

UC Berkeley

UC Berkeley Previously Published Works

Title

A hybrid approach to excited-state-specific variational Monte Carlo and doubly excited states

Permalink

<https://escholarship.org/uc/item/0rw1r6mt>

Journal

The Journal of Chemical Physics, 153(23)

ISSN

0021-9606

Authors

Otis, Leon
Craig, Isaac M
Neuscamman, Eric

Publication Date

2020-12-21

DOI

10.1063/5.0024572

Peer reviewed

A hybrid approach to excited-state-specific variational Monte Carlo and doubly excited states

Leon Otis,[†] Isabel Craig,[‡] and Eric Neuscamman^{*,‡,¶}

[†]*Department of Physics, University of California Berkeley, CA 94720, USA*

[‡]*Department of Chemistry, University of California Berkeley, CA 94720, USA*

[¶]*Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*

E-mail: eneuscamman@berkeley.edu

Abstract

We extend our hybrid linear-method/accelerated-descent variational Monte Carlo optimization approach to excited states and investigate its efficacy in double excitations. In addition to showing a superior statistical efficiency when compared to the linear method, our tests on small molecules show good energetic agreement with benchmark methods. We also demonstrate the ability to treat double excitations in systems that are too large for a full treatment by selected configuration interaction methods via an application to 4-aminobenzonitrile. Finally, we investigate the stability of state-specific variance optimization against collapse to other states' variance minima and find that symmetry, ansatz quality, and sample size all have roles to play in achieving stability.

1 Introduction

Accurate predictions about doubly excited states remain a significant challenge for electronic structure methods. Although they are rare in the low-lying spectra of simple organic molecules, states with significant or near-total doubly excited character are not uncommon in aromatic systems, other π -conjugated settings, and transition metal compounds. Their difficulty can be understood by considering three factors. First, they tend to be strongly multi-reference in character, which frustrates traditional weakly-correlated quantum chemistry approaches. Second, predicting accurate excitation energies requires a method to also capture weak correlation effects, and capturing both strong and weak correlation in medium to large systems remains an open challenge in electronic structure, despite the recent progress in smaller systems.¹⁻¹¹ Third, they are excited states, which remain harder to treat than ground states. When faced with a doubly excited state, many excited state approaches are either not appropriate or difficult to afford. Time-dependent density functional theory (TDDFT), at least within the adiabatic approximation, is famously incapable of predicting double excitations,¹² while equation of motion coupled cluster with singles and doubles (EOM-CCSD) is much less accurate in these states than in single excitations.¹³ Higher-level coupled cluster is more effective,¹⁴ but the $O(N^8)$ and $O(N^{10})$ cost scalings of EOM-CCSDT and EOM-CCSDTQ make them difficult to use outside of small molecules. Similarly, selected configuration interaction (sCI) methods can establish benchmark results for double excitations in small molecules,¹¹ but their exponential scaling makes them difficult to extend to larger systems. For these reasons, more traditional multi-reference methods like CASPT2^{15,16} have long been favored when dealing with double excitations.^{11,17-28} These methods have their limitations as well — intruder states, smaller active space sizes than sCI, sometimes-sensitive state-averaging choices — and so the development of alternative high-accuracy approaches to double excitations, especially ones that rely on very different approximations that could cross-validate current methods’ predictions, remains an important priority.

Quantum Monte Carlo (QMC) methods are, in principle, a promising alternative for doubly excited states and for difficult excited states more generally. Thanks to an ability to employ correlation factors and projector Monte Carlo methods to impart weak correlation effects on top of modest determinant expansions that capture the primary strong correlations, QMC approaches offer one route towards an integrated treatment of weak and strong correlation. Moreover, the ability of variational Monte Carlo (VMC) to work with excited states, either in a state-averaged^{29–33} or a state-specific^{34–36} manner, offers a clear route to extending these advantages to studies of doubly excited states. However, these advantages are only useful in practice if these sophisticated wave function forms can be optimized successfully for excited states. In ground states, recent years have seen significant improvements in the size and complexity of wave functions that can be treated by VMC wave function optimization,^{37–43} from the development of the linear method (LM) to the adoption of accelerated descent (AD) methods and even the combination of the two. Excited state wave function optimization is less well developed, and variance-based state-specific approaches — which are useful when dealing with high-lying states⁴⁴ or in cases where large dipole changes raise concerns for state-averaging⁴⁵ — have been shown to face stability issues in some surprisingly simple settings.³²

To help improve excited-state-specific wave function optimization in VMC, this study extends our hybrid LM/AD optimization approach⁴³ to a variance-based excited state objective function and tests both its stability and its efficacy for double excitations. As in the ground state, we find that AD provides useful support to the Blocked LM, and that their combination is more statistically efficient than the LM alone. In terms of stability against collapse to other states’ variance minima, we find that the difficulty that AD faces in making large changes to the wave function⁴³ actually prevents it from causing stability issues, but that the LM part of the algorithm faces the same challenges as have been seen³² in Newton optimizations of the variance. We therefore focus our stability testing on the LM, and find that achieving stable optimizations is greatly aided by enforcing symmetries, improving

the sophistication of the Jastrow factor, and using a sufficiently large sample size to avoid stochastically jumping out of a shallow minimum. In terms of efficacy for double excitations, we find the hybrid approach to be superior to the LM in all cases, usually in terms of statistical efficiency, but in one case also in terms of successfully finding the minimum at all. Comparisons against benchmark methods confirm VMC’s accuracy, while a demonstration in a doubly excited state of 4-aminobenzonitrile shows that the approach can be expanded well beyond the reach of sCI methods.

2 Theory

2.1 Excited-State-Specific Variational Monte Carlo

While VMC has historically been used for the study of ground states, multiple functionals have been developed for the targeting and variational optimization of excited states. We focus on the recently developed functional,³⁵

$$\Omega(\Psi) = \frac{\langle \Psi | (\omega - H) | \Psi \rangle}{\langle \Psi | (\omega - H)^2 | \Psi \rangle} \quad (1)$$

which is minimized when Ψ is the eigenstate of lowest energy greater than ω . The formulation of VMC for this excited state functional proceeds analogously to the ground state case. Just as ground state VMC seeks to minimize the expectation value of the energy, written as

$$E(\Psi) = \frac{\langle \Psi | H | \Psi \rangle}{\langle \Psi | \Psi \rangle} = \frac{\int d\mathbf{R} \Psi(\mathbf{R}) H \Psi(\mathbf{R})}{\int d\mathbf{R} \Psi(\mathbf{R})^2} = \frac{\int d\mathbf{R} \Psi(\mathbf{R})^2 E_L(\mathbf{R})}{\int d\mathbf{R} \Psi(\mathbf{R})^2} = \int d\mathbf{R} \rho(\mathbf{R}) E_L(\mathbf{R}) \quad (2)$$

using the local energy $E_L(\mathbf{R}) = \frac{H\Psi(\mathbf{R})}{\Psi(\mathbf{R})}$ and probability density $\rho(\mathbf{R}) = \frac{\Psi(\mathbf{R})^2}{\int d\mathbf{R} \Psi(\mathbf{R})^2}$, we can write Ω as

$$\Omega(\Psi) = \frac{\int d\mathbf{R} \Psi(\mathbf{R}) (\omega - H) \Psi(\mathbf{R})}{\int d\mathbf{R} \Psi(\mathbf{R}) (\omega - H)^2 \Psi(\mathbf{R})} = \frac{\int d\mathbf{R} \rho(\mathbf{R}) (\omega - E_L(\mathbf{R}))}{\int d\mathbf{R} \rho(\mathbf{R}) (\omega - E_L(\mathbf{R}))^2}. \quad (3)$$

However, while the probability distribution ρ has a useful zero-variance property,⁴⁶ it may be less useful than other importance sampling functions when estimating quantities such as the energy variance and matrix elements used in the optimization algorithms for minimizing Ω .^{45,47–49} Within this work, we use the importance sampling function

$$|\Phi|^2 = |\Psi|^2 + c_1 \sum_i |\Psi^i|^2 + c_2 \sum_j |\Psi^j|^2 + c_3 \sum_k |\Psi^k|^2 \quad (4)$$

where c_1, c_2, c_3 are weights on sums of squares of wave function derivatives Ψ^i, Ψ^j, Ψ^k for Jastrow, CI, and orbital parameters, respectively. Effective choices for (c_1, c_2, c_3) may be system-dependent and require some experimentation in practice. In our results, we use $(0.0, 0.0001, 0.0)$ for our stability tests on a model cyanine dye, $(0.0004, 0.0002, 0.0)$ for the carbon dimer, $(0.0001, 0.0001, 0.0)$ for nitroxyl, glyoxal, and acrolein, and $(0.0001, 0.0, 0.0)$ for cyclopentadiene and 4-aminobenzonitrile. We have found these values enabled optimization to lower target function values than we could achieve using $|\Psi|^2$ and that using nonzero c_3 for the orbital parameter derivatives also resulted in poorer target function values. However, our exploration of these importance function parameters was not exhaustive and different values may also result in equally good target function values. The general intuition is to include wave function derivative terms, as importance sampling a linear combination of the current wave function probability density and that of its parameter derivatives allows us to obtain a better statistical estimate of the LM Hamiltonian. Whenever the parameters have large effects on the nodal surface, importance sampling $|\Psi|^2$ will lead to a large statistical uncertainty in this Hamiltonian. Importance sampling $|\Phi|^2$ allows us to better sample areas of parameter space where the current Ψ has negligible probability density, but linear updates to Ψ may not, obtaining a better Hamiltonian estimate and therefore better parameter updates in the process. Because this guiding function also has considerably fewer nodes, it can further help avoid the divergence in the variance of the variance^{47,48} encountered with $|\Psi|^2$ while also keeping the distribution close to $|\Psi|^2$.⁴⁹ We also employ clipping of the samples based

on the value of the local energy.^{50,51} We compute a deviation of the form $\frac{1}{N} \sum |E_L(\mathbf{R}) - \bar{E}|$ where \bar{E} is the mean local energy of all initial N samples. Samples with local energies more than 5 times the deviation from \bar{E} are discarded and only the remainder are used in computations for the VMC optimization. We find that clipping can improve optimization performance, especially as more flexible ansatzes with larger numbers of parameters are considered, by reducing the statistical uncertainty in the matrix elements and derivatives used by the algorithms. Essentially, this guards against large parameter updates.

The use of Ω for excited states also requires careful handling of the input ω . For a generic fixed choice of ω , Ω will not be size-consistent and so ω must be updated to transform Ω into state-specific variance minimization, which is size-consistent.³⁶ There are multiple strategies for achieving this transformation to have both state-specificity and size-consistency, such as a linear interpolation between the initial fixed value of ω and the floating value of $E - \sigma$,^{32,36} or, as in this work, a series of fixed- ω optimizations with ω updated between each one until self-consistency between it and $E - \sigma$ is reached. The details of how ω is varied in a calculation is one potential source of instability in the optimization as the target function being minimized with respect to wave function parameters is now changing and there is the possibility of slipping outside the basin of convergence, particularly if ω is varied rapidly.

2.2 Linear Method

The LM^{37,52–54} is based on a first order Taylor expansion of the wave function

$$\Psi(\mathbf{p}) = \Psi_0 + \sum_i \Delta p_i \Psi_i \quad (5)$$

using $\Psi_i = \frac{\partial \Psi(\mathbf{p})}{\partial p_i}$ for first order parameter derivatives of the wave function and Ψ_0 for the wave function at the current parameter values \mathbf{p} . Seeking to minimize the target function Ω with respect to \mathbf{p} leads to the generalized eigenvalue problem

$$(\omega - \mathbf{H}) \mathbf{c} = \lambda (\omega - \mathbf{H})^2 \mathbf{c} \quad (6)$$

which can be solved to produce an update vector $\mathbf{c} = (1, \Delta \mathbf{p})$. The matrices are constructed in the basis of the initial wave function and its first order parameter derivatives, so we have matrix elements of the form

$$\langle \Psi_i | \omega - H | \Psi_j \rangle \quad (7)$$

and

$$\langle \Psi_i | (\omega - H)^2 | \Psi_j \rangle. \quad (8)$$

These matrix elements are evaluated within VMC using parameter derivatives of the wave function and the local energy, and we have found that employing the modified guiding function $|\Phi|^2$ can be crucial for obtaining accurate estimates and effective LM optimizations. In practice, the Hamiltonian matrix of the LM is modified with the addition of shift matrices^{43,55} which help prevent incautiously large parameter changes in the optimization and can noticeably influence the LM's stability and performance. Our implementation of the LM uses by default an adaptive scheme where three sets of shift values are used to produce candidate parameter updates and a correlated sampling procedure is used to determine which one, if any, should be taken to improve the target function value. This approach allows the value of the shifts to vary over the course of the optimization, with the aim of allowing the LM to safely take large steps in parameter space early on while automatically becoming more cautious when close to the minimum, where statistical uncertainty will eventually prevent steps from being able to resolve downhill.

Much recent work on VMC for excited states^{30,35,44,45,49,56-58} has relied on the LM for the task of wave function optimization. However, the LM has multiple limitations, particularly a memory cost that increases with the square of the number of optimizable parameters, and applications with more than 10,000 parameters are rare. The quadratic growth in the number of LM matrix elements exacerbates the nonlinear bias of the LM and eventually

leads to the underdetermined regime where there are too few samples to effectively estimate the LM matrices and the step uncertainty is large. To help address these issues, a variant of the LM known as Blocked LM has been developed⁴⁰ that divides the parameter set into N_b blocks and performs a LM-style matrix diagonalization for each block. Some number N_k of the eigenvectors from each block are retained and combined together along with N_o other directions in parameter space for a second LM diagonalization to obtain an update in the full parameter space. The lower dimension of the matrices constructed by the Blocked LM alleviates the issues faced by the standard LM at the cost of having to run over the samples twice instead of only once. Further details on the Blocked LM can be found in the original publication⁴⁰ and we describe our choices on the number of blocks and retained directions in the supplementary material. A recent study⁴³ of optimization methods for ground state optimization indicates that tighter and more statistically efficient convergence can be obtained through a hybrid combination of the Blocked LM and AD than when using the LM alone.

2.3 Hybrid Optimization

There are multiple flavors of AD optimization methods that require only first order parameter derivatives. These methods are widespread in the machine learning community and have been increasingly used in the context of VMC.^{41–43,59,60} They offer some appealing advantages compared to the LM, including a memory cost linear in parameter number, a lower per-iteration sampling cost to estimate derivatives, and a reduced nonlinear bias from the stochastic evaluation of those derivatives. However, comparisons⁴³ between AD methods and the LM indicate that the former may struggle to reach the minimum in parameter space at comparable computational effort, especially when the wave function contains many challenging nonlinear parameters.

An alternative approach that shows the potential to benefit from the strengths of both classes of methods is to take a hybrid combination that alternates between them.⁴³ Sections

of gradient descent optimization can be used to identify important directions in parameter space, which can be provided as input to the Blocked LM to help it account for coupling between blocks of parameters. The inclusion of Blocked LM steps allows the hybrid method to more efficiently reach the vicinity of the minimum, while the sections of AD enable tighter convergence by correcting poor parameter updates by the Blocked LM due to step uncertainty. In addition, the use of AD and Blocked LM enables the optimization of larger parameter sets that are beyond the reach of the standard LM. In energy minimization we found it especially useful to follow the hybrid optimization with a final section of optimization using only descent, which has been found to more efficiently achieve lower final statistical uncertainties than optimization based solely on the LM.⁴³

The hybrid scheme can be used with any of the variety of AD methods, but in this work, we use a combination of Nesterov momentum with the RMSprop algorithm⁴¹ that was found to work well for energy minimization. It is specified by the following recurrence relations.

$$p_i^{k+1} = (1 - \gamma_k e^{-(\frac{1}{d})(k-1)})q_i^{k+1} - \gamma_k e^{-(\frac{1}{d})(k-1)}q_i^k \quad (9)$$

$$q_i^{k+1} = p_i^k - \tau_k \frac{\partial \Omega(\mathbf{p})}{\partial p_i} \quad (10)$$

$$\lambda_0 = 0 \quad \lambda_k = \frac{1}{2} + \frac{1}{2} \sqrt{1 + 4\lambda_{k-1}^2} \quad \gamma_k = \frac{1 - \lambda_k}{\lambda_{k+1}} \quad (11)$$

$$\tau_k = \frac{\eta}{\sqrt{E[(\frac{\partial \Omega}{\partial p_i})^2]^{(k)} + \epsilon}} \quad (12)$$

$$E[(\partial \Omega)^2]^{(k)} = \rho E \left[\left(\frac{\partial \Omega}{\partial p_i} \right)^2 \right]^{(k-1)} + (1 - \rho) \left(\frac{\partial \Omega}{\partial p_i} \right)^2 \quad (13)$$

Equations 9 through 11 describe how a parameter p_i^k of the wave function on step k of the optimization is updated using knowledge of both the current and previous values of the target function derivatives. The step size τ_k is adaptively adjusted using a running average of the square of parameter derivatives according to equations 12 and 13, where ρ sets the

relative weighting between the past average $E \left[\left(\frac{\partial \Omega}{\partial p_i} \right)^2 \right]^{(k-1)}$ and the current value $\left(\frac{\partial \Omega}{\partial p_i} \right)^2$. The quantities d, η, ρ , and ϵ are hyperparameters whose values are chosen by the algorithm’s user. For all our results, we have used $d = 100$, $\rho = 0.9$ and $\epsilon = 10^{-8}$. The value of η sets an overall scale for step size and can have a significant influence on optimization performance depending on the choices made for different types of parameters. Our choices for η are discussed in the supplementary material.

2.4 Optimization Stability

When working with an exact ansatz and an infinite sample size, any non-degenerate Hamiltonian eigenstate possesses its own variance minimum,³⁴ and an optimization with a guess sufficiently close to that minimum will be stable. In practice, sample sizes are finite, and the use of an approximate ansatz may lead some states’ variance minima to be artificially shallow or to disappear entirely, and so there is a real possibility that state-specific variance minimization will be unstable.³² If the minimum has indeed disappeared, then the only remedy is to improve the ansatz quality, as no amount of statistical precision will allow an optimization to find a minima that is not present. On the bright side, advances in VMC trial functions^{61–69} offer a strong toolkit for improving ansatz quality, albeit one whose use can increase optimization difficulty by increasing the number of variational parameters. If anything, this reality further motivates the development of optimizers, like the hybrid approach tested here, that are designed to handle large, challenging trial functions. In Section 3.2, we will explore the issue of optimization stability by starting with a very simple ansatz for which some states lack variance minima and then making improvements, either by enforcing symmetries or by enhancing the ansatz in order to make the variance minima re-emerge.

If the variance minimum is present but shallow, either due to an approximate ansatz or simply to a near-degeneracy in the spectrum, then the nature of the optimization update steps, and especially their statistical uncertainty, will determine whether an optimization is stable. Strictly speaking, finite-sample-size VMC optimizations always contain at least

some risk of selecting a step in the tail of the distribution of steps that, if taken, will move the optimization out of the basin of convergence for one variance minimum and into the basin for another. To help guard against this issue, a common practice in VMC, and one that we do use for our LM steps, is to verify on a new independent sample (usually using correlated sampling) that the step about to be taken does indeed lower the objective function. Even without statistical uncertainty, step size control is important in nonlinear optimization, an issue that is typically addressed by trust radius methods that guard against single steps making overly large parameter changes.⁷⁰ In the LM, one approach to this issue is a diagonal shift,³⁷ and our LM implementation employs both this shift and an overlap-based shift⁵⁵ for these purposes. In Section 3.2, we will see an example of an optimization that only becomes stable with a large enough sample size due to a shallow minimum created by a near-degeneracy.

2.5 Wave Functions

For demonstrating the hybrid method’s effectiveness within our overall QMC methodology for excited states, we consider multiple types of parameters in our trial wave functions. Our ansatz is the Multi-Slater Jastrow wave function, which has the form

$$\Psi = \psi_{MS}\psi_J \tag{14}$$

$$\psi_{MS} = \sum_{i=0}^{N_D} c_i D_i \tag{15}$$

$$\psi_J = \exp \left\{ \sum_i \sum_j \chi_k(|r_i - R_j|) + \sum_k \sum_{l>k} u_{kl}(|r_k - r_l|) \right\} \tag{16}$$

where D_i are Slater determinants with coefficients c_i and the Jastrow factor ψ_J is constructed from splines that make up the functions χ_k and u_{kl} for the respective one- and two-body terms.⁵⁵ The MSJ ansatz is a common choice in QMC, but it can be augmented further to describe more correlation at the price of a more challenging optimization. Two means for

doing so are to add a more complicated Jastrow factor and to optimize molecular orbital shapes.

A variety^{34,62,67–69,71–76} of many-body Jastrow factors have been considered for improving QMC ansatzes, but in this work we will limit ourselves to adding a three-body term and a number-counting factor. The three-body term is constructed from polynomials of interparticle distances and can be written in the following form.⁶²

$$u(r_{\sigma I}, r_{\sigma' I}, r_{\sigma\sigma'}) = \sum_{l=0}^{M_{eI}} \sum_{m=0}^{M_{eI}} \sum_{n=0}^{M_{ee}} \gamma_{lmn} r_{\sigma I}^l r_{\sigma' I}^m r_{\sigma\sigma'}^n (r_{\sigma I} - \frac{r_c}{2})^3 (r_{\sigma' I} - \frac{r_c}{2})^3 \Theta(r_{\sigma I} - \frac{r_c}{2}) \Theta(r_{\sigma' I} - \frac{r_c}{2}) \quad (17)$$

The maximum polynomial orders are set by M_{eI} and M_{ee} for the electron-ion and electron-electron distances respectively. The γ_{lmn} are the set of optimizable parameters in this polynomial and are subject to constraints for ensuring the Jastrow satisfies symmetry under exchange and cusp conditions, which can be found in the original publication.⁶² Finally, the Theta functions require the three-body term become zero for electron-ion distances more than half a chosen cutoff distance r_c (10 bohr in our case).

Number-counting Jastrow factors are a recently developed^{63,64} ansatz component and can be thought of as a many-body Jastrow factor in real space that aims to recover both strong and weak correlation. They are based on a Voronoi partitioning of space, where the population of electrons in each region, N_I , is given by a sum of local counting functions C_I at each electron coordinate.

$$N_I = \sum_i C_I(\mathbf{r}_i) = \sum_i \frac{g_I(\mathbf{r})}{\sum_j g_j(\mathbf{r})} \quad (18)$$

where

$$g_j(\mathbf{r}) = \exp((\mathbf{r} - \boldsymbol{\mu})^T \mathbf{A}(\mathbf{r} - \boldsymbol{\mu}) + K) \quad (19)$$

are Gaussian basis functions about a center $\boldsymbol{\mu}$. With these populations N_I , we can construct

the Jastrow factor

$$\psi_C = \exp\left(\sum_{IJ} F_{IJ} N_I N_J + \sum_K G_K N_K\right) \quad (20)$$

which can be tacked on to our overall expression for Ψ . The F_{IJ} and G_K are variational parameters though the latter is in practice eliminated with a basis transformation of the region populations and so we only optimize F_{IJ} .⁶⁴

State-specific orbital optimization is useful for obtaining accurate VMC results on particular excited state phenomena including charge transfer⁴⁵ and core excitations⁴⁴ and can avoid some of the pitfalls of state-averaged approaches.⁷⁷ Recent work^{65,66} with the table method has enabled the efficient calculation of orbital parameter derivatives in MSJ wave functions even for large expansions and we refer the reader to the original publications for details. However, we do note that, despite these advances, orbital optimization remains the most challenging part of the optimization, likely due to its inherently high degree of nonlinearity and the fact that it alters the nodal surface. One potential alternative would be to obtain state-specific orbitals from another method⁷⁷ while optimizing only Slater coefficients and Jastrow factors, and we make a preliminary exploration of this idea in our results.

2.6 Variance Matching

While optimization of the parameters improves the absolute quality of the wave functions, the results are still approximate due to the limited ansatz and accurate determination of excitation energy differences requires cancellation of errors. This relies on a balanced treatment of both the ground and the excited state, which we attempt to obtain using variance matching. As shown in previous work,^{45,49} this approach improves predicted excitation energies by optimizing ansatzes of different CI expansion lengths for the two states so that their variances are approximately equal. To facilitate interpolation, the variances for a series of excited state calculations at different expansion lengths can be fit to an analytic form such

as the power law decay

$$\sigma^2(N) = c + \frac{d}{N^\alpha} \tag{21}$$

to determine parameters, c , d , and α . For a given expansion length for the ground state and a resulting variance σ_g^2 , we can estimate the expansion length N^* that will yield a matching variance for the excited state and take the corresponding energy when computing our prediction for the excitation energy.⁴⁹ In practice, some additional varying of N by hand can be performed to find an explicit variance match. All our reported excitation energies were obtained using this explicit variance matching procedure.

3 Results

3.1 Computational Details

All our VMC calculations used an implementation of the described optimization algorithms within a development version of QMCPACK.^{55,78} A recently developed set of pseudopotentials⁷⁹ and associated basis sets were used for all molecules. For constructing our ansatzes, we have employed Molpro⁸⁰ and PySCF⁸¹ for CASSCF calculations to generate Slater determinant expansions. In one case, we instead use CASSCF to provide orbitals for a selected CI calculation in Dice.^{6,7} CASPT2 calculations in Molpro⁸⁰ were used alongside other methods' literature values in benchmarking our VMC results. Specific active space and basis set choices are given in each system's section. Molecular geometries, absolute energies for our doubly excited state calculations, and additional optimization details can be found in the supplementary material.




Figure 1: Examples of unstable and stable LM optimizations of Ω for CN5. We consider the singlet near $-40.5 E_h$, which is dominated by a HOMO to LUMO excitation, as well as the next triplet and singlet above it, which are dominated by a HOMO-2 to LUMO excitation. Initial wave function guesses were obtained from a diagonalization of the Hamiltonian in a 26 determinant space with pre-optimized one- and two-body Jastrow factors present. These Jastrows were optimized further along with the 26 determinant coefficients. For each state, the value of ω was set to the $E-\sigma$ value from a one million sample evaluation using the initial wave function and held fixed throughout the optimization. The LM shifts were kept constant at 0.1 for the diagonal shift and 1 for the overlap-based shift throughout the optimizations with 50,000 samples per iteration in all cases. For the unstable cases, we found that using increased sampling effort alone did not ensure stability. A proposed parameter update was rejected if our correlated sampling assessment predicted it would raise the target function value, but all optimizations contain hundreds of accepted steps. Horizontal lines show the lowest 7 eigenenergies from our diagonalization, with solid lines for singlet states and dashed for triplets.

3.2 Stability in a Model Cyanine Dye

To explore the stability of Ω -based variance minimization, we have performed a series of LM tests on the model cyanine dye $C_3H_3(NH_2)_2^+$ (denoted hereafter as CN5) in which Filippi and coworkers discovered variance optimization instabilities for some choices of the trial function.³² In particular, they showed that while a CAS(6e,5o) wave function was stable,

other ansatzes from CAS(6e,10o) and CIPSI were not. They also showed that a small CSF expansion, constructed of HOMO-to-virtual excitations of B_1 symmetry, was especially prone to optimize to a different state than the one targeted by the initial guess due to the absence of variance minima in this simpler case. Here we perform similar tests to study the absence of variance minima, confirming that instabilities are present when the sample size is small and the trial function is simple, but also demonstrating that improvements in the trial function and sampling effort can overcome these instabilities. To start, we note that virtual orbitals outside the CASSCF active space are often not physical in their shapes, and so for the states we target, we make sure that the primary orbitals involved are within the (6e,5o) CASSCF active space. In this active space, we performed an equal-weight state-averaged CASSCF optimization of the lowest four B_1 singlet states using an aug-cc-pVDZ basis set in Molpro,⁸⁰ after which we imported all determinants with weights above 0.05 in any state into our VMC ansatz. This resulted in a 26-determinant ansatz, or, in the cases where we enforced singlet spin symmetry in VMC, a 13-CSF ansatz. Note that this trial function differs from the one used previously,³² which is intentional, as the previous approach examined some states whose dominant orbitals were virtual in the quantum chemistry calculations and so may not have been as well optimized as orbitals containing electrons. Here, we try to ensure that orbital quality is balanced between states by ensuring that the dominant orbitals of the states we test are within our CASSCF active space. Nonetheless, we can still find optimization instabilities when using this simple ansatz, which confirms that ansatz components beyond a small determinant expansion can be necessary to achieve stable variance optimization.

As seen in Figure 1, the stability of an optimization using our 26-determinant ansatz depends on which state is being targeted and in some cases on whether or not spin symmetry is enforced. First, consider the optimization that guesses the singlet near $-40.5 E_h$ and then collapses to the triplet below it. This failure shows that, at least at this level of statistical resolution, the singlet in question lacks a local variance minimum in the variable space of the 26 different determinant coefficients. However, after enforcing spin symmetry by instead




Figure 2: Fractions of the CI expansion corresponding to the third singlet’s dominant CSF over the course of LM optimizations. When starting from the energy eigenstate with only 1- and 2-body Jastrows are present, the weight of this CSF quickly collapses as the optimization drifts to a higher singlet. This optimization used 50,000 samples per iteration and starting from solely the dominant CSF with an additional 3-body Jastrow delayed, but did not prevent, the drift at this sampling effort. Upon increasing the sample size to 500,000 samples per iteration, the 3-body Jastrow wave function optimization became stable, as shown by the solid line. Note that the CI expansion fraction is simply the sum of the squares of the determinant coefficients within the given CSF divided by the sum of all the squared determinant coefficients.

optimizing the 13 coefficients for the singlet CSFs, the variance minimum for this state reappears. Next, we note that when we optimize the 26 determinant coefficients starting with a guess for the triplet just below $-40.4 E_h$, the optimization is stable. Thus, in these cases, moving to a more sophisticated ansatz is not necessary so long as symmetries are enforced. The same cannot be said for the most difficult case we consider, in which a guess for the singlet just above $-40.35 E_h$ is unstable and optimizes to a higher singlet even when spin symmetry is enforced. With no more symmetries to make use of, we must conclude that variance minimization is not stable for this state when using this simple ansatz.

In principle, variance minimization should become stable as the ansatz is improved and

the sample size increased. The question, of course, is whether stability can be achieved in practice, especially given that overly aggressive embellishments of the ansatz may lead to an impossibly difficult optimization. First, let us improve the ansatz by adding the 3-body Jastrow factor introduced by Needs and coworkers,⁶² which we first optimize for the guessed CSF expansion with the CSF coefficients held fixed before finally optimizing all variables together. As seen in Figure 2, which shows the key CSF weights during the final optimization in which all parameters are free to vary, this ansatz improvement does lead to a stable optimization, but only when a larger sample size is used, suggesting that the variance minimum is now present but shallow. Given how close in energy this state is to the singlet state above it (about 15 mE_h), a shallow minimum makes sense, which is a good reminder that optimization objective functions based on the energy and variance are at a disadvantage when states are nearly degenerate. For low-lying states, state-averaged energy minimization can often deal with this type of situation,³² but near-degeneracies remain a challenge for higher-lying states or in situations where state-averaging introduces its own challenges. In future, it may therefore be worth considering ways to involve other properties in the objective function, possibly via a VMC analogue of a generalized variational principle.⁸² Even with this near-degeneracy, which we stress involves states significantly closer together in energy than any of those tested in the small-determinant-expansion case in the previous study,³² we do see that a modest improvement in ansatz quality produces a stable Ω -based LM optimization.

While this finding is reassuring, we must stress that difficult optimization cases remain a serious challenge. For example, Filippi and coworkers found³² that highly sophisticated wave functions derived from a (6,10) active space faced optimization instabilities even though wave functions derived from a smaller (6,5) active space did not, showing that in practice it is not always easy to predict when these instabilities will arise. This in mind, we have checked carefully for signs of optimization instability in the double excitations we now turn to, and while we did not observe any such issues in these states, finding a more complete

resolution to this problem is clearly an important direction for future research.

3.3 Carbon Dimer

For our doubly excited state applications, we first consider the carbon dimer, a very heavily studied system for the testing and development of theoretical methods.^{2,11,83,84} We use the $2\ ^1\Sigma_g$ state, which is characterized by a HOMO to LUMO squared double excitation, as a simple starting test case for validating the hybrid method’s results against the LM’s and assessing the accuracy of VMC. In this case, selected CI methods are able to achieve millions of determinants and provide high quality benchmark data^{11,85} on the vertical excitation energy that we can use to assess our results with more compact wave functions. We consider only an equilibrium bond length of 1.248 Å.

To construct our ansatzes, we use a (8e,8o) CASSCF calculation in Molpro⁸⁰ with a carbon pseudopotential and the corresponding cc-pVTZ basis.⁷⁹ The resulting CI expansions are used in our variance matching procedure and we consider cases both for standard MSJ ansatzes using only one- and two-body Jastrows, and with an additional NCJF and orbital optimization. The NCJF was produced using a set of 16 counting regions composed of 8 octants for each carbon atom.⁶⁴




Figure 3: Excitation energy of the $2\ ^1\Sigma_g$ state in C_2 for LM and hybrid method on both MSJ and all parameter ansatzes. See also Table 1. 20 determinants were used for all ground state optimizations. The excited state optimizations used 23 determinants for both the LM and Hybrid MSJ cases, 71 for the LM All case, and 50 for Hybrid All case. The benchmark value is taken from a selected CI calculation using the CIPSI algorithm.¹¹

Table 1: Excitation energies and uncertainties for the $2\ ^1\Sigma_g$ state in C_2 . Literature values for selected CI, CCSDT, CC3, and CASPT2 using an aug-cc-pVQZ basis set are included for comparison.¹¹

Method	Excitation Energy (eV)	Uncertainty (eV)	Total Samples
LM MSJ VMC	2.34	0.04	100,000,000
LM All VMC	2.36	0.03	100,000,000
Hybrid MSJ VMC	2.32	0.02	138,000,000
Hybrid All VMC	2.34	0.01	138,000,000
Selected CI	2.40		
CCSDT	2.87		
CC3	3.24		
CASPT2	2.50		

Figure 3 shows the predicted excitation energy achieved by the LM and the hybrid method on both the simpler MSJ ansatzes and on all parameter wave functions that include the NCJF and orbital rotations. It is reassuring to find that the hybrid method’s results agree with the LM’s to within statistical uncertainty. In terms of accuracy, the VMC results are within about 0.05 eV of the selected CI value with the all parameter optimizations offering some improvement over the MSJ results. This has been achieved with very modest CI expansions, using less than 100 determinants in all cases, compared to the 5 million used to produce the benchmark energy.¹¹




Figure 4: Example optimization of target function Ω for all parameters using the hybrid method. Points for AD correspond to an average over 50 iterations while those for the Blocked LM are individual iterations.

We also explore the potential benefits of the hybrid approach over a pure LM optimization. Figure 4 shows an example optimization by the hybrid method of all parameters including NCJF and orbitals. With these more challenging parameters, the Blocked LM is more prone to step uncertainty and upward fluctuations of the target function, which in this case we see the AD sections correct. The hybrid method also exhibits a statistical advantage

over the LM based on the uncertainty and sampling data in Table 1. For slightly greater sampling effort, the hybrid method provides about a factor of 3 improvement in the uncertainty on the excitation energy, which would require a factor of 9 more samples to obtain solely with the LM. This advantage in computational efficiency persists in our results for more difficult systems.

3.4 Nitroxyl, Glyoxal, and Acrolein

One nuance to the study of doubly excited states is distinguishing between different types of such states. For instance, the state we have considered in the carbon dimer is multi-determinantal while some other molecules' states can be viewed as single reference, enabling higher order coupled cluster¹¹ or orbital optimized DFT⁸⁶ to obtain accurate excitation energies. The work of Loos and coworkers has categorized a set of doubly excited states using the percentage of singles amplitudes in CC3¹¹ and we briefly consider several of the same systems to compare our methodology against their benchmark results. Specifically, we consider the $2^1A'$ state in nitroxyl, the 2^1A_g state in glyoxal, and the $3^1A'$ state of acrolein. These systems exhibit some of the diversity of doubly excited states, with acrolein's state containing a high percentage of single excitations, while nitroxyl and glyoxal have almost none. For all systems, we use cc-pVTZ basis sets with pseudopotentials⁷⁹ and generate determinants for MSJ ansatzes from CASSCF calculations in Molpro.⁸⁰ Our active spaces are (12e,9o) for nitroxyl, (8e,6o) for glyoxal, and (10e,10o) for acrolein. Each CASSCF calculation was state-averaged over four singlet states.

The excitation energy predictions from the LM and the hybrid method are shown in Figures 5 through 7 with precise values given in Tables 2 through 4. We compare our results to the theoretical best estimates (TBE) from Loos and coworkers as well as their coupled cluster and CASPT2 values.¹¹ For this trio of systems, we find that we can obtain good accuracy with very modest wave functions and consistent results between the LM and hybrid method. For nitroxyl, we come within 0.03 eV of the TBE, while our excitation energies in

glyoxal and acrolein are within about 0.2 eV. This level of accuracy outperforms the coupled cluster approaches as well as some versions of CASPT2 in the case of acrolein. As in the carbon dimer, these calculations use less than 100 determinants in all cases compared to the millions used in the benchmark selected CI calculations, which are restricted to smaller basis sets in glyoxal and acrolein. We note that our methodology allows for further systematic improvement through more sophisticated ansatzes and in order to consider its performance on larger systems beyond the reach of selected CI, we now turn to some molecules outside the benchmark set of Loos and coworkers.

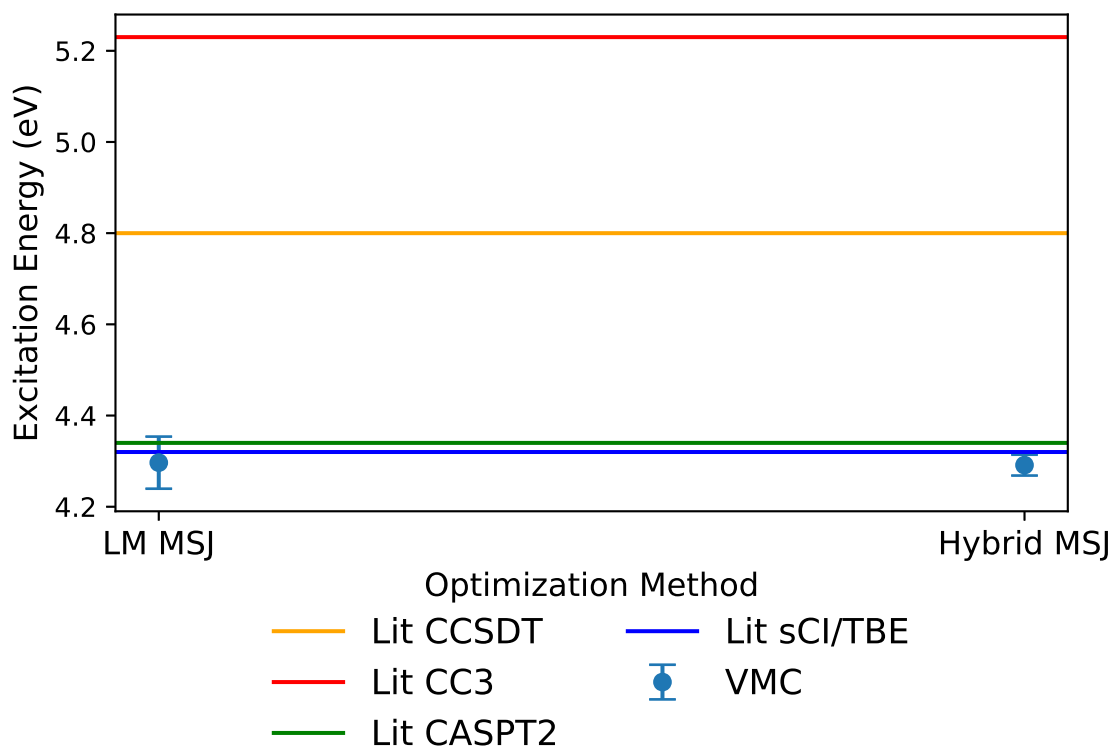


Figure 5: Excitation energy of the 2 ¹A' state in nitroxyl for LM and hybrid method on MSJ ansatzes. See also Table 2. Both methods used 20 determinants the ground state and 40 for the excited state. Reference values for coupled cluster, CASPT2, and selected CI are taken from the work of Loos and coworkers.¹¹ They take the selected CI value as the theoretical best estimate.

Table 2: Excitation energies and uncertainties for the $2^1A'$ state in nitroxyl. The included literature values for CCSDT, CC3, CASPT2, and selected CI all use an aug-cc-pVQZ basis set.¹¹

Method	Excitation Energy (eV)	Uncertainty (eV)	Total Samples
LM MSJ VMC	4.30	0.06	100,000,000
Hybrid MSJ VMC	4.29	0.01	138,000,000
TBE/Selected CI	4.32		
CCSDT	4.8		
CC3	5.23		
CASPT2	4.34		




Figure 6: Excitation energy of the 2^1A_g state in glyoxal for LM and hybrid method on MSJ ansatzes. See also Table 3. Both methods used 50 determinants the ground state and 20 for the excited state. Reference values for coupled cluster, CASPT2, and selected CI are taken from the work of Loos and coworkers.¹¹ Their theoretical best estimate is obtained by adding a 0.06 eV basis set correction to the selected CI result.

Table 3: Excitation energies and uncertainties for the 2^1A_g state in glyoxal. The included literature values use an aug-cc-pVDZ basis set for selected CI and aug-cc-pVQZ for CCSDT, CC3, and CASPT2.¹¹

Method	Excitation Energy (eV)	Uncertainty (eV)	Total Samples
LM MSJ VMC	5.90	0.09	140,000,000
Hybrid MSJ VMC	5.75	0.03	138,000,000
TBE	5.54		
Selected CI	5.48		
CCSDT	6.35		
CC3	6.76		
CASPT2	5.43		

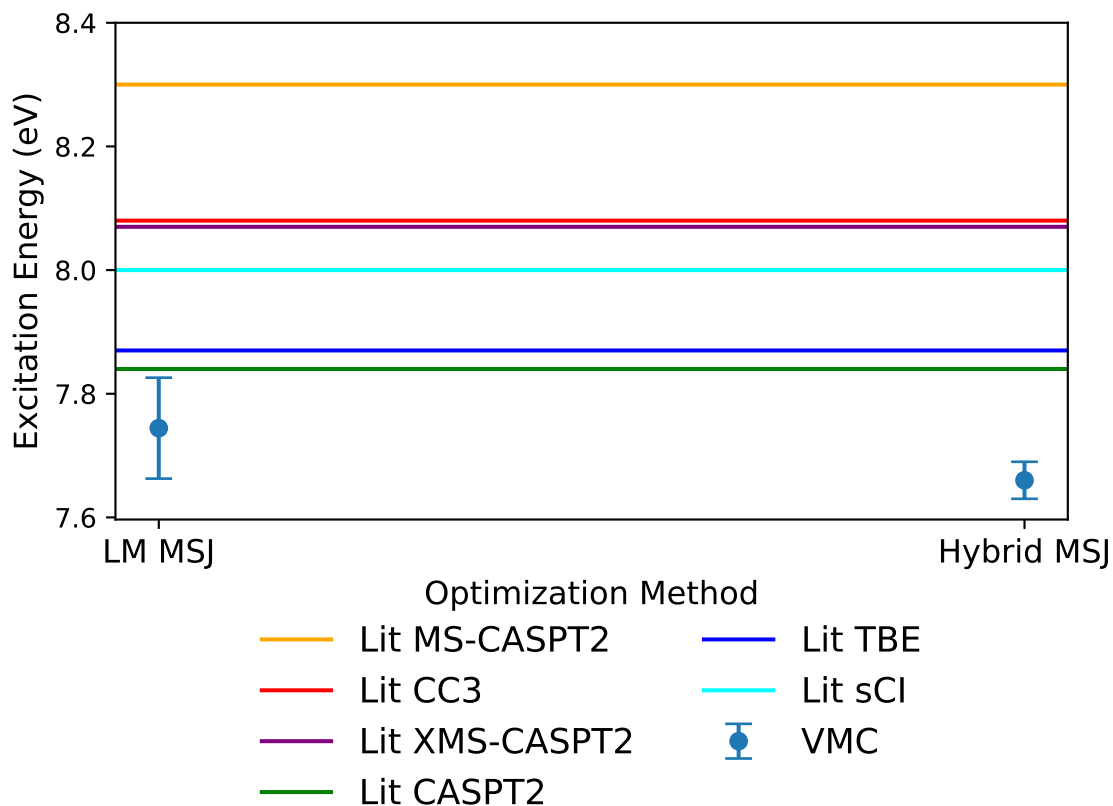


Figure 7: Excitation energy of the $3^1A'$ state in acrolein for LM and hybrid method on MSJ ansatzes. See also Table 4. Both methods used 20 determinants the ground state and 70 for the excited state. Reference values for coupled cluster, CASPT2, and selected CI are taken from the work of Loos and coworkers.¹¹ Their theoretical best estimate is obtained by adding a -0.13 eV basis set correction to the selected CI result.

Table 4: Excitation energies and uncertainties for the $3^1A'$ state in acrolein. The included literature values use a 6-31+G(d) basis set for selected CI, an aug-cc-pVTZ basis CC3, and an aug-cc-pVQZ basis for the three versions of CASPT2.¹¹

Method	Excitation Energy (eV)	Uncertainty (eV)	Total Samples
LM MSJ VMC	7.74	0.08	140,000,000
Hybrid MSJ VMC	7.66	0.03	138,000,000
TBE	7.87		
Selected CI	8.00		
CC3	8.08		
CASPT2	7.84		
MS-CASPT2	8.3		
XMS-CASPT2	7.84		

3.5 Cyclopentadiene

For a more challenging test of our methodology, we consider the doubly excited 3^1A_1 state of cyclopentadiene (CPD). This state has been repeatedly studied in theoretical benchmark investigations^{22,23,87-90} and in some experimental investigations.^{91,92} As before, we construct multi-Slater wave functions and add traditional 1 and 2-body Jastrow factors. For this larger molecule, we use the heatbath selected CI (HCI) method^{6,7} in the Dice code to produce our CI expansions by correlating 26 electrons in the lowest 46 orbitals from a (6e,5o) CASSCF in Molpro with pseudopotentials and cc-pVTZ basis sets.⁷⁹

In this system, we find that the LM fails to optimize the Ω functional as well as the hybrid method and leads to an inferior energy prediction. However, for relatively simple variance-matched multi-Slater Jastrow ansatzes of 20 and 500 determinants for the ground and excited states respectively, the hybrid method is able to achieve an excitation energy within about 0.1 to 0.2 eV of CASPT2 as seen in Figure 8 and Table 5.




Figure 8: Excitation energy of 3^1A_1 state in CPD for LM and hybrid method. See also Table 5. Both the LM and hybrid optimizations used 20 and 500 determinants for the ground and excited states, respectively. Literature values are taken from Schreiber et al.²² for CASPT2 and CCSD LR and from Shen and Li⁸⁹ for MR-CISD+Q.

Table 5: Excitation Energies and uncertainties for 3^1A_1 state in CPD.

Method	Excitation Energy (eV)	Uncertainty (eV)	Total Samples
LM MSJ VMC	7.59	0.06	140,000,000
Hybrid MSJ VMC	8.27	0.03	210,000,000
CASSCF	10.29		
CASPT2	8.39		
Lit. CASPT2 ²²	8.52		
Lit. CCSD LR ²²	8.95		
Lit. MR-CISD+Q ⁸⁹	9.02		

To elaborate on the failure of the LM in this case, we present data in Table 6 on a head to head comparison of the LM and the hybrid method on the same excited state wave function

with a determinant expansion length of 50 and using the same fixed value of ω in the target function. We find that the hybrid method achieves a lower value of the target function, which corresponds to a significantly higher energy for the excited state. This results in its substantially better prediction for the excitation energy as both optimization methods give comparable results for the ground state calculation. We note that extensive experimentation with the technical details of the LM, including choices for the shifts, sampling effort, and guiding function failed to bring it into agreement with the hybrid method. When we started the LM with the hybrid method’s *optimized* wave function, it remained at roughly those parameter values with the same target function value and energy as found by the hybrid method. This test indicates that the LM agrees that the hybrid method has found an optimal location in parameter space if it starts close enough, but is apparently unable to find it itself when starting from the unoptimized wave function.

Table 6: Head to head comparison of LM and hybrid for CPD on the same excited state wave function at fixed ω .

Method	Energy (a.u.)	Energy Uncertainty (a.u.)	Target Function $\Omega(\Psi)$ (a.u.)	Target Uncertainty (a.u.)
LM MSJ VMC	-31.541	0.0013	-0.747	0.0023
Hybrid MSJ VMC	-31.517	0.0007	-0.752	0.0004

3.6 4-Aminobenzonitrile

Our final system, 4-aminobenzonitrile (ABN), has been heavily studied as an example of intramolecular charge transfer (ICT) with many attempts to determine the geometry of the ICT state.^{93–96} Here we test our ability to treat a doubly excited state at the ICT geometry. We selected this system for an initial exploration of possible benefits of using state-specific orbitals within VMC while forgoing orbital optimization. These orbitals were obtained from a recent state-specific CASSCF approach^{77,96} that employs a root tracker based on combination of an excited state variational principle and density matrices. For a twisted geometry of ABN⁹⁶ (coordinates in supplementary material), we construct Multi-Slater wave functions from both state-averaged (over four states) and state-specific CASSCF calculations that use

a (12e,11o) active space along with pseudopotentials and cc-pVDZ basis sets.⁷⁹ Both types of CASSCF calculations were performed in a development version of PySCF.⁸¹ The excited state with double excitation character that we consider appears as the fourth CASSCF state in energy, directly above the ICT state.

Figure 9 shows the excitation energies obtained by the LM and hybrid method for the cases where we use state-averaged and state-specific CASSCF orbitals. In this instance, we find that the optimization methods agree with each other and that there is no clear difference between using state-averaged and state-specific orbitals within our VMC ansatzes. There is about a 0.4 eV difference between our VMC results and CASPT2, but in the absence of an experimental result or higher level benchmark, it is not obvious which is more accurate. The agreement between the state-averaged and state-specific VMC excitation energies may not be too surprising given that we also find little difference at the CASSCF level. While this is a null result for the usefulness of state-specific orbitals for this state in ABN, other cases may perform differently, including the ICT state. In terms of optimization, the across the board agreement in ABN offers further evidence that the hybrid method is at least as accurate as the LM, while continuing to provide better statistical efficiency.




Figure 9: Excitation energy of the doubly excited state in ABN for LM and hybrid method. See also Table 7. Excitation energies for all other methods were obtained in Molpro. 20 determinants were used in all ground state optimizations. The excited state optimizations used 70, 100, 500, and 300 determinants for the LM SA, Hybrid SA, LM SS, and Hybrid SS cases respectively. These numbers of determinants were chosen to achieve an explicit variance match between the ground and excited states for each method.

Table 7: Excitation Energies and uncertainties for ABN.

Method	Excitation Energy (eV)	Uncertainty (eV)	Total Samples
LM SA MSJ VMC	6.73	0.10	140,000,000
Hybrid SA MSJ VMC	6.58	0.04	174,000,000
LM SS MSJ VMC	6.78	0.08	140,000,000
Hybrid SS MSJ VMC	6.61	0.04	192,000,000
SA CASSCF	6.10	-	-
SS CASSCF	6.15	-	-
SA CASPT2	6.23	-	-
EOM-CCSD	7.16	-	-

4 Conclusion

We have presented an extension of a hybrid LM/AD optimization approach to the case of excited-state-specific variance optimization and tested its efficacy for doubly excited states. As in energy minimization, we find the hybrid method to be more statistically efficient than the linear method, and, in one case, we were surprised to find it to be more effective at finding the variance minimum. Thanks to VMC’s ability to combine a linear combination of determinants (for capturing strong correlation) with sophisticated correlation factors (for weak correlation) and its ability to explicitly balance wave function quality between different states, we find it to be highly accurate compared to theoretical benchmarks in our tests on doubly excited states. As it relies on far more modest determinant expansions than sCI methods, it can also be used to treat both strong and weak correlation in system sizes where capturing both through sCI is not currently possible. We have also performed some simple tests on the stability of state-specific variance minimization and found that using symmetry, increasing sample sizes, and improving the quality of the wave function approximation all play important roles in preventing collapse to other states.

Looking forward, there are multiple avenues for further improvement. Our optimization stability testing was limited to a relatively simple example, and while it does suggest steps that can be taken to alleviate stability concerns, it remains to be seen how effective these steps will be in general. The demonstration by Filippi and coworkers³² that optimization instabilities can persist even when highly sophisticated CIPSI-based VMC trial functions are employed makes clear that there is more work to be done to resolve this issue. As variance-based approaches are particularly ill-suited to degenerate or near-degenerate states, it would be quite interesting to explore whether generalized variational principles that incorporate properties beyond the energy can be usefully adapted for VMC optimization. Another priority is improving user-accessibility, as the recent improvements in VMC optimization methodology have in many cases brought with them a significant increase in the methodological complexity. Finding ways to robustly automate choices for stability shifts,

what balance to strike between linear method and accelerated descent steps, what hyperparameters to choose for the descent, and how to arrange variables into blocks within the blocked linear method would significantly simplify the practical application of these tools. The systematic study of effective importance sampling functions is another supporting step for excited state optimizations, where their use is more crucial than in the ground state. A third improvement would be to automatically stage the optimization of different parameters according to the statistical significance of their gradients, which would allow the noisiest and most difficult parameters to be handled at the end of the optimization without any direct human decision-making. Finally, although VMC optimization is becoming increasingly capable, it will likely be profitable to map out areas where difficult parameters like orbital shapes can be safely kept at their quantum chemistry values, whether from state-averaged or state-specific CASSCF. Given the high accuracy that VMC can offer for very challenging excited states, providing easy-to-use incarnations of the best available VMC optimization methods is a high priority.

Supplementary Material

See the supplementary material for molecular geometries, absolute energies, and optimization details for doubly excited states.

Data Availability

The data that support the findings of this study are available from the corresponding author upon reasonable request. Input and output files for QMCPACK are available via the Materials Data Facility with the DOI: <https://doi.org/10.18126/kduf-y6cp>

Acknowledgement

This work was supported by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, Materials Sciences and Engineering Division, as part of the Computational Materials Sciences Program and Center for Predictive Simulation of Functional Materials. Computational work was shared evenly between the Berkeley Research Computing Savio cluster and the LBNL Lawrence cluster.

References

- (1) Booth, G. H.; Thom, A. J. W.; Alavi, A. Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space. *The Journal of Chemical Physics* **2009**, *131*, 054106.
- (2) Blunt, N. S.; Smart, S. D.; Booth, G. H.; Alavi, A. An excited-state approach within full configuration interaction quantum Monte Carlo. *J. Chem. Phys.* **2015**, *143*, 134117.
- (3) White, S. R.; Martin, R. L. Ab initio quantum chemistry using the density matrix renormalization group. *The Journal of Chemical Physics* **1999**, *110*, 4127–4130.
- (4) Chan, G. K.-L.; Head-Gordon, M. Highly correlated calculations with a polynomial cost algorithm: A study of the density matrix renormalization group. *The Journal of Chemical Physics* **2002**, *116*, 4462–4476.
- (5) Kurashige, Y.; Yanai, T. High-performance ab initio density matrix renormalization group method: Applicability to large-scale multireference problems for metal compounds. *The Journal of Chemical Physics* **2009**, *130*, 234114.
- (6) Holmes, A. A.; Tubman, N. M.; Umrigar, C. J. Heat-Bath Configuration Interaction: An Efficient Selected Configuration Interaction Algorithm Inspired by Heat-Bath Sampling. *J. Chem. Theory Comput.* **2016**, *12*, 3674–3680.

- (7) Sharma, S.; Holmes, A. A.; Jeanmairet, G.; Alavi, A.; Umrigar, C. J. Semistochastic Heat-Bath Configuration Interaction Method: Selected Configuration Interaction with Semistochastic Perturbation Theory. *J. Chem. Theory Comput.* **2017**, *13*, 1595–1604.
- (8) Abrams, M. L.; Sherrill, C. D. Important configurations in configuration interaction and coupled-cluster wave functions. *Chemical Physics Letters* **2005**, *412*, 121 – 124.
- (9) Loos, P.-F.; Scemama, A.; Blondel, A.; Garniron, Y.; Caffarel, M.; Jacquemin, D. A Mountaineering Strategy to Excited States: Highly Accurate Reference Energies and Benchmarks. *Journal of Chemical Theory and Computation* **2018**, *14*, 4360–4379, PMID: 29966098.
- (10) Garniron, Y.; Scemama, A.; Giner, E.; Caffarel, M.; Loos, P.-F. Selected configuration interaction dressed by perturbation. *The Journal of Chemical Physics* **2018**, *149*, 064103.
- (11) Loos, P.-F.; Boggio-Pasqua, M.; Scemama, A.; Caffarel, M.; Jacquemin, D. Reference Energies for Double Excitations. *J. Chem. Theory Comput.* **2019**, *15*, 1939–1956.
- (12) Levine, B. G.; Ko, C.; Quenneville, J.; Martínez, T. J. Conical intersections and double excitations in time-dependent density functional theory. *Molecular Physics* **2006**, *104*, 1039–1051.
- (13) Hirata, S.; Nooijen, M.; Bartlett, R. J. High-order determinantal equation-of-motion coupled-cluster calculations for electronic excited states. *Chemical Physics Letters* **2000**, *326*, 255 – 262.
- (14) Watson, M. A.; Chan, G. K.-L. Excited States of Butadiene to Chemical Accuracy: Reconciling Theory and Experiment. *J. Chem. Theory Comput.* **2012**, *8*, 4013–4018, PMID: 26605568.

- (15) Andersson, K.; Malmqvist, P. A.; Roos, B. O.; Sadlej, A. J.; Wolinski, K. Second-order perturbation theory with a CASSCF reference function. *The Journal of Physical Chemistry* **1990**, *94*, 5483–5488.
- (16) Andersson, K.; Malmqvist, P.; Roos, B. O. Second-order perturbation theory with a complete active space self-consistent field reference function. *The Journal of Chemical Physics* **1992**, *96*, 1218–1226.
- (17) Serrano-Andrés, L.; Merchán, M.; Nebot-Gil, I.; Lindh, R.; Roos, B. O. Towards an accurate molecular orbital theory for excited states: Ethene, butadiene, and hexatriene. *The Journal of Chemical Physics* **1993**, *98*, 3151–3162.
- (18) Serrano-Andrés, L.; Merchán, M.; Fülcher, M.; Roos, B. O. A theoretical study of the electronic spectrum of thiophene. *Chemical Physics Letters* **1993**, *211*, 125 – 134.
- (19) Nakayama, K.; Nakano, H.; Hirao, K. Theoretical study of the $\pi \rightarrow \pi^*$ excited states of linear polyenes: The energy gap between 1 1Bu+ and 2 1Ag states and their character. *International Journal of Quantum Chemistry* **1998**, *66*, 157–175.
- (20) Ostojić, B.; Domcke, W. Ab initio investigation of the potential energy surfaces involved in the photophysics of s-trans-1,3-butadiene. *Chemical Physics* **2001**, *269*, 1 – 10.
- (21) Dallos, M.; Lischka, H. A systematic theoretical investigation of the lowest valence- and Rydberg-excited singlet states of trans-butadiene. The character of the 11Bu (V) state revisited. *Theoretical Chemistry Accounts* **2004**, *112*, 16–26.
- (22) Schreiber, M.; Silva-Junior, M. R.; Sauer, S. P. A.; Thiel, W. Benchmarks for electronically excited states: CASPT2, CC2, CCSD, and CC3. *J. Chem. Phys.* **2008**, *128*, 134110.
- (23) Silva-Junior, M. R.; Schreiber, M.; Sauer, S. P. A.; Thiel, W. Benchmarks of electron-

- ically excited states: Basis set effects on CASPT2 results. *J. Chem. Phys.* **2010**, *133*, 174318.
- (24) Duman, S.; Cakmak, Y.; Kolemen, S.; Akkaya, E. U.; Dede, Y. Heavy Atom Free Singlet Oxygen Generation: Doubly Substituted Configurations Dominate S1 States of Bis-BODIPYs. *The Journal of Organic Chemistry* **2012**, *77*, 4516–4527, PMID: 22530939.
- (25) Wen, J.; Han, B.; Havlas, Z.; Michl, J. An MS-CASPT2 Calculation of the Excited Electronic States of an Axial Difluoroborondipyrromethene (BODIPY) Dimer. *Journal of Chemical Theory and Computation* **2018**, *14*, 4291–4297, PMID: 29874458.
- (26) Rauer, C.; Nogueira, J. J.; Marquetand, P.; González, L. Cyclobutane Thymine Photodimerization Mechanism Revealed by Nonadiabatic Molecular Dynamics. *Journal of the American Chemical Society* **2016**, *138*, 15911–15916, PMID: 27682199.
- (27) Ben Amor, N.; Soupart, A.; Heitz, M. C. Methodological CASPT2 study of the valence excited states of an iron-porphyrin complex. *Journal of Molecular Modeling* **2017**, *23*.
- (28) Loos, P.-F.; Lipparini, F.; Boggio-Pasqua, M.; Scemama, A.; Jacquemin, D. A Mountaineering Strategy to Excited States: Highly Accurate Energies and Benchmarks for Medium Sized Molecules. *Journal of Chemical Theory and Computation* **2020**, *16*, 1711–1741, PMID: 31986042.
- (29) Cordova, F.; Doriol, L. J.; Ipatov, A.; Casida, M. E.; Filippi, C.; Vela, A. Troubleshooting time-dependent density-functional theory for photochemical applications: Oxirane. *J. Chem. Phys.* **2007**, *127*, 164111.
- (30) Filippi, C.; Zaccheddu, M.; Buda, F. Absorption Spectrum of the Green Fluorescent Protein Chromophore: A Difficult Case for ab Initio Methods? *J. Chem. Theory Comput.* **2009**, *5*, 2074–2087.

- (31) Guareschi, R.; Zulfikri, H.; Daday, C.; Floris, F. M.; Amovilli, C.; Mennucci, B.; Filippi, C. Introducing QMC/MMpol: Quantum Monte Carlo in polarizable force fields for excited states. *J. Chem. Theory Comput.* **2016**, *12*, 1674–1683.
- (32) Cuzzocrea, A.; Scemama, A.; Briels, W. J.; Moroni, S.; Filippi, C. Variational Principles in Quantum Monte Carlo: The Troubled Story of Variance Minimization. *J. Chem. Theory Comput.* **2020**, *16*, 4203–4212, PMID: 32419451.
- (33) Feldt, J.; Filippi, C. Excited-state calculations with quantum Monte Carlo. *arXiv preprint arXiv:2002.03622* **2020**,
- (34) Umrigar, C. J.; Wilson, K. G.; Wilkins, J. W. Optimized trial wave functions for quantum Monte Carlo calculations. *Phys. Rev. Lett.* **1988**, *60*, 1719–1722.
- (35) Zhao, L.; Neuscamman, E. An Efficient Variational Principle for the Direct Optimization of Excited States. *J. Chem. Theory Comput.* **2016**, *12*, 3436–3440.
- (36) Shea, J. A. R.; Neuscamman, E. Size Consistent Excited States via Algorithmic Transformations between Variational Principles. *J. Chem. Theory Comput.* **2017**, *13*, 6078–6088, PMID: 29140699.
- (37) Umrigar, C. J.; Toulouse, J.; Filippi, C.; Sorella, S.; Hennig, R. G. Alleviation of the fermion-sign problem by optimization of many-body wave functions. *Phys. Rev. Lett.* **2007**, *98*, 110201.
- (38) Sorella, S.; Casula, M.; Rocca, D. Weak binding between two aromatic rings: Feeling the van der Waals attraction by quantum Monte Carlo methods. *J. Chem. Phys.* **2007**, *127*, 014105.
- (39) Neuscamman, E.; Umrigar, C. J.; Chan, G. K. L. Optimizing large parameter sets in variational quantum Monte Carlo. *Phys. Rev. B* **2012**, *85*, 045103.

- (40) Zhao, L.; Neuscamman, E. A Blocked Linear Method for Optimizing Large Parameter Sets in Variational Monte Carlo. *J. Chem. Theory Comput.* **2017**, *13*, 2604–2611.
- (41) Schwarz, L. R.; Alavi, A.; Booth, G. H. Projector Quantum Monte Carlo Method for Nonlinear Wave Functions. *Phys. Rev. Lett.* **2017**, *118*, 176403.
- (42) Sabzevari, I.; Sharma, S. Improved Speed and Scaling in Orbital Space Variational Monte. *J. Chem. Theory Comput.* **2018**, *14*.
- (43) Otis, L.; Neuscamman, E. Complementary first and second derivative methods for ansatz optimization in variational Monte Carlo. *Phys. Chem. Chem. Phys.* **2019**, *21*, 14491–14510.
- (44) Garner, S. M.; Neuscamman, E. A variational Monte Carlo approach for core excitations. *arXiv preprint arXiv:2007.00850* **2020**,
- (45) Flores, S. D. P.; Neuscamman, E. Excited State Specific Multi-Slater Jastrow Wave Functions. *J. Phys. Chem. A* **2019**, *123*, 1487–1497.
- (46) Assaraf, R.; Caffarel, M. Zero-variance principle for monte carlo algorithms. *Phys. Rev. Lett.* **1999**, *83*, 4682–4685.
- (47) Trail, J. R. Heavy-tailed random error in quantum Monte Carlo. *Phys. Rev. E* **2008**, *77*, 016703.
- (48) Trail, J. R. Alternative sampling for variational quantum Monte Carlo. *Phys. Rev. E* **2008**, *77*, 016704.
- (49) Robinson, P. J.; Pineda Flores, S. D.; Neuscamman, E. Excitation variance matching with limited configuration interaction expansions in variational Monte Carlo. *J. Chem. Phys.* **2017**, *147*, 164114.
- (50) Umrigar, C. J.; Nightingale, M. P.; Runge, K. J. A diffusion Monte Carlo algorithm with very small time-step errors. *J. Chem. Phys.* **1993**, *99*, 2865–2890.

- (51) Pfau, D.; Spencer, J. S.; de G. Matthews, A. G.; Foulkes, W. M. C. Ab-Initio Solution of the Many-Electron Schrödinger Equation with Deep Neural Networks. *arXiv preprint arXiv:1909.02487* **2019**,
- (52) Nightingale, M. P.; Melik-Alaverdian, V. Optimization of ground- and excited-state wave functions and van der waals clusters. *Phys. Rev. Lett.* **2001**, *87*, 43401.
- (53) Toulouse, J.; Umrigar, C. J. Optimization of quantum Monte Carlo wave functions by energy minimization. *The Journal of Chemical Physics* **2007**, *126*, 084102.
- (54) Toulouse, J.; Umrigar, C. J. Full optimization of Jastrow–Slater wave functions with application to the first-row atoms and homonuclear diatomic molecules. *The Journal of Chemical Physics* **2008**, *128*, 174101.
- (55) Kim, J.; Baczewski, A.; Beaudet, T. D.; Benali, A.; Bennett, M. C.; Berrill, M. A.; Blunt, N. S.; Borda, E. J. L.; Casula, M.; Ceperley, D. M.; Clark, B. K.; Iii, R. C. C.; Delaney, K. T.; Dewing, M.; Esler, K. P.; Hao, H.; Heinonen, O.; Kent, P. R. C.; Kroegel, J. T.; Kylanpaa, I.; Li, Y. W.; Lopez, M. G.; Luo, Y.; Malone, F. D.; Martin, R. M.; Mathuriya, A.; Mcminis, J.; Melton, C. A.; Mitas, L.; Morales, M. A.; Neuscamman, E.; Parker, W. D.; Flores, S. D. P.; Romero, N. A.; Rubenstein, B. M.; Shea, J. A. R.; Shin, H.; Shulenburger, L.; Tillack, A.; Townsend, J. P.; Tubman, N. M.; Van Der Goetz, B.; Vincent, J. E.; Yang, D. C.; Yang, Y.; Zhang, S.; Zhao, L. QMCPACK : An open source ab initio Quantum Monte Carlo package for the electronic structure of atoms, molecules, and solids. *J. Phys. Condens. Matter* **2018**, *30*, 195901.
- (56) Zhao, L.; Neuscamman, E. Variational Excitations in Real Solids: Optical Gaps and Insights into Many-Body Perturbation Theory. *Phys. Rev. Lett.* **2019**, *123*, 036402.
- (57) Zimmerman, P. M.; Toulouse, J.; Zhang, Z.; Musgrave, C. B.; Umrigar, C. J. Excited states of methylene from quantum Monte Carlo. *The Journal of Chemical Physics* **2009**, *131*, 124103.

- (58) Send, R.; Valsson, O.; Filippi, C. Electronic Excitations of Simple Cyanine Dyes: Reconciling Density Functional and Wave Function Methods. *J. Chem. Theory Comput.* **2011**, *7*, 444–455, PMID: 26596164.
- (59) Luo, D.; Clark, B. K. Backflow Transformations via Neural Networks for Quantum Many-Body Wave Functions. *Phys. Rev. Lett.* **2019**, *122*, 226401.
- (60) Mahajan, A.; Sharma, S. Symmetry-Projected Jastrow Mean-Field Wave Function in Variational Monte Carlo. *J. Phys. Chem. A* **2019**, *123*, 3911–3921.
- (61) Filippi, C.; Umrigar, C. J. Multiconfiguration wave functions for quantum Monte Carlo calculations of first-row diatomic molecules. *J. Chem. Phys.* **1996**, *105*, 213–226.
- (62) Drummond, N. D.; Towler, M. D.; Needs, R. J. Jastrow correlation factor for atoms, molecules, and solids. *Phys. Rev. B* **2004**, *70*, 235119.
- (63) Goetz, B. V. D.; Neuscamman, E. Suppressing Ionic Terms with Number-Counting Jastrow Factors in Real Space. *J. Chem. Theory Comput.* **2017**, *13*, 2035–2042.
- (64) Van Der Goetz, B.; Otis, L.; Neuscamman, E. Clean and Convenient Tessellations for Number Counting Jastrow Factors. *J. Chem. Theory Comput.* **2019**, *15*, 1102–1121.
- (65) Filippi, C.; Assaraf, R.; Moroni, S. Simple formalism for efficient derivatives and multi-determinant expansions in quantum Monte Carlo. *J. Chem. Phys.* **2016**, *144*, 194105.
- (66) Assaraf, R.; Moroni, S.; Filippi, C. Optimizing the Energy with Quantum Monte Carlo: A Lower Numerical Scaling for Jastrow-Slater Expansions. *J. Chem. Theory Comput.* **2017**, *13*, 5273–5281.
- (67) Casula, M.; Sorella, S. Geminal wave functions with Jastrow correlation: A first application to atoms. *J. Chem. Phys.* **2003**, *119*, 6500–6511.
- (68) Casula, M.; Attaccalite, C.; Sorella, S. Correlated geminal wave function for molecules: An efficient resonating valence bond approach. *J. Chem. Phys.* **2004**, *121*, 7110–7126.

- (69) Marchi, M.; Azadi, S.; Casula, M.; Sorella, S. Resonating valence bond wave function with molecular orbitals: Application to first-row molecules. *J. Chem. Phys.* **2009**, *131*, 154116.
- (70) Sorensen, D. C. Newton's Method with a Model Trust Region Modification. *SIAM Journal on Numerical Analysis* **1982**, *19*, 409–426.
- (71) López Ríos, P.; Seth, P.; Drummond, N. D.; Needs, R. J. Framework for constructing generic Jastrow correlation factors. *Phys. Rev. E* **2012**, *86*, 036703.
- (72) Lüchow, A.; Sturm, A.; Schulte, C.; Haghighi Mood, K. Generic expansion of the Jastrow correlation factor in polynomials satisfying symmetry and cusp conditions. *J. Chem. Phys.* **2015**, *142*, 084111.
- (73) Huang, C. J.; Umrigar, C. J.; Nightingale, M. P. Accuracy of electronic wave functions in quantum Monte Carlo: The effect of high-order correlations. *J. Chem. Phys.* **1997**, *107*, 3007–3013.
- (74) Sterpone, F.; Spanu, L.; Ferraro, L.; Sorella, S.; Guidoni, L. Dissecting the Hydrogen Bond: A Quantum Monte Carlo Approach. *J. Chem. Theory Comput.* **2008**, *4*, 1428–1434, PMID: 26621429.
- (75) Beaudet, T. D.; Casula, M.; Kim, J.; Sorella, S.; Martin, R. M. Molecular hydrogen adsorbed on benzene: Insights from a quantum Monte Carlo study. *J. Chem. Phys.* **2008**, *129*, 164711.
- (76) Zen, A.; Luo, Y.; Mazzola, G.; Guidoni, L.; Sorella, S. Ab initio molecular dynamics simulation of liquid water by quantum Monte Carlo. *J. Chem. Phys.* **2015**, *142*, 144111.
- (77) Tran, L. N.; Shea, J. A. R.; Neuscamman, E. Tracking Excited States in Wave Function Optimization Using Density Matrices and Variational Principles. *J. Chem. Theory Comput.* **2019**, *15*, 4790–4803, PMID: 31393725.

- (78) Kent, P. R. C.; Annaberdiyev, A.; Benali, A.; Bennett, M. C.; Landinez Borda, E. J.; Doak, P.; Hao, H.; Jordan, K. D.; Krogel, J. T.; Kylänpää, I.; Lee, J.; Luo, Y.; Malone, F. D.; Melton, C. A.; Mitas, L.; Morales, M. A.; Neuscamman, E.; Reboredo, F. A.; Rubenstein, B.; Saritas, K.; Upadhyay, S.; Wang, G.; Zhang, S.; Zhao, L. QMCPACK: Advances in the development, efficiency, and application of auxiliary field and real-space variational and diffusion quantum Monte Carlo. *J. Chem. Phys.* **2020**, *152*, 174105.
- (79) Bennett, M. C.; Melton, C. A.; Annaberdiyev, A.; Wang, G.; Shulenburger, L.; Mitas, L. A new generation of effective core potentials for correlated calculations. *J. Chem. Phys.* **2017**, *147*, 224106.
- (80) Werner, H.-J.; Knowles, P. J.; Knizia, G.; Manby, F. R.; Schütz, M., et al. MOLPRO, version 2019.1, a package of ab initio programs. 2019; see <http://www.molpro.net>.
- (81) Sun, Q.; Berkelbach, T. C.; Blunt, N. S.; Booth, G. H.; Guo, S.; Li, Z.; Liu, J.; McClain, J. D.; Sayfutyarova, E. R.; Sharma, S.; Wouters, S.; Chan, G. K.-L. PySCF: the Python-based simulations of chemistry framework. *WIREs Computational Molecular Science* **2018**, *8*, e1340.
- (82) Shea, J. A. R.; Gwin, E.; Neuscamman, E. A Generalized Variational Principle with Applications to Excited State Mean Field Theory. *J. Chem. Theory Comput.* **2020**, *16*, 1526–1540, PMID: 32017562.
- (83) Abrams, M. L.; Sherrill, C. D. Full configuration interaction potential energy curves for the $X1\Sigma_g^+$, $B1\Delta_g$, and $B'1\Sigma_g^+$ states of C_2 : A challenge for approximate methods. *J. Chem. Phys.* **2004**, *121*, 9211–9219.
- (84) Booth, G. H.; Cleland, D.; Thom, A. J. W.; Alavi, A. Breaking the carbon dimer: The challenges of multiple bond dissociation with full configuration interaction quantum Monte Carlo methods. *J. Chem. Phys.* **2011**, *135*, 084104.

- (85) Holmes, A. A.; Umrigar, C. J.; Sharma, S. Excited states using semistochastic heat-bath configuration interaction. *J. Chem. Phys.* **2017**, *147*, 164111.
- (86) Hait, D.; Head-Gordon, M. Excited State Orbital Optimization via Minimizing the Square of the Gradient: General Approach and Application to Singly and Doubly Excited States via Density Functional Theory. *Journal of Chemical Theory and Computation* **2020**, *16*, 1699–1710, PMID: 32017554.
- (87) Watts, J. D.; Gwaltney, S. R.; Bartlett, R. J. Coupled-cluster calculations of the excitation energies of ethylene, butadiene, and cyclopentadiene. *J. Chem. Phys.* **1996**, *105*, 6979–6988.
- (88) Silva-Junior, M. R.; Schreiber, M.; Sauer, S. P. A.; Thiel, W. Benchmarks for electronically excited states: Time-dependent density functional theory and density functional theory based multireference configuration interaction. *J. Chem. Phys.* **2008**, *129*, 104103.
- (89) Shen, J.; Li, S. Block correlated coupled cluster method with the complete active-space self-consistent-field reference function: Applications for low-lying electronic excited states. *J. Chem. Phys.* **2009**, *131*, 174101.
- (90) Piecuch, P.; Hansen, J. A.; Ajala, A. O. Benchmarking the completely renormalised equation-of-motion coupled-cluster approaches for vertical excitation energies. *Molecular Physics* **2015**, *113*, 3085–3127.
- (91) Frueholz, R. P.; Flicker, W. M.; Mosher, O. A.; Kuppermann, A. Electronic spectroscopy of 1,3-cyclopentadiene, 1,3-cyclohexadiene and 1,3-cycloheptadiene by electron impact. *J. Chem. Phys.* **1979**, *70*, 2003–2013.
- (92) McDiarmid, R.; Sabljic, A.; Doering, J. P. Valence transitions in 1,3-cyclopentadiene, 1,3-cyclohexadiene, and 1,3-cycloheptadiene. *J. Chem. Phys.* **1985**, *83*, 2147–2152.

- (93) Gómez, I.; Reguero, M.; Boggio-Pasqua, M.; Robb, M. A. Intramolecular Charge Transfer in 4-Aminobenzonitriles Does Not Necessarily Need the Twist. *J. Am. Chem. Soc.* **2005**, *127*, 7119–7129, PMID: 15884954.
- (94) Segado, M.; Gómez, I.; Reguero, M. Intramolecular charge transfer in aminobenzonitriles and tetrafluoro counterparts: fluorescence explained by competition between low-lying excited states and radiationless deactivation. Part I: A mechanistic overview of the parent system ABN. *Phys. Chem. Chem. Phys.* **2016**, *18*, 6861–6874.
- (95) Segado, M.; Mercier, Y.; Gómez, I.; Reguero, M. Intramolecular charge transfer in aminobenzonitriles and tetrafluoro counterparts: fluorescence explained by competition between low lying excited states and radiationless deactivation. Part II: influence of substitution on luminescence patterns. *Phys. Chem. Chem. Phys.* **2016**, *18*, 6875–6884.
- (96) Tran, L. N.; Neuscamman, E. Improving excited-state potential energy surfaces via optimal orbital shapes. *J. Phys. Chem. A* **2020**, *124*, 8273–8279.

Supporting Information:

A hybrid approach to excited-state-specific variational Monte Carlo and doubly excited states

Leon Otis,[†] Isabel Craig,[‡] and Eric Neuscamman^{*,‡,¶}

[†]*Department of Physics, University of California Berkeley, CA 94720, USA*

[‡]*Department of Chemistry, University of California Berkeley, CA 94720, USA*

[¶]*Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*

E-mail: eneuscamman@berkeley.edu

Molecular Geometries

Table S1: Structure of CN5. Coordinates in Å.

C	0.000 000	0.000 000	0.340 620
C	0.000 000	1.188 120	-0.363 814
C	0.000 000	-1.188 120	-0.363 814
H	0.000 000	0.000 000	1.424 530
H	0.000 000	1.160 200	-1.449 150
H	0.000 000	-1.160 200	-1.449 150
N	0.000 000	2.389 120	0.164 413
N	0.000 000	-2.389 120	0.164 413
H	0.000 000	2.531 430	1.161 940
H	0.000 000	-2.531 430	1.161 940
H	0.000 000	3.211 380	-0.414 923
H	0.000 000	-3.211 380	-0.414 923

Table S2: Structure of nitroxyl. Coordinates in Å.

O	0.111 655	0.000 000	1.140 178
N	-0.236 949	0.000 000	-0.018 994
H	0.625 294	0.000 000	-0.621 185

Table S3: Structure of glyoxal. Coordinates in Å.

C	0.642 211	0.401 329	0.000 000
C	-0.642 211	-0.401 329	0.000 000
O	1.722 903	-0.139 984	0.000 000
O	-1.722 903	0.139 984	0.000 000
H	0.508 726	1.491 662	0.000 000
H	-0.508 726	-1.491 662	0.000 000

Table S4: Structure of acrolein. Coordinates in Å.

C	-0.590 800	-0.361 686	0.000 000
C	0.638 441	0.442 999	0.000 000
C	1.835 351	-0.152 787	0.000 000
O	-1.712 769	0.101 534	0.000 000
H	-0.426 590	-1.453 900	0.000 000
H	0.522 297	1.516 693	0.000 000
H	2.756 647	0.409 814	0.000 000
H	1.910 074	-1.232 987	0.000 000

Table S5: Structure of cyclopentadiene. Coordinates in Å.

H	-0.879 859	0.000 000	1.874 608
H	0.879 859	0.000 000	1.874 608
H	0.000 000	2.211 693	0.612 518
H	0.000 000	-2.211 693	0.612 518
H	0.000 000	1.349 811	-1.886 050
H	0.000 000	-1.349 811	-1.886 050
C	0.000 000	0.000 000	1.215 652
C	0.000 000	-1.177 731	0.285 415
C	0.000 000	1.177 731	0.285 415
C	0.000 000	-0.732 372	-0.993 420
C	0.000 000	0.732 372	-0.993 420

Table S6: Structure of 4-aminobenzonitrile. Coordinates in Å.

C	0.031 570	-0.001 678	1.801 388
C	0.046 873	-1.225 280	1.069 477
C	0.035 483	-1.235 146	-0.304 583
C	-0.115 039	0.000 908	-1.046 464
C	0.031 512	1.235 978	-0.302 121
C	0.042 951	1.223 413	1.071 925
C	0.047 593	-0.003 141	3.231 727
H	0.052 819	-2.162 843	1.609 066
H	0.044 455	-2.181 757	-0.831 777
H	0.037 413	2.183 658	-0.827 425
H	0.045 866	2.159 919	1.613 382
N	0.059 850	-0.006 281	4.396 999
N	0.660 760	0.003 409	-2.284 953
H	0.225 584	0.003 579	-3.196 767
H	1.677 891	0.005 163	-2.289 493

Absolute Energies and Optimization Details for Doubly Excited States

Our reported excitation energies in the main text are the energy differences between ground and excited state optimizations of approximately matching variances, and the total numbers of samples we report are the sum of the amounts used for the two optimizations. We note that obtaining variance matched values in practice requires optimizations at multiple determinant expansion lengths beyond pairs of ground and excited state results presented here. For our double excitations results, all LM optimizations used our adaptive three shift scheme described in Section 2.2 of the main text with 1,000,000 samples per iteration. Hybrid optimization consisted of multiple macro-iterations each composed of 100 AD iterations and 3 Blocked LM iterations before a descent finalizing section of 1100 iterations. In all cases over all systems, each Blocked LM iteration used 1,000,000 samples, which were doubled in our calculation of the total sampling effort as the Blocked LM runs over its sample twice and each AD iteration used 30,000 samples, both within the hybrid method and the pure

descent finalizing section. Energies and uncertainties for each state were computed from an average over the last 10 steps for LM optimizations and from the last 500 descent finalizing steps for hybrid optimizations.

We took a staged approach to our excited state optimizations by first optimizing only a short CI expansion alongside the one- and two-body Jastrow factors, and then adding additional determinants until the variance matched that of the ground state. In the cases where they were used, NCJF and orbital parameters were only optimized starting from a pre-optimized MSJ wave function. In the hybrid method, AD sections were used to identify 5 vectors in parameter space that were provided to the Blocked LM steps. The Blocked LM sections divided parameters into 5 blocks in all hybrid optimizations. As discussed in section 2.2 of the main text, the Blocked LM first performs a LM-style matrix diagonalization for each block and retains some of the eigenvectors. For the carbon dimer MSJ case, nitroxyl, glyoxal, and acrolein, 10 eigenvectors were retained from each block and in all other cases, 30 eigenvectors were retained. In each hybrid optimization, these retained eigenvectors were combined with the 5 vectors from the preceding AD section for another diagonalization to obtain the Blocked LM’s parameter update. On AD sections, we lowered initial step sizes for the later optimization stages to avoid kicking the optimization out of the target function minimum. Table S7 shows the initial step sizes used for different parameter types when they were unoptimized. For later stages, the step sizes for the Jastrow factors and CI coefficients were reduced by factors of 10 and 2 respectively. The NCJF and orbital parameters required relatively small changes in their values compared to the other parameter types and we found that significantly smaller step sizes for them were effective in stably improving the target function. Based on our experience from this work, these choices for step sizes and Blocked LM parameters within our staged optimizations form a reasonable default option for applying our methodology in similar studies. However, adjusting these settings may be necessary in challenging cases to either achieve a better optimization or at least verify that one is not easily obtainable.

Table S7: Initial step sizes for completely unoptimized parameters of different types.

Parameter Type	Step Size
One-Body Jastrow	0.1
Two-Body Jastrow	0.1
CI Coefficient	0.01
NCJF F Matrix	0.0005
Orbital Parameter	0.0005

In tables S8 through S13, we list the total number of optimizable parameters and the absolute energies for each of the systems we consider in our doubly excited state calculations. The parameter numbers are for results with CI expansion lengths chosen to obtain an explicit variance match between the ground and excited state. In some cases, the LM and hybrid method needed different expansion lengths to variance match their respective ground state results leading to different total numbers of excited state parameters. Our statistical uncertainties on the last digit of the absolute energies are shown in parentheses. In table S14, we list the number of iterations in each optimization for the different cases in all our molecules.

Table S8: Absolute energies and parameter numbers for the carbon dimer.

	GS Parameters	ES Parameters	GS Energy (a.u.)	ES Energy (a.u.)
LM MSJ	50	53	-11.0232(7)	-10.9372(7)
Hybrid MSJ	50	53	-11.0230(3)	-10.9377(4)
LM All	1080	1251	-11.0338(6)	-10.9469(6)
Hybrid All	1080	1170	-11.0343(2)	-10.9483(1)

Table S9: Absolute energies and parameter numbers for nitroxyl.

	GS Parameters	ES Parameters	GS Energy (a.u.)	ES Energy (a.u.)
LM MSJ	70	90	-26.387(1)	-26.229(1)
Hybrid MSJ	70	90	-26.3858(4)	-26.3851(4)

Table S10: Absolute energies and parameter numbers for glyoxal.

	GS Parameters	ES Parameters	GS Energy (a.u.)	ES Energy (a.u.)
LM MSJ	70	90	-44.445(2)	-44.228(2)
Hybrid MSJ	70	90	-44.4440(6)	-44.2324(5)

Table S11: Absolute energies and parameter numbers for acrolein.

	GS Parameters	ES Parameters	GS Energy (a.u.)	ES Energy (a.u.)
LM MSJ	70	120	-35.319(2)	-35.319(1)
Hybrid MSJ	70	120	-35.3169(6)	-35.3169(5)

Table S12: Absolute energies and parameter numbers for cyclopentadiene.

	GS Parameters	ES Parameters	GS Energy (a.u.)	ES Energy (a.u.)
LM MSJ	60	540	-31.830(1)	-31.551(1)
Hybrid MSJ	60	540	-31.8287(4)	-31.5253(6)

Table S13: Absolute energies and parameter numbers for 4-aminobenzonitrile.

	GS Parameters	ES Parameters	GS Energy (a.u.)	ES Energy (a.u.)
LM SA MSJ	70	120	-62.855(2)	-62.608(2)
Hybrid SA MSJ	70	150	-62.8453(8)	-62.6037(6)
LM SS MSJ	70	550	-62.855(2)	-62.606(2)
Hybrid SS MSJ	70	350	-62.8463(9)	-62.6035(6)

Table S14: Iteration numbers for all systems. The hybrid method values are the number of macro-iterations before the descent finalizing section.

	GS Iterations	ES Iterations
LM MSJ C2	50	50
Hybrid MSJ C2	4	4
LM All C2	50	50
Hybrid All C2	4	4
LM MSJ nitroxyl	50	50
Hybrid MSJ nitroxyl	4	4
LM MSJ glyoxal	70	70
Hybrid MSJ glyoxal	4	4
LM MSJ acrolein	70	70
Hybrid MSJ acrolein	4	4
LM MSJ CPD	50	90
Hybrid MSJ CPD	8	8
LM SA MSJ ABN	70	70
Hybrid SA MSJ ABN	7	5
LM SS MSJ ABN	70	70
Hybrid SS MSJ ABN	7	7