

PAPER • OPEN ACCESS

Quantum eigenvector continuation for chemistry applications

To cite this article: Carlos Mejuto-Zaera and Alexander F Kemper 2023 *Electron. Struct.* 5 045007

View the [article online](#) for updates and enhancements.

You may also like

- [Reactions of electronically excited molecular nitrogen with H₂ and H₂O molecules: theoretical study](#)
Alexey V Pelevkin and Alexander S Sharipov
- [Resolving rainbows with superimposed diffraction oscillations in NO + rare gas scattering: experiment and theory](#)
Jolijn Onvlee, Sjoerd N Vogels, Ad van der Avoird et al.
- [Sub-laser-cycle control of coupled electron–nuclear dynamics at a conical intersection](#)
Maria Richter, Foudhil Bouakline, Jesús González-Vázquez et al.

Electronic Structure

**OPEN ACCESS****PAPER**

Quantum eigenvector continuation for chemistry applications

Carlos Mejuto-Zaera^{1,*} and Alexander F Kemper² **RECEIVED**
28 April 2023**REVISED**
1 October 2023**ACCEPTED FOR PUBLICATION**
9 October 2023**PUBLISHED**
10 November 2023¹ Scuola Internazionale Superiore di Studi Avanzati (SISSA), Trieste, TS, Italy² Department of Physics, North Carolina State University, Raleigh, NC 27695, United States of America

* Author to whom any correspondence should be addressed.

E-mail: cmejutoz@sissa.it**Keywords:** subspace expansion, quantum computing, quantum chemistry, potential energy surface, binding energy curve

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.

**Abstract**

A typical task for classical and quantum computing in chemistry is finding a potential energy surface (PES) along a reaction coordinate, which involves solving the quantum chemistry problem for many points along the reaction path. Developing algorithms to accomplish this task on quantum computers has been an active area of development, yet finding all the relevant eigenstates along the reaction coordinate remains a difficult problem, and determining PESs is thus a costly proposal. In this paper, we demonstrate the use of a eigenvector continuation—a subspace expansion that uses a few eigenstates as a basis—as a tool for rapidly exploring PESs. We apply this to determining the binding PES or torsion PES for several molecules of varying complexity. In all cases, we show that the PES can be captured using relatively few basis states; suggesting that a significant amount of (quantum) computational effort can be saved by making use of already calculated ground states in this manner.

Glossary

We summarize the main choices of notation and nomenclature used throughout the paper in the following list.

Point in parameter space λ, ℓ	A vector containing all values defining a single point in the parameter space of the system studied. In this paper, these correspond to all the relative atomic coordinates describing the molecular geometry. In a spin Hamiltonian, these would correspond to the spin–spin couplings.
Local atomic orbitals (AOs) $\{\phi\}_a$	The basis for the quantum chemistry problem at a particular atomic configuration λ . These are typically not orthogonal.
Fock operator F_a	Fock operator in the basis local AOs at a particular λ_a . Plays the role of the Hamiltonian in the non-linear, generalized eigenvalue problem of the Hartree–Fock approximation.
Local overlap S_a	Overlap integrals of a single local AO basis set at a particular λ_a .
Local molecular orbitals (MOs)	The (orthogonal) MOs found by solving the Hartree–Fock generalized eigenvalue equation.
Local orbital rotation matrix U	The unitary rotation matrix that diagonalizes the Hartree–Fock Hamiltonian, and rotates from AOs to MOs.
FCI Hamiltonian H_a^{FCI}	Full configuration interaction Hamiltonian in the basis of MOs at a particular λ_a .
FCI rotation matrix Q	The unitary rotation matrix that diagonalizes the FCI Hamiltonian.

EC training vector $ v_i^{(n)}\rangle$	n th eigenstate of the FCI Hamiltonian at a particular training point λ_i .
EC overlap $\mathcal{H}_{ij}(\ell)$	Hamiltonian matrix element evaluated with the training state vectors $\mathcal{H}_{ij}(\ell) := \langle v_i H_\ell v_j \rangle$.
EC overlap \mathcal{C}_{ij}	Overlap matrix between the training state vectors $\mathcal{C}_{ij} := \langle v_i v_j \rangle$.
Atomic Metric g^{ab}	A matrix of overlap integrals between two sets of local AO bases $\{\phi\}_a$ and $\{\phi\}_b$ for two different atomic configurations λ_a and λ_b . Note that a and b are a label for the matrix, and not the matrix indices.
EC eigenstate $ \tilde{v}_\ell^{(n)}\rangle$	Approximation of the n th eigenstate of the Hamiltonian at ℓ within the EC representation.

1. Introduction

A central motive of quantum physics and chemistry is the accurate determination of the low lying energy eigenstates of a Hamiltonian describing a system of interest. Whether for finding the ground state of a highly degenerate spin system, such as in spin liquids, or for studying a pathway for a particular chemical reaction, one inevitably uses the eigenstates and expectation values computed with them. Unfortunately, finding the energies and eigenstates is a computationally challenging problem, lying in either the NP-hard (classical) or QMA (quantum) complexity class [1]. A number of classical and quantum algorithms for finding low-lying eigenstates have been developed [2–14], addressing specific difficulties of the problem, but an approach working for all systems and computational demands is so far non-existing.

A common instance of the Hamiltonian eigenstate problem concerns the situation in which one requires the ground and/or excited states as a function of some Hamiltonian parameter(s) λ . For example, such a parameter could be a reaction coordinate in a chemical process, or an interaction strength coefficient when investigating the phase transitions in a spin model in condensed matter physics. In either case, the quantities of interest lie along some path through parameter space. Given that finding the eigenstates of a single Hamiltonian is already complex, it is easy to imagine that doing so with consistent accuracy along a Hamiltonian path is a daunting endeavor. In some reasonably common cases, this problem can be simplified. In particular, unless there is a symmetry-protected level crossing, the eigenvectors and eigenvalues along a given Hamiltonian path are continuous [15]. This allows making use of previously computed eigenstates at a set of parameters $\Lambda = \{\lambda_1, \dots, \lambda_N\}$ along the path as an efficient subspace basis in which to represent the Hamiltonian at some new $\ell \notin \Lambda$. As long as N is not exponentially large, this subspace is a much smaller problem to handle than the full diagonalization at ℓ , and can thus be performed classically at negligible computational cost. This approach, first introduced in the context of nuclear physics [15], is named ‘eigenvector continuation’ (EC), and it has been further extended to condensed matter physics [16, 17] and quantum computing [18].

Since it requires computing eigenstates at a limited number of points, EC may be particularly fruitful in situations where doing so is computationally expensive. Quantum computing may be such a case, as the currently viable algorithms for finding the ground state—including the variational quantum eigensolver (VQE), adiabatic state preparation (ASP), or quantum approximate adiabatic optimization—are expensive hybrid quantum–classical algorithms that are difficult to converge. And yet, to compute e.g. the binding curve of a molecule [19–23], or the phase diagram of a quantum magnet [24, 25], a new iterative loop is started at each new parameter point. While the initial guess for the loop may be improved, no further information is carried forward between iterations. Given the high cost of each eigenstate calculation, being able to reuse previously obtained eigenvectors could be a great advantage [16–18]. The main strategy is thus to perform the costly, exact eigenstate determination in only a small number of parameter points, and then reconstruct the eigenstates in the full path using a reduced, effective Hamiltonian representation in the basis of these selected points.

EC has been successfully employed to model Hamiltonians for solid state systems on a quantum computer [18], but potential applications in the scope of computational chemistry mostly unexplored. In this work we address this issue and demonstrate that EC can be readily applied to computing the binding curves of a number of chemical compounds, with the eye towards applying this to quantum computing. We study singly-bonded (F_2 , HF) doubly-bonded (H_2CO , O_2) and triply-bonded (N_2 , CO) molecules, as well as more strongly correlated examples (C_6H_8 -torsion, Cr_2). Within the context of these molecules, we investigate the use of EC and details of its implementation, particularly the special considerations that are unique to the *ab initio* setting. We evaluate the problems entirely on a classical simulator, however, the method pertains to quantum computation in the same sense that all subspace methods do [26–28]. It is relatively expensive to find the ground state of a model at a single parameter point [29] due to a combination of a large variational

search space [30, 31] barren plateaus [32–34], and deep circuits that are not amenable to today’s hardware, etc. Thus, it is beneficial to reduce the number of times that this task needs to be performed. We also note that EC can be used regardless of the particular circuit ansatz (such as the Hamiltonian Variational Ansatz [35], unitary coupled cluster theory [36] used for the ground state wavefunction, or other variational methods such as ADAPT-VQE [37]), making it quite a general method.

2. EC for *ab initio* calculations

The basic goal of EC, also referred to as the reduced basis method in the linear algebra community [16–18, 38], is accessing the lowest energy solution of a family of time-independent Schrödinger equations which share a parametrized Hamiltonian H_ℓ

$$H_\ell |v_\ell^{(n)}\rangle = E_\ell^{(n)} |v_\ell^{(n)}\rangle. \quad (1)$$

Given H_ℓ , the aim is to access the energies $E_\ell^{(0)}$ (as well as other observables) of the ground state wave functions $|v_\ell^{(0)}\rangle$ in some subset of the parameter phase space. This is to be done without actually undertaking the exponentially expensive exact solution of equation (1) for all parameter points of interest. Instead, the aim is to approximate the ground states for an arbitrary parameter choice ℓ inside the region of interest as a linear combination of a small number of selected parameter points $\lambda_i \in \Lambda$. Hence, after the exact ground state wave functions $|v_i^{(0)}\rangle$ are determined, the problem shifts to finding a set of expansion coefficients $a_i(\ell)$ such that

$$|v_\ell^{(0)}\rangle \approx |\tilde{v}_\ell^{(0)}\rangle = \sum_{i \in \Lambda} a_i(\ell) |v_i^{(0)}\rangle. \quad (2)$$

These coefficients can be variationally optimized by solving the corresponding generalized eigenvalue equation

$$\mathcal{H}_\ell |\tilde{v}_\ell^{(0)}\rangle = \tilde{E}_\ell^{(0)} \mathcal{C} |\tilde{v}_\ell^{(0)}\rangle \quad (3)$$

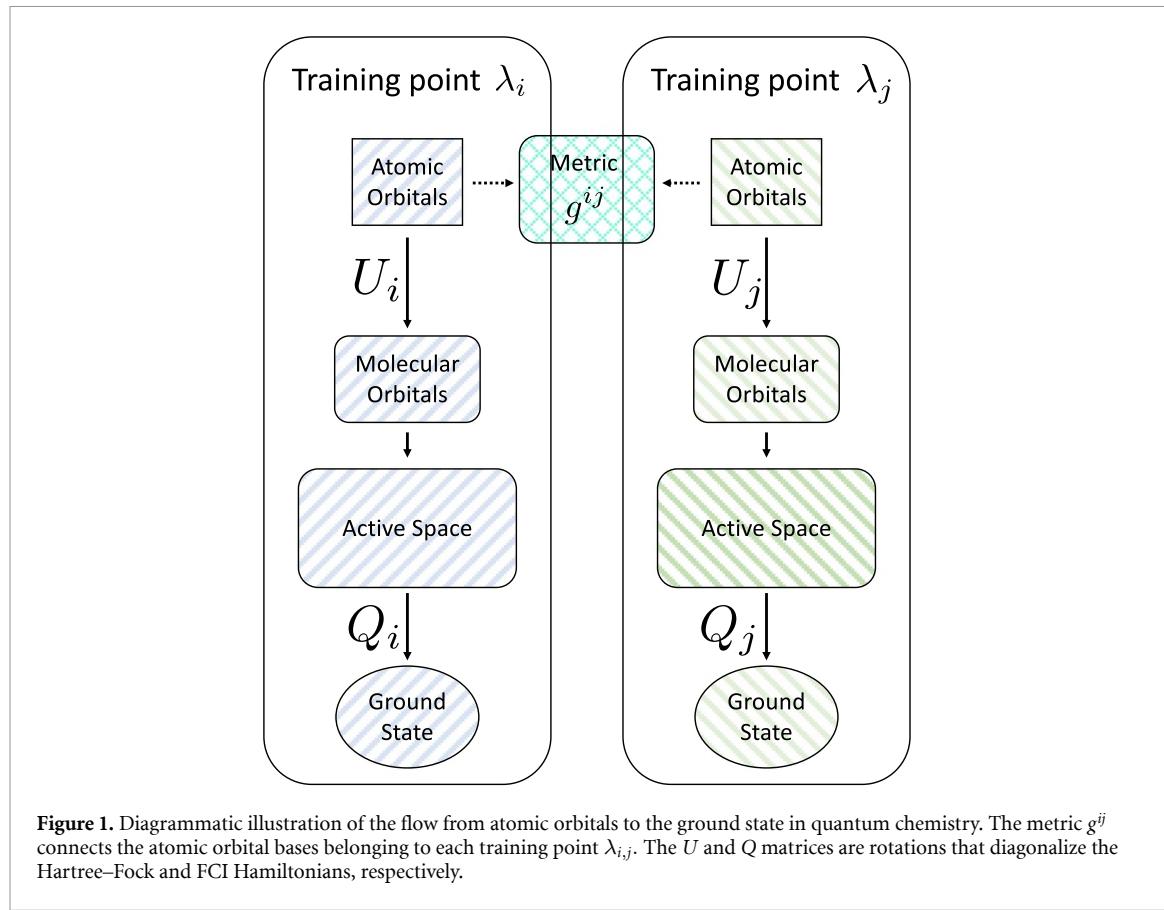
where the Hamiltonian and overlap matrix elements are

$$\begin{aligned} \mathcal{H}_{ij}(\ell) &= \langle v_i^{(0)} | H_\ell | v_j^{(0)} \rangle, \\ \mathcal{C}_{ij} &= \langle v_i^{(0)} | v_j^{(0)} \rangle. \end{aligned} \quad (4)$$

In the above equation, \mathcal{C}_{ij} is in general not the identity matrix since the states $|v_i^{(0)}\rangle$ are eigenvectors of different Hamiltonians. In order to implement EC on a quantum computer, the Hamiltonian and overlap matrix elements need to be measured. This can be straightforwardly achieved via Hadamard test based circuits, as discussed in the literature [18, 26, 28]. Following that, the generalized eigenvalue problem, which is the size of the number of ℓ values, can be diagonalized classically. As is the case in other subspace expansion methods based on a generalized eigenvalue problem, noise and measurement errors may lead to the condition number of the measured overlap matrix \mathcal{C} growing unfavorably large. This issue can be alleviated by performing a singular value decomposition of \mathcal{C} and filtering out all singular values below some threshold [26]. In this work, we choose a singular value threshold of $1E-4$.

An accurate representation of the ground states at all ℓ of interest can be achieved with a judicious choice of a small number of expansion points λ_i , which we will refer to as training points or EC points [15, 18]. Such a compact representation can be of great value for phase space screening of a Hamiltonian, e.g. to characterize the existing phases and transitions. For how to perform this ‘judicious’ choice of training points, we refer to the existing literature, such as the residue estimation method presented in [16, 17]; however, we note that a common ingredient in these approaches is the natural assumption that all Hamiltonians in the phase space of parameters λ share a single Hilbert space. This is not generally true in computational chemistry, as we discuss in detail below.

Within the context of quantum computation, the EC scheme allows a natural and potentially attractive approach to investigate the phase diagram of complex systems, where solving the Schrödinger equation (1) on the full set of parameter points to a desired accuracy is prohibitively expensive. Indeed, if preparing the expansion states $|v_i^{(0)}\rangle$ on a quantum register is feasible, one can then measure the Hamiltonian and overlap matrix elements in equation (4) and solve the generalized eigenvalue problem classically. This strategy has been successfully demonstrated on simple spin and chemical models in [18]. Of particular interest would be the application of EC to problems in *ab initio* computational chemistry, where the phase space studied can be



a parameterized chemical reaction. However, a particular complication arises in the *ab initio* context that needs to be addressed: the fact that the Hamiltonians for different parameter points λ_i will in general live, for realistic applications, in different Hilbert spaces.

2.1. The Hilbert space problem in *ab initio* computational chemistry

When considering EC in the context of quantum chemistry, a particular complication that arises is the fact that the atomic basis is not necessarily consistent for the set of training points. For example, consider the typical task of finding a binding energy curve of a diatomic molecule. At each separation R , the atomic orbitals (AOs) are centered at different points in space, which has to be handled in computing the EC Hamiltonian (\mathcal{H}) and overlap (\mathcal{C}) matrices—we discuss this procedure in detail in this section. A similar issue arises in performing EC in finite volume calculations where the volume is not consistent [39].

To set the stage for our discussion, we first outline the quantum chemistry process to obtaining a ground state for a correlated problem (see figure 1). Each training point λ_i comes with a set of AOs $\{\phi\}_i$. The initial step of finding the ground state of the interacting problem $|\psi_i^{(0)}\rangle$ is usually to solve the Hartree–Fock (HF) problem. This is done solving the generalized eigenvalue problem determined by the Fock-matrix (F) and overlap (S) for that set of AOs. This yields a rotation matrix U which mixes the AOs into a set of molecular orbitals (MOs). In turn, these MOs can be used as single-particle orbital basis to define a Fock space of many-electron states in their occupation number representation. In principle, one can then exactly solve the problem by projecting the Hamiltonian operator into this basis of Fock states, resulting in the so-called full configuration interaction (FCI) Hamiltonian H^{FCI} . The FCI Hamiltonian is diagonalized via another rotation matrix Q in the exponentially large Fock space to finally obtain the ground state $|\psi_i^{(0)}\rangle$. Note that in this diagonalization, one typically does not need to consider an overlap matrix, since the MOs are typically orthonormal.

One complexity arises in the EC when the inner product must be taken between two vectors $|\psi_i^{(0)}\rangle$ and $|\psi_j^{(0)}\rangle$ that arose from distinct sets of AOs. This is already clear from the overlap matrix element \mathcal{C}_{ij} at the HF level. Neglecting other complexities for now, the overlap between MOs $|\alpha(\lambda_i)\rangle$ and $|\beta(\lambda_j)\rangle$ that arise from the *different* Hilbert spaces at training points λ_i and λ_j is given by

$$\begin{aligned}\langle \alpha(\lambda_i) | \beta(\lambda_j) \rangle &= \sum_{mn} (U^i)_{n,\alpha}^* (U^j)_{\beta,m} \langle \phi_n^i | \phi_m^j \rangle \\ &= \sum_{mn} (U^i)_{n,\alpha}^* (U^j)_{\beta,m} (g^{ij})_{nm},\end{aligned}\quad (5)$$

where in the last line we have introduced the *metric* g^{ij} between the two training points i and j , which is a matrix containing the inner product between the two sets of AOs (cf figure 1). This already suggests that in following the EC strategy, some care will need to be taken to account for this difference in the orbital basis between different training points. An additional step, matching the orbitals of different training points, will be necessary to evaluate the expectation values in equation (4).

To the above considerations concerning the overlap of the AOs and MOs between different geometries one has to add an additional complication which arises frequently in realistic *ab initio* calculations: the notion of an active space. Even after the massive dimensionality decrease from the uncountable real-space basis, required to describe continuous space, to the finite number of AOs, the exponential scaling of the many-body Hilbert space as a function of the number of orbitals and electrons makes it computationally impossible to perform all-orbital, all-electron calculations except in the smallest molecules with the most modest basis sets. In all other cases, one typically restricts the post-HF (mean-field) determination of correlation effects to a subset of all orbitals, i.e. those orbitals deemed to be the most relevant for the electronic properties of the system. These are typically chosen to be the first N_o orbitals around the Fermi-level (the HOMO/LUMO frontier in the single reference description) containing the first N_e electrons in the mean-field reference determinant. These N_o orbitals with N_e electrons constitute the so called active space. An effective Hamiltonian for the active space can be formulated, in which all occupied orbitals outside the active space appear only as a constant shift in energy and as modified one-body terms. Post-HF correlated methods can then be applied to the active space alone, and additionally feedback correlation effects between the active space and the non-active orbitals can also be taken into account at different levels [2, 8].

This notion of active space adds another layer of inconsistency between the training FCI vectors on each geometry: Since the simplest way to define active space orbitals is in reference to the mean-field orbitals, and these change between each geometry, there is no guarantee that any given subset of them (such as the active space) span the same region of the one-body Hilbert space at each geometry. In an extreme example, if the mean-field orbitals close to the Fermi-level have a completely different AO character between two given parameters, then the FCI vectors obtained from the corresponding effective Hamiltonians will be essentially orthogonal.

These notions of orbital matching has to be included in the evaluation of the matrix elements in equation (4), where a transformation between the now active orbital basis in $\{|v_i^{(0)}\rangle, |v_j^{(0)}\rangle\}$ and $H(\ell)$ has to be performed. If the mismatch between the active orbital basis between the EC points and the target point ℓ is large, then this transformation will result in a reduction of the norm of the training vectors in the new basis. This is detrimental to the information contained in the overlap matrix, and thus to the conditioning of the generalized eigenvalue problem, although the issue can in principle be remedied by incorporating more training points to faithfully model the parameterized ground states in the parameter range of interest.

In order to minimize this effect, it is necessary to ensure that the active spaces in the parameter range to be studied with EC are spanned by AOs of the same character. This can be achieved by choosing large active spaces, such that all the relevant AOs for all parameter points will always be included; alternatively, one can choose the nature of the active space orbitals by more systematic means than proximity to the Fermi-level, e.g. using complete active space self-consistent field (CASSCF) orbital optimization (cf chapter 12 in [2]). In the result section below, we exemplify the first strategy for weakly-correlated molecules and employ the second for the strongly correlated Cr_2 . On a molecular torsion example, we will show a case in which this orbital mismatch is harder to solve, and consequently an increased number of training points is needed to cover the full parameter space of interest.

2.2. Possible orbital matchings

As discussed above (figure 1), finding the EC Hamiltonian and overlap matrices involves a local rotation from AOs to MOs U , a local rotation from MOs to FCI eigenvectors Q , and an inner product between two separate sets of AOs (which we encode in the metric g). In principle, to capture the proper inner products, each of these must be taken into account; in practice, however, it can be beneficial to neglect one or more of these.

In figure 2 we show the results of applying EC to the dissociation of F_2 using up to 4 training points and different orbital matching strategies. We compare the binding energy to the HF results, as well as a reference complete active space (CAS)-FCI result. The calculations are performed in the cc-pVDZ basis set, with an (8o, 14e) active space. The panels show three possible orbital matching approaches:

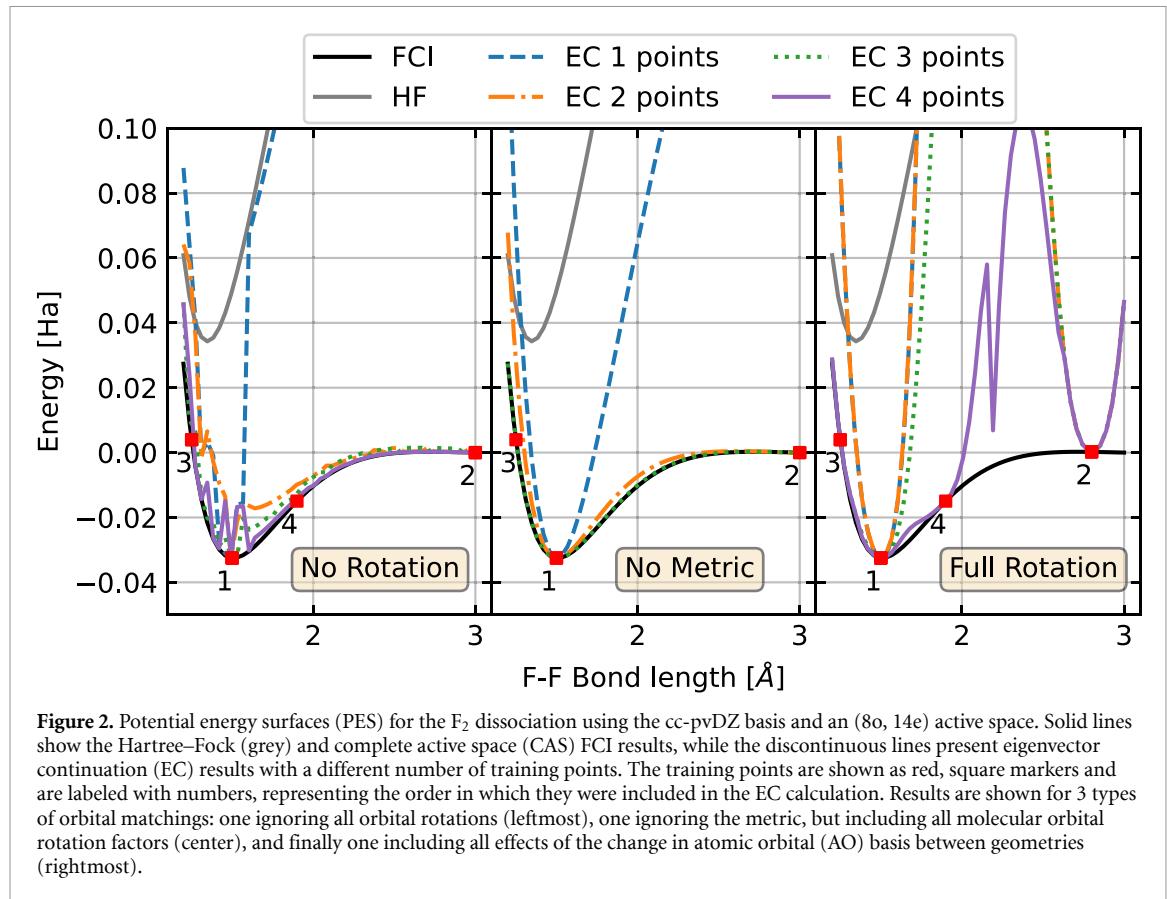


Figure 2. Potential energy surfaces (PES) for the F_2 dissociation using the cc-pvDZ basis and an (80, 14e) active space. Solid lines show the Hartree–Fock (grey) and complete active space (CAS) FCI results, while the discontinuous lines present eigenvector continuation (EC) results with a different number of training points. The training points are shown as red, square markers and are labeled with numbers, representing the order in which they were included in the EC calculation. Results are shown for 3 types of orbital matchings: one ignoring all orbital rotations (leftmost), one ignoring the metric, but including all molecular orbital rotation factors (center), and finally one including all effects of the change in atomic orbital (AO) basis between geometries (rightmost).

1. Full rotation: the \mathcal{H} and \mathcal{C} matrices are determined as discussed above, incorporating the U and Q rotations, as well as the metric.
2. No metric: the U and Q rotations from the FCI vectors to the AOs are kept, but the metric is neglected.
3. No rotation: the FCI vectors are treated entirely without reference to their origin. The U and Q rotations are neglected, as well as the metric.

Explicit mathematical expressions for the implementation of these three orbital matching choices are presented in the [appendix](#). Somewhat counter-intuitively, the best results are obtained when the metric is neglected, and the method still works when the metric and rotations are neglected (although with limited success). On the other hand, the notionally correct calculation which incorporates the rotations and the metric performs quite poorly.

The poor behavior of the calculation with full rotations can be understood by considering the metric. In our calculations we use localized AOs; while their highly localized nature is desirable from a quantum chemistry perspective, it also leads to a rapidly decaying metric. In essence, the overlap between atomic bases at different training point R tends exponentially to zero as R increases. We illustrate this in figure 3 where we plot the vector norm of one of a training state at $R = 1.5 \text{ \AA}$ in the basis corresponding to a range of R . When the metric is included, the norm drops and nearly vanishes for $R > 2.5 \text{ \AA}$. The inner product is thus not well captured, and the EC fails.

When the metric is neglected, however, EC performs quite well. In particular, when the local rotations U and Q are kept, 3 training points are sufficient to obtain the full binding energy curve. Intuitively, this can be understood as follows. The Q and U rotations describe how the final FCI eigenvector is composed of the MOs, and how the molecules orbitals are composed of the AOs. In other words, the FCI eigenvector at a given parameter ℓ is a vector in the space spanned by the basis of AOs at that parameter point. As the atomic separation is varied the FCI eigenvector rotates in the space spanned by the local AOs. However, it is in fact irrelevant that the local atomic basis is now shifted in real space. For EC, it suffices that the FCI eigenvector expressed in *its own* local basis can be spanned by the training points in *their own* local basis. This information is encoded in Q and U , and thus keeping those is sufficient. Putting this together with the issues with the metric, we conclude that neglecting the metric is a better choice than keeping it. In figure 3, we show that the previous issues with the vanishing overlap due to the metric do not arise here.

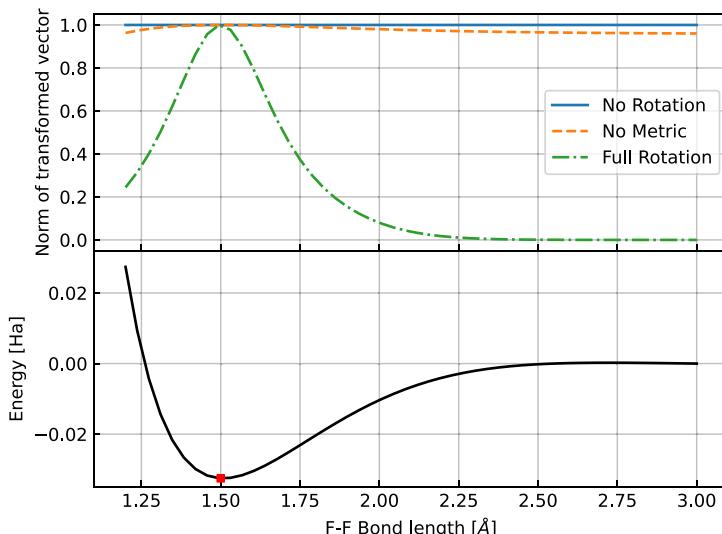


Figure 3. Norm of the transformed vector used as basis in the eigenvector continuation (EC) calculation, for the three different orbital matchings (upper panel), in an F_2 cc-pvDZ calculation with an (8o, 12e) active space. The original vector is the CAS-FCI solution at $R = 1.5$ Å (red marker in lower panel). For reference, the FCI potential energy surface is shown as a black solid curve (lower panel).

Finally, we can choose to neglect all rotations, and treat the FCI eigenvector as a vector divorced from any basis information. Here, a more straightforward linear algebra perspective is insightful. The FCI eigenvector simply needs to be spanned by a sufficient number of linearly independent basis vectors; the basis vectors need to be sufficiently expressive in order to be able to orthogonalize the ground state with respect to any other states in the subspace. Thus, this method works, but a larger number of training states may be required. Figure 2 shows that the 4 number of training points considered here are not sufficient to achieve agreement with the reference FCI result.

It is worth mentioning that the EC framework remains variational regardless of the orbital matching condition employed. Indeed, the orbital matching conditions just correspond to different effective expansion bases for a given point in parameter space, but the structure of the generalized eigenvalue problem in equation (4) remains the same for any of these choices. Hence, choosing the orbital matching condition minimizing the energy, besides that which generates continuous potential energy surface (PES), is a perfectly valid variational strategy.

The main goal of this work concerns showing the applicability of the EC framework to *ab initio* systems. As discussed in this section, the main ingredient that need to be added to previous implementation strategies [16–18, 38] is the orbital matching between different points in parameter space. Here, we implement this orbital matching as full orbital rotations with different rotation matrices (see [appendix](#)) on the FCI training vectors. This is of course not a realistic approach to target large molecular systems, since rotating the FCI vectors formally scales exponentially with system size, just like solving the FCI problem. Future work should be dedicated to developing other, perhaps approximate orbital matching implementations circumventing the explicit FCI vector rotation.

3. Analyzing the performance of EC

In this section we turn to the analysis of the reliability of EC as a compact and accurate approximation to characterize the ground state of *ab initio* molecular problems in simple one-dimensional parameter spaces. Arguably the most relevant parameter entering the molecular Hamiltonian, within the Born–Oppenheimer approximation, is the molecular geometry. The electronic energy eigenvalues as a function of the nuclear positions are commonly referred to as PESs. Hence, we can reformulate our goal as the study of how many EC points are necessary for accurately reconstructing one-dimensional PES in a few cases of chemical interest, namely stretching and torsion of covalent bonds.

While a one-dimensional PES for a particular molecule and bond is a fundamentally well defined target, the different approximations typically invoked in a computational chemistry calculation limit the ultimate accuracy of even a hypothetical and exact FCI simulation. Indeed, choices like the atomic basis set, the single-particle orbitals and the correlated active space all affect the FCI reference, and a careful analysis of the convergence of the observables of interest with respect to these factors is a necessary step in an electronic

Table 1. Table summarizing the computational details of the molecular potential energy surfaces (PES) studied in this work.

Molecule	Basis set	Orbital basis	Active space
F_2	cc-pVDZ	RHF	(8o, 14e)
O_2	cc-pVDZ	RHF	(8o, 12e)
N_2	cc-pVDZ	RHF	(8o, 10e)
HF	cc-pVDZ	ROHF	(7o, 10e)
H_2CO	cc-pVDZ	ROHF	(8o, 10e)
CO	cc-pVDZ	ROHF	(8o, 10e)
Cr_2	cc-pVTZ-dk	CASSCF	(12o, 12e)
C_6H_8	cc-pVDZ	RHF	(6o, 6e)

structure investigation. However, these considerations fall beyond the scope of this work, as we concern ourselves with examining how well EC can reproduce a given FCI reference with a small number of training points. We will thus choose a single, reasonable, but by no means final, FCI reference for each molecular case study, but make no claims as to its ultimate relevance towards the accurate description of the real physical system. We compensate this simplification by examining molecular examples of different degree of electronic correlation and computational complexity, in order to keep the validity of our conclusions as broad as possible.

A brief description of the FCI references for all molecular systems follows, with a subsequent presentation and discussion of the numerical results. All calculations were performed using the PYSCF package for electronic structure [40–42].

3.1. Molecular systems

3.1.1. Bond stretching of weakly correlated molecules

The majority of the PESs studied in this work fall under the category of bond stretching of ‘weakly correlated’ molecules. By this, we mean that the nature of the electronic correlation in the equilibrium geometry is well captured by single-reference methods. Nonetheless, in the bond stretching process the ground state naturally becomes multi-reference (to some degree strongly-correlated), making the accurate description of bond dissociation energies a challenge for effective single-particle theories even in these comparatively simple molecules. In addition, the study of bond stretching is of relevance to the quantum computing community, where the bond stretching and dissociation problem is a *drosophila* [11, 13, 43–47].

We consider bonds of different chemical character. We take into account the common heuristic distinction of single, double and triple bonds derived mostly from Lewis structures, and distinguish between symmetric and asymmetric bonds, i.e. bonds between chemically equivalent and inequivalent atoms. As examples of single bonds, we perform EC calculations for F_2 and HF, while we consider O_2 and the CO bond in H_2CO for the double bond category; N_2 and CO are our triple bond representatives. For all these systems, we used a cc-pVDZ basis set, in which at each geometry we perform a restricted HF calculation. For the asymmetric bond stretchings, in order to generate smooth PES, all possible spin states within restricted open-shell HF were considered, and the lowest in energy for each bond length was used as the MO basis for the subsequent FCI calculations. Even in this small molecules and moderate basis set, performing FCI on all electrons and orbitals is computationally prohibitive on a single processor. Hence, we performed instead CAS calculations including the 2p and 2s AO manifolds involved in the bond breaking. We summarize the active space sizes in table 1. We considered bond lengths up to 3.5 times the FCI equilibrium bond length.

For all the 6 molecules presented in figure 4, the PES curve obtained by EC is in good agreement with the FCI reference. To ease the comparison between different molecules, we rescale the bond length axis by the FCI equilibrium distance of each molecule, and the energy axis by the bonding energy, taking the energy at 3.5 times the equilibrium bond length as the dissociated asymptote. For each molecule, we present the minimal number of EC training points that leads to an acceptable result compared to FCI. None of the molecules require a particularly large number, but some variation does exist between the molecules. We note in particular that F_2 , N_2 and CO exhibit better agreement; while the other EC PES curves have some departure from the FCI result, these could be readily improved by the addition of more EC points (cf figure 2). In the case of F_2 , it is remarkable that just 3 points are enough to recover the full PES faithfully. These can be interpreted as the three distinct physical regions in the bond dissociation process: the bound region, the dissociated region, and the Coulson-Fisher point where a mean-field description would start breaking translational and/or spin symmetry. While for all other bonds shown in figure 4 more than 3 training points are needed, these typically agglomerate around the Coulson-Fisher region, where the system is arguably more strongly correlated. In this sense, a qualitative relationship can be established between the variability of the eigenstate character in a bond-length region and the number of EC points needed to sample

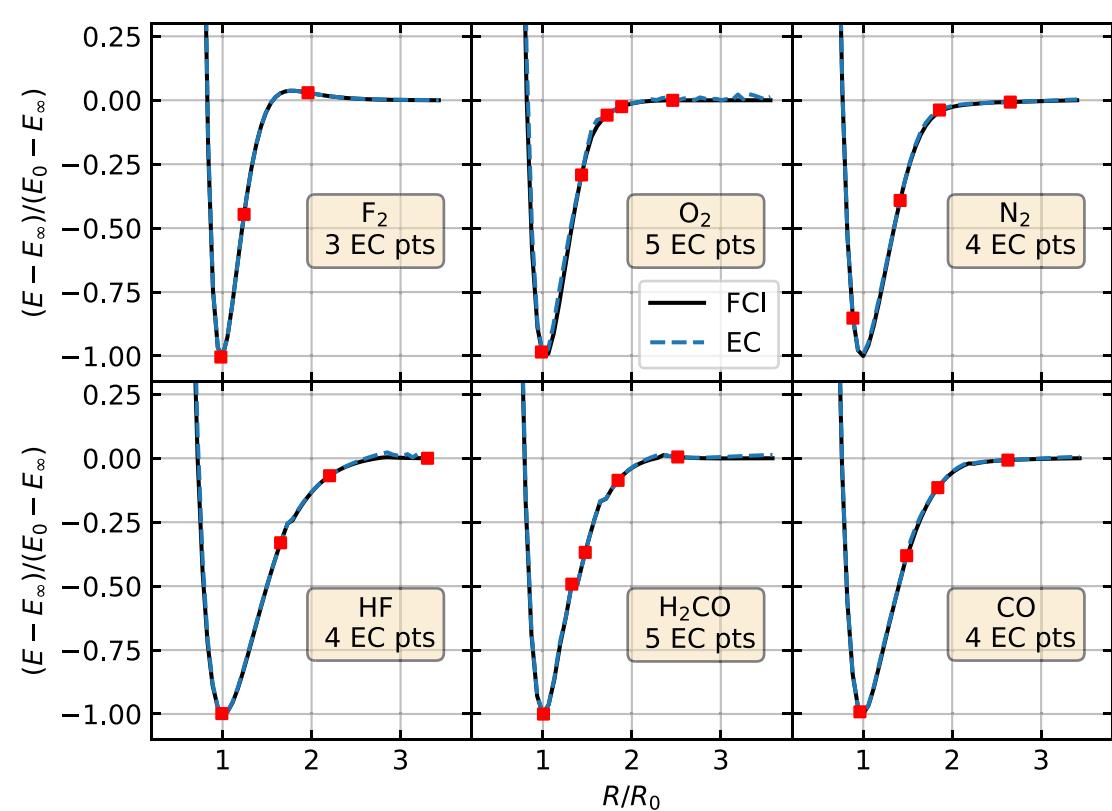


Figure 4. Bond stretching potential energy surfaces (PES) for small molecules, comparing FCI and eigenvector continuation (EC). The x -axis is the bond length rescaled with the equilibrium value for the given molecule, and the y -axis is the ground state energy rescaled by the minimum value and shifted by the large distance asymptotic (i.e. the bond energy). Symmetric bonds are shown in the upper row, while asymmetric bonds are found in the lower row.

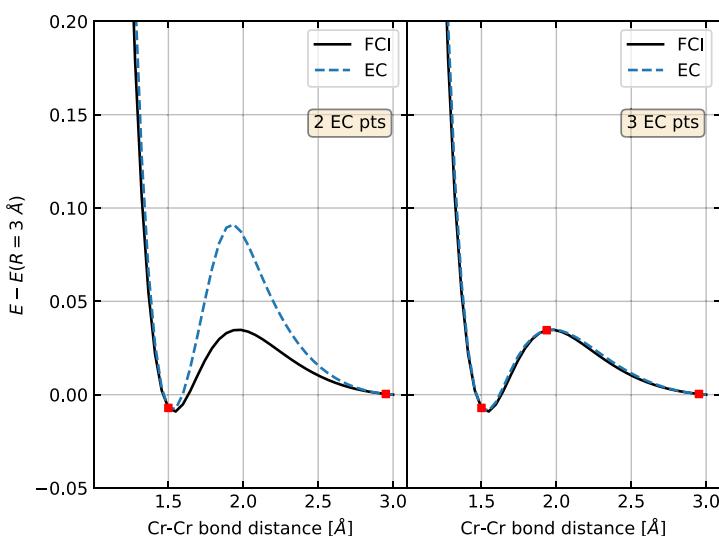


Figure 5. Potential energy surface (PES) for Cr_2 dimer in cc-pVTZ-dk basis. The FCI results correspond to a CASSCF (12o, 12e) calculation. Shown are eigenvector continuation (EC) for two different numbers of training points.

that zone accurately. This matches well the observations using EC in spin models [16–18]. We note that there are some unusual kinks in the PES curves for the asymmetric bond breakings, which is due to the limitations of the FCI calculations, rather than an artifact introduced by EC.

3.1.2. Cr_2 and bond torsion

To test the performance of EC for intrinsically strongly correlated molecules, we consider the bond stretching of a Cr_2 dimer, where we used a cc-pVTZ-dk basis set. Besides the restricted HF calculation at each bond length, a further orbital optimization was performed at the CASSCF level of theory, with a (12o, 12e) active

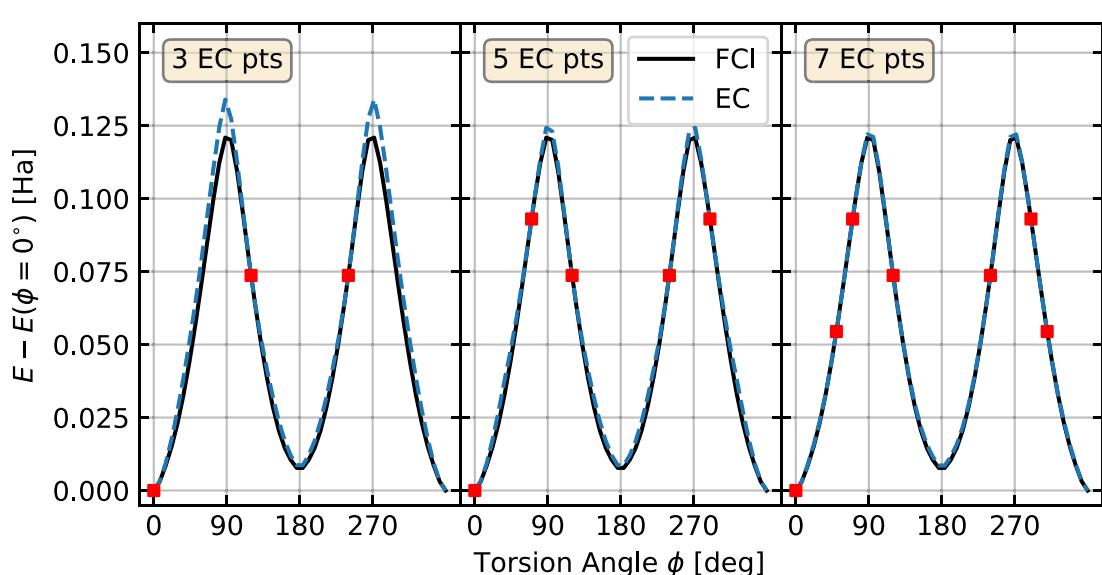


Figure 6. Potential energy surface (PES) for hexatriene in the cc-pvDZ basis as a function of the torsion angle ϕ around the central C–C double bond. $\phi = 0^\circ$ corresponds to the trans configuration, $\phi = 180^\circ$ to cis. The FCI results correspond to a complete active space (6o, 6e) calculation involving only the π orbital manifold. Shown are eigenvector continuation (EC) for three different numbers of training points, always symmetrically chosen around $\phi = 180^\circ$.

space. The multi-reference orbital optimization was necessary to obtain a homogeneous 3d and 4s orbital character in the active space through the bond dissociation. The (12o, 12e) CASSCF energy served then also as FCI reference for the EC. Figure 5 shows the results of the EC calculation for 2 and 3 training points. As was observed in the weakly correlated molecules, in Cr_2 as well a sparse sampling of the bound, dissociated and Coulson-Fisher regions is sufficient to recover the full PES.

Considering all the bond stretching examples, it is remarkable that the EC scheme seems to be fairly insensitive to the chemical nature of the PES modelled. Indeed, regardless of the chemical complexity of the bond, represented by single, double, triple, symmetric, asymmetric and correlated bonds, as well as the computational complexity of the FCI reference, based on either RHF or CASSCF orbitals, the EC representation shows a relatively homogeneous convergence in terms of the training points. A handful (up to five) points along the bond stretching, typically including at least the bound, dissociated and Coulson-Fisher region, are enough to obtain a visually accurate representation of the ground state PES along the full reaction path.

Finally, we considered the bond torsion of trans-hexatriene around the central CC double bond in figure 6. This PES was evaluated in the cc-pVQZ basis, using a minimal active space including all π orbitals, namely (6o, 6e), on restricted HF orbitals. The $\phi = 0^\circ$ geometry was taken from [48]. The orbital mismatch problem is more severe in this case as the rotation mixes the p -orbital manifold. By rotating around the bond, the atomic p_z -orbitals of one half of the molecule become eventually completely orthogonal to the p_z orbitals of the other half, and consequently the AO character of the frontier MOs changes drastically from 0 torsion to 90° . As a result, the PES requires a larger number of training points (7) to capture the full surface properly. Nonetheless, this is still modest sampling with which to recover the full PES.

3.2. Choosing the EC training points

As mentioned in section 2, how to judiciously choose the EC training points to maximize the compactness of the approximation without compromising accuracy has been discussed by Herbst *et al* [16], who suggest the use of a residual estimate for determining what points to add. In essence, given a previous EC approximation, the next point to be chosen is the one that maximizes the additional information in the basis as measured by the accuracy of the EC eigenvalue at that point. Here, we briefly exemplify how the residue estimate proposed therein satisfactorily adapts to the aim of achieving ‘chemical accuracy’ in *ab initio* simulations. By chemical accuracy, one refers to a maximal error of 1.6 mHa of a computational estimate with the true or reference value. When controlled experimental results are not available, one often uses as reference a computational result from a more accurate theoretical model. For our purposes, we can use the FCI reference to determine the error of our EC results at each molecular geometry.

If the average error in the energy of a given PES curve calculated via EC with m training points (m -EC) is above chemical accuracy, a natural choice for the $m + 1$ th training point is to pick a molecular geometry in

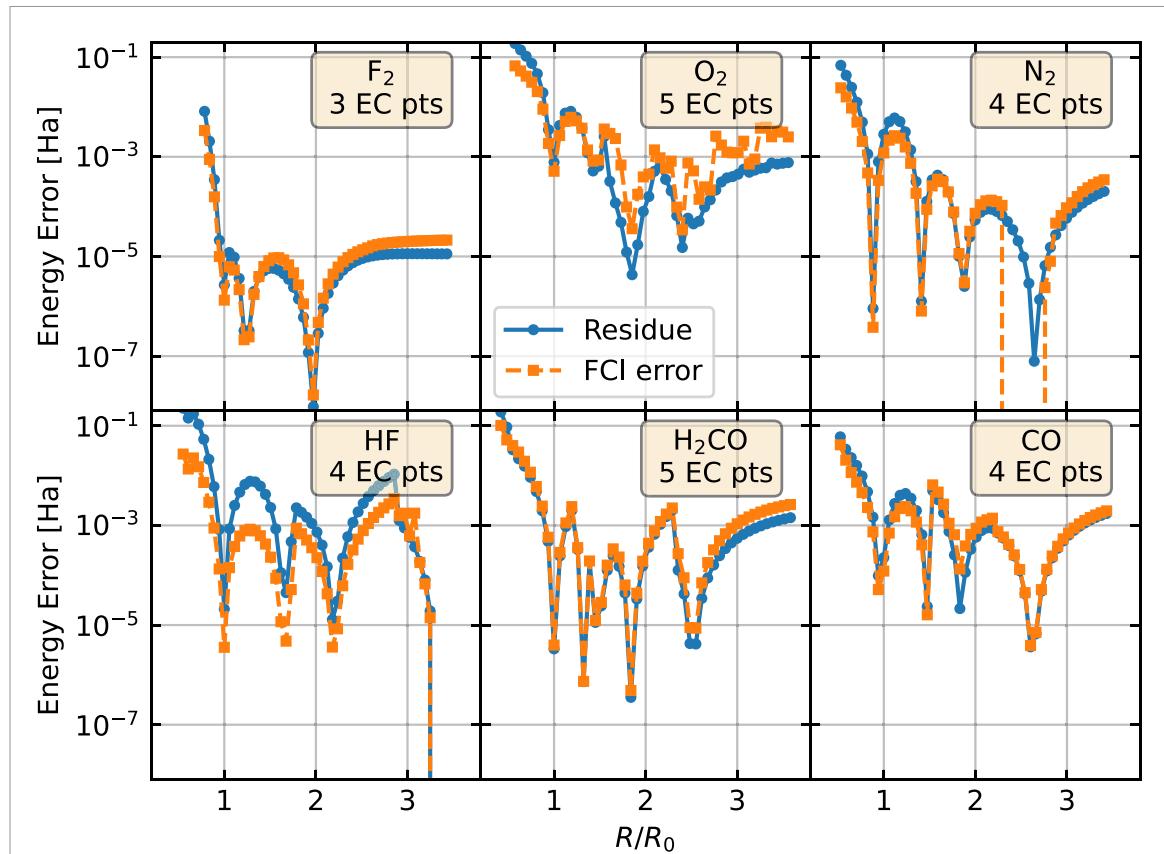


Figure 7. Two measures of the error on the potential energy surfaces (PES) for the bond stretchings examples in figure 4. The exact error with respect to FCI is shown with square markers, and the residue estimate in equation (6) with round markers. The approximate residue follows the exact error closely.

the region of maximal deviation. The FCI reference is in general not available, and it is necessary to obtain an estimate of the error using exclusively the data available within the EC calculation. Following [16], one can evaluate the residue of the EC approximation at each geometry of interest. Given a geometry ℓ , for which the m -EC simulation provides with a ground state wave function approximant $|\psi_\ell^{(0),m}\rangle$ with estimate energy $\tilde{E}_\ell^{(0),m}$, the residue is defined as

$$r_\ell^{(m)} = \left| H_\ell |\psi_\ell^{(0),m}\rangle - \tilde{E}_\ell^{(0),m} |\psi_\ell^{(0),m}\rangle \right|^2. \quad (6)$$

This residue $r_\ell^{(n)}$ can be written in terms of the Hamiltonian and overlap matrices, the eigenvector from the generalized eigenvalue problem in equation (3), and the matrix elements of the squared Hamiltonian in the EC training basis $(H^2)_{ij} = \langle \psi_i^{(0)} | H_\ell^2 | \psi_j^{(0)} \rangle$. The only additional cost on top of the EC calculation is thus the measurement of the squared Hamiltonian matrix elements.

In figure 7, we present the residues of the EC calculations for the bond stretchings in figure 4, keeping the same number of training points. We compare these to the error in the energy with respect to the FCI energies in the corresponding active spaces. As can be seen in figure 7, the residue estimate closely follows the actual error with respect to FCI, and thus offers an effective indicator to choose the EC training points in order to ultimately reach chemical accuracy with respect to the FCI reference. Of course, this does not guarantee an excellent agreement with experiment, as several approximations enter the FCI reference chosen in each case. However, the compactification offered by the EC approximation enables FCI-quality results within a small fraction of the cost of actually performing an FCI-level calculation (be it using a classical or quantum algorithm) at each point on the PES.

3.3. Accessing excited states with EC

In principle, the EC formalism is not limited to the approximation of ground states. As long as more than one training point is used, the generalized eigenvalue problem in equation (3) will have multiple eigenvectors, some of which could be accurate approximations of some excited state in the full Hamiltonian. Indeed, there are two scenarios in which it is natural to expect EC to provide a compact representation of excited states.

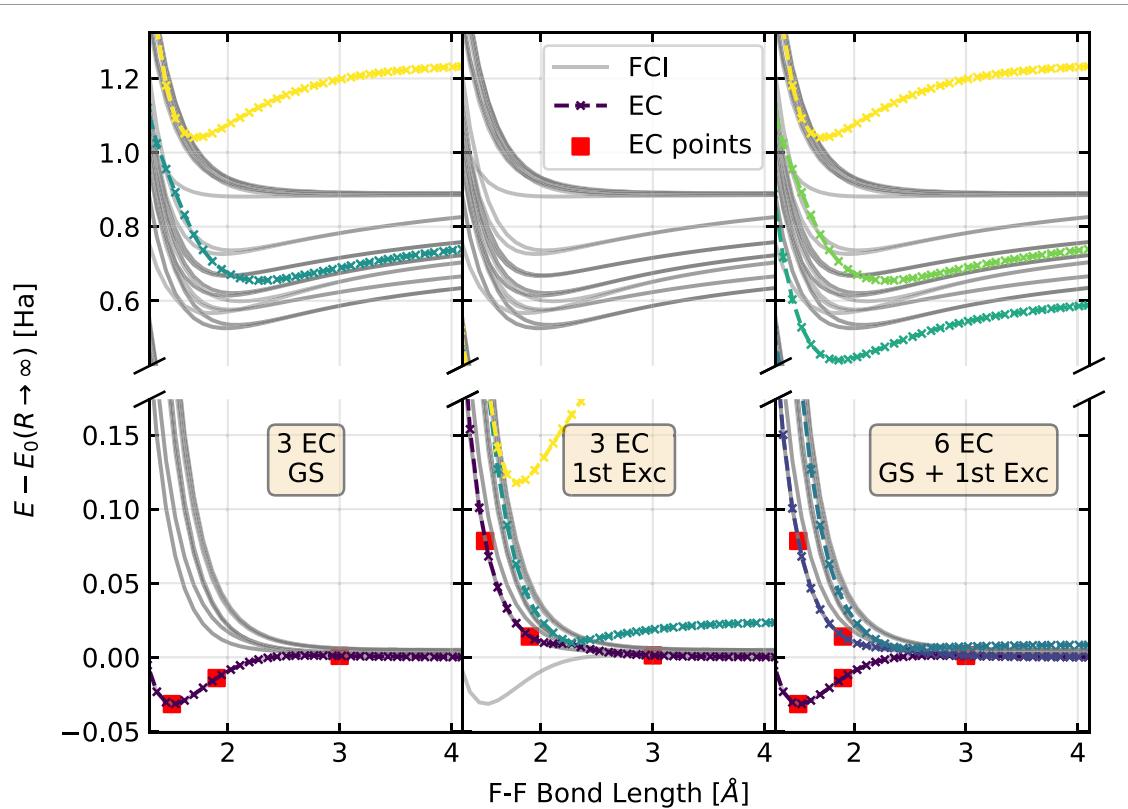


Figure 8. Results of eigenvector continuation (EC) for excited state potential energy surfaces (PES) in F_2 using the cc-pVDZ basis set. The FCI PES in the (8o, 14e) active space for the first few excited states are presented in grey. Results are shown for three different EC simulations. Left panel: EC with 3 training points using always FCI ground state vectors. Middle panel: EC with 3 training points using always FCI first excited state vectors. Right panel: EC with 6 training points in 3 different geometries, using both the ground state and 1st excited state of the FCI Hamiltonian in each point.

First, consider Hamiltonians that have avoided level crossings, such that the ground and first excited states switch character continuously across some path in the parameter space. In this case, using only ground states as training points along a path through the level crossing should also potentially result in an acceptable representation of the first excited state. Second, there is no fundamental need to use exclusively FCI ground states to build the EC training sets. If excited states are used, this should produce an EC approximation targeting the corresponding excited state PES. Simultaneously, in the presence of the aforementioned level crossings, having a mixed EC training set containing ground and excited states can lead to accurate representations for both. Here we consider these different possibilities in the example of the F_2 dimer.

We present excited state PES for the F_2 molecule in the cc-pVDZ basis, with a (8e, 14o) active space in figure 8. The FCI surfaces, shown as grey lines, do not show a level crossing between the ground and first excited state along the dimer bond dissociation. Nonetheless, these states become degenerate in the dissociation limit, and hence a complete decoupling between both states is not obvious *a priori*. Figure 8 shows the results for three different EC calculations in three panels. For all of these, the training points were obtained from the same molecular geometries, matching those already shown in the upper left panel of figure 4.

In the left-most panel of figure 8, we compare all three eigenstates obtained from the effective Hamiltonians of a ground state based EC to the first few lying FCI eigenstates. While the exact ground state PES is perfectly reproduced, as was already discussed in figure 4, the excited states of the effective EC Hamiltonian do not match well any of the exact excited state PESs. Moreover, these EC excited state PESs come after the first group of close-lying FCI excited state PESs, appearing at ~ 0.8 Ha above the ground state. This suggests that the subspace that captures the ground state PES could be orthogonal to the first bundle of excited states.

A similar picture occurs in the middle panel of figure 8, where the three training points in the EC simulation were all first excited states. Consequently, a relatively faithful approximation of the FCI first excited state PES is obtained, although a noticeable deviation appears at ~ 2.2 Å through an artificial, i.e. not present in the FCI reference, avoided crossing with the second excited state. Similarly to the first case, neither of the two higher lying excited states from the EC calculation approximate any of the FCI PESs well.

Surprisingly, all the EC curves in this case have some degree of bonding character (a minimum), even when all excited states in the corresponding energy window are dissociating, including the one used to obtain the training points. Still, since there is no overlap with the results from the EC calculation using ground state training vectors, it seems that the ground state and first excited eigenstate manifolds are mostly decoupled.

To further confirm this, an EC calculation was performed using 6 training points in the same 3 molecular geometries, shown in the right-most panel of figure 8. For each of the 3 bond lengths, the ground and first excited states at the FCI level were used as independent training vectors for the EC simulation. The three PES already obtained in the EC calculation using only ground state training points (cf left panel in figure 8) are found again in this larger simulation. However, the three PES that would correspond to the EC calculation based on excited states only (cf middle panel in figure 8) are significantly changed. The lowest in energy of the three—the overall first excited state of the EC simulation—follows the exact result better than in the middle panel, missing the deviation at ~ 2.2 Å. Furthermore, the next excited state PES is significantly shallower, closer to the expected non-binding behavior. Finally, the highest excited state among the central three is now pushed in energy much closer to the excited state manifold at ~ 0.8 Ha, better justifying its bound character. Despite these noticeable improvements, we still find that the only PES that accurately reproduces the FCI reference are those which are sampled explicitly, namely the FCI ground and first excited state in this case.

4. Conclusions

The spirit of the EC approach is proposing low-dimensional effective models to accurately reproduce targeted eigenstates of a parameterized Hamiltonian in some region of the parametric phase space. This is done by sampling a small number of points in said region, i.e. performing a computationally expensive, accurate determination of the eigenstates of interest in these few points, and then using their information to reconstruct the eigenstates inexpensively in the rest of phase space. The computation at the training points may be exact FCI when feasible [16], based on highly accurate matrix-product state Ansätze [17], or even the result of a quantum computation [18] for systems beyond the current reach of classical approaches. With a modest number of training points, the accurate results of comparable quality to these expensive methods can be recovered in the full parameter phase space at a fraction of the computational complexity. This becomes especially attractive for studying chemical reactions, which involves the accurate determination of the PESs of ground and excited states along the reaction coordinates. Therefore, here we have investigated the applicability and effectiveness of EC in the *ab initio* setting.

One of the major hurdles in applying EC to *ab initio* quantum chemistry is the mismatch in basis that arises from disparate molecular geometries for the subspace basis point, which we discussed extensively in the text. One significant conclusion from this work is that parts of this mismatch may be entirely neglected; specifically, the mismatch between the most basic level of the calculations, i.e. the AO overlap between different training states. After doing so, we have shown that the PES can be captured with remarkably few subspace basis vectors for a number of chemically distinct molecules single, double and triple bonds between chemically equivalent and inequivalent atoms in weakly correlated molecules, bond stretching of the intrinsically strongly correlated Cr_2 , and the bond torsion of trans-hexatriene around the central CC double bond. The associated error as compared to the FCI reference calculations is quite low.

Several aspects of the results that go beyond simple ground state manifolds are worth highlighting. First, EC can correctly handle level crossings in the ground state spectrum in chemical molecules, as long as training points are chosen on both sides of the crossing; this extends to any situation where multiple orthogonal subsectors are of interest. Second, we have shown that the use of EC is not limited to the ground state. Excited state manifolds can also be captured by inclusion of representatives of the excited states into the subspace. We exemplify this in F_2 by sampling with excited states instead of ground states.

There are two promising directions of future work on the EC framework worth mentioning at this point: First, as discussed in section 2.2, the current implementation of EC, which involves rotating the exponentially large FCI vectors, is not suitable for large calculations. Thus, in order to extend its impact to complex PES in larger molecules, there is need to develop an alternative approach to evaluate the expectation values in equation (4), heeding the issues with orbital matching presented here but avoiding the explicit rotation of the FCI vectors. The other direction concerns the use of approximate solutions instead of exact FCI for the training points. Indeed, any approximate ansatz giving access to the expectation values in equation (4), upon performing an orbital matching, can be used to perform the EC scheme. Moreover, the use of such approximate states does not affect the variational property of the resulting PES, just the obtained accuracy. When using EC as a classical algorithm, one could consider employing coupled cluster based approximations [49], while in a quantum algorithm ASP [50, 51] or other subspace expansion algorithms

[20, 21, 26] could be employed. In either case, it is interesting to note that the PES obtained using EC at the training points themselves can be variationally more accurate than the approximate solution it is built from. In this sense, EC can be a way of not only extracting the most information out of a small number of accurate PES samples, but also of improving the accuracy of said samples.

In short, EC is a promising tool for *ab initio* calculations in any situation where the eigenstates are difficult to obtain. This is in particular true on quantum computers, where finding ground states is a primary target and yet remains elusive; the current state of the art is plagued with issues in the optimization. It is thus quite difficult to find a ground state, and when this feat is accomplished, it should be used to maximum effect. EC is one way to achieve this goal.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Acknowledgments

C M Z acknowledges financial support from the European Research Council (ERC), under the European Union's Horizon 2020 research and innovation programme, Grant Agreement No. 692670 'FIRSTORM'. A F K acknowledges financial support from the US National Science Foundation under Grant No. NSF DMR-1752713.

Appendix A. Equilibrium geometries for H₂CO and trans-hexatriene

The equilibrium geometry used for H₂CO in this paper was optimized using the PYSCF interface to PyBerry [52] at the restricted HF level, using the cc-pVDZ basis. The obtained geometry is presented in table 2. The bond stretched in figure 4 is the CO double bond.

Table 2. Equilibrium geometry for H₂CO as used for the bond stretching in figure 4 and figure 7.

Atom	X	Y	Z
C	0.000 000	0.000 000	0.000 000
O	0.000 000	0.000 000	1.181 970
H	0.000 000	-0.932 542	-0.586 845
H	0.000 000	0.932 542	-0.586 845

The equilibrium geometry of trans-hexatriene (C₆H₈), i.e. the geometry corresponding to $\phi = 0^\circ$ in figure 6, was taken from [48]. We reproduce it in table 3 for completeness. The rotation in the paper is performed around the CC-bond between the carbon atoms in the first and third lines of table 3.

Table 3. Equilibrium (i.e. $\phi = 0^\circ$) geometry for trans-hexatriene, as reported in [48].

Atom	X	Y	Z
C	0.5987 833	0.2969 975	0.0000 000
H	0.6520 887	1.3822 812	0.0000 000
C	-0.5987 843	-0.2970 141	0.0000 000
H	-0.6520 904	-1.3822 967	0.0000 000
C	-1.8607 210	0.4195 548	0.0000 000
H	-1.8010 551	1.5036 080	0.0000 000
C	-3.0531 867	-0.1693 136	0.0000 000
H	-3.9685 470	0.4053 361	0.0000 000
H	-3.1479 810	-1.2485 605	0.0000 000
C	1.8607 264	-0.4195 599	0.0000 000
H	1.8010 777	-1.5036 141	0.0000 000
C	3.0531 816	0.1693 296	0.0000 000
H	3.9685 551	-0.4052 992	0.0000 000
H	3.1479 561	1.2485 793	0.0000 000

Appendix B. Derivation of the generalized eigenvalue equation

In this section, we briefly review the mathematics that leads to the generalized eigenvalue problem in equation (3). We start with the usual formulation of an eigenvalue problem—for a linear operator \mathcal{H} , an eigenvector $|\nu\rangle$ satisfies

$$\mathcal{H}|\nu\rangle = \lambda|\nu\rangle. \quad (\text{B1})$$

We can turn this into a matrix equation by expanding $|\nu\rangle$ in an orthonormal basis:

$$|\nu\rangle = \sum_j \nu_j |j\rangle. \quad (\text{B2})$$

Then,

$$\langle i|\mathcal{H}|\nu\rangle = \lambda \langle i|\nu\rangle \quad (\text{B3})$$

$$\sum_j \langle i|\mathcal{H}|j\rangle \nu_j = \lambda \sum_j \nu_j \langle i|j\rangle \quad (\text{B4})$$

$$\sum_j \mathcal{H}_{ij} \nu_j = \lambda \nu_i \quad (\text{B5})$$

which is the usual matrix eigenvalue equation and we have used the fact that $\langle i|j\rangle = \delta_{ij}$. A generalized eigenvalue problem arises when the basis $\{|i\rangle\}$ is not orthogonal. Then $\langle i|j\rangle$ is the (i,j) th element of an overlap matrix \mathcal{C} , and we obtain

$$\sum_j \mathcal{H}_{ij} \nu_j = \lambda \sum_j \mathcal{C}_{ij} \nu_j. \quad (\text{B6})$$

Appendix C. Mathematical expressions for orbital matching conditions

In this section, we briefly motivate the different orbital matchings presented in section 3.2 and figure 2, and give explicit mathematical expressions for the transformation matrices. For the sake of completeness, we start reviewing orbital bases. Given a basis of N_{AO} non-orthogonal AOs at the parameter space point λ_i (e.g. a given molecular geometry), denoted by $\{|\phi_m^i\rangle\}_{m=1}^{N_{AO}}$, their local overlap integrals are collected in the matrix S_i as

$$S_{i;mn} = \langle \phi_m^i | \phi_n^i \rangle. \quad (\text{C1})$$

Usually, the first step in a quantum chemistry computation involves finding an optimal MO basis (e.g. in the mean-field sense for HF), which can be expressed as a linear combination of AOs, such as

$$|\alpha(\lambda_i)\rangle = \sum_m U_{\alpha m}^i |\phi_m^i\rangle. \quad (\text{C2})$$

These MOs are typically chosen orthonormal, i.e. $\langle \alpha(\lambda_i) | \beta(\lambda_i) \rangle = \delta_{\alpha\beta}$, such that the matrices U^i and S_i fulfill

$$U^i S_i (U^i)^\dagger = \mathbb{I}. \quad (\text{C3})$$

Hence, the inverse of U^i is the matrix $\sqrt{S_i} (U^i)^\dagger$ is unitary. Given two different MO bases at one parameter point λ_i , with corresponding ‘AO to MO’ matrices U^i and W^i respectively, then the orbital transformation between these two MOs $T_{U \rightarrow W}^i$ is given by

$$T_{U \rightarrow W}^i = W^i S_i (U^i)^\dagger = W^i \sqrt{S_i} \sqrt{S_i} (U^i)^\dagger, \quad (\text{C4})$$

which is unitary since both MO bases are orthonormal. In order to transform a many-electron object between MO bases, such as the full molecular wave function, one has to apply the exponential operator $\exp [(\vec{c})^\dagger T_{U \rightarrow W}^i \vec{c}]$, where $(\vec{c})^\dagger$ is a vectorized representation of the list of creation operators in the MO basis of U^i .

Now, we consider two MOs corresponding to *different* parameter points λ_i and λ_j , with their ‘AO to MO’ matrices U^i and W^j respectively. Each parameter point has further its own AOs with their own local overlap

matrices S_i and S_j . The non-local overlap between the AOs at both parameter points also become relevant in this scenario. These form the metric g^{ij} introduced in the main text, namely

$$(g^{ij})_{nm} = \langle \phi_n^i | \phi_m^j \rangle. \quad (C5)$$

This metric enters the overlap between MOs from the two different parameter points, such as

$$\langle \alpha(\lambda_i) | \beta(\lambda_j) \rangle = \sum_{mn} (U_{\alpha m}^i)^* (g^{ij})_{mn} W_{\beta n}^j. \quad (C6)$$

The proper transformation between the U^i and W^j orbital bases in this case thus includes the metric explicitly through

$$T_{U \rightarrow W}^{i \rightarrow j} = W^j g^{ij} (U^i)^\dagger, \quad (C7)$$

which then enters the exponential transformation operator for many-body objects, $\exp \left[(\vec{c})^\dagger T_{U \rightarrow W}^{i \rightarrow j} \vec{c} \right]$. Unitary transformation matrices of this form enter the computation of the overlap and Hamiltonian matrices in equation (4). In this work, we evaluate this expectation values in the MO basis of the Hamiltonian of equation (4), which will in general be different than that of U^i and W^j .

The previous considerations including the metric would not be a major issue, if we kept all orbitals in our calculation. This is however computationally unachievable in general, and instead an active space is chosen. Then, the presence of g^{ij} can lead to an important problem in the computational chemistry setting: even if the atomic character (s, p, d, \dots) of the orbitals in the active space does not change between the U^i and W^j bases, the fact that the AOs are typically localized (e.g. as Gaussian orbitals) results in the inner products in equation (C5) naturally decreasing exponentially when λ_i and λ_j represent different molecular geometries. As a consequence, the generalized eigenvalue problem in equation (4) becomes ill-conditioned exponentially quickly along the PES, and one would formally need an exponentially dense grid of sample points to recover the full PES within the EC Ansatz.

This will happen if we try to perform the EC with the exact orbital matching, which includes the metric g^{ij} in equation (C7). Instead, we can alleviate the orbital mismatch due to the exponential decrease of the metric within the active space by making the substitution $g^{ij} \rightarrow \sqrt{S_j} \sqrt{S_i}$ in equation (C7). This corresponds to the orbital matching ignoring the metric described in section IIIB. Intuitively, this simplification ignoring the spatial displacement of the AOs with the change in orbital geometry, while still accounting for the changes in the MO composition in terms of those AOs. An even more insensitive approximation would be to assume $\langle \alpha(\lambda_i) | \beta(\lambda_j) \rangle = \delta_{\alpha\beta}$, which indeed would correspond to not rotating the MO basis at all between parameter points λ_i , which is also briefly presented in section IIIB. In our current implementation in pyscf, we use the function `fci.addons.transform_ci_for_orbital_rotation` to apply the exponential transformation operator the FCI ground states for the given choice of rotation matrix $T_{U \rightarrow W}^{i \rightarrow j}$. In summary, these choices are

$$T_{U \rightarrow W}^{i \rightarrow j} = \begin{cases} \mathbb{I} & \text{No Rotation} \\ W^j \sqrt{S_j} \sqrt{S_i} (U^i)^\dagger & \text{No metric} \\ W^j g^{ij} (U^i)^\dagger & \text{Full rotation} \end{cases}. \quad (C8)$$

ORCID iDs

Carlos Mejuto-Zaera  <https://orcid.org/0000-0001-5921-0959>

Alexander F Kemper  <https://orcid.org/0000-0002-5426-5181>

References

- [1] Bharti K *et al* 2022 *Rev. Mod. Phys.* **94** 015004
- [2] Helgaker T, Jorgensen P and Olsen J 2000 *Molecular Electronic Structure Theory* (Wiley)
- [3] Onida G, Reining L and Rubio A 2002 *Rev. Mod. Phys.* **74** 601
- [4] Martin R M 2004 *Electronic Structure: Basic Theory and Practical Methods* (Cambridge University Press)
- [5] Schollwöck U 2005 *Rev. Mod. Phys.* **77** 259
- [6] Kotliar G, Savrasov S Y, Haule K, Oudovenko V S, Parcollet O and Marianetti C 2006 *Rev. Mod. Phys.* **78** 865
- [7] Schollwöck U 2011 *Ann. Phys., NY* **326** 96
- [8] Szabo A and Ostlund N S 2012 *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Courier Corporation)
- [9] Martin R M, Reining L and Ceperley D M 2016 *Interacting Electrons* (Cambridge University Press)
- [10] Becca F and Sorella S 2017 *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press)
- [11] Cao Y *et al* 2019 *Chem. Rev.* **119** 10856

[12] Bauer B, Bravyi S, Motta M and Chan G K 2020 arXiv:[2001.03685](https://arxiv.org/abs/2001.03685)

[13] McArdle S, Endo S, Aspuru-Guzik A, Benjamin S C and Yuan X 2020 *Rev. Mod. Phys.* **92** 015003

[14] Motta M and Rice J E 2022 *Wiley Interdiscip. Rev.-Comput. Mol. Sci.* **12** e1580

[15] Frame D, He R, Ipsen I, Lee D, Lee D and Rrapaj E 2018 *Phys. Rev. Lett.* **121** 032501

[16] Herbst M F, Stamm B, Wessel S and Rizzi M 2022 *Phys. Rev. E* **105** 045303

[17] Brehmer P, Herbst M F, Wessel S, Rizzi M and Stamm B 2023 arXiv:[2304.13587](https://arxiv.org/abs/2304.13587)

[18] Francis A, Agrawal A A, Howard J H, Kökcü E and Kemper A 2022 arXiv:[2209.10571](https://arxiv.org/abs/2209.10571)

[19] Kandala A, Mezzacapo A, Temme K, Takita M, Brink M, Chow J M and Gambetta J M 2017 *Nature* **549** 242

[20] Parrish R M and McMahon P L 2019 arXiv:[1909.08925](https://arxiv.org/abs/1909.08925)

[21] Stair N H, Huang R and Evangelista F A 2020 *J. Chem. Theory Comput.* **16** 2236

[22] Andersson S *et al* 2020 Learn quantum computation using qiskit (available at: <http://community.qiskit.org/textbook>)

[23] Cohn J, Motta M and Parrish R M 2021 arXiv:[2104.08957](https://arxiv.org/abs/2104.08957)

[24] Nishimoto S, Shibata N and Hotta C 2013 *Nat. Commun.* **4** 2287

[25] Broholm C, Cava R, Kivelson S, Nocera D, Norman M and Senthil T 2020 *Science* **367** eaay0668

[26] Klymko K *et al* 2022 *PRX Quantum* **3** 020323

[27] Shen Y, Klymko K, Sud J, Williams-Young D B, de Jong W A and Tubman N M 2022 arXiv:[2208.01063](https://arxiv.org/abs/2208.01063)

[28] Cortes C L and Gray S K 2022 *Phys. Rev. A* **105** 022417

[29] Tilly J *et al* 2022 *Phys. Rep.* **986** 1

[30] Bittel L and Kliesch M 2021 *Phys. Rev. Lett.* **127** 120502

[31] Anschuetz E R and Kiani B T 2022 *Nat. Commun.* **13** 7760

[32] McClean J R, Boixo S, Smelyanskiy V N, Babbush R and Neven H 2018 *Nat. Commun.* **9** 1

[33] Cerezo M, Sone A, Volkoff T, Cincio L and Coles P J 2021 *Nat. Commun.* **12** 1

[34] Larocca M, Czarnik P, Sharma K, Muraleedharan G, Coles P J and Cerezo M 2022 *Quantum* **6** 824

[35] Wiersema R, Zhou C, de Sereville Y, Carrasquilla J F, Kim Y B and Yuen H 2020 *PRX Quantum* **1** 020319

[36] Anand A, Schleich P, Alperin-Lea S, Jensen P W, Sim S, Díaz-Tinoco M, Kottmann J S, Degroote M, Izmaylov A F and Aspuru-Guzik A 2022 *Chem. Soc. Rev.* **51** 1659

[37] Grimsley H R, Economou S E, Barnes E and Mayhall N J 2019 *Nat. Commun.* **10** 3007

[38] Quarteroni A, Manzoni A and Negri F 2015 *Reduced Basis Methods for Partial Differential Equations: an Introduction* vol 92 (Springer)

[39] Yapa N and König S 2022 *Phys. Rev. C* **106** 014309

[40] Sun Q *et al* 2018 *WIREs Comput. Mol. Sci.* **8** e1340

[41] Sun Q *et al* 2020 *J. Chem. Phys.* **153** 024109

[42] Sun Q 2015 *J. Comput. Chem.* **36** 1664

[43] Lee J, Huggins W J, Head-Gordon M and Whaley K B 2019 *J. Chem. Theory Comput.* **15** 311

[44] Arute F *et al* Google AI Quantum and Collaborators 2020 *Science* **369** 1084

[45] Huggins W J, O'Gorman B A, Rubin N C, Reichman D R, Babbush R and Lee J 2022 *Nature* **603** 416

[46] Baek U, Hait D, Shee J, Leimkuhler O, Huggins W J, Stetina T F, Head-Gordon M and Whaley K B 2022 Say no to optimization: a non-orthogonal quantum eigensolver (arXiv:[2205.09039](https://arxiv.org/abs/2205.09039) [quant-ph])

[47] Clary J M, Jones E B, Vigil-Fowler D, Chang C and Graf P 2023 *Int. J. Quantum Chem.* **123** e27097

[48] Huggins W J, Lee J, Baek U, O'Gorman B and Whaley K B 2020 *New J. Phys.* **22** 073009

[49] Bartlett R J and Musial M 2007 *Rev. Mod. Phys.* **79** 291

[50] Veis L and Pittner J 2014 *J. Chem. Phys.* **140** 214111

[51] Kremenetski V, Mejuto-Zaera C, Cotton S J and Tubman N M 2021 *J. Chem. Phys.* **155** 234106

[52] Hermann J 2016 Pyberny Documentation (available at: <https://jan.hermann.name/pyberny/index.html>)