obtained as linear combinations. The symmetric Σ and the antisymmetric Δ matrices are generated in the following way: the former is obtained by generating a random matrix and then multiplying it by its transpose, to ensure that it is symmetric and positive-definite, while the

latter antisymmetric matrix is obtained by generating a random matrix and subtracting its transpose.

As a first test example, to investigate the elapsed diagonalization times, we consider a generalized eigenvalue problem of size $2n \times 2n$, where $n = 10\,000$ is the size of the matrices A , B , Σ , and Δ . We solved the generalized eigenvalue problem by applying the three algorithms discussed in this work: the original procedure proposed by Olsen and co-workers (LR_{std}),¹⁷ where the reduced problem is solved for $1/\omega$ to have a symmetric, positive-definite metric, the algorithm recently introduced by Helmich-Paris (LR_{HP}),²⁰ and our implementation shown in Algorithm 1 (SMO-GD). We seek 10–100 eigenpairs with increments of 10. Convergence is achieved when the root-mean-square norm of the residual is smaller than 10^{-6} and its maximum absolute value is smaller than 10^{-5} . We use a subspace dimension of 20, i.e., we keep in memory up to 20 vectors per eigenvector in the history. For all of the algorithms, multiple eigenvectors are expanded simultaneously in the same expansion space, and we exploit a locking procedure for the converged eigenvectors. In Figure 1, we report the cumulative elapsed time required to solve the

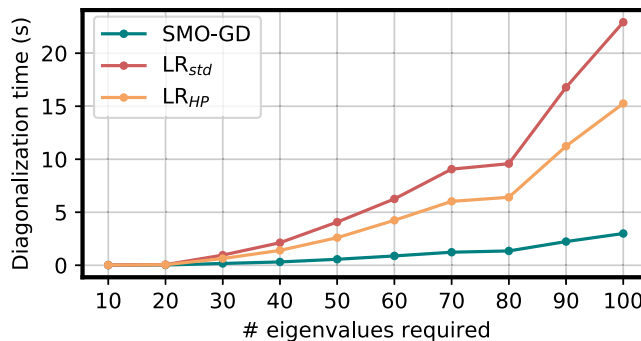


Figure 1. Time analysis of the three methods discussed for a $2n \times 2n$ problem, where $n = 10\,000$. Diagonalization time was calculated with respect to the number of eigenvalues required.

reduced eigenvalue problem for the three algorithms. Such timings are the sum over all the iterations; note that by design, the iterations produced by all algorithms are equivalent, and therefore the number of iterations is shared by all methods. For the algorithm presented here (SMO-GD), the reported timings include the time spent for the additional orthogonalization with respect to the metric (line 20 of Algorithm 1). We observe that the proposed algorithm significantly outperforms both the original Olsen’s algorithm and the new method presented by Helmich-Paris. The differences become more significant as the size of the reduced problem increases, i.e., when more eigenvalues are sought. Note that in our simple-minded example, the overall cost is still dominated by the dense matrix–vector multiplications. Nevertheless, the difference in timings is sizable and can be expected to have an impact on large-scale applications when many states are computed. To better appreciate the overall difference between the algorithms, we plot, in Figure 2, the total time required to compute 100 eigenvalues and eigenvectors for increasingly larger systems with n ranging from 1000 to 10 000. In all cases, our SMO-GD algorithm is faster than the others, and in particular, the time difference increases as the dimension decreases, consistently with what was observed in Figure 1.

These results can be clearly rationalized by looking at the dense linear algebra operations performed by the various

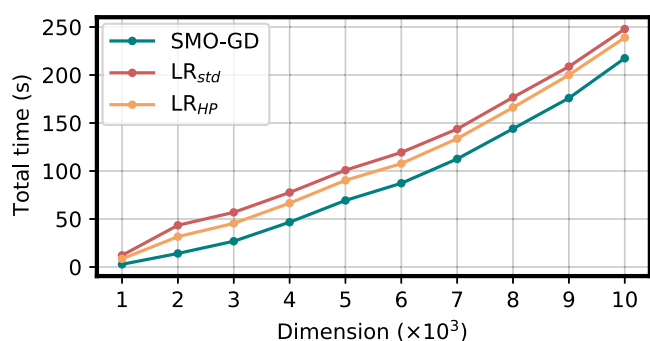


Figure 2. Total time (s) of the three methods for different dimensions n and 100 eigenvectors required.

algorithms. Focusing on the leading $O(m^3)$ operations, the original algorithm by Olsen et al. requires the solution to a generalized eigenvalue problem of size $2m$, which is performed in our implementation using the DSYGV LAPACK routine, which performs a Cholesky decomposition of the metric to transform the problem into a standard eigenvalue problem and then solves the latter using the DSYEV LAPACK routine. Both operations require $O((2m)^3)$ floating-point operations. The algorithm recently presented by Helmich-Paris²⁰ avoids $O((2m)^3)$ by reducing the subspace problem to a m -sized one, by performing two singular value decompositions (LAPACK routine DGESVD), two Cholesky factorizations (LAPACK routine DPOTRF), and eight matrix–matrix multiplications. In contrast, our algorithm requires only one matrix–matrix multiplication (to assemble $S^T S$ and one symmetric diagonalization (LAPACK routine DSYEV), resulting in a significantly reduced computational cost. The two additional Λ -orthonormalizations require only $O(n_{\text{eig}}^3)$ operations, which is by definition smaller than m .

Table 1. Number and Type of $O(m^3)$ Dense Linear Algebra Operations Required by the Algorithm SMO-GD, Compared with the Strattmann-Scuseria-Frisch (SSF) and Helmich-Paris (HP) Ones^a

	SMO-GD	SSF	HP
EV	1	2	0
MM	1	5	8
SVD	0	0	2
CD	0	0	2

^aHere, m is the size of the expansion subspace. EV is the symmetric diagonalization, MM the matrix–matrix multiplication, SVD the singular value decomposition, and CD the Cholesky factorization.

We conclude these remarks by comparing our algorithm to the method proposed by Strattmann, Scuseria, and Frisch (SSF)⁴ for the specific case of LR DFT. The SSF algorithm solves the non-Hermitian problem

$$(A - B)(A + B)(\mathbf{y} + \mathbf{z}) = \omega^2(\mathbf{y} + \mathbf{z}) \quad (28)$$

introducing a nonfaithful representation of eq 28 in the expansion subspace

$$E^- E^+ \mathbf{u}^+ = \omega^2 \mathbf{u}^+ \quad (29)$$

In their work, they show that the two problems become equivalent when convergence is achieved. Again, the problem

in eq 29 is m -sized and can be solved by transforming it to the symmetric eigenvalue

$$(E^-)^{1/2} E^+ (E^-)^{1/2} \mathbf{u}' = \omega^2 \mathbf{u}' \quad (30)$$

where $\mathbf{u}' = (E^-)^{-1/2} \mathbf{u}^+$. In practice, the implementation requires diagonalization of E^- to compute $(E^-)^{1/2}$ (one symmetric diagonalization and one matrix–matrix multiplication), which is then used to assemble the symmetric matrix in eq 30 (two matrix–matrix multiplications), which is then diagonalized (a second symmetric diagonalization). Two further matrix–matrix multiplications are used to recover the \mathbf{u}^+ and \mathbf{u}^- eigenvectors, the latter being the left eigenvectors to eq 29. The left and right eigenvectors $\mathbf{y} + \mathbf{z}$ and $\mathbf{y} - \mathbf{z}$ are then biorthogonalized. A summary of all the dense linear algebra operations performed by the SSF, HP, and SMO-GD algorithms is reported in Table 1. The SSF algorithm generates a different expansion space, with respect to the Olsen algorithm (which generates the same expansion space as the HP algorithm and the one presented in this communication), making a one-to-one comparison somewhat harder, since, in general, the number of iterations may be different. However, results presented by Strattmann, Scuseria, and Frisch in their paper⁴ show that their procedure generates subspaces that are of the same size as the ones generated by the Olsen algorithm, making thus a qualitative comparison possible. We have thus employed both algorithms to solve TD-DFT like equations, obtained by setting $\Sigma = 1$ and $\Delta = 0$ in eq 1. We use the same A and B matrices used in previous tests; however, as these are very well-behaved, starting the iterations from an optimal guess (the canonical basis vectors corresponding to the lowest diagonal elements of $A + B$) results in almost immediate convergence of both algorithms and offers little material to compare results. To offer a more realistic comparison, we perturb the optimal guess by adding a random vector to each guess vector, the elements of which are random numbers chosen in the interval $[0, 0.01]$. The results are reported in Figures 3 and 4. As can be seen from the figure, SMO-GD

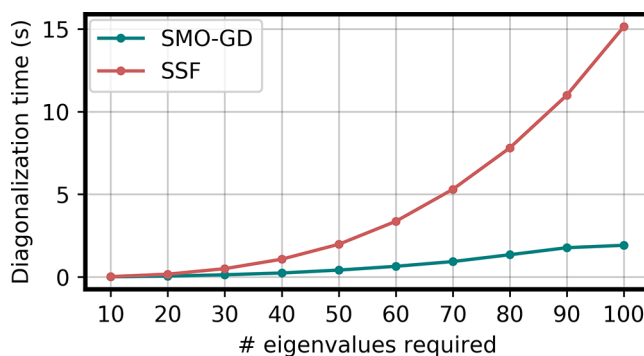


Figure 3. Comparison of the SMO-GD and SSF algorithms for the solution of a $2n \times 2n$ TD-DFT like generalized eigenvalue equations, with $n = 10\,000$ as a function of the number of required eigenvectors. The total time for the dense linear algebra operations required to solve the problem in the subspace is reported for both methods.

outperforms SSF in this well-behaved case. The better performances of SMO-GD are due not only to the smaller number of dense linear algebra operations but also to its faster convergence. All the calculations reported in Figures 3 and 4 required in fact 10 SMO-GD and 9 SSF iterations to achieve convergence. As the SSF algorithm performs twice as many




Figure 4. Comparison of the SMO-GD and SSF algorithms for the solution of a $2n \times 2n$ TD-DFT like generalized eigenvalue equations, with n ranging from 1000 to 10 000, when 100 eigenvalues are sought. The total execution time is reported.

matrix-vector multiplications than SMO-GD per iteration, the total number of matrix–vector multiplications is smaller in SMO-GD, which explains the sizable difference in the total elapsed time. We note, in passing, that we have tested the algorithms using randomly generated matrices, where the diagonal was then shifted to make the problem more diagonally dominant and could not achieve convergence using the SSF implementation. On the other hand, SMO-GD always managed to converge. Similar results were obtained by using a very tight (10^{-10} RMS of the residual) convergence threshold. In the latter case, SSF approached convergence and then started to stagnate, while SMO-GD exhibited no problem. Both cases as summarized in the [Supporting Information](#). While neither describes a particularly realistic application scenario, as the response equations in TD-DFT are typically strongly diagonally dominant and thus well-conditioned, we believe that such tests further confirm the good numerical stability of the SMO-GD algorithm.

In conclusion, we have presented a new algorithm to solve the MCSCF linear response equations and also the related TD-SCF equations that is not only more efficient than what was previously reported in the literature, but also conceptually simple and easy to implement. Thanks to the robust orthogonalization procedures described in the [Supporting Information](#), it is also numerically robust and stable. If the expansion space becomes ill-conditioned, which is bound to happen near convergence, then the metric in the reduced space can exhibit small (i.e., numerically zero) or even negative eigenvalues, independent of whether the actual metric is ill-conditioned. This can make the overall procedure fail. By choosing expansion vectors that are orthogonal with respect to the scalar product induced by the metric, we avoided this problem from the beginning. Furthermore, by limiting the linear algebra operations to a symmetric diagonalization in the subspace, we avoid further propagation of possible numerical instabilities. The combination of robustness and efficiency makes the new algorithm therefore an ideal strategy to tackle the solution to the linear response equations.

■ ASSOCIATED CONTENT

SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.3c00989>.

Pseudocode of the orthogonalization primitives and additional tests and results ([PDF](#))

■ AUTHOR INFORMATION

Corresponding Authors

Tommaso Nottoli – Dipartimento di Chimica e Chimica Industriale, Università di Pisa, 56124 Pisa, Italy;
Email: tommaso.nottoli@phd.unipi.it

Filippo Lipparini – Dipartimento di Chimica e Chimica Industriale, Università di Pisa, 56124 Pisa, Italy;
orcid.org/0000-0002-4947-3912;
Email: filippo.lipparini@unipi.it

Authors

Riccardo Alessandro – Dipartimento di Chimica e Chimica Industriale, Università di Pisa, 56124 Pisa, Italy

Ivan Gianni – Dipartimento di Chimica e Chimica Industriale, Università di Pisa, 56124 Pisa, Italy

Federica Pes – Dipartimento di Chimica e Chimica Industriale, Università di Pisa, 56124 Pisa, Italy;
orcid.org/0000-0001-9064-7876

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acs.jctc.3c00989>

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

The authors acknowledge financial support from ICSC-Centro Nazionale di Ricerca in High Performance Computing, Big Data, and Quantum Computing, funded by the European Union – Next Generation EU – PNRR, Missione 4 Componente 2 Investimento 1.4. F.L. and F.P. further acknowledge funding from the Italian Ministry of Research, under Grant No. 2020HTSXMA_002 (PSI-MOVIE). F.P. is a member of the GNCS group of INdAM.

■ REFERENCES

- (1) Norman, P.; Ruud, K.; Saue, T. *Principles and Practices of Molecular Properties: Theory, Modeling, And Simulations*; John Wiley & Sons, 2018.
- (2) Bauernschmitt, R.; Ahlrichs, R. Treatment of electronic excitations within the adiabatic approximation of time dependent density functional theory. *Chem. Phys. Lett.* **1996**, *256*, 454–464.
- (3) Casida, M. E.; Jamorski, C.; Casida, K. C.; Salahub, D. R. Molecular excitation energies to high-lying bound states from time-dependent density-functional response theory: Characterization and correction of the time-dependent local density approximation ionization threshold. *J. Chem. Phys.* **1998**, *108*, 4439–4449.
- (4) Stratmann, R. E.; Scuseria, G. E.; Frisch, M. J. An efficient implementation of time-dependent density-functional theory for the calculation of excitation energies of large molecules. *J. Chem. Phys.* **1998**, *109*, 8218–8224.
- (5) Monkhorst, H. J. Calculation of properties with the coupled-cluster method. *Int. J. Quantum Chem.* **1977**, *12*, 421–432.
- (6) Sekino, H.; Bartlett, R. J. A linear response, coupled-cluster theory for excitation energy. *Int. J. Quantum Chem.* **1984**, *26*, 255–265.
- (7) Christiansen, O.; Jørgensen, P.; Hättig, C. Response functions from Fourier component variational perturbation theory applied to a time-averaged quasienergy. *Int. J. Quantum Chem.* **1998**, *68*, 1–52.
- (8) Koch, H.; Jørgensen, P. Coupled cluster response functions. *J. Chem. Phys.* **1990**, *93*, 3333–3344.
- (9) Yeager, D. L.; Jørgensen, P. A multiconfigurational time-dependent Hartree-Fock approach. *Chem. Phys. Lett.* **1979**, *65*, 77–80.
- (10) Yeager, D. L.; Olsen, J.; Jørgensen, P. The evaluation of MCRPA (MCTDHF) electronic excitation energies, oscillator

strengths, and polarizabilities: Application to O₂. *Int. J. Quantum Chem.* **1981**, *20*, 151–162.

(11) Sasagane, K.; Mori, K.; Ichihara, A.; Itoh, R. The multi-configuration time-dependent Hartree–Fock method based on a closed-shell-type multiconfiguration self-consistent field reference state and its application to the LiH molecule. *J. Chem. Phys.* **1990**, *92*, 3619–3632.

(12) Helgaker, T.; Coriani, S.; Jørgensen, P.; Kristensen, K.; Olsen, J.; Ruud, K. Recent advances in wave function-based methods of molecular-property calculations. *Chem. Rev.* **2012**, *112*, 543–631.

(13) Norman, P. A perspective on nonresonant and resonant electronic response theory for time-dependent molecular properties. *Phys. Chem. Chem. Phys.* **2011**, *13*, 20519–20535.

(14) Bai, Z.; Li, R.-C. Minimization principles for the linear response eigenvalue problem I: Theory. *SIAM J. Matrix Anal. Appl.* **2012**, *33*, 1075–1100.

(15) Bai, Z.; Li, R.-C. Minimization principles for the linear response eigenvalue problem II: Computation. *SIAM J. Matrix Anal. Appl.* **2013**, *34*, 392–416.

(16) Bai, Z.; Li, R.-C. Minimization principles and computation for the generalized linear response eigenvalue problem. *BIT* **2014**, *54*, 31–54.

(17) Olsen, J.; Jensen, H. J. A.; Jørgensen, P. Solution of the large matrix equations which occur in response theory. *J. Comput. Phys.* **1988**, *74*, 265–282.

(18) Jørgensen, P.; Jensen, H. J. A.; Olsen, J. Linear response calculations for large scale multiconfiguration self-consistent field wave functions. *J. Chem. Phys.* **1988**, *89*, 3654–3661.

(19) Saue, T.; Jensen, H. A. Linear response at the 4-component relativistic level: Application to the frequency-dependent dipole polarizabilities of the coinage metal dimers. *J. Chem. Phys.* **2003**, *118*, 522–536.

(20) Helmich-Paris, B. CASSCF linear response calculations for large open-shell molecules. *J. Chem. Phys.* **2019**, *150*, 174121.

(21) Davidson, E. R. The Iterative Calculation of a Few of the Lowest Eigenvalues and Corresponding Eigenvectors of Large Real-Symmetric Matrices. *J. Comput. Phys.* **1975**, *17*, 87–94.

(22) Olsen, J.; Jørgensen, P. Linear and nonlinear response functions for an exact state and for an MCSCF state. *J. Chem. Phys.* **1985**, *82*, 3235–3264.

(23) Nottoli, T.; Gianni, I.; Levitt, A.; Lipparini, F. A robust, open-source implementation of the locally optimal block preconditioned conjugate gradient for large eigenvalue problems in quantum chemistry. *Theor. Chem. Acc.* **2023**, *142*, 1–13.

(24) Hylleraas, E. A.; Undheim, B. Numerische berechnung der 2 S-terme von ortho-und par-helium. *Z. Phys.* **1930**, *65*, 759–772.

(25) MacDonald, J. Successive approximations by the Rayleigh-Ritz variation method. *Phys. Rev.* **1933**, *43*, 830.