

© 2024 William A Wheeler

ADVANCING MANY-BODY METHODS FOR ELECTRONIC STRUCTURE
CALCULATIONS

BY

WILLIAM A WHEELER

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Materials Science and Engineering
in the Graduate College of the
University of Illinois Urbana-Champaign, 2024

Urbana, Illinois

Doctoral Committee:

Associate Professor André Schleife, Chair
Assistant Professor Lucas K. Wagner, Director of Research
Associate Professor Daniel P. Shoemaker
Associate Professor Nicola H. Perry

ABSTRACT

Ab initio calculations of materials complement experimental techniques by enabling prediction of properties of materials, as well as accessing quantities that are difficult or impossible to measure in the lab. In strongly correlated materials, properties depend strongly on the coordination of many electrons together, requiring many-body descriptions that explicitly capture electron interactions. This dissertation advances the methods available for *ab initio* many-body calculations in several areas.

In this work, I analyze the capabilities and limitations of a new formula for calculating bulk electric quadrupole moments of many-body wave functions. Interpreting the quadrupole as the result of dipoles flowing through the material highlights subtleties of defining this quantity. With the aim of streamlining development of new formulas and algorithms, I introduce PyQMC, a new implementation of quantum Monte Carlo (QMC) in Python. PyQMC reduces the human time in algorithm development and workflow design through a modular, easy-to-modify codebase. The capabilities of PyQMC are brought to my next project, characterizing spin fluctuations in the ground state of unconventional superconductor CaCuO_2 . Multi-determinant QMC trial functions are shown to capture spin fluctuations in the undoped parent compound and orbital optimization is demonstrated to be viable on these wave functions, establishing the basis for analyzing spin fluctuations in the less understood doped material. Moving towards excited states, I analyze an algorithm for calculating excited states using many-body wave function based techniques. The approach introduces the first variational principle for an ensemble of trial states, where an objective function is minimized only by the ensemble of the lowest N eigenstates. My work opens the path towards constructing an effective model of the cuprates at QMC level accuracy, which requires high-accuracy ground and excited states and reliable characterization of the spin and charge correlations in the materials. The tools and methods I present will transfer to accurate characterization of other systems and support future efforts to understand and model strongly correlated materials.

ACKNOWLEDGMENTS

I would like to thank the staff in both the Physics and Materials Science and Engineering departments for making the various administrative processes of being a graduate student as easy as I can imagine.

I would like to thank all my collaborators and group members who helped me progress in my research through their stimulating conversations, encouragement, and advice, including (in no particular order) Scott Jensen, Chun Yu Chow, Kevin Kleiner, Shivesh Pathak, Alex Munõz, João Rodrigues, Mick Krongchon, Yueqing Chang, Brian Busemeyer, Kiel Williams, Scott McCormack, Nirvaan Khara, Sonali Joshi, Billy Passias, Kevin Ly, Shubhang Goswami, Prof. Taylor Hughes, and Prof. David Ceperley, among others. I am also grateful to many other friends and family members for support and encouragement along the way.

I would like to thank André Schleife, who supervised me through several rewarding semesters as a TA and who continued to offer me periodic guidance throughout my time here.

Finally, I thank my advisor, Lucas Wagner, for his mentorship throughout my Ph.D. He taught me valuable research skills ranging from code development to organization to communication, and helped shape my perspective on research. I particularly value his emphasis on careful reasoning and quality writing, and valuing sound research over rapid output. The experience working with him would not have been complete without the many far-fetched analogies, almost funny jokes, and irrelevant tangents on food, languages, and life that made meetings entertaining and long lasting.

I gratefully acknowledge support in funding and computational resources that made my work possible. I thank the Materials Science and Engineering department for support through the Hamer Fellowship. I also thank the U.S. National Science Foundation for support via Award No. 1931258. This work used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This work made use of the Illinois Campus Cluster, a computing resource that is operated by the Illinois Campus Cluster Program (ICCP) in conjunction with the National Center for Supercomputing Applications (NCSA) and which is supported by funds from the University of Illinois Urbana-Champaign.

TABLE OF CONTENTS

Chapter 1	Introduction	1
1.1	Calculating Materials' Properties from the Atomic Scale	1
1.2	Thesis Outline	2
Chapter 2	Background	5
2.1	Describing Systems of Many Electrons	5
2.2	Quantum Monte Carlo	10
2.3	Observables	20
2.4	Polarization and Quadrupole Moments in Bulk Systems	20
Chapter 3	Bulk electric quadrupole moments	27
3.1	Electric Moments as Integrated Current	28
3.2	Adiabatic Evolution	31
3.3	Practical Evaluation	32
3.4	Example Calculations	34
3.5	Analysis of constraints and finite size errors	37
3.6	Connection to Recent Work	47
3.7	Conclusion	50
Chapter 4	PyQMC code	51
4.1	All-Python Code for Human Efficiency	52
4.2	Code Structure and Implementation	53
4.3	Diverse Workflow Support	63
4.4	Acceleration Strategies	65
4.5	Conclusion	68
4.6	Implementation Details	69
Chapter 5	Describing Spin Fluctuations in CaCuO_2	71
5.1	The Challenge of Describing Cuprates	71
5.2	Methods: Multideterminant QMC Approach	72
5.3	Calculation Details	74
5.4	Results	77
5.5	Conclusion	83

Chapter 6	Ensemble VMC for excited states	84
6.1	Introduction	85
6.2	Penalty-Based Cost Functional	86
6.3	Application to <i>Ab Initio</i> System	93
6.4	Conclusion	95
Chapter 7	Conclusion	97
References		99

CHAPTER 1

INTRODUCTION

1.1 Calculating Materials' Properties from the Atomic Scale

Materials derive many of their properties from the organization of the electrons that make them up. Examples of such properties include electrical conductivity, chemical bonding, stiffness, absorption of light, magnetism, and superconductivity. Understanding how these properties interact and respond to external stimuli enables the design and control of materials for applications. One example I will discuss later on is the unconventional superconductivity found in cuprate materials at temperatures up to 130K – higher than any other materials known so far. A complete description of how this “high-temperature” superconductivity arises is unknown, but could aid in designing materials that superconduct closer to room temperature.

Calculating materials' properties from first principles (the Schrödinger equation) enables us to interpret experimental observations and make predictions about the properties of new materials. Although the Schrödinger equation is fairly simple conceptually, the dimensionality of the problem with many particles is so high that it can only be solved through approximations. An electronic state, called a wave function, is a function of the positions of all electrons in the system – if there are N electrons, it is a $3N$ -dimensional function. One of the most widely-used approaches is the mean-field approximation, where only one electron is considered at a time, and the effects of all other electrons are averaged together. Such models are able to describe a surprisingly wide range of phenomena, given their simplicity. However, accurate calculations on systems where correlation between electrons is important, such as in the cuprates, often require many-body methods (treating many electrons simultaneously). Transition metals like copper have electrons in d orbitals, which means many electrons are crowded together close to the nucleus and need to coordinate their positions to avoid one another. In other words, the many-body probability distribution favors configurations where electrons are spaced apart over those where electrons collide. Describing these kinds of situations is the strength of *ab initio* many-body methods, the topic of this work.

The main approach used in this work is quantum Monte Carlo (QMC), a high-accuracy many-body method

that scales up to systems with thousands of electrons on modern computers. Traditionally, QMC targets the electronic ground state of a system. Although consistently quite accurate, its accuracy is limited by the quality of trial wave functions (essentially parameterized guesses for the target function). The contributions of this work expand QMC capabilities; in particular, improving the quality of QMC wave functions, advancing excited state approaches, and enabling calculation of new quantities in QMC. An application of improved trial functions is demonstrated to describe spin fluctuations in the strongly-correlated cuprate CaCuO_2 .

1.2 *Thesis Outline*

The work in this dissertation consists of four projects. The document is organized as follows.

Chapter 2, reviews the computational methods implemented in the PyQMC code and those used in the calculations on CaCuO_2 . It also reviews the methods for computing bulk polarization and for computing bulk quadrupole moments of non-interacting systems to provide background for the work on many-body quadrupole moments.

Chapter 3 is focused on calculating bulk electric quadrupole moments in periodic systems. The many-body formula we introduced extends an existing one-body formalism to a form that can be applied to the wave functions used in QMC and other many-body methods. We successfully tested the many-body formula on several non-interacting models from prior work on bulk quadrupole moments. We also gave a physical interpretation of the quadrupole as resulting from dipoles flowing through the material, analogous to charge currents that generate polarization. However, there are two notable differences: (1) in contrast to charges, dipole moments are not conserved, so there is no clear way to separate dipoles at the surface from those in the bulk; (2) dipole moments are not quantized, and small fluctuations in dipole moment pose a problem, since quadrupole moments are only defined when the dipole moment is zero. These problems are fundamental to the quadrupole moment, independent of our formula, and constrain the existence of such a quantity in real materials.

Chapter 4 covers our new implementation of quantum Monte Carlo (QMC) in Python, called PyQMC. PyQMC is designed to reduce the human time in code development and workflow design, while maintaining reasonably high computational performance. It overcomes traditional challenges of developing and implementing new methods by using Python, a simple programming language that enables easy and efficient modification of code. Among QMC software packages, PyQMC is uniquely flexible in enabling users to interface it with their own custom code without requiring any modifications to the code base. Integration with the quantum chemistry package PySCF streamlines the process of importing data from other methods to generate trial functions for PyQMC. Additionally, by using PySCF objects to set up calculations, PyQMC

results can be directly compared to results of other methods for benchmarking.

Chapter 5 is focused on testing new trial wave functions for the cuprate material CaCuO_2 , achieving a better description of the physics that leads to unconventional superconductivity in this family of materials. The movement of holes through the material is affected by spin fluctuations – superpositions of different spin configurations around the hole. We use high quality multi-determinant trial functions, enabled by PyQMC, to measure the importance of spin fluctuations in the ground state of cuprates. Optimizing the orbitals of the determinants – also enabled by PyQMC – is necessary to combine the trial function with the additional short-range correlations captured by QMC. We study a 2×2 unit cell of the simplest cuprate, CaCuO_2 , to compare the performance of different trial wave functions in the undoped material. The 2×2 Heisenberg model, a good approximation to the undoped cuprates, is known to have spin fluctuations in the ground state, so we expect good *ab initio* wave functions to describe them as well. Since QMC is variational, wave function energies are upper bounds to the true ground state energy, so a decrease in energy indicates a more accurate estimate of the ground state. A few determinants with large weights represent the fluctuations between different spin configurations. Surprisingly, we find that these more sophisticated wave functions do not improve the final energies over the simpler single-determinant wave function. In the *ab initio* calculations, the fluctuations between different spin orderings are dwarfed by the random fluctuations of the many electrons in the system; however, when the large background is subtracted, the spin fluctuations are remarkably similar to the Heisenberg ground state. Although the multi-determinant wave functions do not outperform traditional ones in this study, they are shown to effectively capture the physics known to be relevant for holes in the superconducting state.

Chapter 6 analyzes a method for computing several excited states in a single QMC optimization. Traditional ground-state QMC simply minimizes the energy of a single trial function. A recent excited state approach computes the n -th excited state by minimizing its energy and penalizing its overlap with the first $n - 1$ eigenstates – the n -th excited state is the lowest energy state orthogonal to the first $n - 1$ states. A drawback of this approach is the necessity to optimize each trial function in a separate calculation. In the new approach, an ensemble of trial functions is optimized simultaneously. The optimizer minimizes a weighted sum of energies plus a penalty to ensure all the states are orthogonal. This work derives a condition on the penalty and weight parameters to ensure that the optimal ensemble reaches the true eigenstates of the system. There is generally a critical penalty value, above which the eigenstates are reached. The algorithm is carried out on a toy model as well as realistic small molecules to demonstrate practical application. We find that when the penalty is only slightly above the critical value, the objective function has a shallow minimum, and optimizing to the exact minimum is difficult. A practical critical value must include a buffer to account for the stochastic noise of Monte Carlo evaluations of the cost function, in which case the states are orthogonal and are the

Table 1.1: Summary of projects in the dissertation.

Chapter	Problem	Solution	Publication
3	How to define bulk quadrupole moment in many-body systems	Propose a quadrupole formula, give physical interpretation, benchmark on previous TB models, and analyze limitations.	Wheeler, Wagner, Hughes, PRB 100, 245135 (2019).
4	Accelerate QMC algorithm development	Develop software package PyQMC for QMC calculations in Python.	Wheeler et al., JCP 158, 114801 (2023).
5	Capture spin fluctuations in <i>ab initio</i> calculations of cuprates	Enable spin fluctuations in QMC trial functions and compare calculations on CaCuO_2 with and without.	In progress
6	Optimize multiple excited states simultaneously	Prove that an appropriate objective function optimizes an ensemble exactly to the excited states.	Wheeler, Kleiner, Wagner, Electron. Struct. 6, 025001 (2024)

best possible approximation to the true eigenstates (under the given ensemble parameterization).

Chapter 7 summarizes the results of the dissertation and provides an outlook on their impact on future work in the field. An overview of the projects is presented in Table 1.1.

CHAPTER 2

BACKGROUND

Ab initio, or first principles, calculations provide an avenue for probing materials in parallel to experiment. The Schrödinger equation governs the electrons in the material subject to Coulomb interactions with other electrons and with fixed atomic nuclei. By solving this equation, properties of the material can be computed without any experimental inputs, giving *ab initio* methods predictive power.

This chapter reviews the background of the various *ab initio* methods used throughout this work. Additionally, Section 2.4 introduces the many-body formulation of bulk polarization and the one-body formulation of bulk quadrupole moments as background for the discussion of the many-body bulk quadrupole moment.

2.1 Describing Systems of Many Electrons

Systems of N electrons are governed by the Schrödinger equation

$$\hat{H}\Psi = i\frac{\partial}{\partial t}\Psi, \quad (2.1)$$

where the many-body Hamiltonian $\hat{H}$ acts on wave functions $\Psi(\mathbf{R})$, functions of all $3N$ electron coordinates $\mathbf{R} = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N\}$. In the atomic environment of a material, the Hamiltonian is given by

$$\hat{H} = \underbrace{-\frac{1}{2} \sum_i \nabla_i^2}_{\hat{T}} - \underbrace{\sum_i \sum_I \frac{Z_I}{|\mathbf{r}_i - \mathbf{r}_I|}}_{\hat{V}_{\text{ext}}} + \underbrace{\sum_{i < j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}}_{\hat{V}_{\text{ee}}}, \quad (2.2)$$

where lowercase indices i, j index electrons, the uppercase index I indexes the ions (atomic nuclei), $\mathbf{r}$ represents a particle's position, and Z an ionic charge. The Hamiltonian consists of three components: the kinetic energy $\hat{T}$, the external potential $\hat{V}_{\text{ext}}$ due to the atomic nuclei, and the Coulomb potential $\hat{V}_{\text{ee}}$ between electrons. The equation is expressed in Hartree units, used throughout this dissertation, where the electron charge

$e = 1$, electron mass $m_e = 1$, Planck constant $\hbar = 1$, and Coulomb constant $4\pi\epsilon_0 = 1$ are the fundamental units. We are often interested in finding the low energy eigenstates of $\hat{H}$. Many materials properties can be determined just from calculating the ground state of the electrons. For example, fixing atoms at different positions can identify stable atomic structures and phonon excitations. Similarly, finding the ground state energies under different spin constraints can identify magnetic structures and magnetic excitations. Insulating or metallic behavior corresponds to the localization of electrons in the ground state. Low energy excited states have relevant information as well. Often electronic degrees of freedom interact with spin or lattice degrees of freedom. To describe an exchange of energy between a spin state and an electronic state, excited electronic states enter the picture. Including electronic states can be crucial to determining good quality models for systems with many kinds of interactions.

2.1.1 Hartree-Fock, the Uncorrelated Approximation

All wave functions of electrons must satisfy fermion anti-symmetry,

$$\Psi(\mathbf{r}_1, \mathbf{r}_2, \dots) = -\Psi(\mathbf{r}_2, \mathbf{r}_1, \dots).$$

The Hartree-Fock approximation (see *e.g.* [1]) models wave functions as Slater determinants – the antisymmetrized product of a set of orbitals $\{\phi_j\}$,

$$\Psi_S = \frac{1}{\sqrt{N!}} \begin{vmatrix} \phi_1(\mathbf{r}_1) & \phi_2(\mathbf{r}_1) & \dots & \phi_N(\mathbf{r}_1) \\ \phi_1(\mathbf{r}_2) & \phi_2(\mathbf{r}_2) & \dots & \phi_N(\mathbf{r}_2) \\ \vdots & & \ddots & \\ \phi_1(\mathbf{r}_N) & \phi_2(\mathbf{r}_N) & \dots & \phi_N(\mathbf{r}_N) \end{vmatrix}. \quad (2.3)$$

Each orbital in the determinant is considered to be “occupied” by an electron. The Hartree-Fock ground state is the Slater determinant that minimizes the energy $\langle \Psi_S | \hat{H} | \Psi_S \rangle$, which provides a variational upper bound to the true many-body ground-state energy. The method can be thought of as finding the best set of one-particle orbitals to occupy for approximating the electronic ground state of a system. In this model, an excitation is thought of as an electron moving from an occupied orbital ϕ_i to an “unoccupied” orbital $\phi_{i'}$ – one that is not in the set $\{\phi_j\}$ of orbitals in Ψ_S – while keeping the remaining electrons fixed.

A single determinant is called uncorrelated, since it encodes the smallest amount of correlation allowed by anti-symmetry (that it approaches zero at coincidence of two or more electrons). The antisymmetry constraint does not apply to electrons having different spins, and the $1/r$ Coulomb interaction is “long

ranged”, contributing significantly even when electrons are not near coincidence. For these reasons, wave functions that describe additional correlation are necessary for accurate descriptions of most systems.

2.1.2 Determinant Expansion / Configuration Interaction

Correlation can be introduced into the wave function ansatz by a linear combination of multiple determinants, enabling more accurate approximations.[1] For example, in a two-site system with one up and one down electron, a state with low energy will delocalize the electrons across the sites. A “determinant” in this system is just a product of the up and down orbitals. A low-energy single determinant will put both electrons into an orbital spread across the two sites, but in this state both electrons occupy the same site with 50% probability. Increasing the probability of the two configurations where each electron occupies a different site (such as the singlet state) requires two determinants, and has lower energy than the single determinant by reducing the probability of both electrons occupying the same site. In this sense of sets of occupied orbitals, determinants are often called configurations in quantum chemistry, and linear combinations are referred to as configuration interaction (CI).

Realistic systems consist of many atoms with many occupied and unoccupied orbitals that are relevant for arbitrary electronic excitations. In principle, the set of all possible occupation configurations forms a complete many-body basis to describe any wave function and can give exact results. In practice, there are two big limitations on this expansion. In a computation, orbitals are represented in a finite basis of three-dimensional functions, and results must be extrapolated to the infinite limit in that basis. Even for a fixed basis, describing all possible occupation configurations requires a combinatorially large number of determinants. These constraints make the exact solution, called full CI, too expensive to scale to large systems. There are several approaches to reducing the number of determinants that need to be considered; this work uses the complete active space approach, discussed in the sections that follow.

Complete Active Space Configuration Interaction (CASCI)

When constructing a CI wave function, a strategy must be chosen for selecting which determinants to include, where each determinant is a different set of occupied orbitals. Effective models of materials reduce the full geometry of atoms and electrons to a small set of sites with a small subset of electrons. These sites typically correspond to some orbitals in the original atomic system, which are observed to host most of the electronic fluctuations. In the complete active space (CAS) approach,[2]–[4] such a set of orbitals is selected to be the “active space:” these are the orbitals whose occupation is allowed to change across different determinants, sometimes occupied, sometimes not. Every possible determinant from occupying active space orbitals with the

specified number of active space electrons is included in the CAS expansion, and the coefficients are optimized to minimize the energy. The “core” orbitals are not in the active space, and are present in every determinant. Any virtual orbitals not in the active space do not contribute to any determinant in the expansion. By fixing the core and virtual orbitals, the number of determinants (and the cost of the calculation) is greatly reduced compared to full CI.

The CASCI approach takes advantage of the fact that orbital occupations in low-energy states tend to fluctuate among orbitals within a narrow energy window. A good description of many-body ground states (as well as low-energy excitations) is thus achievable by including only orbitals that contribute to these fluctuations. This set of “important” orbitals depends on the accuracy needed; describing more detail requires more active space orbitals. Often, the active space can be judiciously chosen based on orbital energies and physical intuition. A systematic approach increases the size of the active space through a series of calculations until the energy converges within a desired tolerance.

Complete Active Space Self-Consistent Field (CASSCF)

A disadvantage of the CASCI approach above is the fixed shape of the orbitals. When multiple configurations are superposed, the electrons are better described by orbitals shifted from their mean-field shapes. While the full expansion can correctly describe these effects, the right choice of orbitals can result in a much smaller number of determinants to achieve the same accuracy.

Similar to CASCI, the CASSCF method[2]–[4] generates a determinant expansion where each determinant is composed of different orbitals in the active space (along with all of the core orbitals). A self-consistent field (SCF) equation is then iterated to find optimal orbital shapes, using CASCI to optimize the determinant coefficients at each iteration. The orbital optimization allows CASSCF to perform better than CASCI even with a smaller active space, since the orbitals can adapt to mix with orbitals outside the original active space.

In Chapter 5, a CAS approach is used to capture correlations between electrons on copper $3d_{x^2-y^2}$ orbitals in copper-oxygen plane of CaCuO_2 . Four copper orbitals are considered relevant; however they are strongly hybridized with oxygen $2p$ orbitals, and the mean-field description is not very good. While CASCI struggles even with many more than the four desired orbitals, but CASSCF successfully extracts the relevant four-orbital space.

2.1.3 Density Functional Theory (DFT)

Density functional theory (DFT) is an approach to diagonalizing the Hamiltonian similar to Hartree-Fock (HF), but with an aim to capture electron correlations that HF neglects. The foundation of the approach

is a theorem by Hohenberg and Kohn[5] asserting that the external potential $\hat{V}_{\text{ext}}$ of a Hamiltonian is uniquely determined by its ground-state density. Since $\hat{V}_{\text{ext}}$ specifies the full many-body Hamiltonian, all the information of the Hamiltonian, its eigenstates, and eigenvalues (energy spectrum), is contained in the ground-state density.

In addition, Hohenberg and Kohn proved that an energy functional of the density can be defined, whose minimum is at the true energy and density of the ground state. The expectation value of the energy and all terms contributing to it are functionals of the ground state, and therefore also functionals of the density,

$$E[\rho] = T[\rho] + V_{\text{ext}}[\rho] + V_{\text{Coulomb}}[\rho] + V_{\text{xc}}[\rho], \quad (2.4)$$

where $V_{\text{xc}}[\rho]$ represents the exchange (present in HF) and correlation (not present in HF) contributions to the energy. This exchange-correlation functional represents all the effects that are not included in the other terms. It is universal[5] – the functional is the same for any atoms and geometry – but its exact form is unknown.

Most DFT calculations use a local exchange-correlation functional, where evaluation at a point $\mathbf{r}$ depends only on the value and gradients of the density at $\mathbf{r}$ [6], [7]. These local functionals have been parameterized numerically based on accurate quantum Monte Carlo calculations of the homogeneous electron gas[8], and perform surprisingly well for most systems. Importantly, DFT (like HF) scales N_{elec}^3 , allowing it to be applied to materials containing thousands of atoms.

The key difference between DFT and HF is the approximate exchange-correlation term V_{xc} included in the DFT Hamiltonian. While HF tends to spread orbitals out and grossly increase energy gaps, the correlation effects in DFT compensate these errors, typically underestimating gaps instead. Although consistently underestimating gaps, DFT is often much more accurate than HF for solid calculations. Qualitatively, the orbitals are much better descriptions of the electron distributions. Although not the orbitals that minimize the energy of a Slater determinant, along with the exchange-correlation contribution, they result in a low-energy wave function and are designed to reproduce the ground state density very accurately.

While accounting for the energy due to electron correlation, DFT does not include it explicitly in the calculation. This approximation is quite good when the correlations have a weak renormalization effect on the effective one-body orbital description, but becomes unreliable when the correlation effects are central to the behavior in a material, such as with the highly localized d electrons in transition metal atoms. Determinant expansions are one way of explicitly including correlation, but scale poorly to large systems. Quantum Monte Carlo, on the other hand, can achieve the same N_{elec}^3 scaling as DFT – albeit with a large prefactor – and include many-body correlations explicitly.

2.2 Quantum Monte Carlo

2.2.1 Slater-Jastrow Ansatz

An alternative way to add correlation to a Slater determinant, aside from determinant expansion, is by adding Jastrow factor,[\[9\]](#) which reweights the probability of different spatial configurations of electrons. The Slater-Jastrow function takes the form

$$\Psi(\mathbf{R}) = \Psi_S(\mathbf{R})e^{J(\mathbf{R})}, \quad (2.5)$$

where the Jastrow factor $e^{J(\mathbf{R})}$ is positive everywhere, reweighting the probabilities of different electron configurations. The function $J(\mathbf{R})$ is a function of all N electron positions, but can be classified by the number of particles contributing to each term. For example, the (homogenous) two-body Jastrow is a sum over particle pairs

$$J_2(\mathbf{R}) = \sum_{(i,j)} u_{ee}(\mathbf{r}_i - \mathbf{r}_j), \quad (2.6)$$

where $u_{ee}(\mathbf{r})$ is some function of the pair displacements, usually spherically symmetric (depending only on the scalar $|\mathbf{r}|$). In an atomic environment, the two-body correlation allows electrons to reside closer to the nuclei, which is accommodated by an electron-ion pair factor, called the one-body Jastrow (involving one electron),

$$J_1(\mathbf{R}) = \sum_{(i,I)} u_{ei}(\mathbf{r}_i - \mathbf{r}_I). \quad (2.7)$$

A (confusingly termed) three-body Jastrow augments the homogeneous two-body Jastrow by additionally depending on the electrons' distances to an ion

$$J_3(\mathbf{R}) = \sum_{(i,j,I)} u(\mathbf{r}_i - \mathbf{r}_j, \mathbf{r}_i - \mathbf{r}_I, \mathbf{r}_j - \mathbf{r}_I). \quad (2.8)$$

In practice, this function is typically represented as a product of pair terms, $u(r_{ij}, r_{iI}, r_{jI}) = u_{ee}(r_{ij})u_{ei}(r_{iI})u_{ei}(r_{jI})$.

Many other forms of Jastrow factors $J(\mathbf{R})$ exist, and do not necessarily need to fall under the n-body classification; the only constraint is that they are symmetric under particle exchange, so that the overall function $\Psi(\mathbf{R})$ is anti-symmetric.

The Slater-Jastrow ansatz can be combined with the determinant expansion concept in [Sec. 2.1.2](#), called a multi-Slater-Jastrow (MSJ) trial function,

$$\Psi(\mathbf{R}) = \sum_i c_i D_i(\mathbf{R}) e^{J(\mathbf{R})}, \quad (2.9)$$

where D_i represents the i -th determinant in the expansion and c_i are the expansion coefficients. The Jastrow factor captures short-range correlations that typically correspond to large numbers of low-weight determinants, allowing a much smaller number of determinants to be used in the MSJ function to capture the longer-range correlations.

2.2.2 QMC Expectation Values

The expectation value of an operator $\hat{O}$ over a wave functions Ψ is defined to be

$$\langle O \rangle_{\mathbf{R} \sim |\Psi|^2} = \frac{\int d\mathbf{R} \Psi^*(\mathbf{R}) \hat{O} \Psi(\mathbf{R})}{\int d\mathbf{R} |\Psi(\mathbf{R})|^2}. \quad (2.10)$$

Note that the normalization is included explicitly in this expression, rather than assuming that Ψ is normalized. Eq. (2.10) can be rearranged to have the form of an integral over a probability distribution,

$$\langle O_L \rangle_{\mathbf{R} \sim |\Psi|^2} = \int d\mathbf{R} \frac{|\Psi(\mathbf{R})|^2}{\int d\mathbf{R} |\Psi(\mathbf{R})|^2} \frac{\hat{O} \Psi(\mathbf{R})}{\Psi(\mathbf{R})}. \quad (2.11)$$

The integration variable $\mathbf{R}$ is $3N$ -dimensional, making standard numerical grid integration techniques prohibitively expensive; however, Monte Carlo sampling can be very effective for evaluating such integrals. Values of $\mathbf{R}$ are randomly sampled with probability $p_\Psi(\mathbf{R}) \propto |\Psi(\mathbf{R})|^2$, and values of $O_L(\mathbf{R}) = \frac{\hat{O} \Psi(\mathbf{R})}{\Psi(\mathbf{R})}$ are averaged over the samples, yielding the expectation. We refer to $O_L(\mathbf{R})$ as the local quantity or estimator for the operator $\hat{O}$, e.g. the local energy for $\hat{H}$.

In certain situations, one may wish to sample alternative probability distributions. This can be useful for sampling low-probability regions that are unusually important, or to more efficiently bridge separated regions of high probability. In this work, it is used for correlated sampling, where multiple wave functions are evaluated on the same set of samples. Correlated samples result in low-variance differences, where differences in estimators on the same sample $\mathbf{R}$ are small in comparison to the fluctuations, which would dominate the difference of independent averages. When sampling a distribution $p(\mathbf{R})$, the estimator is modified to $f(\mathbf{R}) = O_L(\mathbf{R}) p_\Psi(\mathbf{R}) / p(\mathbf{R})$.

The Metropolis Algorithm

The integral of Eq. (2.11) is sampled by the Metropolis method.[10] Given a configuration of electrons $\mathbf{R}$, a move is proposed to a new configuration $\mathbf{R}'$. An acceptance ratio is determined by the probability of proposing the move $T(\mathbf{R} \rightarrow \mathbf{R}')$ and the probability density of the target distribution $|\Psi(\mathbf{R})|^2$. The acceptance ratio is

defined as

$$A = \frac{|\Psi(\mathbf{R}')|^2}{|\Psi(\mathbf{R})|^2} \frac{T(\mathbf{R}' \rightarrow \mathbf{R})}{T(\mathbf{R} \rightarrow \mathbf{R}')}. \quad (2.12)$$

If $A > 1$, the move $\mathbf{R} \rightarrow \mathbf{R}'$ is accepted. If $A < 1$, the move is accepted with probability A ; otherwise, it is rejected, and the walker remains at $\mathbf{R}$ for the next step.

A common, simple choice for the move proposal $T(\mathbf{R} \rightarrow \mathbf{R}')$ is to sample moves from a Gaussian distribution $p(\mathbf{R}'|\mathbf{R}) \propto e^{-|\mathbf{R}-\mathbf{R}'|^2/2\Delta\tau^2}$, where $\Delta\tau$ is a parameter called the “time step” that determines the average distance of the proposed moves. This choice of move proposal has the same probability forwards and backwards, so the acceptance ratio depends only on the distribution, $A = |\Psi(\mathbf{R}')/\Psi(\mathbf{R})|^2$. In systems with many electrons, sampling is more efficient when $\mathbf{R} \rightarrow \mathbf{R}'$ changes the position of only a single electron at a time; i.e., the move is really $\mathbf{r}_e \rightarrow \mathbf{r}'_e$ for electron e . The choice of $\Delta\tau$ affects the efficiency of the sampling: choosing $\Delta\tau$ too large results in a high rejection rate; too small, and walkers move very slowly through configuration space.

This tradeoff can be improved by better choices of the proposal distribution $T(\mathbf{R} \rightarrow \mathbf{R}')$, preferentially proposing moves that increase $|\Psi(\mathbf{R})|$. Instead of centering the Gaussian move of electron e at its current position $\mathbf{r}_e$, the move is “biased”, shifting the center along the gradient of $|\Psi(\mathbf{R})|$. The resulting proposal is $\mathbf{r}'_e = \mathbf{r}_e + \Delta\tau \mathbf{F}_e(\mathbf{R}) + \chi$, where χ is sampled from a Gaussian distribution, and $\mathbf{F}_e = \nabla_e \ln \Psi(\mathbf{R})$ is related to the force used in classical molecular dynamics sampling.^[11] The move probabilities are therefore

$$T(\mathbf{R} \rightarrow \mathbf{R}') = e^{-|\chi|^2/2\Delta\tau} \quad (2.13)$$

$$T(\mathbf{R}' \rightarrow \mathbf{R}) = e^{-|[\mathbf{F}_e(\mathbf{R}) + \mathbf{F}_e(\mathbf{R}')] \Delta\tau + \chi|^2/2\Delta\tau}. \quad (2.14)$$

This procedure is called smart Monte Carlo in classical simulations.^[12] It should be noted that the timestep $\Delta\tau$ in these equations corresponds to time squared in molecular dynamics, and in these equations has effective units of length squared; however, it is nothing more than an adjustable parameter to control the size of the proposed moves.

2.2.3 Variational Monte Carlo (VMC)

The trial functions of the form in Eq. (2.9) have many adjustable parameters. To approximate the ground state, the parameters are variationally optimized by minimizing the energy expectation

$$E[\Psi] = \frac{\langle \Psi | \hat{H} | \Psi \rangle}{\langle \Psi | \Psi \rangle} = \left\langle \frac{\hat{H} \Psi(\mathbf{R})}{\Psi(\mathbf{R})} \right\rangle_{R \sim |\Psi|^2}. \quad (2.15)$$

The gradients of $E[\Psi]$ with respect to the parameters p are computed through the Monte Carlo sampling as

$$\frac{\partial}{\partial p} E[\Psi] = 2\text{Re} \left[\frac{\langle \Psi | \hat{H} | \partial_p \Psi \rangle}{\langle \Psi | \Psi \rangle} - E[\Psi] \frac{\langle \Psi | \partial_p \Psi \rangle}{\langle \Psi | \Psi \rangle} \right], \quad (2.16)$$

where ∂_p is shorthand for $\partial/\partial p$. Expressing the derivative as a “real part” assumes that the operator ∂_p is Hermitian, which is not guaranteed; however, the expression seems to be accurate in practice. The energy can then be optimized with any standard or specialized stochastic optimization methods.[13]–[15]

2.2.4 Diffusion Monte Carlo

Diffusion Monte Carlo (DMC) is a method for projecting out the ground state Φ_0 starting from any trial function Ψ_T , provided $\langle \Phi_0 | \Psi_T \rangle \neq 0$. In principle, it draws samples from the exact ground state wave function, but has the drawback that Φ_0 cannot be evaluated at a pre-selected position $\mathbf{R}$. In practice, the process scales exponentially unless an approximation is made to restrict the sign of the wave function: the fixed-node or fixed-phase approximation, which are used in practice, and which typically give highly accurate results. The details of the algorithm are discussed below.

Projection

The ground state of a Hamiltonian $\hat{H}$ can be extracted from a trial function by repeated application of an appropriate projector $\hat{G}$, [11], [16]

$$|\Phi_0\rangle = \lim_{N \rightarrow \infty} \hat{G}^N |\Psi_T\rangle, \quad (2.17)$$

where Ψ_T is a trial wave function having nonzero overlap with the ground state, and $\hat{G}$ has the same eigenstates Φ_i as $\hat{H}$ with eigenvalues $\lambda(E_i)$, such that $|\lambda(E_i)| < |\lambda(E_0)|$ for any $i > 0$. To see why this yields the ground state, we expand Ψ_T in the eigenstate basis Φ_i ,

$$\Psi_T = \sum_i c_i \Phi_i. \quad (2.18)$$

The operator acting repeatedly on an eigenstates is equivalent to multiplying by a power of the eigenvalue, resulting in

$$\hat{G}^N |\Psi_T\rangle = c_0 \lambda(E_0)^N |\Phi_0\rangle + c_1 \lambda(E_1)^N |\Phi_1\rangle + c_2 \lambda(E_2)^N |\Phi_2\rangle \dots \quad (2.19)$$

Normalizing by the zeroth coefficient,

$$\frac{\hat{G}^N |\Psi_T\rangle}{c_0 \lambda(E_0)^N} = |\Phi_0\rangle + \frac{c_1}{c_0} \left(\frac{\lambda(E_1)}{\lambda(E_0)} \right)^N |\Phi_1\rangle + \frac{c_2}{c_0} \left(\frac{\lambda(E_2)}{\lambda(E_0)} \right)^N |\Phi_2\rangle \dots \quad (2.20)$$

it becomes clear that in the limit as $N \rightarrow \infty$, all but the first term are suppressed to zero, since $|\lambda(E_i)/\lambda(E_0)| < 1$, leaving just the ground state.

The operator $\hat{G}^N$ is applied stochastically in the basis of real-space electron configurations as an integral over paths $\{\mathbf{R}_0 \rightarrow \mathbf{R}_1 \rightarrow \dots \rightarrow \mathbf{R}_N\}$,

$$\hat{G}^N |\Psi_T\rangle = \int d\mathbf{R}_0 \dots d\mathbf{R}_N |\mathbf{R}_N\rangle G_{\mathbf{R}_N, \mathbf{R}_{N-1}} \dots G_{\mathbf{R}_2, \mathbf{R}_1} G_{\mathbf{R}_1, \mathbf{R}_0} \langle \mathbf{R}_0 | \Psi_T \rangle, \quad (2.21)$$

where $G_{\mathbf{R}, \mathbf{R}'} = \langle \mathbf{R} | \hat{G} | \mathbf{R}' \rangle$. Expectation values over the ground state have the form

$$\frac{\langle \Phi_0 | \hat{O} | \Phi_0 \rangle}{\langle \Phi_0 | \Phi_0 \rangle} = \frac{\langle \Psi_T | \hat{G}^N \hat{O} \hat{G}^N | \Psi_T \rangle}{\langle \Psi_T | \hat{G}^{2N} | \Psi_T \rangle} \quad (2.22)$$

In diffusion Monte Carlo, the projector takes the form

$$G_{\mathbf{R}, \mathbf{R}'} = e^{-\tau(\hat{H} - E_T)} = e^{-\tau\hat{H}} e^{\tau E_T}, \quad (2.23)$$

where τ is a time step and E_T is an estimate of E_0 (whose value is unknown until the calculation is finished). The term containing E_T is a constant multiplier to correct the normalization, but does not affect the sampling.

The Trotter decomposition[17] is used to separate the kinetic ($\hat{T}$) and potential ($\hat{V}$) terms of the Hamiltonian,

$$e^{-\tau\hat{H}} \approx (e^{-\tau\hat{H}}) \approx (e^{-\tau\hat{T}} e^{-\tau\hat{V}}), \quad (2.24)$$

becoming exact in the limit $\tau \rightarrow 0$. The Coulomb potential is diagonal in the position basis (i.e. it is just a function of $\mathbf{R}$), so it contributes a factor to the weight w , but does not affect the sampling. The path is generated by choosing each successive position $\mathbf{R}'$ from the distribution

$$\langle \mathbf{R}' | e^{-\tau\hat{T}} | \mathbf{R} \rangle = e^{-\tau|\mathbf{R} - \mathbf{R}'|^2} \quad (2.25)$$

at the current position $\mathbf{R}$, and the weight factor is

$$w = \prod_{i=0}^N e^{-\tau V(\mathbf{R}_i)}. \quad (2.26)$$

As we saw with variational Monte Carlo, sampling using Gaussian moves is not very efficient, spending a lot of time sampling low-probability regions of configuration space. Additionally, the weights fluctuate enormously as $V(\mathbf{R})$ diverges at particle coincidences, causing high variance in the estimates of observables. These issues are improved by importance sampling, concentrating samples in regions of higher probability.

Importance Sampling

Instead of sampling the ground state distribution $|\Phi_0|^2$, diffusion Monte Carlo samples the mixed distribution $f(\mathbf{R}) = |\Phi_0(\mathbf{R})\Psi_T(\mathbf{R})|$. Expectation values take the form

$$E_{\text{DMC}}[\Psi_T] = \frac{\langle \Psi_T | \hat{H} | \Phi_0 \rangle}{\langle \Psi_T | \Phi_0 \rangle} = \left\langle \frac{\hat{H}\Psi_T(\mathbf{R})}{\Psi_T(\mathbf{R})} \right\rangle_{R \sim |\Psi_T \Phi_0|}. \quad (2.27)$$

The wave functions Ψ_T, Φ_0 are assumed to be real-valued in this section; the complex case is considered in Sec. 2.2.4.

We use the derivatives of the trial function $\Psi_T(\mathbf{R})$ to generate stochastic moves that evolve $f(\mathbf{R}, t)$. The time evolution of $f(\mathbf{R}, t)$ is given by the imaginary-time Schrödinger equation

$$\frac{\partial_t f(\mathbf{R}, t)}{f(\mathbf{R}, t)} = \frac{\partial_t \Phi(\mathbf{R}, t)}{\Phi(\mathbf{R}, t)} = -\frac{1}{2} \frac{\nabla^2 \Phi(\mathbf{R}, t)}{\Phi(\mathbf{R}, t)} + V(\mathbf{R}) \quad (2.28)$$

Substituting in the local energy E_{loc} of Ψ_T replaces the potential (thus reducing the large weight fluctuations when particles are near each other),

$$\frac{\partial_t f(\mathbf{R}, t)}{f(\mathbf{R}, t)} = -\frac{1}{2} \frac{\nabla^2 \Phi(\mathbf{R}, t)}{\Phi(\mathbf{R}, t)} + E_{\text{loc}} + \frac{1}{2} \frac{\nabla^2 \Psi_T(\mathbf{R})}{\Psi_T(\mathbf{R})}. \quad (2.29)$$

The Laplacian terms are expressed in terms of f as

$$\frac{\nabla^2 \Phi(\mathbf{R}, t)}{\Phi(\mathbf{R}, t)} - \frac{\nabla^2 \Psi_T(\mathbf{R})}{\Psi_T(\mathbf{R})} = \frac{\nabla^2 f}{f} - 2\nabla \mathbf{v}_D - 2\mathbf{v}_D \frac{\nabla f}{f}, \quad (2.30)$$

where $\mathbf{v}_D(\mathbf{R}) = \text{Re}(\nabla \Psi_T(\mathbf{R})/\Psi_T(\mathbf{R}))$ is called the drift velocity. Combining this with Eq. (2.29) yields the drift-diffusion equation for f ,

$$\partial_t f(\mathbf{R}, t) = -\frac{1}{2} \nabla^2 f + \nabla(\mathbf{v}_D f) + E_{\text{loc}} f. \quad (2.31)$$

In this final form, the diffusion part is equivalent to the force-biased moves in VMC, resulting in more efficient sampling. The weight factor becomes

$$w = \prod_{i=0}^N e^{-\tau E_{\text{loc}}(\mathbf{R}_i)}, \quad (2.32)$$

now using E_{loc} , which fluctuates very little compared to $V(\mathbf{R})$, leading to more stable evolution of the population of walkers.

Branching

Although the weight factor fluctuates less using importance sampling, it is still a product of N exponentials, and the weights of different walkers will exponentially vary. Continuing the random walk as described above, one walker will inevitably have the largest weight, and all other walkers will have negligible weight in comparison. Regardless of how many walkers are used, only one will contribute to the averages. To make efficient use of parallel computing resources, we limit the weights through a step called branching.

In its simplest form, branching is dividing a single walker into two walkers each having half the original weight, and both at the same configuration $\mathbf{R}_i$. All statistical averages remain unchanged by this procedure, but now each of the new walkers can propagate independently, better sampling the distribution. An obvious problem with simply branching walkers that have large weights is that the number of walkers increases exponentially, and the computer quickly runs out of resources. To offset this, some walkers also need to be removed. To preserve the statistical averages, a walker is kept with probability proportional to its weight w , and the new weight is set to one: the expectation value is therefore $\langle w \rangle = 1 \cdot w + 0 \cdot (1 - w) = w$. Removing walkers in this way introduces a small bias, but with a large enough population of walkers, the bias is insignificant.

Through the birth and death processes of branching, the number of walkers can be kept stable, but is not necessarily constant. For simpler implementation, these processes are combined to keep the number of walkers constant by the stochastic comb strategy.[18]–[20] In this strategy, the unit interval is divided into subintervals: the interval for walker n is

$$\left[\frac{\sum_{i < n} w_i}{\sum_j w_j}, \frac{\sum_{i \leq n} w_i}{\sum_j w_j} \right).$$

For N_w walkers, the comb is a set of N_w evenly spaced points on $[0, 1]$ shifted by a uniform random number $x \in [0, 1/N_w)$. Each point in the comb selects a new walker based on which interval it lands in; some walkers (typically larger weights) will be selected multiple times, and others (typically having smaller weights) not at all.

Given two walkers with coordinates $(\mathbf{R}_0, \mathbf{R}_1)$ and weights (w_0, w_1) , the expectation value of an operator O_L is

$$\langle O_L \rangle = \frac{1}{w_0 + w_1} (w_0 O_L(\mathbf{R}_0) + w_1 O_L(\mathbf{R}_1)). \quad (2.33)$$

Supposing $w_0 < w_1$, walker 0 will be kept with probability $2w_0/(w_0 + w_1)$; otherwise, R_1 will be copied to R_0 . This process is represented by the diagram in Fig. 2.2.4. In both cases, the weights are reset to $w_0 = w_1 = 1$.




Figure 2.1: Branching of two walkers into two new walkers. The vertical lines indicate the two walkers that are selected as the new walkers, based on the uniform random shift $x \in [0, 0.5)$. They are spaced 0.5 apart. The blue region (left) represents walker 0 and the green region (right) represents walker 1.



Figure 2.2: Branching of a population of walkers into the same number of new walkers. The vertical lines indicate the two walkers that are selected as the new walkers, based on the uniform random shift $x \in [0, 0.5)$. They are spaced 0.5 apart. The blue region (left) represents walker 0 and the green region (right) represents walker 1.

This results in the new expectation value

$$\langle O_L \rangle = \frac{2w_0}{w_0 + w_1} \frac{1}{2} (O_L(\mathbf{R}_0) + O_L(\mathbf{R}_1)) + \frac{w_1 - w_0}{w_0 + w_1} O_L(\mathbf{R}_1) = \frac{1}{w_0 + w_1} (w_0 O_L(\mathbf{R}_0) + w_1 O_L(\mathbf{R}_1)), \quad (2.34)$$

which is the same as before the branching.

With more than two walkers, the same stochastic comb strategy can be employed, illustrated in Fig. 2.2.4.

Fixed-Node Approximation

The expectation values from equation Eq. (2.27) are evaluated by integrals of the form

$$\frac{\langle \Phi_0 | \hat{O} | \Phi_0 \rangle}{\langle \Phi_0 | \Phi_0 \rangle} = \frac{\int d\mathbf{R} d\mathbf{R}' \Psi_T(\mathbf{R}') \hat{O} \hat{G}^N \Psi_T(\mathbf{R})}{\Psi_T(\mathbf{R}') \hat{G}^N \Psi_T(\mathbf{R})}, \quad (2.35)$$

where $\hat{G}_N$ is defined in Eq. (2.21). Fermion symmetry requires that $\Psi_T(\mathbf{R}')$ and $\Psi_T(\mathbf{R})$ have opposite signs when $\mathbf{R}'$ has two electrons' positions exchanged relative to $\mathbf{R}$. For real-valued trial functions Ψ_T , the numerator and denominator in Eq. 2.35 must have both positive and negative contributions (since only positive distributions can be sampled, the sign may be incorporated into the weights) and approach zero. Reducing the stochastic noise in this calculation scales exponentially; an alternative is to fix the nodal surface of the trial function Ψ_T and approximate the ground state Φ_0 by the lowest energy wave function sharing the nodal surface $\Psi_T(\mathbf{R}) = 0$. To implement this approximation, moves $\mathbf{R} \rightarrow \mathbf{R}'$ that result in a sign change,

$\Psi_T(\mathbf{R}')/\Psi_T(\mathbf{R}) < 0$, are rejected in the random walk, ensuring that the contribution from the trial function is always positive. This avoids the exponentially scaling sign problem at the cost of approximating the nodal surface, which can nevertheless be variationally improved by comparing different trial nodal surfaces.

Fixed-phase Approximation

The fixed-node approximation is extended to complex wave functions by fixing the phase at every point to match the trial function. The Schrödinger equation for complex wave functions is actually two equations, one for the real and one for the imaginary parts,

$$-\frac{1}{2} \left[\frac{\nabla^2 |\Psi(\mathbf{R})|}{|\Psi(\mathbf{R})|} - [\nabla \phi(\mathbf{R})]^2 \right] + V(\mathbf{R}) = E \quad (2.36)$$

$$\frac{1}{2} \nabla^2 \phi(\mathbf{R}) + \frac{\nabla |\Psi(\mathbf{R})|}{|\Psi(\mathbf{R})|} \nabla \phi(\mathbf{R}) = 0, \quad (2.37)$$

where the wave function is represented as $\Psi(\mathbf{R}) = |\Psi(\mathbf{R})|e^{i\phi(\mathbf{R})}$. The local energy $E_{\text{loc}}(\mathbf{R})$ is real-valued, given by Eq. (2.36),

$$E_{\text{loc}}(\mathbf{R}) = \text{Re} \left(\frac{\hat{H} \Psi(\mathbf{R})}{\Psi(\mathbf{R})} \right). \quad (2.38)$$

Eq. 2.36 has the same form as the usual Schrödinger equation, with an effective potential $V(\mathbf{R}) + [\nabla \phi(\mathbf{R})]^2/2$, and the wave function replaced by its magnitude $\Psi(\mathbf{R}) \rightarrow |\Psi(\mathbf{R})|$. Note that for real-valued wave functions, ratios such as $\nabla^2 \Psi(\mathbf{R})/\Psi(\mathbf{R})$ are unchanged by the absolute values. Thus, the importance sampling derivation follows the same steps as in Sec. 2.2.4.

These steps are briefly repeated here for the complex case. The time evolution given by the imaginary-time Schrödinger equation now includes the phase gradient,

$$\frac{\partial_t f(\mathbf{R}, t)}{f(\mathbf{R}, t)} = \frac{\partial_t \Phi(\mathbf{R}, t)}{\Phi(\mathbf{R}, t)} = -\frac{1}{2} \frac{\nabla^2 |\Phi(\mathbf{R}, t)|}{|\Phi(\mathbf{R}, t)|} + V(\mathbf{R}) + \frac{1}{2} \nabla \phi(\mathbf{R})^2. \quad (2.39)$$

Substituting in the local energy E_{loc} from Eq. (2.36) replaces both the phase gradient and the potential,

$$\frac{\partial_t f(\mathbf{R}, t)}{f(\mathbf{R}, t)} = -\frac{1}{2} \frac{\nabla^2 |\Phi(\mathbf{R}, t)|}{|\Phi(\mathbf{R}, t)|} + E_{\text{loc}} + \frac{1}{2} \frac{\nabla^2 |\Psi|}{|\Psi|}. \quad (2.40)$$

In Eq. 2.30 for the Laplacian terms, the drift velocity is defined as the real part of the gradient, $\mathbf{v}_D(\mathbf{R}) = \text{Re}(\nabla \Psi(\mathbf{R})/\Psi(\mathbf{R}))$. The final equation for the evolution of $f(\mathbf{R}, t)$ is the same as Eq. 2.31.

The difference in implementation for complex functions (fixed phase) compared to real functions (fixed node) has two parts: (1) taking the real part of the gradient for computing the drift, and (2) removing the sign-change rejection step. The local energy E_{loc} requires no modification to include the additional phase

term $\nabla\phi(\mathbf{R})^2/2$, since the kinetic energy Laplacian is evaluated directly on the complex wave function.

2.2.5 Effective Core Potentials (ECPs)

All-electron calculations (including all the electrons of a molecule or material in a calculation) are computationally expensive. One reason is that on a system with N_e electrons, all methods scale at least as N_e^3 , since that is the cost of evaluating a determinant. DFT and QMC typically achieve this N_e^3 scaling, while the other methods mentioned above scale even more quickly. In QMC, an additional expense is associated with the electrons close to a nucleus of charge Z . Since both the Coulomb interaction and kinetic energy diverge near $r = 0$, the variance in local energy is quite large. The short length scales near the nucleus also require a very small time step. These effects lead to a computational cost scaling as $Z^{5.5}$ for all-electron QMC calculations.[21]

To make calculations with heavy atoms feasible, the core electrons of an atom are replaced by an effective core potential (ECP) – a nonlocal function centered around the nucleus that represents the effect of the core electrons without having to include them explicitly in the calculation.[22] ECPs smooth out the potential energy surface for electrons near a nucleus, reducing the variance in QMC local energy in addition to reducing the total number of electrons. They are conventionally split into a local part $V_{\text{local}}^{\text{ECP}}(r)$ and a nonlocal part $V_{\text{nl},l}^{\text{ECP}}$ that depends on angular momentum l . The ECP operator for each atom acts on a function $f(\mathbf{r})$, with the nucleus at the origin, as

$$\hat{V}^{\text{ECP}}f(\mathbf{r}) = V_{\text{local}}^{\text{ECP}}(r)f(\mathbf{r}) + \sum_l V_{\text{nl},l}^{\text{ECP}}(r)Y_{lm}(\Omega)f_l(r), \quad (2.41)$$

where the l -component of $f(r, \Omega)$ is $f_l(r) = \int_{4\pi} Y_{lm}^*(\Omega')f(r, \Omega')d\Omega'$, and (r, Ω) are the radius and angle in the spherical representation of $\mathbf{r}$. The functions $V_{\text{local}}^{\text{ECP}}(r)$, $V_{\text{nl},l}^{\text{ECP}}(r)$ are optimized to match energies between all-electron and ECP calculations on single atoms and small molecules. These calculations may be done with DFT; however for correlated calculations, the ECPs need to be optimized for a correlated method. Since ECPs are localized in the cores of the atoms, they are generally transferable to larger systems with many atoms.

The many-body ECP operator $V^{\text{ECP}}(\mathbf{R}, \mathbf{R}')$ is a sum of independent terms given by Eq. 2.41 for each electron and each atom. Is a nonlocal operator, depending on two many-body configurations, $\mathbf{R}$ and $\mathbf{R}'$. Importantly for QMC approaches, this form is nonlocal only in the angular channel, making the local evaluation feasible: it only requires an integral over the angular coordinates of $\mathbf{r}_{ea}$, not over the radial coordinate r_{ea} .

2.3 Observables

2.3.1 Structure Factor $S(\mathbf{q})$

The structure factor $S(\mathbf{q})$ is a property of bulk systems measuring two-particle correlations. The structure factor is related to the Fourier transform of the pair correlation function, giving information about how the electrons scatter radiation as a function of momentum $\mathbf{q}$. In diffraction experiments, which measure $S(\mathbf{q})$, $\mathbf{q}$ is the momentum exchanged between the material and the diffracted particles.

$S(\mathbf{q})$ is defined as[11]

$$S(\mathbf{q}) = \frac{1}{N} \langle \Psi | \sum_{i,j} e^{i\mathbf{q} \cdot (\mathbf{r}_i - \mathbf{r}_j)} | \Psi \rangle, \quad (2.42)$$

To efficiently evaluate $S(\mathbf{q})$, it is rewritten as

$$S(\mathbf{q}) = \frac{1}{N} \langle \Psi | \sum_i e^{i\mathbf{q} \cdot \mathbf{r}_i} \sum_j e^{-i\mathbf{q} \cdot \mathbf{r}_j} | \Psi \rangle \quad (2.43)$$

$$= \frac{1}{N} \langle \Psi | \sum_i e^{i\mathbf{q} \cdot \mathbf{r}_i} |^2 | \Psi \rangle, \quad (2.44)$$

which only depends on the coordinates $\mathbf{r}_i$, not the pair distances $(\mathbf{r}_i - \mathbf{r}_j)$.

The spin structure factor is a related quantity weighted by spin, measuring the spin-spin correlations of electrons,

$$S(\mathbf{q}) = \frac{1}{N} \langle \Psi | \sum_{i \neq j} s_i s_j e^{i\mathbf{q} \cdot \mathbf{r}_{ij}} | \Psi \rangle = \frac{1}{N} \langle \Psi | \sum_i s_i e^{i\mathbf{q} \cdot \mathbf{r}_i} |^2 | \Psi \rangle, \quad (2.45)$$

where s_i is the spin of electron i .

2.4 Polarization and Quadrupole Moments in Bulk Systems

Like the structure factor, polarization is a property of a bulk material, usually represented with periodic boundary conditions in calculations. Although polarization had been known and studied for decades, it was not until the 1990s that a consistent approach was developed to calculate it for periodic systems using noninteracting methods[23], [24] and later for many-body methods as well.[25] A related quantity, the quadrupole moment, is a two-dimensional analog of one-dimensional polarization. An approach to calculating bulk quadrupole moments was recently introduced for noninteracting periodic systems.[26], [27] An open question was to develop a formalism for this quantity that can also be used in many-body methods, the subject of Chapter 3. This section provides the background for that work, summarizing the development of the modern theory of polarization and its many-body representation as well as the formalism for computing

the bulk quadrupole moment in noninteracting systems.

Polarization is traditionally defined as the dipole moment per unit volume. Conceptually, the simplest dipole moment is formed by a pair of equal and opposite charges separated by a small distance. In a finite system of classical charges, the dipole moment $\mathbf{p}$ is straightforward to calculate: it is the sum of the positions of the particles, $\mathbf{r}$, weighted by their charge values, q_i ,

$$\mathbf{p} = \sum_i q_i \mathbf{r}_i. \quad (2.46)$$

The dipole moment, defined this way, is independent of the choice of origin as long as the total charge is zero. If the origin is shifted from $(0, 0, 0)$ to $\mathbf{x}_0$, every position $\mathbf{r}$ is shifted to $\mathbf{r} - \mathbf{x}_0$, changing the sum by $-\sum_i q_i \mathbf{x}_0$. This term is proportional to the total charge $\sum_i q_i$; if the total charge is not zero, the dipole moment is not well defined.

In the periodic simulation cell, the position $\mathbf{r}$ is not continuous at the cell boundaries, since particles moving across the boundary at one end enter the cell from the other end, instantaneously jumping across the cell. In addition to being independent of the choice of origin of the coordinate system, the dipole moment should be independent of the choice of unit cell; however, it is easy to see that is not the case for the definition above. Suppose a positive and negative charge are arranged in a unit cell so that there is a small positive dipole moment. If a new cell boundary is chosen to be between the two particles, they suddenly appear separated on opposite sides of the cell, and the formula yields a very large negative dipole moment. In a similar example, if a dipole is created by moving a positive charge halfway across the cell, the final configuration is the same as if the charge had been moved in the opposite direction, generating a dipole moment of opposite sign. The effects of such a process on the bulk simulation cell would appear at the surface, where a positive or negative charge builds up, and identifies the direction of the polarization.

The solution to the problem comes from reframing polarization as the result of integrating a current. Polarization in physical systems is the result of charge moving through an insulator, i.e. current flowing. Such a polarization current can be thought of as the rearrangement of the ground state charge density in response to a slow change in the external potential. Niu and Thouless[28], [29] took this perspective and showed that a periodically varying potential can induce a current that pumps charge across a unit cell over once cycle. In particular, they showed that the charge pumped across one cell of an infinite crystal must be quantized – consistent with the intuition that it must be a whole number of electrons. Thouless’s formula was the foundation for a more general formula introduced by King-Smith and Vanderbilt[23] that did not require the time-varying potential to return to its starting point. They showed that the integral of polarization

current over a time-dependent process can be written as a pure difference between initial and final values

$$P_\alpha = P_\alpha^{(1)} - P_\alpha^{(0)} \quad (2.47)$$

of the quantity

$$P_\alpha^{(\lambda)} = \frac{2i}{8\pi^3} \sum_{n=1}^M \int_{BZ} d\mathbf{k} \langle u_{\mathbf{k}n}^{(\lambda)} | \frac{\partial}{\partial k_\alpha} | u_{\mathbf{k}n}^{(\lambda)} \rangle, \quad (2.48)$$

where λ takes values of 0 and 1 for the initial and final states, respectively, $u_{\mathbf{k}n}$ is a Bloch function with reciprocal wavevector $\mathbf{k}$ and band index n , and α is a coordinate index. That formula, involving an integral of Bloch functions over the Brioullin zone, allows the bulk polarization of a solid to be computed in a single band structure calculation, rather than requiring the full time-dependent process to be simulated.

In practice, a discretized version of Eq. 2.48 is used that can be computed over a finite mesh of $\mathbf{k}$ points. The discretized formula is used in the Wilson loop formalism by Benalcazar *et al* [26] discussed below. The clearest explanation comes from considering the formula applied to a many-body wave function, as done by Souza *et al* [30] and letting $\Phi_{\mathbf{k}}$ be the cell-periodic part of a Slater determinant of Bloch functions. The many-body version of Eq. 2.48 is

$$P_\alpha^{(\lambda)} = \frac{2i}{8\pi^3} \int_{BZ} d\mathbf{k} \langle \Phi_{\mathbf{k}}^{(\lambda)} | \frac{\partial}{\partial k_\alpha} | \Phi_{\mathbf{k}}^{(\lambda)} \rangle, \quad (2.49)$$

where the only difference is no longer summing over a band index. The discretization only needs to be considered along the axis α , since the derivative in the integrand is along α ; the perpendicular integrals can be computed by the usual sum. Discretizing the derivative along α to a resolution of Δk has the form

$$P_\alpha^{(\lambda)} = \text{Im} \sum_{\mathbf{k}} \langle \Phi_{\mathbf{k}}^{(\lambda)} | \Phi_{\mathbf{k}+\Delta\mathbf{k}}^{(\lambda)} \rangle - 1 \quad (2.50)$$

$$\approx \text{Im} \sum_{\mathbf{k}} \ln \langle \Phi_{\mathbf{k}}^{(\lambda)} | \Phi_{\mathbf{k}+\Delta\mathbf{k}}^{(\lambda)} \rangle \quad (2.51)$$

$$= \text{Im} \ln \prod_{\mathbf{k}} \langle \Phi_{\mathbf{k}}^{(\lambda)} | \Phi_{\mathbf{k}+\Delta\mathbf{k}}^{(\lambda)} \rangle. \quad (2.52)$$

In the non-interacting case, $\Phi_{\mathbf{k}}$ is a Slater determinant of cell-periodic Bloch functions $u_{\mathbf{k}n}$, so it satisfies the identity

$$\langle \det u_{\mathbf{k}n}(\mathbf{r}_i) | \det u_{\mathbf{k}m}(\mathbf{r}_j) \rangle = \det \langle u_{\mathbf{k}n} | u_{\mathbf{k}m} \rangle. \quad (2.53)$$

This leads to the discretized formula of King-Smith and Vanderbilt on a mesh of k points $\mathbf{k}_j$ in one dimension,

$$P^{(\lambda)} = \text{Im} \ln \prod_j \det F_{\mathbf{k}_j}^{mn}, \quad (2.54)$$

where the overlap matrix $F_{\mathbf{k}_j}^{mn} = \langle u_{\mathbf{k}_j m} | u_{\mathbf{k}_{j+1} n} \rangle$ is the overlap between single particle states in bands m and n at adjacent k points in the mesh. The logarithm in Eq. 2.54 makes the formulas work out, but will be given physical interpretation in the discussion of the many-body formula below.

With a formula available for non-interacting calculations (including DFT), it remained to find a formulation compatible with many-body methods. An important insight was made by Ortiz and Martin[31]: since charges are localized in an insulator, the system can be divided into bulk and surface regions that do not exchange charge. This charge conservation allows for the existence of bulk polarization in insulators, but breaks down in metals. They used this fact to extend the non-interacting formulas to the many-body case, effectively replacing the integral over the Brillouin zone with an average over boundary conditions of different twists. A few years later, Resta[25] showed another formula to be equivalent in the thermodynamic limit, one with a big advantage: requiring only a single calculation, instead of many with different twisted boundary conditions. Resta's formula involves putting the traditional dipole, $\hat{\mathbf{X}} = \sum_i \mathbf{x}_i$, into a complex exponential

$$\hat{U}_\alpha = \exp \left[\frac{2\pi i \hat{X}^\alpha}{L_\alpha} \right], \quad (2.55)$$

where α is the axis to compute along. The exponential is averaged over the many-body wave function, $z_\alpha = \langle \Psi_0 | \hat{U}_\alpha | \Psi_0 \rangle$, and the polarization is extracted by restoring units to the complex phase,

$$P^\alpha = \frac{L_\alpha}{2\pi \mathcal{V}} \text{Im} \ln z_\alpha, \quad (2.56)$$

where $\mathcal{V}$ is the volume of the system. Treating the dipole as a complex phase naturally wraps the cell into a circle, removing the discontinuity in position. It also adds an additional quantity to the problem: the magnitude of the complex average. If all samples X are localized to a single point in the unit cell, they map onto a single point of the unit circle, and the average has magnitude one – this is the infinitely localized case, where the wave function is a delta function. As points spread out, the average moves into the interior of the unit circle, and in the maximally spread out case, the average is clearly exactly at the center, i.e. has magnitude zero. Supporting this intuition, it was shown that the magnitude of z provides a quantitative measure of localization length[32],

$$\lambda_\alpha^2 = - \left(\frac{L_\alpha}{2\pi} \right)^2 \ln |z_\alpha|^2, \quad (2.57)$$

where λ_α^2 is the variance of the position along axis α . In the delocalizing case $|z| \rightarrow 0$, the localization length $\lambda \rightarrow \infty$ diverges, indicating a metallic state. The phase of z becomes undefined, correctly indicating that the polarization is undefined in a metal.

There is a connection between this formula and the one suggested by King-Smith and Vanderbilt. The exponential operator of position used in the many-body polarization formula shifts the k index of the wave function by the coefficient $\mathbf{a}$,

$$\langle \Phi_{\mathbf{k}} | e^{i\mathbf{a} \cdot \hat{\mathbf{X}}} | \Phi_{\mathbf{k}} \rangle = \langle \Phi_{\mathbf{k}} | \Phi_{\mathbf{k}+\mathbf{a}} \rangle. \quad (2.58)$$

Here, $\Phi_{\mathbf{k}}(\mathbf{R}) = e^{-i\mathbf{k} \cdot \mathbf{R}} \Psi_{\mathbf{k}}(\mathbf{R})$ is the cell-periodic part of the wave function, as in the discussion of the discretized formula above. The overlap of the functions $\Phi_{\mathbf{k}}$ is exactly the one in the integrand of Eq. 2.50. This connection explains the appearance of a logarithm in that formula: it is extracting a complex phase. A thorough discussion and interpretation of these formulas is given by Souza, Wilkins, and Martin.^[30]

About a decade after these polarization formulas were established, Bernevig, Hughes, and Zhang considered how symmetry limits possible polarization states.^[33] In a molecule, mirror symmetry forbids a dipole moment from forming across the mirror plane. However, in an infinite solid, both the unpolarized state and the $\pi/2$ polarized state satisfy mirror symmetry. There is a caveat, namely, a real material always has boundaries, and the surfaces can't be left with asymmetric charges. Imagining the polarization as a pumping process, if a charge is pumped halfway across the cell, then half a charge must be left at each surface. While this may sound odd, it is actually realized in models having time-reversal and mirror symmetry, as well as real materials.^[33] When spin is considered, the “half charge” manifests as a half-filled surface state, each end having opposite spin. The half-filled states, required by the symmetry of the system, lead to a surprising phenomenon: an insulator that conducts on the surface. These materials were dubbed “topological” insulators, since their polarization is quantized by the Chern number of the band structure – the same integer quantity computed by Thouless in his quantized charge pump – and is insensitive to any perturbations that preserve the bulk insulating gap. (If the gap is closed, bulk polarization becomes undefined, and there is no guarantee of preserving the surface states.) The discovery of symmetry-protected topological phases of matter led to interest in the community in exploring new possible topological phases.^{[34]–[38]}

One idea pursued by Benalcazar, Bernevig, and Hughes^[26] was to consider higher-order generalizations of the dipole, i.e. quadrupole and octupole moments. This overview will focus only on the quadrupole case, since all the work on the octupole is a natural analog to the quadrupole results. Analogous to the dipole, the quadrupole can be thought of as two equal and opposite dipoles separated by a distance. The quadrupole

moment in an open system of point charges is defined as

$$Q_{xy} = \sum_i q_i x_i y_i, \quad (2.59)$$

where q_i is the charge of particle i , as before. $Q_{\alpha\beta}$ is a tensor, since α, β each range over x, y, z , and each component depends on the coordinate frame it is measured in. The condition for origin independence can be seen by imagining a shift in the origin in the y -direction by y_0 ,

$$Q'_{xy} = \sum_i q_i x_i (y_i - y_0) = Q_{xy} - y_0 \sum_i q_i x_i. \quad (2.60)$$

The extra term appearing from the shift is exactly the dipole moment along the x -direction. A similar term appears from shifting the origin in the x -direction. The dipole moment (and therefore also the total charge) of the system must be zero for the quadrupole to be well-defined.

Like the dipole, the quadrupole formula is a function of position operators and under periodic boundary conditions faces the problem of discontinuous position across the boundaries. Benalcazar *et al* used an approach of nested Wilson loops to solve the problem.[26] A Wilson loop $\mathcal{W}_{x,\mathbf{k}}$ is related to the discretized polarization formula of King-Smith and Vanderbilt from Eq. 2.54,

$$\mathcal{W}_{x,\mathbf{k}} = \prod_j F_{x,\mathbf{k}_j} = \prod_j \langle u_{\mathbf{k}_j m} | u_{\mathbf{k}_{j+1} n} \rangle, \quad (2.61)$$

where the x indicates that the set of $\mathbf{k}$ points is taken along the x direction. The difference in this case, compared to computing polarization, is that the result is still a matrix rather than a determinant. The eigenvalues of $\mathcal{W}_{x,\mathbf{k}}$ have complex phases proportional to the Wannier centers of the wave function. With two electrons per cell, two Wannier bands can be computed in this way. The polarizations of these bands are each be computed by integrating over k_y , which they call “Wannier-sector polarizations” $p_y^{v_x^\pm}$. Performing the integrals along y and x in the opposite order gives corresponding Wannier-sector polarizations $p_x^{v_y^\pm}$. The relation of the bulk quadrupole moment q_{xy} to these Wannier-sector polarizations depends on the symmetries of the system; for example, under mirror and inversion symmetries $M_x, M_y, \mathcal{I}$, the quadrupole moment is $q_{xy} = 2p_y^{v_x^\pm} p_x^{v_y^\pm} = 0$ or $1/2$.

Although the calculation is somewhat complicated, the basic idea of this formulation is computing the Wannier centers along one direction for two particles per unit cell using the existing non-interacting framework for polarization. Ensuring that the Wannier centers are separated at every point in $\mathbf{k}$ space, a separate Wannier-sector polarization is computed for each of them. These Wannier-sector polarizations determine

the bulk quadrupole moment. Analogous to Thouless's charge pump for polarization, the generation of Wannier-sector polarizations is interpreted as pumping dipoles across the system by transporting charges along Wannier bands. While the Wilson loop approach can calculate quadrupole and higher multipole moments of non-interacting systems both in quantized and non-quantized regimes, the formulas are not directly applicable to many-body interacting systems. For polarization, one can consider the phase of a many-body ground state parameterized by twisted periodic boundary conditions[30], [31] discussed above, but it is not obvious how this approach can be practically extended to the nested Wilson loop approach to calculating quadrupole and higher multipole moments.

CHAPTER 3

BULK ELECTRIC QUADRUPOLE MOMENTS

This chapter is based on my published work “Many-body electric multipole operators in extended systems,” [39] in which we provide a formulation of the quadrupole and octupole moments in terms of ground-state expectation values of many-body operators, analogous to Resta’s formulation of the charge polarization [25].

The modern theory of electric polarization in crystals has had a wide-ranging impact over the past 25 years [23], [24], [31], [40]–[42]. Aspects of the theory are not only useful for predictions of observable properties of real materials, but are tightly connected to a variety of topological insulator phenomena. For example, insulating crystals having a symmetry that would naively forbid the existence of a dipole moment, may instead have non-vanishing, but quantized, polarization if they are in a topological insulator phase [43]–[46]. Furthermore, by tracking changes in polarization in an adiabatic cycle, one can observe quantized charge pumping characterized by a Chern number topological invariant [28]. More recently, the theory of quantized electronic polarization in topological insulators was extended to account for higher electric multipole moments [26], [27]. Therein it was shown that two and three dimensional, non-interacting topological crystalline phases of matter can support quantized electric quadrupole or octupole moments.

We define a many-body quadrupole operator and discuss its connection to the flow of adiabatic dipole currents. In addition, we test our operator for some explicit model systems: (1) We show that our formula reproduces the phase diagrams of the models from Ref. [26]; (2) we use our formula to track the changes in quadrupole moment during the higher order pumping process in Ref. [27], and we show that the results match those from Ref. [27]; (3) we test our formula on a model with corner charges but no bulk quadrupole moment, and show that our operator indicates zero quadrupole moment despite the presence of corner charges.

Higher electric multipole operators, such as the quadrupole $\hat{q}^{ij} = \sum_n \hat{x}_n^i \hat{x}_n^j$ or octupole $\hat{o}^{ijk} = \sum_n \hat{x}_n^i \hat{x}_n^j \hat{x}_n^k$, are problematic in periodic systems for the same reason as the many-body polarization operator discussed in Sec. 2.4 – the discontinuous position operators. We follow the same line of reasoning that led to Resta’s polarization formula, [25] examining exponentiated versions of these operators. Consider a lattice with primitive vectors $\mathbf{a}_b$ and reciprocal lattice vectors $\mathbf{b}_a$, satisfying $\mathbf{b}_a \cdot \mathbf{a}_b = \delta_{ab}$. The exponentiated operators

(including the polarization operator here for reference) have the form

$$\hat{\mathcal{U}}_a^P = \exp \left[\frac{2\pi i \sum_i X^i b_a^i}{N_a} \right], \quad (3.1)$$

$$\hat{\mathcal{U}}_{ab}^Q = \exp \left[\frac{2\pi i \sum_{ij} \hat{q}^{ij} b_a^i b_b^j}{N_a N_b} \right], \quad (3.2)$$

$$\hat{\mathcal{U}}_{abc}^O = \exp \left[\frac{2\pi i \sum_{ijk} \hat{o}^{ijk} b_a^i b_b^j b_c^k}{N_a N_b N_c} \right], \quad (3.3)$$

where N_a is the number of unit cells in the a -th lattice direction. The bulk moments are then defined as

$$P^a = \lim_{N_a \rightarrow \infty} \frac{eL_a}{2\pi\mathcal{V}} \text{Im} \ln \langle \Psi_0 | \hat{\mathcal{U}}_a^P | \Psi_0 \rangle \quad (3.4)$$

$$Q^{ab} = \lim_{N_a, N_b \rightarrow \infty} \frac{eL_a L_b}{2\pi\mathcal{V}} \text{Im} \ln \langle \Psi_0 | \hat{\mathcal{U}}_{ab}^Q | \Psi_0 \rangle \quad (3.5)$$

$$O^{abc} = \lim_{N_a, N_b, N_c \rightarrow \infty} \frac{eL_a L_b L_c}{2\pi\mathcal{V}} \text{Im} \ln \langle \Psi_0 | \hat{\mathcal{U}}_{abc}^O | \Psi_0 \rangle. \quad (3.6)$$

As mentioned above, I will focus mainly on the quadrupole formulas, since the octupole case is analogous.

3.1 Electric Moments as Integrated Current

We will start by showing that our definition of the quadrupole moment (Eqs. 3.2, 3.5) matches the expected physical observables by generalizing Resta's arguments from Ref. [25]. Consider a generic many-body Hamiltonian ¹

$$H = \sum_{n=1}^{N_{elec}} \sum_{i=x,y,z} \frac{(\hat{p}_n^i)^2}{2m} + V(\hat{\mathbf{X}}). \quad (3.7)$$

We will examine how the unitary operators in Eqs. 3.1-3.3 transform the Hamiltonian. These operators all involve different functions of positions $\hat{\mathbf{X}}$; to generalize, let us consider an arbitrary real differentiable function $\lambda(\hat{\mathbf{X}})$ and the corresponding unitary transformation $\hat{U}_\lambda(\alpha) = \exp \left(i\alpha \lambda(\hat{\mathbf{X}}) \right)$. This operator transforms each momentum operator of the Hamiltonian as

$$\hat{U}_\lambda(\alpha) \hat{p}_n^i \hat{U}_\lambda^\dagger(\alpha) = \hat{p}_n^i - \hbar \alpha \frac{\partial \lambda(\hat{\mathbf{X}})}{\partial \hat{x}_n^i}, \quad (3.8)$$

and leaves the potential energy operator unchanged. With this, we can define a transformed Hamiltonian $H(\alpha) = \hat{U}_\lambda(\alpha) H \hat{U}_\lambda^\dagger(\alpha)$. The difference between $H(0)$ and $H(\alpha)$ can be expressed as an electromagnetic vector

¹The superscript i on the momentum operator $\hat{p}_n^i$ indicates the x , y , or z component; we also use numerals 1, 2, and 3 to refer to these components (i.e. p^1 , p^2 , p^3 or x^1 , x^2 , x^3).

potential acting on each electron, e.g.

$$A^i(\hat{\mathbf{x}}_n) = -\frac{\hbar\alpha}{e} \frac{\partial \lambda(\hat{\mathbf{X}})}{\partial \hat{x}_n^i}$$

for the n -th electron, since a vector potential shows up in the Hamiltonian as a shift in the momentum operator. We are interested in evaluating expectation values on the ground state $|\Psi_0\rangle$ of H , as in Resta's formula. If we assume $|\Psi_0\rangle$ is non-degenerate (which is usually a reasonable assumption), then $U_\lambda(\alpha)|\Psi_0\rangle$ is an eigenstate of $H(\alpha)$,

$$H(\alpha)U_\lambda(\alpha)|\Psi_0\rangle = E_0U_\lambda(\alpha)|\Psi_0\rangle,$$

which comes from rearranging the definition of $H(\alpha)$.

Returning to something more concrete, we will reproduce Resta's polarization argument by considering $\lambda(\hat{\mathbf{X}}) = \hat{X}^1$. This generates a constant vector potential $A^x = -\frac{\hbar\alpha}{e}$, $A^y = A^z = 0$, which is equivalent to an electric field – the quantity that couples to polarization. To have $U_{\hat{X}^x}(\alpha)|\Psi_0\rangle$ satisfy periodic boundary conditions, we can choose $\alpha = 2\pi/L_x$. This way, if an electron moves a distance L_x in the x direction, the phase changes by exactly 2π , and the operator becomes

$$e^{i\frac{2\pi}{L_x}(\hat{\mathbf{X}}+L_x)} = e^{i\frac{2\pi}{L_x}\hat{\mathbf{X}}}e^{i2\pi} = \hat{U}_{\hat{\mathbf{X}}^x}\left(\frac{2\pi}{L_x}\right),$$

returning to the same value it started with. When approaching the thermodynamic limit, L_x is large and we can treat α as a small parameter. Thus we can treat $U_{\hat{X}^x}(2\pi/L_x)|\Psi_0\rangle$, an eigenstate of $H(2\pi/L_x)$, as a perturbed version of the initial ground state $|\Psi_0\rangle$, and so we can expand $U_{\hat{X}^x}(2\pi/L_x)|\Psi_0\rangle$ in terms of the eigenstates of $H(0)$ using perturbation theory (this is why we need $U_{\hat{\mathbf{X}}^x}(2\pi/L_x)|\Psi_0\rangle$ to be periodic – expanding it in terms of eigenstates of $H(0)$ requires it to obey the same boundary conditions). The perturbation term in the Hamiltonian is $H' = -\frac{2\pi\hbar}{mL_x} \sum_{n=1}^{N_{elec}} \hat{p}_n^x$, and the leading order correction to the unperturbed ground state is

$$U_{\hat{X}^x}(2\pi/L_x)|\Psi_0\rangle \approx e^{i\gamma_{px}} \left(|\Psi_0\rangle - \frac{2\pi\hbar}{mL_x} \sum_{j \neq 0} |\Psi_j\rangle \frac{\langle \Psi_j | \hat{\mathcal{P}}^x | \Psi_0 \rangle}{E_0 - E_j} \right), \quad (3.9)$$

where $|\Psi_j\rangle$ are the excited states of $H(0)$ with energies E_j , and $\hat{\mathcal{P}}^a = \sum_{i \in \text{electrons}} \hat{p}_i^a$ is the many-body momentum operator of all the electrons. In addition to the usual form, we have allowed for a phase factor γ_{px} which determines the x -component of the electric polarization.^[25] As we stated earlier, we now see explicitly that in the thermodynamic limit

$$\lim_{L_x \rightarrow \infty} U_{\hat{X}^x}\left(\frac{2\pi}{L_x}\right)|\Psi_0\rangle = e^{i\gamma_{px}}|\Psi_0\rangle.$$

To connect to the physical polarization we can show that the time-derivative of this polarization definition matches the physically expected result of an electric current. Indeed, Ref. [25] shows from the perturbation theory calculation above that in 1D

$$\frac{dP^1}{dt} \approx \frac{ie\hbar}{mL_1} \sum_{j \neq 0} \langle \dot{\Psi}_0 | \Psi_j \rangle \frac{\langle \Psi_j | \hat{\mathcal{P}}^1 | \Psi_0 \rangle}{E_0 - E_j} + \text{c.c.}, \quad (3.10)$$

which is the adiabatic electric charge current.[28]

Now consider the quadrupole moment $\hat{Q}^{12}$. We will consider a unitary transformation $\hat{U}_\lambda(\alpha)$ with $\lambda(\hat{\mathbf{X}}) = \sum_{n=1}^{N_{elec}} \hat{x}_n^1 \hat{x}_n^2$ and with $\alpha = 2\pi/L_1 L_2$. Using the same arguments as the case for polarization, we can use perturbation theory to calculate

$$U_{\hat{Q}^{12}}(2\pi/L_1 L_2) |\Psi_0\rangle \approx e^{i\gamma_{q^{12}}} \left(|\Psi_0\rangle - \frac{2\pi\hbar}{L_1 L_2} \sum_{j \neq 0} |\Psi_j\rangle \frac{\langle \Psi_j | \hat{\mathcal{J}}_D^{12} | \Psi_0 \rangle}{E_0 - E_j} \right) \quad (3.11)$$

where $\hat{\mathcal{J}}_D^{12} = \frac{1}{m} \sum_{n=1}^{N_{elec}} (\hat{p}_n^1 \hat{x}_n^2 + \hat{p}_n^2 \hat{x}_n^1)$ is the *dipole-current* operator. Note that the velocity $d\hat{x}_n^i/dt = (i/\hbar)[H, \hat{x}_n^i] = \hat{p}_n^i/m$ is like the time derivative of the dipole operator, while the time derivative of the quadrupole operator is $d(\hat{x}_n^1 \hat{x}_n^2)/dt = (i/\hbar)[H, \hat{x}_n^1 \hat{x}_n^2] = (1/m)(\hat{p}_n^1 \hat{x}_n^2 + \hat{p}_n^2 \hat{x}_n^1)$, hence our notion of a dipole current operator. From this result we can further calculate

$$\frac{dQ^{12}}{dt} \approx \frac{ie\hbar}{L_1 L_2} \sum_{j \neq 0} \langle \dot{\Psi}_0 | \Psi_j \rangle \frac{\langle \Psi_j | \hat{\mathcal{J}}_D^{12} | \Psi_0 \rangle}{E_0 - E_j} + \text{c.c.} \quad (3.12)$$

which relates the time derivative of our definition of the quadrupole to a dipole-current as we would physically expect. We note that the above expression for $\hat{\mathcal{J}}_D^{12}$ (and thus Eqs. 3.11 and 3.12) cannot be evaluated in a periodic system in its current form because it contains position operators. Nonetheless, the notion of a dipole current is useful for understanding the origin of a bulk quadrupole moment as the result of a process that rearranges charge moments, even in extended systems. The problem with periodicity is reminiscent of the polarization operator $\hat{\mathbf{X}}$, which cannot be evaluated directly in a periodic system, but can be computed from the operator in Eq. 3.1. Finding an alternate expression for $\hat{\mathcal{J}}_D^{12}$ that is compatible with periodic boundary conditions is an interesting direction for future work.

3.2 Adiabatic Evolution

The charge pumping process that generates polarization is driven by a uniform electric field over the system. An electric field can be represented by a time-varying vector potential, which lets us track how the system has changed after the electric field has been applied. Here, we review the “adiabatic evolution” description – gradually changing some parameter during a process – of polarization under uniform electric field and extend it to describe the evolution of quadrupole moments under electric field gradients.

The action of $\hat{U}_{\hat{X}^1}(2\pi/L_1)$ on a state can be treated as adiabatic evolution from a system with a vanishing vector potential to one with $A^1 = -h/eL_1$. The process can be accomplished through a time-dependent vector potential of the form $A^1 = -\frac{ht}{eL_1T}$, and which is 0 for $t < 0$, and $-h/eL_1$ for $t > T$. The process turns on a uniform electric field $E^1 = -\frac{h}{eL_1T}$ during the time interval $t \in [0, T]$. We choose the duration T to be very large so that the strength of E^1 is small, approaching the adiabatic limit. We are interested in the evolution of the non-degenerate ground state of a neutral insulator, and so, in the adiabatic limit, the presence of the uniform electric field will only generate an additional phase factor $\gamma_{p1} = -\hbar^{-1} \int_0^T \mathbf{d} \cdot \mathbf{E} dt$, where $\mathbf{d}$ is the total dipole moment of the material. Besides the weak, time-varying vector potential, we assume the insulator is otherwise static so that $\mathbf{d}$ is time-independent. Hence, we find the phase (cf. Eq. 3.9)

$$\begin{aligned} \gamma_{p1} &= \hbar^{-1} d^1 \int_0^T \partial_t A^1 dt = d^1 \cdot (A^1(T) - A^1(0)) \\ &= -2\pi d^1 / eL_x \end{aligned} \quad (3.13)$$

which exactly produces the right value of $P^1 = \frac{e}{2\pi} \text{Im} \ln \langle \Psi_0 | \hat{U}_{\hat{X}^1}(2\pi/L_1) | \Psi_0 \rangle$ in the thermodynamic limit.

The essential idea in the previous paragraph is that, to detect a polarization in the material, we adiabatically turn on a small electric field and track how the phase of the ground state responds during this process. For the quadrupole, instead of turning on a uniform electric field, we turn on a uniform electric field *gradient*. The quadrupolar operator $U_{\hat{Q}^{12}}(2\pi/L_1L_2)$ can be interpreted as adiabatically evolving our ground state from a vanishing vector potential at $t = 0$ to the vector potential $A^i = -\frac{h\sigma^{ij}x^j}{eL_1L_2}$ at $t = T$. This process can be carried out using a time-dependent vector potential

$$A^i = -\frac{ht\sigma^{ij}x^j}{eL_1L_2T} \quad i, j = 1, 2, \quad (3.14)$$

where $\sigma^{12} = \sigma^{21} = 1$, $\sigma^{11} = \sigma^{22} = 0$, and T is large. This vector potential represents a constant electric field gradient. Thus, a natural way to interpret the proposed quadrupole operator is as an evolution process where a small, uniform electric field gradient is turned on for a finite amount of time. During this time, the

ground state will develop the phase factor γ_{q12} shown in Eq. 3.11, in the thermodynamic limit. Analogous to polarization, we have to consider a non-degenerate ground state in a neutral insulator that is also unpolarized and has a gap to dipole excitations.² Since the Q^{12} quadrupole moment couples to an electric field gradient $\frac{1}{2}(\partial_1 E^2 + \partial_2 E^1)$, its contribution to the phase will be $\gamma_{q12} = -\frac{1}{2\hbar} \int_0^T (q^{12} \partial_1 E^2 + q^{21} \partial_2 E^1) dt$. We assume that the quadrupole moment of our system is static so we can simplify this expression to

$$\begin{aligned} \gamma_{q12} &= \frac{q^{12}}{2\hbar} \int_0^T (\partial_1 \partial_t A^2 + \partial_2 \partial_t A^1) dt \\ &= \frac{q^{12}}{2\hbar} (\partial_1 (A^2(T) - A^2(0)) + \partial_2 (A^1(T) - A^1(0))) \\ &= -\frac{2\pi q^{12}}{eL_1 L_2}, \end{aligned} \tag{3.15}$$

which confirms our operator definition of the quadrupole moment.

3.3 Practical Evaluation

Having shown that $\hat{\mathcal{U}}_{ab}^Q$ can be used to determine dipole currents in the bulk, and that it has a physical interpretation via adiabatic evolution in an electric field gradient, we need to verify that it can have a well-defined, non-vanishing expectation value $z_{ab}^{(Q)}$ in periodic systems. The following argument from Ref. [47] shows that the expectation value $z_j^{(P)}$ for polarization is non-vanishing only when the expectation value of $\hat{\mathcal{U}}_j^P$ transforms in the trivial representation of the translation group. Assuming translation invariance, the ground state $|\Psi_0\rangle$ is an eigenstate of the translation operator carrying many-body momentum $\mathbf{K}$, i.e. $\hat{T}_{\mathbf{a}_c}^{-1} |\Psi_0\rangle = e^{-i\mathbf{K} \cdot \mathbf{a}_c} |\Psi_0\rangle$. We find

$$\begin{aligned} \langle \Psi_0 | \hat{\mathcal{U}}_a^P | \Psi_0 \rangle &= \langle \Psi_0 | \hat{T}_{\mathbf{a}_c}^{-1} \hat{T}_{\mathbf{a}_c} \hat{\mathcal{U}}_a^P \hat{T}_{\mathbf{a}_c}^{-1} \hat{T}_{\mathbf{a}_c} | \Psi_0 \rangle \\ &= e^{i\mathbf{K} \cdot \mathbf{a}_c} e^{-i\mathbf{K} \cdot \mathbf{a}_c} \exp(2\pi i \delta_{ac} n_{0a}) \langle \Psi_0 | \hat{\mathcal{U}}_a^P | \Psi_0 \rangle, \end{aligned} \tag{3.16}$$

²This last constraint, of being gapped to formations of dipoles, is strong, and unlikely to be strictly satisfied in real materials. It is analogous to requiring that a material be insulating to define the polarization, which is an important requirement.

where we translated all the electrons in the c -th primitive direction using $\hat{T}_{\mathbf{a}_c}$ to evaluate the transformed operator,

$$\hat{T}_{\mathbf{a}_c} \hat{\mathcal{U}}_a^P \hat{T}_{\mathbf{a}_c}^{-1} = \exp \left[2\pi i \frac{\sum_i (X^i + a_c^i) b_a^i}{N_a} \right] \quad (3.17)$$

$$= \exp \left[2\pi i \frac{\sum_i X^i b_a^i + N_e \delta_{ac}}{N_a} \right] \quad (3.18)$$

$$= \hat{\mathcal{U}}_a^P \exp [2\pi i \delta_{ac} n_{0a}], \quad (3.19)$$

and assumed $|\Psi_0\rangle$ is an eigenstate of the total number operator. We see that if $n_{0a} = N_e/N_a$ is an integer for N_a unit cells in the a -th direction, then the expectation value does not have to vanish. We can relax n_{0a} to be rational with a suitable modification of $\hat{\mathcal{U}}_a$ [47].

Repeating the same transformation for the quadrupole, we find

$$\begin{aligned} \langle \Psi_0 | \hat{\mathcal{U}}_{ab}^Q | \Psi_0 \rangle &= \exp [2\pi i n_{0ab} \delta_{ac} \delta_{bc}] \times \langle \Psi_0 | \hat{\mathcal{U}}_{ab}^Q \exp \left[2\pi i \hat{X}^j (b_a^j \delta_{bc} + \delta_{ac} b_b^j) / N_a N_b \right] | \Psi_0 \rangle \\ &= e^{i\Gamma} \langle \Psi_0 | \hat{\mathcal{U}}_{ab}^Q | \Psi_0 \rangle + O(1/L), \end{aligned} \quad (3.20)$$

where $n_{0ab} = N_e/N_a N_b$ is assumed to be an integer, and we have applied the results of Ref. [25] to evaluate $e^{2\pi i \hat{X}^j (b_a^j \delta_{bc} + \delta_{ac} b_b^j) / N_a N_b} |\Psi_0\rangle = e^{i\Gamma} |\Psi_0\rangle + O(1/L)$. Hence, in the thermodynamic limit we find the relation $\langle \Psi_0 | \hat{\mathcal{U}}_{ab}^Q | \Psi_0 \rangle = e^{i\Gamma} \langle \Psi_0 | \hat{\mathcal{U}}_{ab}^Q | \Psi_0 \rangle$, and $z_{ab}^{(Q)}$ must therefore vanish unless $\Gamma = 2\pi p$ for some integer p . The phase factor Γ is different than the phase generated by translations of the polarization operator $\hat{\mathcal{U}}_a^P$, which depends on the particle number n_{0a} . Instead, Γ depends on the polarization $\hat{X}^j$. Specifying that the planar filling factor n_{0ab} is an integer (or rational fraction for suitably generalized $\hat{\mathcal{U}}_{ab}^Q$ [47]) will not be enough in this case, and one must also specify that the polarization vanishes (up to a quantum).

However, even after satisfying this constraint, the expectation value of $\hat{\mathcal{U}}_{ab}^Q$ is not invariant under translations in general. Constraining the polarization to vanish (which should ensure $\Gamma = 2\pi p$) will not fix the issue in finite-sized systems due to the non-vanishing *fluctuations* of the dipole moment. The dipole fluctuations are a measure of the Wannier function localization length, or alternatively, the non-vanishing correlation length controlled by the insulating gap. When the dipole fluctuations are non-vanishing, the state $|\Psi_0\rangle$ is not an eigenstate of the $\hat{X}^j$ operator, so even with no average dipole, evaluating the expectation value of the quadrupole operator can be problematic. If we have single particle orbitals in a zero-correlation length limit where Wannier functions are point-localized, then the dipole fluctuations vanish and the many-body ground state will be an exact eigenstate of the $\hat{\mathcal{U}}_a^P$ operators. For this situation the expectation value of $\hat{\mathcal{U}}_{ab}^Q$ is well-defined even in finite-size as long as the polarization vanishes, and the particle filling is integer. Indeed,

for many-body ground states that are an exact eigenstate of $\hat{U}_a^P$ the magnitude of the expectation value of $\hat{U}_{ab}^Q$ will tend to unity in the thermodynamic limit. To evaluate the expectation value of the quadrupole operator more generally we must be more careful.³ The effect of dipole fluctuations on $\hat{U}_{ab}^Q$ is discussed further in Sec. 3.5.3.

In this work we replace the quadrupole operator by an approximation: we evaluate the expectation value of $\hat{U}_{ab}^Q$ on a finite supercell of size $L_x \times L_y$, then extend the truncated operator periodically outside the supercell. We then evaluate the expectation value of this periodic operator. The error introduced by this approximation depends on the size of the supercell we use, but improves as the ratio of the correlation/localization length and the characteristic length of the supercell goes to zero. From our analytic evaluation of the expectation value of $\hat{U}_{ab}^Q$ for Gaussian charge configurations in Sec. 3.5.3, we expect that the magnitude of the expectation value will approach a constant in a thermodynamic limit where both N_x and N_y are taken to infinity together. This constant will not generically be unity, and it will depend on the size of the dipole fluctuations in the x and y directions as well as the aspect ratio N_x/N_y . In this limit the phase factor that determines the quadrupole moment converges to the correct value as well. Discussions of finite size error and more details about the effects of dipole fluctuations are shown in Secs. 3.5.2 and 3.5.3. We will leave a systematic study of this approximation and alternatives to future work.

3.4 Example Calculations

3.4.1 Multipole Models with Mirror or C_4 Symmetry

Now we provide some examples in which we calculate the multipole moments using our operators in tightbinding models. Our first example will be focused on the insulating phases of a particular tightbinding model[26], [27] on a square lattice with four degrees of freedom per unit cell, which, for simplicity, we will treat as spinless orbitals (we also provide calculations for a related octupole model[26], and in the subsequent subsection we provide calculations for another tightbinding model with a non-vanishing quadrupole moment, but with C_4T symmetry[48]). The unit cell basis and tunneling terms are illustrated in Fig. 3.1, and the

³Similar problems would occur for the polarization operator if the ground state was not an eigenstate of the number operator, but instead had fluctuations in particle number.




Figure 3.1: Tightbinding model having four spinless orbitals per cell with inter-cell tunneling $\lambda_{x,y}$, intra-cell tunneling $\gamma_{x,y}$, and onsite potential δ that takes opposite signs on the filled and empty circles. There is a flux ϕ between each unit cell and within each unit cell in a gauge where the dotted lines have a relative phase factor of $e^{i\Phi}$.

Bloch Hamiltonian for a system with periodic boundary conditions is given by

$$H(\mathbf{k}, \Phi) = \begin{pmatrix} \delta & \Gamma_x(k_x) & \Gamma_y(k_y) & 0 \\ \Gamma_x(-k_x) & -\delta & 0 & \Gamma_y(k_y) \\ \Gamma_y(-k_y) & 0 & -\delta & e^{i\Phi}\Gamma_x(k_x) \\ 0 & \Gamma_y(-k_y) & e^{-i\Phi}\Gamma_x(-k_x) & \delta \end{pmatrix} \quad (3.21)$$

where $\Gamma_x(k_x) = \gamma_x + \lambda_x e^{ik_x}$, $\Gamma_y(k_y) = \gamma_y + \lambda_y e^{ik_y}$, $\delta, \gamma_{x/y}, \lambda_{x/y}$ are real parameters representing onsite energies, intra-cell tunneling, and inter-cell tunneling respectively, and Φ is the flux in each (intra- and inter-cell) plaquette in our chosen gauge.

This model can be tuned to a variety of insulator phases, and we will consider several test cases. For case (i) we explore this model in the quadrupolar phase protected by M_x and M_y mirror symmetries where $\Phi = \pi$, $\delta = 0$, and we fix $\lambda_x = \lambda_y \equiv \lambda$ to plot a phase diagram as a function of γ_x/λ and γ_y/λ in the interval $[-2, 2]$. We find that the calculation of the quadrupole moment using our operator (see Fig. 3.2a), matches the expected phase diagram from Ref. [26] where the system has a quantized quadrupole moment with magnitude

$Q^{xy} = e/2$ when $\gamma_x/\lambda, \gamma_y/\lambda$ are both in the interval $[-1, 1]$.

Let us make some technical remarks about this calculation. Although difficult to see by eye, there is some deviation from the expected phase diagram very close to the phase boundaries where the correlation length is increasing without bound, but this deviation is just a finite-size effect as it decreases rapidly with the size of the supercell. Additionally, while the phase of the operator $\hat{U}_{ab}^Q$ yields the quadrupole moment of the ground state in the thermodynamic limit, it can only be extracted if the magnitude of the expectation value is non-vanishing. According to Eq. 3.11, one would expect its magnitude to approach one in the thermodynamic limit; however, we find that, for our choice of boundary conditions and our choice for the approach to the thermodynamic limit, the magnitude does not go to unity in the presence of dipole fluctuations. In Sec. 3.5.2, we show the scaling of the magnitude of the operator as a function of system size for two points in the phase diagram shown in Fig. 3.2a. Furthermore, we discuss the complications associated with dipole fluctuations in more detail in Sec. 3.5.3.

For case (ii) we consider a variation of the quadrupole phase protected by C_4 symmetry where we tune $\gamma_x = \gamma_y \equiv \gamma$, $\lambda_x = \lambda_y \equiv \lambda$ and allow for $\Phi \in [0, \pi]$. At $\Phi = 0$ the model is gapless at half filling, but for any finite Φ in this interval we expect to find a quantized quadrupole of $Q^{xy} = e/2$ when $|\lambda| > |\gamma|$. We show this calculation in the extended planes along the diagonals of the phase diagram in Fig. 3.2a, and find that our operator correctly reproduces the phase diagram.

For case (iii) we consider the adiabatic dipole pumping process introduced in Refs. [26], [27] where the quadrupole moment is continuously tuned from the topological phase where $Q^{xy} = e/2$ to the trivial phase where $Q^{xy} = 0$. We can parameterize this pumping process by $\lambda_x = \lambda_y = 1$, $\delta \rightarrow -\sin(\theta(t))$, and $\gamma = 1/2(1 - \cos(\theta(t)))$. We show the result in Fig. 3.2b, in which we plot the result of our operator compared to calculations of the corner charge (for an open, square geometry) and the edge polarization on a cylinder, and we find they all match exactly as expected for a continuously varying bulk quadrupole moment Q^{xy} .

For case (iv) we tune the system to $\Phi = \delta = 0$. This model has M_x and M_y mirror symmetries so we expect the quadrupole moment to be quantized, though that does not mean it takes the non-trivial topological value. Indeed this model is not expected to have a bulk quadrupole moment despite having insulating phases with fractional $e/2$ corner charge, since this charge is associated with edge polarization rather than a bulk quadrupole moment. Here we show a phase diagram for fixed $\lambda_y = 1, \gamma_x = 0$, where we vary λ_x and γ_y . We show that our formula, while yielding a quantized value for Q^{xy} , does not distinguish between the insulating phases with and without fractional corner charge. This is a clear and important indication that our invariant is only sensitive to the bulk quadrupole moment, and is not just testing for the possible existence of corner charge alone.

Finally, in Fig. 3.2d we show a calculation for the octupole moment in the phase diagram of the mirror-

symmetric octupole model of Ref. [26] where we have fixed the inter-cell hopping λ and varied the intra-cell hoppings $\gamma_{x,y,z}$ in the interval $[0, 2]$. We see that the phase diagram matches the expected result with a topological octupole moment $O^{xyz} = e/2$ when $|\gamma_i/\lambda| < 1$ for all i .

3.4.2 Quadrupole Model with C_4T Symmetry

In this subsection we perform a calculation of the quadrupole moment in a 2D version of the chiral hinge-insulator model presented in Ref. [48]. This model has a quantized quadrupole moment protected by the product of rotation and time-reversal symmetry C_4T . The Bloch Hamiltonian we consider is

$$\begin{aligned} H(\mathbf{k}) = & \sin k_x \Gamma^1 + \sin k_y \Gamma^2 + (2 - m - \cos k_x - \cos k_y) \Gamma^0 \\ & + \Delta (\cos k_x - \cos k_y) \Gamma^3 \end{aligned} \quad (3.22)$$

where $\Gamma^0 = \tau^z \otimes \mathbb{I}$, $\Gamma^1 = \tau^x \otimes \sigma^x$, $\Gamma^2 = \tau^x \otimes \sigma^y$, $\Gamma^3 = \tau^x \otimes \sigma^z$. When $\Delta = 0$ this model has both C_4 and time-reversal symmetry T , and is a model representing a 2D time-reversal invariant topological insulator with helical edge states when, e.g., $0 < m < 2$. When $\Delta \neq 0$ the system breaks both C_4 and T but preserves the product C_4T . The non-vanishing Δ acts to gap the helical edge states and generate corner modes for an open square geometry that respects global C_4 symmetry. Thus we expect to find an ill-defined quadrupole moment when $\Delta = 0$, and a non-vanishing topological quadrupole moment $Q^{xy} = e/2$ when $0 < m < 2, \Delta \neq 0$. When $m < 0, \Delta \neq 0$ we expect to find a well-defined, but vanishing quadrupole moment. We show the calculated phase diagram as a function of m and Δ in Fig. 3.3. We note that we look at a range of Δ that is small enough not to destroy the bulk topology and we see that our operator correctly reproduces the phase diagram.

3.5 Analysis of constraints and finite size errors

This section works out detailed calculations of constraints on the evolution of the quadrupole moment and finite size errors relating to dipole fluctuations.

3.5.1 Constraints on Charge Configurations for the Quadrupole and Octupole Operators

Here we will evaluate the conditions under which one will generically find a non-vanishing expectation value of multipole operators if the electron orbitals are point-charge localized. For example, in a limit where all dipole fluctuations vanish, we expect that the expectation values of the quadrupole operators are well-behaved in extended systems. More precisely, to evaluate the quadrupole moment we want to consider a many-body state that is an exact eigenstate of the $\mathcal{U}_x$ and $\mathcal{U}_y$ operators. This is not necessarily an unphysical situation,




Figure 3.2: (a) Phase diagram for model in Eq. 3.21 with $L_x = L_y = 40$ as a function of $(\gamma_x/\lambda, \gamma_y/\lambda, \Phi)$. The green and purple regions have a values of Q^{xy} that differ by $e/2$, as calculated using Eq. 3.2. Dependence on Φ is only shown along the diagonal lines that maintain C_4 symmetry. (b) The evolution of the quadrupole moment (Eq. 3.2) during a pumping process as compared with the corner charge and edge polarization. All three match exactly, as expected. (c) Phase diagram for Eq. 3.21 with $\Phi = 0$ as a function of $(\gamma_y/\lambda_y, \lambda_x/\lambda_y)$. Eq. 3.2 does not distinguish two different phases that have corner charge and a trivial insulator with no corner charge. This is as expected as in this regime the model has a vanishing quadrupole moment even in the phases with corner charge/modes.[27] (d) Phase diagram for an octupole model[26] with $L_x = L_y = L_z = 10$ as a function of $(\gamma_x/\lambda, \gamma_y/\lambda, \gamma_z/\lambda)$. The green and purple regions have values of O^{xyz} that differ by $e/2$ as determined by Eq. 3.3.




Figure 3.3: Phase diagram for model in Eq. 3.22 with $L_x = L_y = 40$ as a function of (m, Δ) . The green and purple regions have values of Q^{xy} that differ by $e/2$ according to Eq. 3.2. The data along the line $\Delta = 0$ is not shown because the values of Q^{xy} are highly fluctuating, and the quadrupole moment is not well-defined since the edges are gapless.

as we know it will be approximately true as we approach the thermodynamic limit. That is, if we consider our insulating system to be large compared to the characteristic size of individual Wannier functions of the occupied bands, then we can treat the charge distribution as a collection of point charges located at the Wannier centers that are periodically repeated throughout the crystal lattice[24]. For a many-body state represented by a product of single-particle orbitals, vanishing dipole fluctuations imply that the orbitals are δ -function localized, so we will consider the expectation values of the multipole operators in such a state. Even in this situation, there are constraints under which the expectation values of the quadrupole or octupole operators are well-defined in the thermodynamic limit, i.e, when they effectively transform in the trivial representation of the translation group, for such point-charge product states. We will determine those conditions now.

For simplicity, let us consider our electron configuration to be a rectangular lattice of point charges and focus on the quadrupole moment Q^{xy} associated to the operator

$$\hat{U}_{xy}^Q = \exp\left(\frac{2\pi i \sum_j x_j y_j}{L_x L_y}\right). \quad (3.23)$$

Now let us transform this operator by a translation $\mathbf{R} = h_1 a_x \hat{x} + h_2 a_y \hat{y}$ where h_1, h_2 are integers. We find that

$$\begin{aligned} T_{\mathbf{R}} \hat{U}_{xy}^Q T_{\mathbf{R}}^{-1} &= \hat{U}_{xy}^Q \exp\left[2\pi i \sum_j h_1 h_2 / N_x N_y\right] \exp\left[2\pi i h_1 \sum_j y_j / N_x L_y\right] \exp\left[2\pi i h_2 \sum_j x_j / L_x N_y\right] \\ &= \hat{U}_{xy}^Q \exp\left[2\pi i \sum_{\mathbf{R}} \sum_{\alpha=1}^{\nu} h_1 h_2 / N_x N_y\right] \exp\left[2\pi i h_1 \sum_{\mathbf{R}} \sum_{\alpha=1}^{\nu} y_{\mathbf{R},\alpha} / N_x L_y\right] \exp\left[2\pi i h_2 \sum_{\mathbf{R}} \sum_{\alpha=1}^{\nu} x_{\mathbf{R},\alpha} / L_x N_y\right] \\ &= \hat{U}_{xy}^Q \exp[2\pi i \nu h_1 h_2] \exp\left[2\pi i h_1 \sum_{n_2} \sum_{\alpha=1}^{\nu} (\bar{y}_{\alpha} + n_2 a_y) / L_y\right] \exp\left[2\pi i h_2 \sum_{n_1} \sum_{\alpha=1}^{\nu} (\bar{x}_{\alpha} + n_1 a_x) / L_x\right] \\ &= \hat{U}_{xy}^Q \exp[2\pi i \nu h_1 h_2] \exp\left[\frac{2\pi i h_1}{a_y} \sum_{\alpha=1}^{\nu} \bar{y}_{\alpha}\right] \exp\left[\frac{2\pi i h_2}{a_x} \sum_{\alpha=1}^{\nu} \bar{x}_{\alpha}\right] \exp[\pi i h_1 \nu (N_y + 1)] \exp[\pi i h_2 \nu (N_x + 1)] \end{aligned} \quad (3.24)$$

where the sum over α runs over all the electrons in a single unit cell and ν is the electron filling.

$$\exp[2\pi i \nu] \exp\left[\frac{2\pi i}{a_y} \sum_{\alpha=1}^{\nu} \bar{y}_{\alpha}\right] \exp\left[\frac{2\pi i}{a_x} \sum_{\alpha=1}^{\nu} \bar{x}_{\alpha}\right] \exp[\pi i \nu (N_y + N_x + 2)] = 1. \quad (3.25)$$

To satisfy this constraint in a manner that is independent of the evenness/oddness of N_x, N_y it is sufficient to choose $\nu = 2\mathbb{Z}$ and the x and y components of the dipole moment per unit cell must be integer multiples of a_x, a_y respectively. We can relax the constraint of ν being integer valued if we modify our starting operator Eq. 3.23 according to a prescription analogous to that presented in Ref. [47], but in that case we will still

require $\nu = 2\mathbb{Q}$. While we have derived these constraints assuming the charge distribution is point-like we expect this result to apply to more complicated distributions in the thermodynamic limit where the system size is much larger than the spatial extent of the Wannier functions that compose our insulating ground state. The coordinates $\bar{x}_\alpha, \bar{y}_\alpha$ will then be associated to the centers of the Wannier functions in each unit cell.

We can repeat this calculation for the octupole moment, e.g, O^{xyz} , and translating by a vector $\mathbf{R} = a_x\hat{x} + a_y\hat{y} + a_z\hat{z}$. We find the constraint that the dipole and quadrupole moments per unit cell must be integer valued as well as the constraint

$$\exp \left[\frac{2\pi i \nu (N_x + 1)(N_y + 1)(N_z + 1)}{4} \right] = 1. \quad (3.26)$$

In order for this constraint to be satisfied independently of the system size we thus need $\nu = 4\mathbb{Z}$ for the octupole moment.

3.5.2 Finite Size Effects of the Quadrupole Operator

The evaluation of the expectation value of the quadrupole operator $\hat{U}_{xy}^Q$ (Eq. 3.23) exhibits some challenges. We show how this manifests in our calculation as a function of system size. Fig. 3.4 shows the γ_x - γ_y phase diagram of the quantized, mirror-symmetric quadrupole model Eq. 3.21 from Ref. [27] for $L_x = L_y = 10, 20, 40$. There is an error in a thin region near the phase boundary, where only one of γ_x and γ_y is close to (but less than) one. The expectation value predicts the wrong topological phase in this region; however, this region decreases rapidly with system size. In the thermodynamic limit, the error disappears. This type of error is due to the large correlation length that exceeds the system size for parameters near the phase boundaries.

The expectation value of the operator $\hat{U}_{xy}^Q$ has a phase and a magnitude. While the phase gives the quadrupole moment, the magnitude reflects its fluctuations. Fig. 3.5 shows how the magnitude depends on our finite simulation size for selected points in the γ_x - γ_y phase diagram shown in Fig. 3.4. The magnitudes are close to zero and decrease as the system gets larger. This does not mean that the physical quantity is not meaningful; in comparison, for a 3D system, Resta's polarization operator also decreases to zero magnitude if $N_x = N_y = N_z \rightarrow \infty$. Exploring efficient ways to evaluate the expectation value for large systems is still an open challenge. In Sec. 3.5.3 we will show how this magnitude decrease can be traced back to fluctuations in the dipole moments.

3.5.3 Scaling of the Magnitude of the Quadrupole Expectation Value

Finally, we discuss the scaling properties of the expectation values of the dipole and quadrupole operators using a test configuration of continuum charge distributions arranged on a crystalline lattice. We will




Figure 3.4: Phase diagrams in γ_x - γ_y of $\text{Im} \ln \langle \Phi_0 | \hat{\mathcal{U}}_{xy}^Q | \Phi_0 \rangle$ for $N_x = N_y = 10$ (a), 20 (b), 40 (c). Teal lines at $|\gamma_i| = 1$ indicate the phase boundary of the model. Our operator gives a quadrupole moment of zero in the green regions and 0.5 in the purple regions. In larger systems, our operator's quadrupole phase more closely approaches the theoretical phase boundary.




Figure 3.5: Log of the magnitude $|\langle \Phi_0 | \hat{\mathcal{U}}_{xy}^Q | \Phi_0 \rangle|$ at different points in the γ_x - γ_y phase diagram. The magnitude goes to zero with increasing system size at all points in the phase diagram.

consider the cases of product states of δ -function point charges as well as Gaussian charges, both product and determinant states. In these systems, the dipole fluctuations cause the magnitude of the quadrupole operator to go to a limiting value which is not equal to one.

Polarization

For the calculation of the polarization let us consider a 1D system with a filling of one electron per unit cell at a position x_0 with respect to the origin of each unit cell. The electrons in our system are thus located at the positions $x_0 + n$ in units of the lattice constant a which we set to $a = 1$, and where $n \in \mathbb{Z}$. Now we can calculate the polarization via the expectation value of the operator $\hat{\mathcal{U}}_x = \exp \left[\frac{2\pi i X}{L_x} \right]$ in the many-body ground state consisting of a tensor product of localized δ -function charges. The polarization is given by

$$P_x = \frac{eL_x}{2\pi\mathcal{V}} \text{Im} \ln \left[\langle \hat{\mathcal{U}}_x \rangle \right] \quad (3.27)$$

where $\mathcal{V}$ is the volume of the system, i.e., $L_x = N_x$ in 1D. After integrating over the real-space coordinates of each electron to generate the expectation value, we find the result

$$P_x = ex_0 + e \frac{N_x + 1}{2}, \quad (3.28)$$

which is only well-defined mod e , so the second term either contributes 0 or $1/2$ depending on the parity of N_x .

Now let us take localized charges obeying a Gaussian distribution

$$\psi_{x_0}(x) = \frac{1}{\sqrt{\pi^{1/2}\sigma_x}} \exp \left[-\frac{(x - x_0)^2}{2\sigma_x^2} \right]. \quad (3.29)$$

We find that, for a single charge at x_0

$$\langle \hat{\mathcal{U}}_x \rangle = \exp \left[\frac{2\pi i x_0}{N_x} \right] \exp \left[-\frac{\pi^2 \sigma_x^2}{N_x^2} \right]. \quad (3.30)$$

We see that this result includes a phase factor encoding the position x_0 as well as an exponential factor with magnitude ≤ 1 . The phase factor is identical to what the δ -function charge distribution produces, and when all electrons across the lattice are considered it will exactly reproduce Eq. 3.28 for the value of the polarization. On the other hand, if we have a number of electrons $N_e = N_x$ then the *magnitude* of the expectation value will be $e^{-\pi^2 \sigma_x^2 / N_x}$ which will tend to unity as $N_x \rightarrow \infty$ as expected. Thus, this operator recovers unitarity in the thermodynamic limit.

Let us now illustrate an important subtlety by reconsidering the case of Gaussian charges in a 2D system. Suppose that we have one electron per unit cell on a square lattice, and that the electrons are centered at $(x_0, 0) + \mathbf{R}$ where $\mathbf{R}$ is the set of 2D lattice vectors. For a single-electron, which is Gaussian distributed in x and y we have

$$\langle \hat{\mathcal{U}}_x \rangle = \exp \left[\frac{2\pi i x_0}{N_x} \right] \exp \left[-\frac{\pi^2 \sigma_x^2}{N_x^2} \right] \quad (3.31)$$

which is the same as the 1D case above. The distinction between dimensions appears when we use the fact that $N_e = N_x N_y$. After including all of the electrons we have

$$|\langle \hat{\mathcal{U}}_x \rangle| = \exp \left[-\frac{\pi^2 \sigma_x^2}{N_x^2} \right]^{N_x N_y} = \exp \left[-\frac{\pi^2 \sigma_x^2 N_y}{N_x} \right]. \quad (3.32)$$

The important feature is that the magnitude of this operator approaches unity in a specific thermodynamic limit where we take $N_x \rightarrow \infty$ before N_y , but it vanishes if we use the opposite order, and approaches a constant < 1 if we take the limit while keeping a fixed aspect ratio N_x/N_y . Under the same conditions, but in 3D, we would find a factor $e^{-\pi^2 \sigma_x^2 N_y N_z / N_x}$. Thus, generically to keep the operator expectation value non-vanishing we must always take the thermodynamic limit in the direction of the polarization component of interest before letting the transverse directions approach infinity.

Quadrupole Moment

Now let us consider the quadrupole moment for a 2D system on a square lattice. We can begin with the periodic arrangement of δ -function point charges and calculate the expectation value of $\hat{\mathcal{U}}_{xy} = \exp \left[\frac{2\pi i \hat{Q}_{xy}}{L_x L_y} \right]$. We must be careful in this case to specify that the total dipole moment vanishes (up to an integer). We will satisfy this constraint by taking the system with two electrons per cell with coordinates (x_0, y_0) and $(-x_0, -y_0)$ with respect to the origin of the cell, and we repeat these positions across the entire lattice. For this configuration, and for a lattice indexed from $1 \dots N_x, 1 \dots N_y$, we find

$$Q_{xy} = 2ex_0y_0 + \frac{e(N_x + 1)(N_y + 1)}{2}, \quad (3.33)$$

which depends on the parity of N_x and N_y similarly to the polarization.

Now let us move on to charges with a Gaussian distribution described by

$$\psi_{(x_0, y_0)}(x) = \frac{1}{\sqrt{\pi\sigma_x\sigma_y}} \exp \left[-\frac{(x - x_0)^2}{2\sigma_x^2} - \frac{(y - y_0)^2}{2\sigma_y^2} \right]. \quad (3.34)$$

In this case, the expectation value generated by a single electron with a Gaussian wave function centered at

(x_0, y_0) is

$$\begin{aligned}
\langle \hat{U}_{xy} \rangle &= \frac{2}{\sqrt{4 + \alpha^2 \sigma_x^2 \sigma_y^2}} \\
&\times \exp \left[-\frac{\alpha^2 \sigma_y^2 x_0^2 + \alpha^2 \sigma_x^2 y_0^2 - 4i\alpha x_0 y_0}{4 + \alpha^2 \sigma_x^2 \sigma_y^2} \right] \\
&\equiv \Lambda \exp \left[\frac{i\alpha x_0 y_0}{1 + \alpha^2 \sigma_x^2 \sigma_y^2 / 4} \right]
\end{aligned} \tag{3.35}$$

where $\alpha = 2\pi/L_x L_y$. We can see from this result that both the magnitude *and* the phase of this expectation value depend on the system size when σ_x and/or σ_y are non-vanishing. This is not the case for the polarization, where the imaginary phase is independent of the system size for a Gaussian distribution. Here the difference arises because we allow for fluctuations of the dipole moment, which is equivalent to having non-vanishing σ_i . The dipole moment is vanishing on average, but has non-zero fluctuations that lead to this ambiguity in the phase factor. The same would likely occur for the polarization if we let the charge in each unit cell fluctuate, but we have implicitly assumed that it does not fluctuate. We can calculate the quadrupole moment density from the imaginary part of the log of this expression for the expectation value to find

$$x_0 y_0 \left(1 - \frac{\pi^2 \sigma_x^2 \sigma_y^2}{L_x^2 L_y^2} + \dots \right) \tag{3.36}$$

for a single electron. We find that this value converges to the correct result as $L_x, L_y \rightarrow \infty$, i.e., when the system size is much larger than the dipole fluctuation lengths. When all of the electrons in the lattice are taken into account this will reproduce the value calculated for the δ -function distribution.

Finally, let us consider the magnitude Λ of this expectation value. We find a magnitude for a single electron to be

$$\Lambda = \frac{2}{\sqrt{4 + \alpha^2 \sigma_x^2 \sigma_y^2}} \exp \left[-\frac{\alpha^2 \sigma_y^2 x_0^2 + \alpha^2 \sigma_x^2 y_0^2}{4 + \alpha^2 \sigma_x^2 \sigma_y^2} \right]. \tag{3.37}$$

We want to determine the value of this magnitude for a collection of electrons on the lattice. The magnitude for the full ground state is harder to calculate since it depends on the positions of the electrons, but for a product state including two electrons per unit cell, i.e., $N_e = 2N_x N_y$, and ignoring contributions of the charge

density that spreads from the last cells back into the first cells due to periodic boundary conditions, we find

$$\begin{aligned}\Lambda_{N_e} &\approx \left(\frac{2}{\sqrt{4 + \alpha^2 \sigma_x^2 \sigma_y^2}} \right)^{2N_x N_y} \exp \left[-\frac{2\pi^2(\sigma_y^2(N_x^2/3 + x_0^2) + \sigma_x^2((N_y^2/3 + y_0^2)))}{N_x N_y + \pi^2 \sigma_x^2 \sigma_y^2 / (N_x N_y)} \right] \\ &\approx \left(\frac{1}{\sqrt{1 + \pi^2 \sigma_x^2 \sigma_y^2 / N_x^2 N_y^2}} \right)^{2N_x N_y} \exp \left[-\frac{(2/3)\pi^2(\sigma_y^2 N_x / N_y + \sigma_x^2 N_y / N_x)}{1 + \pi^2 \sigma_x^2 \sigma_y^2 / (N_x^2 N_y^2)} \right]\end{aligned}\quad (3.38)$$

where in the first approximation we have dropped the terms that shift $N_{x,y}$ by 1 coming from a discrete sum over unit cells, and in the second approximation we have dropped the terms proportional to x_0, y_0 since they are coordinates within a unit cell and are less than 1 in units of the lattice constant. We find that as $N_x, N_y \rightarrow \infty$ that the first factor tends to unity:

$$\left(\frac{1}{\sqrt{1 + \pi^2 \sigma_x^2 \sigma_y^2 / N_x^2 N_y^2}} \right)^{2N_x N_y} \approx \left(1 - \frac{\pi^2 \sigma_x^2 \sigma_y^2}{2N_x^2 N_y^2} \right)^{2N_x N_y} \approx \exp(-\pi^2 \sigma_x^2 \sigma_y^2 / N_x N_y) \rightarrow 1. \quad (3.39)$$

If we make an analogy with the polarization, this factor has a dependence on the quadrupole fluctuations $\sigma_x^2 \sigma_y^2 / 4$ and tends to unity as long as those fluctuations are finite.

Now let us consider the second factor

$$\exp \left[-\frac{(2/3)\pi^2(\sigma_y^2 N_x / N_y + \sigma_x^2 N_y / N_x)}{1 + \pi^2 \sigma_x^2 \sigma_y^2 / (N_x^2 N_y^2)} \right]. \quad (3.40)$$

This factor has interesting features: (i) the numerator has a dependence on the dipole fluctuations in the x and y directions, (ii) the denominator has a dependence on the quadrupole fluctuations. If dipole fluctuations in either direction vanish, then there is a consistent way to take the thermodynamic limit such that this term tends to unity. If the dipole fluctuations are non-vanishing in both directions then the best that one can do is to take $N_x = N_y \rightarrow \infty$. In this case the factor tends to the finite value $e^{-2\pi^2(\sigma_y^2 + \sigma_x^2)/3}$, which is non-vanishing, but not unity. We also note that if one takes N_x or N_y to infinity first then this factor generically tends to zero. This result seems to imply that even though $\hat{\mathcal{U}}_{xy}$ may not be strictly a well-defined operator in periodic systems with dipole fluctuations, its expectation value is still meaningful and can recover the correct results for the quadrupole moment.

The above discussion considers a product state of Gaussian orbitals; however, the tightbinding calculations presented in the paper use a Slater determinant, so we also show a comparison with a determinant of the Gaussian orbitals. An analytic calculation of the determinant state is impractical, so we compared the numerical results of the determinant and product states.

We confirm that the determinant state yields similar results to the product state except when the Gaussian spreads σ_x, σ_y grow so large that determinant state exhibits a transition to a conducting state (i.e., the magnitude of the polarization operator goes to zero in the infinite limit), while the product state remains insulating (Fig. 3.6). For smaller values of σ_x, σ_y , the phases and magnitudes of the product state agree with the determinant state. We confirm that the phase of the polarization operator expectation value in our states is always zero, so the quadrupole is well-defined in principle. We also confirm that the magnitude of the quadrupole operator remains fixed when we vary $\sigma = \sigma_x = \sigma_y$ and even values of $N_x = N_y$, and is the same for the determinant and product states. When σ increases, even when the magnitude of the polarization goes to one, the limit of the quadrupole magnitude decreases (it does not always approach one, even when finite). Beyond a certain value of σ , the quadrupole magnitude goes to zero. Based on our numerics, the insulating-conducting transition of the determinant state seems to coincide with the transition where the quadrupole magnitude goes to zero.

Our simple model of Gaussian orbitals can only go so far in explaining the results from our tightbinding calculations; however, we have shown that even with very simple wave functions, the quadrupole magnitude is suppressed by dipole fluctuations, and does not generally tend to one in the thermodynamic limit. We expect that our results would closely match the tightbinding results if we accounted for dipole fluctuations across the periodic boundary conditions.

3.6 Connection to Recent Work

Recent papers[49], [50] have also investigated the evaluation of the quadrupole and octupole operators we presented in Eqs. 2-5. Kang et. al.[49] simultaneously proposed the same definitions for the quadrupole and octupole operators as we do in Eqs. 2-5. Similar to our argument in Eq. 3.20, they investigate the invariance of the operator $\hat{\mathcal{U}}_{xy}$ under translations, and conclude that the polarization must be zero for the operator to be well-defined. We have both pointed out that the ground state is not an eigenstate of polarization $\hat{\mathcal{U}}_j$, which is what complicates the evaluation of $\hat{\mathcal{U}}_{xy}$. They point out a violation of translation invariance of $\mathcal{O}(1/E_{\text{gap}})$, inversely proportional to the excitation gap:

$$\langle \hat{\mathcal{U}}'_{xy} \rangle = \langle \hat{\mathcal{U}}_{xy} \rangle \langle \hat{\mathcal{U}}_y \rangle + \mathcal{O}\left(\frac{1}{E_{\text{gap}}}\right),$$

where $\hat{\mathcal{U}}'_{xy}$ is the operator $\hat{\mathcal{U}}_{xy}$ translated in the x direction by $x \rightarrow x + L_x$. We describe this same error as being a result of dipole fluctuations, which, for non-interacting systems, measure the Wannier localization length. For single-particle orbitals, the error disappears in the limit that the Wannier functions are δ -function




Figure 3.6: Comparison of determinant (solid lines) and product states (dashed lines) of Gaussian orbitals as a function of system size, keeping $N_x = N_y$ (always even) and $\sigma_x = \sigma_y$. For all these calculations, we fixed $x_0 = y_0 = 0.2$. Top left: magnitude of the polarization operator goes to zero for the determinant with larger values of σ_x , but goes to one for the product state, regardless of σ_x . Top right: phase of the polarization operator is always zero. Bottom left: Magnitude of the quadrupole operator plateaus in the infinite limit to a value depending on σ_x for both determinant and product states. For large enough σ , the magnitude from the determinant state tends toward zero in the infinite limit. Bottom right: The phase of the quadrupole operator is consistent for all calculations. For odd values of $N_x = N_y$ (not shown), the quadrupole phase is shifted by 0.5.

localized, equivalent to their limit of $E_{\text{gap}} \rightarrow \infty$. We drew a connection between dipole fluctuations and decreasing magnitude of $\langle \hat{\mathcal{U}}_{xy} \rangle$, which is a signature of the problem with translation invariance. In contrast to their discussion, we point out that zero polarization is not sufficient to have a well-defined quadrupole in finite-sized periodic systems. As such, for systems with dipole fluctuations, evaluating this operator in PBCs necessarily uses an approximate form; we described one approximation that we used in our calculation, which is similar to that used in Ref. [49].

To confirm the definitions of the quadrupole operator $\hat{\mathcal{U}}_{xy}$ as a reliable estimator of the response to electric field gradients, they also give an argument similar to our adiabatic evolution discussion, but in a more field-theoretical approach.

In response to these proposals, Ono et. al. [50] comment on some problems that arise when applying the quadrupole operator $\hat{\mathcal{U}}_{xy}$ to certain tight-binding models, that were not considered in our work [50]. They find that their evaluation of the quadrupole operator has more serious issues than those discussed here, namely that the result seems to depend non-trivially on the correlation length in an insulator phase when the quadrupole moment should nominally be quantized. They also comment that their results can depend on the parity of the number of lattice sites, but as we show in Secs. 3.5.1 and 3.5.3 this latter issue has a simple origin, and similar effects can appear even for the polarization, and even for classical point-charge configurations.

In the models in which we evaluated the quadrupole moment we did not find any dependence on the correlation length besides typical finite-size effects, so it might be useful to comment on the distinctions between the models we considered and one of their models, though fully resolving the discrepancy will require further research. One of their models is a C_4 invariant system with four occupied bands that form a Wannier representable obstructed atomic limit. The electrons in the model sit at maximal Wyckoff positions $1a, 1b, 2c$, i.e., one at the center, one at the corner, and two in the middle of the edges of the unit cell. One possibly important distinction is that this model does not have gapped Wannier bands. While the need for gapped Wannier bands was emphasized in Ref. [27], it has not been proven that this is a generic requirement to define the quadrupole moment (though a gapped, neutral edge *is* required). A second distinction is that in the models we consider, one can always relate the configurations with a non-vanishing quadrupole to a configuration representing a trivial onsite limit through a continuous deformation that preserves a symmetry under which the *dipole moment does not change*. Thus, in our models one can clearly calculate a *change* in quadrupole moment when going from the trivial atomic limit to a non-vanishing quadrupole configuration. In their model, such an interpolation does not seem to exist, so it is not clear how to compare the quadrupole moment in their obstructed atomic limit to the trivial atomic limit. While again, this may not be a requirement to define a quadrupole model, it does seem like a natural consideration.

Ref. [50] also makes a note regarding second-order contributions to the perturbative expansion of the magnitude $|\langle \Psi_0 | \hat{\mathcal{U}}_x | \Psi_0 \rangle|$, and comments that Resta's formula for polarization is not guaranteed to work, i.e., the magnitude may go to zero in higher dimensions. This is a well-known fact for polarization. Indeed to obtain a nonzero magnitude of the expectation value of $\hat{\mathcal{U}}_x$ in higher dimensions one must take the correct thermodynamic limit, namely, where the direction parallel to the polarization component of interest goes to infinity first. We also find that, for periodic boundary conditions, the quadrupole magnitude $|\langle \hat{\mathcal{U}}_{xy} \rangle|$ goes to zero in the thermodynamic limit if dipole fluctuations are finite, seemingly no matter what thermodynamic limit one uses. Despite this we find that this issue does not affect the consistency of the phase in the tight-binding calculations we performed in finite size systems. Based on the calculations in Sec. 3.5.3, we expect that if we truncate the system, as indicated by our approximation scheme, then the magnitude $|\langle \hat{\mathcal{U}}_{xy} \rangle|$ can approach a finite (though likely not unity) value if we take the limit $L_x = L_y \rightarrow \infty$, even in the presence of dipole fluctuations.

In summary, Ono et. al. raise some interesting questions about the evaluation of the many-body quadrupole operator and about what constraints are necessary to enforce in order to have a well-defined quadrupole moment in a crystalline system. The issues they have pointed out will be valuable in specifying careful definitions and treating boundary effects correctly for detecting higher order multipole moments in bulk systems.

3.7 Conclusion

In conclusion, we have proposed definitions for many-body operators whose expectation values determine the quadrupole and octupole moments of insulators. We showed that the change in quadrupole moment corresponds to dipole currents in the material. We proposed a method to evaluate the quadrupole; however, the method is not entirely satisfactory because it suffers from significant finite size errors related to the fluctuations of the dipole moment, even when the total dipole moment is vanishing. However, we showed that on several nontrivial tightbinding models the operators do indeed capture the bulk properties, and can be used to evaluate topological indices. The operator can be evaluated using many-body wave functions as well as using tightbinding wave functions, so it can be used to generate new many-body topological indices for systems with interactions and/or disorder.

CHAPTER 4

PYQMC CODE

This chapter is based on my published work “PyQMC: an all-Python real-space quantum Monte Carlo module in PySCF.”[51] As discussed in previous chapters, quantum Monte Carlo (QMC) is a highly accurate – but computationally expensive – method for computing properties of correlated electron systems. The capabilities of QMC are constantly being extended. This includes improving the accuracy and efficiency of the algorithms, as well as extending the range of what can be computed; for example, the ability to compute electronic excitations in addition to the ground state. Most QMC codes are written in high performance, compiled languages such as C++ and Fortran, which are computationally efficient, but the lengthy and complex code limits the rate at which new algorithms and capabilities can be implemented.

In collaboration with our research group, I developed **PyQMC**, a real-space QMC package written entirely in Python, a language known for its readability and simplicity. Our aim is to accelerate the process of developing new QMC methods by greatly simplifying implementation and testing. **PyQMC** has enabled other projects in our research group on gradient regularization[52], excited states[53], and wave function characterization by density matrix entropy [54], as well as my own projects covered in Chapters 5 and 6. It has already been used by an external group for developing neural network wave functions[55] and will continue to be useful to the QMC community.

One area we have made advances in using **PyQMC** is improving trial wave functions for periodic systems. Periodic QMC calculations have been limited by the assumption that the low energy wave functions are close to single determinants. Long-range correlation between electrons on different atoms cannot be described by a Jastrow factor and requires a superposition of determinants; however, a systematic way of selecting determinants and their weights for use in periodic QMC calculations has been impractical within the structure of previously available QMC packages. (This has been done for molecules, e.g. in Ref. [56].) Long-range correlation occurs for example in partially filled degenerate orbitals, such as an electron equally occupying p_x and p_y orbitals in a system with C_4 symmetry. It also occurs in systems where the eigenstates are superpositions of different spin orderings, such as a spin singlet state between two sites. In the doped cuprates,

it has been shown that multiple spin orderings are close in energy[57], indicating that a multi-determinant wave function may be necessary to describe the ground state. In the PyQMC code presented in this chapter, it is simple and efficient to generate trial wave functions for QMC from determinant expansions optimized using systematic quantum chemistry methods. The seamless software connection from quantum chemistry methods to QMC enabled my study on the cuprates in Chapter 5: to determine whether spin fluctuations (multiple determinants) are an essential ingredient to accurately describing the low-energy physics.

4.1 *All-Python Code for Human Efficiency*

Calculations have an inherent trade-off between accuracy and computational cost: more accurate methods scale more steeply with number of electrons, and exact calculations scale exponentially with system size. Quantum Monte Carlo (QMC) offers a good balance between accuracy and scalability, capable of treating systems with thousands of electrons.[11], [58]–[60] The past few years have seen several advances in QMC methods: new wave functions using machine learning techniques,[55], [61]–[65] new algorithms for optimizing excited states,[53], [66]–[71] complex observables such as energy density[72], [73] and density matrices,[74] a new method to derive effective Hamiltonians from *ab initio* QMC,[75]–[77] and new time-stepping algorithms to reduce timestep error.[78], [79]

Developing new tools for quantum Monte Carlo (QMC) comes with challenges. Computational efficiency can be sensitive to subtle details of implementation, and adding new methods can require significant changes to algorithms. A bottleneck in this development process is the testing and implementation of new ideas in code. Several high-performance real-space QMC packages are written in low-level compiled languages such as C++ or Fortran to achieve performance suitable for large-scale calculations, but come with the challenge that they are cumbersome to modify, add to, or customize.

PyQMC is an all-Python, flexible implementation of real-space QMC for molecules and materials designed to streamline development and teaching of new ideas in QMC. Although it does not currently achieve the performance of other compiled QMC softwares, it remains competitive and significantly reduces human time in code development and workflow design. PyQMC is part of the PySCF ecosystem, a collection of libraries for quantum chemistry methods that achieve performance close to that of compiled languages while implemented in the much more flexible Python language. In this chapter, I will describe the implementation of PyQMC and note some of its advantages: integration with PySCF, fast development, modularity and compatibility with user-modified code, flexibility of parallelization across diverse platforms (traditional desktop, cloud, high performance computing), and a unified codebase for running on graphics processing units or central processing units.

4.2 Code Structure and Implementation

PyQMC carries out the calculations described in Section 2.2. The main objects for executing a calculation in PyQMC, and their corresponding roles in the QMC algorithm, are shown in the table below.

Coordinates (walkers)	$\mathbf{R}$	configurations object
Wave function	Ψ	wave function object
Local observables	$\hat{H}, \hat{O}$	accumulator objects

The physical system – atoms, geometry, electrons, boundary conditions – is defined by a PySCF Mole or Cell object (for molecular or periodic systems, respectively), which is used to construct the appropriate external potential, wave function, and configuration objects in PyQMC. The Mole/Cell object also contains parameters that define the approximations used in the calculation, namely (1) the basis set, which controls the nodal surface, and (2) pseudopotentials, which replace the core electrons.

The configuration objects store the positions of all the electrons, and are used in most other parts of the code. They are essentially wrappers of a NumPy array with several convenience functions. Their main purpose is storing the boundary conditions: `OpenConfigs` and `PeriodicConfigs` objects both function identically in functions such as VMC, DMC, and accumulators. They store information for computing distances – raw or minimal image, respectively – and `PeriodicConfigs` keeps track of when electrons cross cell boundaries for computing twist phases.

Wave functions and observables are described in the next two sections.

4.2.1 Flexible Wave Function Objects

The wave function objects handle evaluation of wave functions and their derivatives. The standard implementation is the multi-Slater Jastrow (MSJ) wave function, having the form

$$\Psi(\mathbf{R}) = e^{J(\mathbf{R}, \boldsymbol{\alpha})} \sum_k c_k D_k(\mathbf{R}, \boldsymbol{\beta}), \quad (4.1)$$

where $\boldsymbol{\alpha}$ (Jastrow), $\mathbf{c}$ (determinant), and $\boldsymbol{\beta}$ (orbital) are variational parameters. The MSJ trial function allows for a compact representation of the wave function by using fewer determinants to represent static correlation and the Jastrow factor to represent dynamical correlation.[80] Fig. 4.1 compares the number of determinants needed with and without a Jastrow factor for a chain of six hydrogen atoms with lattice spacing 3.0 Bohrs, in the strongly correlated regime. The variational Monte Carlo (VMC) calculation uses a two-body Jastrow with electron-electron and electron-ion pair correlation. Fixed-node diffusion Monte Carlo (DMC)




Figure 4.1: Ground state energy of 6-atom hydrogen chain at spacing of 3 Bohr. Because the nodal error is small, DMC performs well with only a single determinant. VMC with a two-body Jastrow achieves a similar result with an order of magnitude fewer determinants than the pure multi-determinant methods. This figure is reproduced from Ref. [54].

can be interpreted as using the best possible Jastrow factor, shifting the wave function distribution without changing the nodal surface. Coupled cluster (CCSD) in the V5Z basis is near the complete basis set limit, and is consistent with the DMC energy. Heat-bath configuration interaction (HCI) approaches the CCSD value as the determinant basis size increases. The pure determinant methods require two orders of magnitude more determinants than the VMC with a two-body Jastrow to converge.

In addition, we have implemented **GeminalJastrow**, the three-body Jastrow proposed by Sorella et. al.[81] Any number of wave functions can be combined through the **MultiplyWF** and **AddWF** objects, enabling mixing and matching of wave function forms. For efficiency, the Slater object includes a linear combination of determinants without the need for combining multiple wave function objects. As a subject of active research, we expect additional wave function forms to be added over time.

New wave functions are easily implemented in the **PyQMC** framework. Any object that conforms to the wave function interface can be used in all **PyQMC** methods. For example, other groups have implemented neural network trial functions[55] and used the algorithm outlined in section 4.2.4 to optimize the wave function parameters. **PyQMC**’s testing framework makes it possible to quickly check for bugs and ensure compatibility of new objects for seamless integration.

A few comments on the implementation: Since most evaluations in **PyQMC** require computing a change of only one electron, many values are computed and stored once at the beginning of a calculation, and updated as electrons are moved one at a time. Most operations are done on the logarithm of the wave function or on

ratios for numerical precision, since wave function values often vary by many orders of magnitude.

4.2.2 Accumulator Objects

Accumulator objects compute expectation values of arbitrary operators by accumulating averages over a QMC run. They are used to compute energy and energy gradients (used in optimization), as well as observables such as density matrices and structure factors. These objects follow a simple API, and are otherwise independent of the rest of the code. An example of the accumulator API is given in Fig. 4.6.

An arbitrary operator $\hat{O}$ is evaluated on wave functions Φ and Ψ as follows

$$\frac{\langle \Phi | \hat{O} | \Psi \rangle}{\langle \Phi | \Psi \rangle} = \frac{\int d\mathbf{R} d\mathbf{R}' \Phi^*(\mathbf{R}) \Psi(\mathbf{R}') O(\mathbf{R}, \mathbf{R}')}{\int d\mathbf{R} \Phi^*(\mathbf{R}) \Psi(\mathbf{R})} \quad (4.2)$$

$$= \frac{\int d\mathbf{R} \Phi^*(\mathbf{R}) \Psi(\mathbf{R}) \int d\mathbf{R}' \frac{\Psi(\mathbf{R}')}{\Psi(\mathbf{R})} O(\mathbf{R}, \mathbf{R}')}{\int d\mathbf{R} \Phi^*(\mathbf{R}) \Psi(\mathbf{R})} \quad (4.3)$$

$$= \frac{\int d\mathbf{R} \Phi^*(\mathbf{R}) \Psi(\mathbf{R}) O_L(\mathbf{R}, \Psi)}{\int d\mathbf{R} \Phi^*(\mathbf{R}) \Psi(\mathbf{R})} \quad (4.4)$$

where the highlighted term is the local evaluation of the operator $\hat{O}$,

$$O_L(\mathbf{R}, \Psi) = \int d\mathbf{R}' \frac{\Psi(\mathbf{R}') O(\mathbf{R}, \mathbf{R}')}{\Psi(\mathbf{R})}. \quad (4.5)$$

In PyQMC, the integral over $\mathbf{R}$ is handled by the VMC algorithm, where $\Psi = \Phi = \Psi_T$, and in the case of DMC, $\Psi = \Psi_T$ is the trial function while $\Phi = \Phi_{FN}$ is the fixed-node projected wave function. We define an accumulator as an object that evaluates $O_L(\mathbf{R}, \Psi)$.

In this section, we summarize the accumulator objects implemented in PyQMC.

Gradient Operators

For semilocal operators such as gradients, the expression in Eq. 4.5 simplifies to

$$O_L(\mathbf{R}) = \frac{[\hat{O}\Psi](\mathbf{R})}{\Psi(\mathbf{R})}. \quad (4.6)$$

In PyQMC, all wave function objects can compute $\frac{\nabla_r \Psi}{\Psi}$, $\frac{\nabla_r^2 \Psi}{\Psi}$, and $\frac{\nabla_p \Psi}{\Psi}$, where ∇_r refers to the gradient with respect to a single electronic coordinate, and ∇_p refers to the gradient with respect to all variational parameters in the wave function.

Effective Core Potentials

PyQMC is compatible with semilocal effective core potentials (ECPs) (nonlocal in the angular part, but local in the radial part). ECP evaluation is implemented as in QWalk[82] using the form described by Mitas et al.[83] PyQMC automatically reads the ECPs from the PySCF Mole or Cell object.

The nonlocal operator takes the form of Eq. 4.5. The ECP operator $H^{\text{ECP}}(\mathbf{R}, \mathbf{R}') = \sum_{e,a} H_{ea}^{\text{ECP}}(\mathbf{R}, \mathbf{R}')$ is a sum of independent terms between electron e and atom a ,

$$H_{ea}^{\text{ECP}}(\mathbf{R}, \mathbf{R}') = \delta(r_{ea} - r'_{ea}) \sum_l \frac{2l+1}{4\pi} v_l(r_{ea}) P_l(\cos \theta'), \quad (4.7)$$

where r_{ea} is the distance between positions of electron e and atom a , v_l is a radial pseudopotential for angular momentum channel l , P_l is a Legendre polynomial, and θ' is the angle between $\mathbf{r}_{ea}$ and $\mathbf{r}'_{ea}$. The angular integral for each (e, a) pair is evaluated using a randomly oriented quadrature rule

$$\begin{aligned} \int d\mathbf{R}' H_{ea}^{\text{ECP}}(\mathbf{R}, \mathbf{R}') \frac{\Psi(\mathbf{R}')}{\Psi(\mathbf{R})} \\ = \frac{4\pi}{N_\Omega} \sum_{\Omega} w_\Omega H_{ea}^{\text{ECP}}(\mathbf{R}, \mathbf{R}') \frac{\Psi(\mathbf{R}'_{ea\Omega})}{\Psi(\mathbf{R})}, \end{aligned}$$

where the auxiliary configurations $\mathbf{R}'_{ea\Omega}$ are generated from $\mathbf{R}$ by moving electron e about ion a by angles $\Omega = (\theta, \phi)$ of the quadrature grid and corresponding weights w_Ω . All the quadrature rules of octahedral and icosahedral symmetries listed by Mitas et. al.[83] are implemented in PyQMC.


Reduced Density Matrices

All one-particle observables can be calculated from the one-particle reduced density matrix (1-RDM), making it a useful quantity to characterize many-body wave functions alongside the energy. In PyQMC, the 1-RDM is represented in a basis of single-particle orbitals $\phi_i(\mathbf{r})$ as

$$\rho_{ij} = \langle \Psi | c_i^\dagger c_j | \Psi \rangle, \quad (4.8)$$

where $c_i^\dagger$ and c_i are creation and annihilation operators for orbital ϕ_i , respectively. Since the reduced density matrices are completely nonlocal, we perform an auxilliary random walk, sampling a conditional probability $P'(\mathbf{R}'|\mathbf{R})$ and evaluating

$$\mathcal{O}(\mathbf{R}) = \left\langle \frac{\Psi(\mathbf{R}') O(\mathbf{R}, \mathbf{R}')}{\Psi(\mathbf{R}) P'(\mathbf{R}'|\mathbf{R})} \right\rangle_{\mathbf{R}' \sim P'(\mathbf{R}'|\mathbf{R})}. \quad (4.9)$$



```

1 import numpy as np
2 import pyqmc.api as pyq
3 import pyscf.tools
4
5 mol, mf = pyq.recover_pyscf("MnO.chk")
6 rdm1_dmc = pyq.read_mc_output(
7     "MnO_DMC.hdf5",
8     warmup=100,
9     reblock=20,
10 )
11 rdm1_ccsd_t = np.load("MnO_CCSD(T).npy")
12 ao_rdm1_dmc = np.einsum(
13     "pi,ij,qj->pq",
14     mf.mo_coeff,
15     rdm1_dmc - rdm1_ccsd_t,
16     mf.mo_coeff.conj()
17 )
18 dens_dmc = pyscf.tools.cubegen.density(
19     mol, "MnO_dmc.cube", ao_rdm1_dmc
20 )

```

Figure 4.2: The integration between PySCF and PyQMC makes it straightforward to compare properties of wave functions between different high level methods. (a) Electron density of the MnO molecule at bond length 1.6477 Å for Hartree-Fock and diffusion Monte Carlo referenced to CCSD(T) calculations. (b) The code used to compute the DMC densities. For the entire code used to generate the plots, see the supplementary information.

The 1-RDM is evaluated in QMC by averaging the quantity[74]

$$\rho_{ij} = \frac{1}{\sqrt{N_i N_j}} \left\langle \sum_{a=1}^N \frac{\Psi^*(\mathbf{R}'_a)}{\Psi^*(\mathbf{R})} \frac{\phi_i(\mathbf{r}'_a) \phi_j^*(\mathbf{r}_a)}{\rho_{\text{aux}}(\mathbf{r}'_a)} \right\rangle_{\substack{\mathbf{R} \sim |\Psi|^2; \\ \mathbf{r}'_a \sim \rho_{\text{aux}}}}, \quad (4.10)$$

where $\mathbf{R}'_a$ is generated from $\mathbf{R}$ by moving electron a , $\mathbf{r}_a \rightarrow \mathbf{r}'_a$, and $\rho_{\text{aux}}(\mathbf{r}) = \sum_i |\phi_i(\mathbf{r})|^2$ is proportional to the one-particle distribution used to sample the auxiliary coordinate $\mathbf{r}'_a$. We use McMillan's method of using the same auxiliary coordinates $\mathbf{r}'_a$ for every electron a in the sum.[84] The normalization factors

$$N_i = \left\langle \frac{|\phi_i(\mathbf{r})|^2}{\rho_{\text{aux}}(\mathbf{r})} \right\rangle_{\mathbf{r} \sim \rho_{\text{aux}}} \quad (4.11)$$

are accumulated during the Monte Carlo run, and are applied as a post-processing step using the function `normalize_obdm`.

The two-particle reduced density matrix (2-RDM)

$$\rho_{ijkl} = \langle \Psi | c_i^\dagger c_k^\dagger c_l c_j | \Psi \rangle \quad (4.12)$$

can be used to calculate all two-body observables, and is analogous to the 1-RDM. Note that in some modules in PySCF and other quantum chemistry codes, $\langle \Psi | c_i^\dagger c_j c_k^\dagger c_l | \Psi \rangle$ is evaluated instead. It is relatively easy to translate between these two representations as

$$\langle \Psi | c_i^\dagger c_k^\dagger c_l c_j | \Psi \rangle = \langle \Psi | c_i^\dagger c_j c_k^\dagger c_l | \Psi \rangle - \delta_{jk} \langle \Psi | c_i^\dagger c_l | \Psi \rangle, \quad (4.13)$$

which is done by the PySCF function `reorder_rdm`. The 2-RDM is evaluated in QMC as

$$\rho_{ijkl} = \left\langle \sum_{a < b} \frac{\Psi(\mathbf{R}'_{ab})}{\Psi(\mathbf{R})} \frac{\phi_j^*(\mathbf{r}'_a) \phi_l^*(\mathbf{r}'_b) \phi_i(\mathbf{r}_a) \phi_k(\mathbf{r}_b)}{\rho_{\text{aux}}(\mathbf{r}'_a) \rho_{\text{aux}}(\mathbf{r}'_b)} \right\rangle_{\substack{\mathbf{R} \sim |\Psi|^2; \\ \mathbf{r}'_a \sim \rho_{\text{aux}}}}, \quad (4.14)$$

PyQMC's implementation can evaluate the RDMs in an arbitrary basis.

PySCF routines can be applied directly to the 1-RDM computed in DMC to compute and plot density or other one-body quantities (Fig. 4.2). Using PySCF's built-in `cubegen.density` function removes the need to write a new script for plotting.

Computing RDMs in the same basis allows for seamless comparison between methods, i.e., by simply subtracting the matrices. Different methods are commonly compared by their energies, a single number. One- and two-particle density matrices capture more of the state and offer better comparison of properties;

methods that result in the same energy may still produce states with different densities. QMC computations of RDMs in a basis have an additional advantage that the statistical noise is much smaller compared to computing on a grid, resulting in smoother density plots. The difference in densities between Hartree-Fock, DMC, and CCSD(T) is shown in Fig. 4.2.

4.2.3 Bulk Systems

Infinite solids are approximated by finite simulation cells with twisted boundary conditions (TBCs)

$$\Psi(\mathbf{r}_1, \dots, \mathbf{r}_i + \mathbf{L}, \dots, \mathbf{r}_N) = e^{i\mathbf{k} \cdot \mathbf{L}} \Psi(\mathbf{r}_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_N), \quad (4.15)$$

where $\mathbf{k}$ is the twist; thus the basis functions are eigenstates of a translation operator. The one-particle part of the Hamiltonian commutes with the translation of a single electron, and thus can be diagonalized using basis functions of definite twist. The total energy per cell is obtained by averaging over all twists in the Brillouin zone.[85] However, the Coulomb operator does not commute with the translation operator of a single electron, and thus causes the energy eigenstates to in general be superpositions of twists.

In PyQMC, practical calculations are performed using a *supercell* approximation, in which a simulation cell larger than the primitive cell is chosen, and the Coulomb operator is truncated to remove matrix elements between different twists. This truncation can be partially corrected using the structure factor,[86] with an error proportional to $\frac{1}{N}$, where N is the number of electrons in three dimensions. Note that this correction should be performed after twist averaging above.

PyQMC contains several features to facilitate extrapolation to infinite system size. First, a PySCF mean-field calculation is performed on the primitive cell. The k -points used in the mean-field calculation determine which twists are available for a given supercell $\mathbf{S}$. The available twists are obtained in PyQMC using the function `available_twist(cell, mf, S)`, where `cell` and `mf` are PySCF cell and mean-field objects. The code then automatically generates the appropriate supercell objects from the primitive cell mean-field object. By averaging over twist, one can remove the kinetic energy finite size correction.[86] The small- k limit of the structure factor gives the approximate Coulomb finite size correction.[86] The structure factor is available as an accumulator in PyQMC.

1. Generate walkers $\mathbf{R}$
2. Compute regularization factor⁵² $f(\mathbf{R})$
 - (a) $d(\mathbf{R}) \leftarrow \frac{1}{r_{\text{cutoff}}} \frac{\Psi(\mathbf{R})}{\sqrt{\sum_e |\nabla_e \Psi(\mathbf{R})|^2}}$
 - (b) $f(\mathbf{R}) \leftarrow 9d(\mathbf{R})^2 - 15d(\mathbf{R})^4 + 7d(\mathbf{R})^6$
3. Stochastic reconfiguration^{53, 54}
 - (a) $G_{\Psi}^i(\mathbf{R}) \leftarrow \frac{\partial_{p_i} \Psi}{\Psi} \Big|_{\mathbf{R}} f(\mathbf{R})$
 - (b) $E_L(\mathbf{R}) \leftarrow \frac{\hat{H}\Psi}{\Psi} \Big|_{\mathbf{R}}$
 - (c) $G_E^i \leftarrow \langle E_L(\mathbf{R}) G_{\Psi}^i(\mathbf{R}) \rangle_{\mathbf{R}} - \langle E_L(\mathbf{R}) \rangle_{\mathbf{R}} \langle G_{\Psi}^i(\mathbf{R}) \rangle_{\mathbf{R}}$
 - (d) $S_{ij} \leftarrow \left\langle \frac{\partial_{p_i} \Psi}{\Psi} \Big|_{\mathbf{R}} G_{\Psi}^j(\mathbf{R}) \right\rangle_{\mathbf{R}} - \langle G_{\Psi}^i(\mathbf{R}) \rangle_{\mathbf{R}} \langle G_{\Psi}^j(\mathbf{R}) \rangle_{\mathbf{R}}$
 - (e) $u_i \leftarrow \sum_j (\mathbf{S})_{ij}^{-1} G_E^j$ (regularized gradient)
4. Line minimization using correlated sampling
 - (a) Select walkers $\mathbf{R}$
 - (b) $\mathbf{p} \leftarrow$ parameters of Ψ
 - (c) **for** x in $[-1, 0, 1, 2, 3]$, **do**
 - i. $\Psi_x \leftarrow$ replace $\mathbf{p}$ with $\mathbf{p} + x\mathbf{u}$
 - ii. $E(x) = \left\langle \frac{\hat{H}\Psi_x}{\Psi_x} \Big| \frac{\Psi_x}{\Psi_0} \right\rangle_{\mathbf{R}}^2$
 - (d) $E_{\text{fit}} \leftarrow$ fit $E(x)$ to cubic function
 - (e) $x_{\min} \leftarrow \arg \min_x E_{\text{fit}}(x)$
 - (f) $\Psi \leftarrow \Psi_{x_{\min}}$

Figure 4.3: Pseudo-code for the wave function optimization routine in PyQMC. The three main parts of each step of the optimization algorithm: variance regularization factor, stochastic reconfiguration, and line minimization.

4.2.4 Methods

Variational Monte Carlo (VMC)

The trial functions in section 2.2 can contain hundreds or thousands of parameters. To approximate the ground state, the parameters of the trial function are variationally optimized by minimizing the VMC energy

$$E[\Psi] = \langle \Psi | \hat{H} | \Psi \rangle = \left\langle \frac{\hat{H}\Psi(\mathbf{R})}{\Psi(\mathbf{R})} \right\rangle_{R \sim |\Psi|^2}. \quad (4.16)$$

The gradient of $E[\Psi]$ is used to determine the updates to the parameters $\mathbf{p}$ during optimization. The gradient estimator $\frac{\hat{H}\Psi}{\Psi} \frac{\partial_p \Psi}{\Psi}$ has infinite variance near the nodes of Ψ , which is removed by including the regularization factor of Ref [52]. Next, the parameter update direction is determined from $\frac{\partial E}{\partial \mathbf{p}}$ using the




Figure 4.4: Optimization of a minimal multi-Slater-Jastrow wave function for the ground and first excited states of a CO molecule, with and without optimizing the orbitals. Better variational estimates are achieved for the excited state by optimizing the orbitals.

stochastic reconfiguration technique of Casula and Sorella.[13], [87] Finally, the magnitude is determined by the minimum energy along the update direction. The parameters corresponding to the minimum are determined by a polynomial fit of correlated samples of the energy along the line. The parameters are updated, and the process is repeated to convergence.

For multi-Slater-Jastrow functions (Eq. 4.1), PyQMC supports optimization of α (Jastrow), c (determinant), and β (orbital) parameters.

VMC for Excited States

A standard approach to computing excited states is to hold orbital coefficients fixed from an excited mean-field determinant.[88] To optimize excited-state orbitals, additional measures are required to keep them from reverting to the orbitals of lower-energy states. Methods such as the state-averaged CASSCF method[89], [90] or other state averaged methods[56], [67], [71], [91], [92] allow orbital shapes to vary, but require the orbitals to be the same for all energy eigenstates. This requirement makes it easier to enforce orthogonality of eigenstates but severely limits the expressiveness of the wave functions.

In PyQMC’s implementation, excited states are kept orthogonal to lower-energy states through an overlap penalty introduced in Ref [53], allowing orbital coefficients to be optimized for each state independently. The objective function for the optimization is given by

$$O[\Psi] = \langle \Psi | \hat{H} | \Psi \rangle + \sum_{i=1}^{n-1} \lambda_i |\langle \Psi_i | \Psi \rangle|^2 \quad (4.17)$$

$$= \left\langle \frac{\hat{H} \Psi(\mathbf{R})}{\Psi(\mathbf{R})} \frac{|\Psi(\mathbf{R})|^2}{\rho} \right\rangle + \sum_{i=1}^{n-1} \lambda_i \left| \frac{N_{in}}{\sqrt{N_{ii}N_{nn}}} \right|^2, \quad (4.18)$$

where

$$N_{ij} = \left\langle \frac{\Psi_i^*(\mathbf{R})\Psi_j(\mathbf{R})}{\rho(\mathbf{R})} \right\rangle, \quad (4.19)$$

is the wave function overlap matrix and $\mathbf{R}$ is sampled from the distribution $\rho(\mathbf{R})$. Typically $\rho(\mathbf{R}) \propto \sum_i |\Psi_i(\mathbf{R})|^2$.

To demonstrate the importance of orbital optimization for excited states, we show optimizations of the ground and first excited states of the CO molecule (Fig. 4.4) using a 400-determinant multi-Slater-Jastrow ansatz, with and without optimizing orbitals. The energy is shown at each iteration over the course of both optimizations. Fixed-orbital wave functions yield an excitation energy of 9.43(5) eV, compared with 6.68(5) eV from optimized-orbital wave functions of the same form. Compared with the experimentally determined vertical excitation energy 4.76 eV,[93] optimizing orbitals results in a 60% improvement at the VMC level.

Diffusion Monte Carlo

Diffusion Monte Carlo (DMC) is implemented in PyQMC using importance sampling and the fixed-node approximation. Samples are drawn from the mixed distribution $f(\mathbf{R}) = |\Psi^*(\mathbf{R})\Phi_0(\mathbf{R})|$, where Φ_0 is the ground state, by stochastically applying a projection operator $\hat{\mathcal{P}}_\tau$ to a trial function Ψ ,

$$\Psi^*(\mathbf{R})\Phi_0(\mathbf{R}) = \lim_{N \rightarrow \infty} \langle \Psi | \mathbf{R} \rangle \langle \mathbf{R} | \hat{\mathcal{P}}_\tau^N | \Psi \rangle. \quad (4.20)$$

The time step τ is a parameter that must be extrapolated to $\tau \rightarrow 0$. Positions and weights $(\mathbf{R}_i, w_i)$ are generated by the projection $\hat{\mathcal{P}}_\tau$ at each Monte Carlo step. The fixed-node approximation is used for real wave functions, rejecting moves $\mathbf{R} \rightarrow \mathbf{R}'$ that change the sign of the trial function Ψ . For complex wave functions, the fixed-phase approximation is used.[94] Because the gradient of the phase enters into the potential, no rejection based on sign is required.

Sampling the mixed distribution results in mixed-estimator averages

$$\langle \hat{O} \rangle = \frac{\langle \Psi | \hat{O} | \Phi_0 \rangle}{\langle \Psi | \Phi_0 \rangle} \quad (4.21)$$

which are computed similarly to Eq. 4.2 as averages over walkers with additional weights w_i ,

$$\left\langle w_i \frac{\int d\mathbf{R}' O(\mathbf{R}_i, \mathbf{R}') \Psi^*(\mathbf{R}')}{\Psi^*(\mathbf{R}_i)} \right\rangle_{\mathbf{R}_i \sim f(\mathbf{R})}. \quad (4.22)$$

Branching is performed every few steps to keep weights balanced, replicating some walkers and removing others depending on their weights. In PyQMC, the branching is implemented by the stochastic comb method:[18]–

[20] walkers are resampled with probability proportional to their weights, the total weight $\sum_j w_j$ is saved, and the new weights are subsequently set equal to one. The expected contribution from each walker is correct on average, and the resulting population bias is small. This approach has the advantage of keeping the number of walkers fixed, which simplifies efficient parallelization on a fixed number of processors.

PyQMC employs two strategies proposed by Anderson and Umrigar[79] to reduce time-step errors: modified weight updates and modified T-moves[95] for nonlocal ECPs.

4.3 Diverse Workflow Support

4.3.1 Integration with PySCF

QMC trial functions are typically based on DFT or quantum chemistry methods. In many QMC codes, converters from other packages make up a large portion of the programming effort. PyQMC uses PySCF objects directly to initialize calculations, eliminating the need for converters. Mole and Cell objects define the Hamiltonian, including geometry, number of electrons, basis set, and pseudopotentials. The use of PySCF’s `eval_gto()` function to evaluate orbitals guarantees compatibility with any basis set supported by PySCF. QMC trial wave function determinants are generated from SCF objects, and there is some compatibility with multireference methods such as CAS, CASSCF, and full CI without requiring conversion steps.

Tight coupling to PySCF enables easy use of analysis routines. A common example is the calculation and plotting of density differences discussed in section 4.2.2 and shown in Fig. 4.2.

PyQMC allows for file-free computation – executing a full calculation from atomic structure to QMC result without saving any intermediate results (Fig. 4.5). Having all objects and data in the workspace streamlines prototyping of new algorithms and workflows. Python workflows using PyQMC avoid the need to repeatedly generate input files and read output files, enabling the user to handle the objects directly. Additionally, PyQMC data is saved and loaded in HDF5 format with minimal code to maintain.

4.3.2 Monkey Patching

PyQMC allows users to add modifications to a calculation locally without changing the package directory, a practice known as “monkey patching.” Although modifying the package directory is certainly possible, it poses a barrier to users in our experience. With PyQMC’s all-Python, modular structure, built-in routines are compatible with objects defined outside the package directory, such as customized wave function and accumulator objects for VMC and DMC; built-in objects can be used in externally-defined customized methods as well. Fig. 4.6 shows code outside of the package defining an accumulator object that is used

```

1 from pyscf import gto, scf
2 import pyqmc.api as pyq
3
4 mol = gto.M(
5     atom=f"Mn 0. 0. 0.; 0 0. 0. 1.6477",
6     basis="ccecp-ccpvtz",
7     ecp="ccecp",
8     spin=5,
9 )
10 mf = scf.UHF(mol)
11 mf.run()
12
13 # QMC
14 configs = pyq.initial_guess(mol, nconfig)
15 wf, to_opt = pyq.generate_wf(mol, mf)
16 pgrad_acc = pyq.gradient_generator(mol, wf, to_opt)
17 wf, optimization_data = pyq.line_minimization(wf, configs, pgrad_acc)
18 configs, dmc_data = pyq.rundmc(wf, configs)

```

Figure 4.5: Single script execution of a QMC calculation from atomic positions to QMC result. PySCF objects are used directly in PyQMC functions. Writing intermediate results to disk is not required.

```

1 import numpy as np
2 import pyqmc.api as pyq
3
4
5 class DipoleAccumulator:
6     def __init__(self):
7         pass
8
9     def __call__(self, configs, wf):
10         return {"electric_dipole": configs.configs.sum(axis=1)}
11
12     def avg(self, configs, wf):
13         avg = {}
14         data = self(configs, wf)
15         for k, it in data.items():
16             avg[k] = np.mean(it, axis=0)
17         return avg
18
19     def shapes(self):
20         return {"electric_dipole": (3,)}
21
22     def keys(self):
23         return self.shapes().keys()
24
25
26 accumulators = {"extra_accumulators": dict(dipole=DipoleAccumulator())}
27 pyq.VMC(
28     "Mn0_scf.hdf5",
29     "Mn0_vmc_dipole.hdf5",
30     load_parameters="Mn0_opt.hdf5",
31     accumulators=accumulators,
32 )

```

Figure 4.6: User code can be injected into PyQMC's QMC routines. In this example, we defined a class on the fly to compute the molecular dipole moment within the script that performs the calculation.




Figure 4.7: User-defined code can be mixed and matched in multiple ways. Externally defined accumulators are input directly into built-in VMC and DMC routines, as in Fig 6. Similarly, externally defined wave functions are simultaneously input directly into VMC and DMC.

directly in `PyQMC`'s VMC and DMC routines, in this case to compute the dipole moment of a molecule. Using custom accumulator objects is depicted in the flowchart in Fig. 4.7.

As an example of the benefits of this platform, we contrast implementation of a new VMC algorithm between Python and C++ (e.g. for sampling the sum of two wave functions in excited state optimizations). In C++, the new algorithm would require adding a file into the package, adding the file into the make system, and recompiling the distribution. In Python, a customized VMC is written, tested, and run at scale without the user modifying the distributed package at all, as depicted in the flowchart in Fig. 4.7. It is completely portable; the new algorithm file(s) can be shared and it will work for another user or machine. Developing new QMC methods and algorithms is often iterative, and by requiring fewer steps, this Python implementation greatly reduces friction for users and developers to explore new ideas.

4.4 Acceleration Strategies

`PyQMC` supports two acceleration strategies: parallel execution, and the use of graphical processing units (GPUs). The strategies work simultaneously: QMC calculations can use multiple GPUs across multiple computational resources.




Figure 4.8: Parallel scaling of VMC on a coronene molecule, with the number of walkers scaled proportionally to the number of cores. Calculations were run on Summit and parallelized with `Dask`.[\[101\]](#)

4.4.1 Parallelization

PyQMC makes use of Python’s standard library `futures` objects for parallelization. For compatibility with PyQMC, a futures object only needs to implement the `submit` function, which distributes work onto a remote process or server. The wave function data and a subset of walkers are sent to each worker process, and the results are collected as the processes finish.

By using `futures` objects, PyQMC can take advantage of many parallelization strategies. The Python standard library `concurrent.futures` provides on-node process-based parallelization. Other packages can be installed and used with the code transparently; for example `mpi4py`[\[96\]](#)–[\[99\]](#) provides futures over the high performance computing standard Message Passing Interface.[\[100\]](#) Similarly, `Dask`[\[101\]](#) provides a futures-based interface using pilot processes that are very flexible, allowing for remote execution on cloud-based resources.

Quantum Monte Carlo methods are often termed “embarrassingly parallel,” meaning that the computational time decreases almost linearly as the number of processors increases. Fig. 4.8 shows the number of Monte Carlo steps executed per second as a function of the number of nodes used for a VMC calculation on a coronene molecule. The parallel efficiency on 64 Summit nodes (2688 cores) is above 99.9%. This scaling is representative of what one should expect in optimization, DMC, and excited state calculations (i.e., all types of calculations). PyQMC’s flexible parallel implementation does not seem to have any disadvantage compared to more standard approaches using MPI.

4.4.2 Graphical Processing Unit Acceleration

PyQMC uses the same code to run on GPUs or on CPUs without any modifications. GPU capability is implemented using the CuPy library,[102] which is used as a drop-in replacement for NumPy. When `gpu.py` is loaded at the start of a calculation, it checks if GPU resources are available. Specifically, if the CuPy library is available and has access to a GPU, it is used as the library for array operations; otherwise, NumPy is used. A few auxiliary functions are used from CuPy (`get_array_module`, `asnumpy`, `fuse`), which are defined manually if CuPy is not used. In the current implementation, PyQMC only looks for a single GPU; if multiple GPUs are available, they need to be divided up so that each process sees only a single GPU. Currently, wave function evaluation and Ewald summation, which are computationally intensive, run on GPU when available and return arrays on CPU. VMC, DMC, and other algorithmic-level functions are coded entirely on CPU; implementing new algorithms does not require any extra interfacing to make use of GPU resources.

GPUs are massively parallel computing devices that can greatly speed up calculations, but only when given a sufficient amount of computational work to perform. Transferring data between CPU and GPU is slow, so the GPU must perform enough accelerated work to compensate for the latency before any speedup can be obtained. As with CPUs, data for a set of walkers is stored together in a single array. The arrays sent to GPU for wave function evaluation – distances and atomic orbital evaluations – additionally scale with the number of electrons. Thus, most GPU work scales approximately linearly with $N_{\text{walker}}N_{\text{elec}}$.

Fig. 4.9 shows the GPU speedup (ratio of CPU time to GPU time) versus $N_{\text{walker}}N_{\text{elec}}$ for a sequence of hydrocarbon molecules: benzene (30 electrons), anthracene (66 electrons), coronene (108 electrons), ovalene (142 electrons), and hexabenzocoronene (186 electrons). The calculations used correlation-consistent effective core potentials and corresponding VDZ basis sets for both H and C atoms from pseudopotentiallibrary.org.[103], [104] Each calculation was carried out on a single node of the Summit supercomputer at Oak Ridge National Laboratory. For sufficiently large numbers of electrons (about 60-100), the speedup collapses onto a single line which only depends on $N_{\text{walker}}N_{\text{elec}}$, approximately the amount of work given to the GPU.

We believe that there could be improvements to the GPU performance of the code by porting more of the code from CPU to GPU. In particular, PyQMC uses PySCF’s functions to evaluate the atomic orbitals on the CPU. For the molecules shown in Fig. 4.9, the atomic orbital evaluation takes up 15-20% of the time, meaning that the GPU speedup in these tests is limited to a maximum of five or six, even if it performed the work in zero time with zero latency. In the future, we are thus targeting this bottleneck to achieve better GPU speedups.




Figure 4.9: GPU speedup ($t_{\text{CPU}}/t_{\text{GPU}}$) versus the product of the number of electrons and number of walkers (amount of work) for a sequence of hydrocarbon molecules: benzene (30 electrons), anthracene (66 electrons), coronene (108 electrons), ovalene (142 electrons), and hexabenzocoronene (186 electrons). Comparisons are run on a single Summit node, using 37 CPUs in both cases, and six GPUs for the GPU case. The calculations were parallelized using Dask. For large enough problem size, the GPU speedup depends only on the amount of work available.

4.5 Conclusion

PyQMC is a production-level, feature-complete, and state-of-the-art QMC implementation linked with PySCF. Because PyQMC is implemented entirely in Python, it is extremely flexible and modular. Similarly to PySCF for standard quantum chemistry methods, PyQMC is aimed at both production level calculations and development of new methods. Just within our group and others, these features have already led to new algorithmic developments.[52]–[55], [105] PyQMC is licensed under the MIT license[106]–[108], and is thus freely available to download and modify. Other groups are free to build on the base implementations laid out here.

Python’s high level of abstraction greatly reduces the human time required to customize implementations and develop new ideas. The library ecosystem is well-developed, including libraries for scientific computing (NumPy,[109] SciPy[110]), data I/O (h5py[111]), parallelization (concurrent, MPI for Python[99], Dask[101]), and GPU execution(CuPy[102]) PyQMC is written in such a way that almost all computationally intensive tasks are actually executed in compiled C or Fortran code provided by one of those libraries, so that the performance is competitive with packages implemented completely in compiled languages while code can be written at high level.

The latest version of the PyQMC source code is available at <https://github.com/WagnerGroup/pyqmc>. PyQMC can also be installed from the Python Package Index (PyPI) via `pip install pyqmc` or `pip install`

`pyscf[pyqmc]`.

4.6 Implementation Details

4.6.1 Optimization of Complex Parameters

QMC wave functions $\Psi(X; p)$ are generally complex functions of electron coordinates X and a set of parameters p , which may be complex. Optimizing the wave function requires the derivative of the energy $E(p)$ with respect to the parameter p .

To compute these derivatives dE/dp , the parameters p must be real. If $p \in \mathbb{C}$ is a complex parameter, and $f(p) : \mathbb{C} \rightarrow \mathbb{C}$ is an analytic function, then f can be expanded around a point p_0 as

$$f(p) \approx f(p_0) + \left(\frac{df}{dp} \right) \bigg|_{p_0} (p - p_0) + \dots \quad (4.23)$$

The second term is generally complex, since the expansion is valid for $(p - p_0)$ of any phase. For that reason, if $E(p) : \mathbb{C} \rightarrow \mathbb{R}$ is a real-valued function of p , it can't be analytic, and $\frac{dE}{dp}$ is not defined. In this case it only makes sense to ask about the derivatives with respect to real parameters, such as $dE/d\text{Re}(p)$ and $dE/d\text{Im}(p)$.

In our case, the energy $E(p) = \int \Psi^*(X; p) H \Psi(X; p) dX / \int \Psi^*(X; p) \Psi(X; p) dX$ is a functional of a complex-valued function $\Psi(X; p)$, so to compute derivatives of E , we need the derivatives of Ψ . Suppose the derivative $\frac{\partial \Psi}{\partial p}$ is available. Assuming Ψ is analytic in p , the derivative is the same from every direction,

$$\frac{\partial \Psi}{\partial \text{Re}(p)} = \frac{\partial \Psi}{i \partial \text{Im}(p)} = \frac{\partial \Psi}{\partial p}. \quad (4.24)$$

For the complex conjugate, the “real derivative” is

$$\frac{\partial \Psi^*}{\partial \text{Re}(p)} = \lim_{h \rightarrow 0} \frac{\Psi^*(p + h) - \Psi^*(p)}{h} = \left(\frac{\partial \Psi}{\partial p} \right)^*, \quad (4.25)$$

and the “imaginary derivative” has a flipped sign,

$$\frac{\partial \Psi^*}{i \partial \text{Im}(p)} = \lim_{h \rightarrow 0} \frac{\Psi^*(p + ih) - \Psi^*(p)}{ih} = - \left(\frac{\partial \Psi}{\partial p} \right)^*, \quad (4.26)$$

because of the imaginary denominator, which flips sign under conjugation.

PyQMC splits complex parameters into real and imaginary parts in the `LinearTransform` object. `LinearTransform` converts parameters and their gradients from dictionaries of arrays, as stored in the wave functions, to a one dimensional list, which are usable in the optimization algorithms. (And converts a parameter list back

into the appropriate dictionary of arrays, of course.) If complex parameters are detected, their real and imaginary parts are added to the list. The serialized gradients list concatenates the original gradient list with the gradients of complex parameters multiplied by i ,

$$\left[\frac{\partial \Psi}{\partial \text{Re}(p)}, \frac{\partial \Psi}{\partial \text{Im}(p)} \right] = \left[\frac{\partial \Psi}{\partial p}, i \frac{\partial \Psi}{\partial p} \right]. \quad (4.27)$$

This procedure correctly converts each complex parameter into two real parameters that can be used to represent energy gradients.

CHAPTER 5

DESCRIBING SPIN FLUCTUATIONS IN CaCuO_2

This chapter is based on unpublished work in collaboration with Chun Yu Chow and Lucas Wagner.

5.1 The Challenge of Describing Cuprates

Identifying a mechanism for the superconductivity observed in cuprates has been a long-standing challenge.[112]–[119] In their pure stoichiometric form, the cuprates exhibit anti-ferromagnetic (AFM) correlations. When hole doped in the range of approximately 12%-25% (depending on the specific material), a superconducting phase is observed.[112] The proposed and debated effective models for cuprates describe the movement of holes through the lattice. An interaction leading to pairing of the holes would explain the observed superconductivity; however, no pairing mechanism has yet been established.

There is evidence that the holes interact with the background AFM spins to form spin polarons that move through the lattice.[57], [115], [120]–[122] The details of the interactions affect the nature of the hole motion and their interactions and are not yet well understood. *Ab initio* quantum Monte Carlo (QMC) calculations have characterized the low-energy spin configurations surrounding a hole in the Cu-O plane;[57] however, these calculations were restricted to single fixed-spin configurations. Multiple configurations close in energy were an indication that spin fluctuations are important to the polaron description – and spin fluctuations have long been known to be relevant to the physics of the cuprates.[120], [123], [124] Capturing spin fluctuations requires investigation with multideterminant trial wave functions.

We assess the performance of multideterminant wave functions in QMC calculations of a simpler system – a 2×2 unit cell of undoped CaCuO_2 . Since we have a good understanding of the spin behavior, characterizing the multideterminant wave functions in the undoped material establishes a baseline for using these trial functions in the more interesting doped material.

5.2 Methods: Multideterminant QMC Approach

In their undoped form, the cuprates exhibit anti-ferromagnetic (AFM) correlations, and correspond well to a Heisenberg Hamiltonian on a square lattice, where each site represents a Cu $3d_{x^2-y^2}$ orbital,

$$\hat{H} = J \sum_{(i,j)} \hat{\mathbf{S}}_i \cdot \hat{\mathbf{S}}_j, \quad (5.1)$$

where (i, j) are nearest-neighbor sites on a square lattice, $\hat{\mathbf{S}}_i$ is the spin operator on site i , and $J > 0$ corresponds to antiferromagnetic coupling. To compare against our *ab initio* results, we consider the same 2×2 unit cell of the Heisenberg model. The lowest energy classical spin configuration, where each site is either up or down, is the Néel state: the spins alternate in a checkerboard pattern. However, the ground state of the model is an equal superposition of the two Néel states *plus* negative contributions from the four possible stripe-ordered configurations (shown in Fig. 5.1). These superpositions are spin fluctuations - they are necessary for describing the 2×2 Heisenberg ground state. This model gives us a reference for assessing how well our *ab initio* wave functions capture the expected spin fluctuations.

5.2.1 Ab Initio Calculations

In *ab initio* cuprate calculations, orbitals are typically set to break spin symmetry in the Néel ordering, so that each site has a net moment up or down. Traditional QMC approaches use a Slater-Jastrow Ansatz with a single Slater determinant constructed from a set of symmetry-broken DFT orbitals. A single determinant corresponds to a single spin configuration in the Heisenberg model – the spins at each site are fixed by the orbitals in the determinant. A state constructed from orbitals with Néel spin ordering is therefore associated with the Néel state of the Heisenberg model. The orbitals additionally determine the final projected state of fixed-node diffusion Monte Carlo (DMC) through the nodal surface of the determinant.

5.2.2 Multideterminant QMC Wave Functions

To establish the potential impact of spin fluctuations in cuprate calculations, we study infinite-layer CaCuO_2 , one of the simplest structures in the cuprate family. We compare the single-determinant approximation to multideterminant, spin-symmetric trial functions. Small determinant expansions capture static correlation – including spin fluctuations associated with the Heisenberg ground state – and enable the nodal surface to be optimized in the presence of short-range Jastrow correlation. Additionally, we assess the quality of the DFT orbitals by optimizing the orbital shapes in the presence of explicit correlation.




Figure 5.1: (a) Ground state and Néel state of the 2×2 Heisenberg model. Upper plots show spin structure factor S_q . All q values beyond the first three are higher harmonic copies, but are shown for later comparison to *ab initio*. Lower plots show the superposition coefficients in the basis of spin configurations. (b) The basis of spin configurations.

5.3 Calculation Details

We use a unit cell containing four Cu atoms, which can realize antiferromagnetic and stripe orderings. Although finite size errors are present, the calculation is large enough to observe the effects of spin fluctuations in the wave functions.

The workflow for the calculation proceeded as follows: density functional theory (DFT) calculations were performed for restricted (RKS) and unrestricted (UKS) spin ansätze. The RKS orbitals were used in a complete active space self-consistent field (CASSCF) calculation with a four-electron, four-orbital active space. The QMC multideterminant wave functions were initialized from the CASSCF results; the single-determinant wave functions were initialized from the symmetry-broken UKS orbitals. Different sets of parameters were variationally optimized and used as trial functions in diffusion Monte Carlo.

5.3.1 Converging DFT

For the single determinant calculation, we obtain orbitals from an unrestricted Kohn-Sham (UKS) DFT calculation. This allows the up-spin and down-spin orbitals to have different shapes and be centered on different atoms, which makes different magnetic orderings possible; however, it does not allow for spin fluctuations. The DFT calculation was carried out in PySCF[125], [126] using the hybrid exchange-correlation functional PBE0.[127] We used correlation consistent effective core potentials (ccECPs) and corresponding triple-zeta Gaussian-type basis elements from pseudopotentiallibrary.org.[103], [104], [128] The calculation cell contains four formula units and the calculation was performed only for the γ point.

To converge the AFM spin order, we first converged a calculation with smearing, which causes orbitals near the Fermi energy E_F to be partially filled, rather than fully filled below E_F and empty above E_F . The resulting spin up and spin down densities were slightly adjusted to create an AFM spin density, and used as the starting point for a calculation with no smearing. The converged result maintained the AFM order.

5.3.2 CASSCF Calculations

For the multideterminant calculation, we obtain the orbitals initially from a restricted Kohn-Sham (RKS) DFT calculation. From these initial RKS orbitals, we select an active space for the CASSCF calculation. Since we expect the low-energy properties to come primarily from the partially filled Cu $3d_{x^2-y^2}$ orbitals, the active space orbitals are selected as the four orbitals closest to the Fermi energy that have strong $d_{x^2-y^2}$ character. The blue curve in Fig. 5.2 indicates the Cu $d_{x^2-y^2}$ character of the orbitals. The peaks at 70, 81, 82, 83 correspond to the four orbitals chosen to be the active space.




Figure 5.2: Partial density of states of RKS orbitals on the copper $d_{x^2-y^2}$ orbitals. The four orbitals with the highest values are used to construct the CASSCF active space.

We performed a CASSCF calculation using a much larger active space of 22 orbitals and 20 electrons to check the active space. The natural orbitals and occupations were computed from the 1-RDM of the ground state. Four natural orbitals had partial occupation, and these exactly matched the four molecular orbitals with $3d_{x^2-y^2}$ character, confirming that the four-orbital active space is adequate.

5.3.3 Parameter Optimization in VMC

We carry out VMC optimization of SJ (single-determinant) and MSJ (multideterminant) trial functions. For both SJ and MSJ forms, we compare optimizations of four different sets of parameters: (1) β ; (2) α, β ; (3) β, g ; (4) α, β, g . The symbols represent the following parameters:

symbol	parameters
β	two-body Jastrow
α	orbitals
g	geminal Jastrow
c	determinant coefficients (MSJ only)

In the MSJ optimizations, c and α are optimized together. The orbital parameters are optimized in batches of one to eight orbitals at a time to avoid optimizing too many parameters simultaneously.




Figure 5.3: Geminal Jastrow basis comparison

The geminal Jastrow has the form

$$J_G(\mathbf{R}) = \sum_{ij} \sum_{mn} g_{mn} \chi_m(\mathbf{r}_i) \chi_n(\mathbf{r}_j), \quad (5.2)$$

where $\chi_m(\mathbf{r})$ are one-particle orbitals chosen from the atomic GTO basis functions. To keep the number of parameters manageable, only s orbitals were used. Different numbers of s orbitals on Cu and O atoms were tested on a CuO molecule (Fig. 5.3). With a few orbitals, there is a noticeable improvement over the VMC energy with only a two-body Jastrow. We chose to include four orbitals on Cu atoms and two orbitals on O atoms in the geminal basis. On the full solid, two Ca orbitals did not show improvement over one, so a single Ca orbital was used.

5.3.4 Diffusion Monte Carlo

Fixed-node diffusion Monte Carlo (DMC) is run on the optimized trial functions to project out the ground state. The fixed-node approximation restricts the projected wave function to have the same nodal surface ($\Psi(\mathbf{R}) = 0$) as the trial function. The error due to the fixed-node approximation is determined by the shape of the nodal surface, which is modified by the orbital parameters α and determinant coefficients c . Since Jastrow factors do not affect the sign of the wave function, the Jastrow parameters β, g do not directly affect the nodal surface. However, the orbitals are optimized in the presence of the Jastrow, so we expect a better Jastrow (i.e., the geminal Jastrow) to result in better orbital optimization, and thus a better nodal surface.




Figure 5.4: VMC (left) and DMC (right) energies of different optimized wave functions. The x-axis labels indicate which parameters were optimized: β two-body Jastrow; α orbital shapes; c determinant coefficients (MSJ only); g geminal Jastrow. Blue crosses are MSJ functions and orange crosses are SJ. VMC energies decrease as more parameters are optimized. DMC energies are lowest for single-determinant symmetry-broken orbitals. DMC energies were extrapolated to zero time step from time steps 0.01, 0.02, and 0.04 Ha^{-1} . Note the difference in energy offset: VMC energies are offset by -1060 Ha; DMC energies are offset by -1061 Ha. Differences between VMC and DMC are much larger than differences between wave functions.

Another source of error is introduced by T-moves[79], [95], used to treat the nonlocal ccECPs. This error depends on the accuracy of the trial function, so we expect that all parameters contribute to reducing this error. The DMC calculations are carried out in PyQMC[129]. Time steps of 0.01, 0.02, and 0.04 Ha^{-1} were used for all wave functions and extrapolated to zero.

5.4 Results

For single-determinant (SJ) and multideterminant (MSJ) wave functions, and each of the four optimized parameter sets, we compute total energies at the VMC and DMC level. We additionally characterize the magnetic properties of the wave functions through spin populations and spin structure factor.

5.4.1 Energies

The final energies after VMC optimization for the eight cases are shown in the left panel of Fig. 5.4. Optimizing orbitals (α) from their original shapes can only lower the variational energy, and indeed, the VMC energy after optimizing orbitals is lower in every case. Similarly, including a geminal Jastrow (g parameters) lowers the energy in every case.

The determinant expansion is based on spin-restricted orbitals (restricted to preserve spin symmetry), so it is not guaranteed to give lower energy than the spin-unrestricted single determinant; and in fact, with Jastrow optimization, the MSJ Ansatz does not achieve lower energy than SJ. In the VMC optimization, up and down orbitals are optimized separately, so in principle the MSJ can achieve at least the accuracy of the SJ wave function when the orbitals are optimized. Reaching the same or lower minimum is not guaranteed, because the trial functions start at very different places in a complicated parameter space with many local minima; however, our results confirm that the MSJ achieves slightly lower energy than the SJ wave function when orbitals are optimized. We note that this would not be possible without the capability to optimize orbitals in QMC, since the CASSCF orbitals need to adapt to the correlations added by the Jastrow factor. Thus, at the VMC level, we are able to achieve the same or better performance while restoring spin symmetry to the wave function.

The DMC results are quite different, shown in the right panel of Fig. 5.4. Surprisingly, optimizing the orbitals of SJ wave functions *increased* the DMC energy, indicating that the nodal surface of the wave function became worse. We attribute this to the two-body Jastrow not being sufficiently close to the DMC projection. To explain this, we imagine a hypothetical “perfect” Jastrow factor. Since DMC does not change the sign of the wave function, the “perfect” Jastrow would result in a wave function equivalent to the DMC projection. In the presence of such a Jastrow, orbital optimization would be guaranteed to improve the DMC energy, since it is being optimized directly. With the aim of getting more accurate orbitals, we added a geminal Jastrow to better describe electron correlations; however, even with the improved Jastrow, orbital optimization worsens the DMC result for the single-determinant wave functions. The orbitals from DFT seem to be very good in this case, and improving the nodal structure through orbital optimization is difficult. A possible reason for this is that the regions near the nodes are sampled with low probability, and at the VMC level the quality of the nodes may be sacrificed for improvements in higher probability regions; however, at the DMC level, the nodal surface is more consequential.

MSJ wave functions perform worse than SJ in all optimization cases. We note that orbital optimization does improve MSJ energies. In CASSCF, the orbitals are determined without any other correlation effects (e.g. the xc-functional in DFT), so we expect them to relax in the presence of a Jastrow factor. Although we optimized the MSJ orbitals to the best of our current capability, with a large number of nonlinear parameters and many local minima, we may find further improvement with better optimization strategies. From our current data, we conclude that spin fluctuations are not a dominant contributor to the energetics of the undoped CaCuO_2 ground state.




Figure 5.5: Unit cell of calculation (solid square) and unit cell for counting spin and charge populations (dashed rectangle). Blue circles represent Cu atoms, red represent O atoms. The Cu plaquettes are the shaded regions defined by the four neighboring O atoms as their corners. Néel, h-stripe, and v-stripe diagrams show spin orderings. Order parameters are evaluated by adding (yellow) and subtracting (cyan) spin populations in each of the Cu plaquettes.

5.4.2 Magnetic Properties

We go on to compare magnetic properties of the different optimized wave functions in reference to the Heisenberg model. Our primary approach to this is to count the spin populations in different regions of the unit cell. Measuring spin population around the Cu atoms gives us an estimate of the magnetic moments that correspond to a model “site” and a measure of order parameters, namely the AFM and stripe ordering. The unit cell is divided into plaquettes surrounding the Cu atoms and plaquettes in the empty spaces, illustrated in Fig. 5.5. In each of the eight regions (plaquettes with and without Cu atoms), we measure the population of up and down electrons, whose sum and difference are the charge and spin populations. The Néel order parameter is computed as the correlation of spin populations on Cu plaquettes,

$$\text{Néel}(\mathbf{R}) = S_0(\mathbf{R}) + S_2(\mathbf{R}) - S_5(\mathbf{R}) - S_7(\mathbf{R}), \quad (5.3)$$

where $S_i(\mathbf{R}) = n_{i\uparrow}(\mathbf{R}) - n_{i\downarrow}(\mathbf{R})$ is the net spin of configuration $\mathbf{R}$ in plaquette i , and $n_{i\sigma}(\mathbf{R})$ counts the number of electrons with spin σ in plaquette i in configuration $\mathbf{R}$. Similarly, the stripe order parameters are defined as

$$\text{h-stripe}(\mathbf{R}) = S_0(\mathbf{R}) - S_2(\mathbf{R}) + S_5(\mathbf{R}) - S_7(\mathbf{R}) \quad (5.4)$$

$$\text{v-stripe}(\mathbf{R}) = S_0(\mathbf{R}) - S_2(\mathbf{R}) - S_5(\mathbf{R}) + S_7(\mathbf{R}). \quad (5.5)$$

Order parameter histograms over configurations are shown in the left plots of Fig. 5.6. The main differences in magnetic properties were observed between single-determinant (SJ) and multideterminant (MSJ) wave

functions. Whether orbitals, a geminal Jastrow, or even a regular two-body Jastrow was optimized had little impact on the measured quantities. For clarity, only a single representative histogram is shown for each of SJ and MSJ wave function groups.

An immediate striking feature is that the signal of spin ordering is dwarfed by fluctuations. These fluctuations come from the many other electrons (not just the $3d_{x^2-y^2}$ orbitals that we consider as the model basis) fluctuating across the boundary of the defined plaquette regions. While the ordered moment has a moment of at most one spin per site, the fluctuations are around five spins per site. To isolate the ordering present in the wave function, we subtract off a reference computed from an unpolarized DFT Slater determinant.

The effects we anticipated from the model are shown by the differences in the right plots of Fig. 5.6. In the Néel order parameter, the SJ function shows a shift to the right as a result of breaking spin symmetry and the imposed order. The MSJ function preserves spin symmetry, but also shifts weight away from zero to equal positive and negative values, a measure of the antiferromagnetic correlation. Both stripe order parameters (horizontal and vertical) have the same distributions, as expected from fourfold rotation symmetry. The SJ functions show negligible stripe ordering, while there are clear stripe correlations in the MSJ functions, about six times smaller than the Néel ordering. The ground state of the 2×2 Heisenberg model has Néel order four times stronger than its stripe order, surprisingly consistent for such a reduced model. The clear Néel and stripe correlations apparent in the MSJ wave functions – and their absence in the SJ wave functions – show that the improved wave functions are effective at capturing the spin fluctuations present in the Heisenberg ground state of the corresponding 2×2 unit cell.

Another standard tool for measuring magnetic correlation is the spin structure factor, S_q . The left plots in Fig. 5.7 show structure factors of the different orbital-optimized wave functions. As seen in the order parameters, the background dominates over the spin signal. The background was approximated by fitting the form

$$(1 - e^{-b|\mathbf{q}|})^2 \quad (5.6)$$

at all the points at which the Heisenberg ground state S_q is zero, and shown in orange on top of S_q . This form satisfies the known properties that $\lim_{q \rightarrow \infty} S_q = 1$ and $S_q \propto q^2$ for small q , and fits the data well. S_q with the background subtracted and scaled by $2N_{\text{elec}}$ is shown in the right plots in blue, overlaid with S_q of the Heisenberg ground state in orange. Overall, the alignment of many peaks shows that the square lattice geometries match. In the Heisenberg data, only the first two peaks are distinct, as a result of having only four sites; the rest are images occurring at harmonics. In the *ab initio* data, the presence of oxygen atoms and electrons in real space distinguishes several more peaks, although the signal still includes the same harmonic




Figure 5.6: Order parameter histograms from spin populations within Cu plaquettes (VMC). Orange distributions are the single determinant (SJ) functions and blue distributions are multideterminant (MSJ). Left plots show the full distribution; right plots show the difference with respect to an unpolarized wave function. Top: Néel order. The SJ distribution is shifted to the positive side, while the MSJ preserves spin symmetry. Both are shifted away from zero relative to the unpolarized state. Middle: horizontal stripe order. Bottom: vertical stripe order. The horizontal and vertical stripe distributions are the same due to symmetry. The MSJ distribution is shifted away from the reference, indicating stripe order. The SJ distribution matches the unpolarized, showing no stripe order.




Figure 5.7: Left: Spin structure factor S_q of optimized SJ and MSJ wave functions (blue). The orange curve is fit to the background. Right: S_q with background subtracted (blue) compared to S_q of the Heisenberg model (orange). MSJ wave functions more closely match the Heisenberg ground state than the SJ wave functions.

$q / \left(\frac{2\pi}{a} \right)$	q	interpretation
1	0.43	AFM lattice vector
$\sqrt{2}$	0.61	Cu diagonal neighbors
2	0.86	Cu nearest neighbors
$\sqrt{5}$	0.96	Cu-O second neighbors
$2\sqrt{2}$	1.22	O-O nearest neighbors
3	1.29	in between Cu-O bonds across Cu or O
4	1.72	Cu-O neighbors

Table 5.1: Amplitudes of reciprocal lattice vectors $\mathbf{q}$ used to compute S_q and descriptions of the real-space distances they correspond to in the Cu-O plane. Only in-plane vectors with $q_z = 0$ are considered.

contributions as the Heisenberg S_q . The length scales of some of the lower q -values are described in Table 5.1.

The SJ wave functions have some noticeable deviations: they miss the first peak at $q \sim 0.43$ entirely, are low at the $q \sim 0.96$ peak and high at $q \sim 1.22$, where there is no Heisenberg peak. The lower peak in all functions at $q = 1.72$ is due to the oxygen bridging neighboring Cu atoms, interfering with the Heisenberg model harmonically than the Heisenberg model. The multideterminant wave functions match the Heisenberg model noticeably better at all q values. Based on the peak mismatches, it seems that the symmetry-broken functions describe oxygen spins differently from the MSJ functions. We regard the SJ-MSJ difference in S_q as improved descriptive capacity that can be useful for characterizing holes in doped CaCuO_2 , which reside primarily on oxygen sites.

5.5 Conclusion

We have demonstrated that spin fluctuations are well described by relatively simple MSJ wave functions in a 2×2 unit cell of undoped CaCuO_2 , indicating promise for applying them to the doped case, where the spin fluctuations are less understood. At the variational level, this is the lowest energy result to date using multideterminant wave functions and orbital optimization, showing that QMC is no longer restricted by the quality of wave functions generated from other techniques; all trial function parameters can consistently be optimized within QMC. With this capability, we can investigate complex wave functions for complex systems, including improved excited states with relaxed orbital shapes. In addition, we have identified a system in which the lowest energy variational wave function does not correspond to the best nodes. This is likely due to the specific shape at the nodes, despite the wave functions being more physically accurate overall. Our attempt to improve the nodes through an improved Jastrow did not resolve the difference; however, there are other algorithms like self-healing DMC that target the nodal surface directly.[130], [131] The test case of CaCuO_2 will be a useful benchmark for further research on wave function optimization algorithms.

CHAPTER 6

ENSEMBLE VMC FOR EXCITED STATES

This chapter is based on my published work “Ensemble variational Monte Carlo for optimization of correlated excited state wave functions.” [132] In this paper, we present a new variational approach to simultaneously finding several low-energy eigenstates of a Hamiltonian in a single calculation.

As discussed in Chapter 2, variational Monte Carlo (VMC) is a powerful method for finding the ground state of a many-body Hamiltonian $\hat{H}$. VMC allows for flexible, many-body wave function parameterizations that are well suited to describing strongly correlated systems. The variational principle is particularly helpful in optimizing these wave functions. It asserts that the average energy of an arbitrary wave function Ψ is an upper bound to the energy of the ground state, $E[\Psi] \geq E[\Phi_0]$. This upper bound property allows us to optimize wave function parameters by minimizing a single number and we can easily identify an increase in accuracy by a decrease in the energy.

While there has been recent work on applying VMC to excited states, [53], [133]–[135] there is to date no general variational principle for an ensemble of N states of the form $O[\{\Psi_i\}] \geq O[\{\Phi_i\}]$, analogous to the ground state variational principle, where $\{\Phi_i\}$ is the set of N lowest-energy eigenstates. We introduce such a variational principle by defining a functional O and showing that it has this upper bound property. This allows us to apply all the same optimization techniques for the ground state to the more general case, and we can compare ensemble accuracies by their average O values, which has not been directly possible through previous methods.

This work is also relevant to the problem of generating low-energy effective models from *ab initio* calculations, such as for the cuprate system CaCuO_2 discussed in Chapter 5. A useful model Hamiltonian should reproduce not just ground state properties, but energy gaps and excited state properties as well, requiring accurate *ab initio* excited state calculations. Chapter 5 discussed using improved wave function forms for more accurate calculations of the ground state of CaCuO_2 . The ensemble variational principle discussed in this chapter provides a robust and systematic way of optimizing these improved wave functions for computing excited states.

6.1 Introduction

Quantum Monte Carlo (QMC) has been established as a highly accurate method for computing ground-states of *ab initio* systems.[58], [59], [136]–[138] Progress has also been made towards efficiently computing accurate excited states in QMC;[139]–[143] however, there are limitations to the methods that have been applied to *ab initio* systems so far. A state-averaging approach has been used to simultaneously optimize a set of ground and excited states, where a common set of orbitals and Jastrow are used to construct all the states.[56] In many systems, optimal orbitals for describing an excited state are quite different from the optimal ground-state orbitals, so large determinant expansions are required to accurately represent these states using a shared set of orbitals. An optimization approach using variance of the local energy has been used to target excited states without state averaging;[66], [144]–[146] however, relying on the variance leads to difficulty optimizing to the correct minimum in some situations.[92] An alternative method avoids state averaging and variance optimization by penalizing the wave function’s overlap with a set of known eigenstates.[53], [147] This method only optimizes one excited state at a time, requiring several expensive calculations in sequence.

We note that there has been some work that optimizes all states simultaneously[133], [135] using a modification of the method in Refs. [53], [147]. Additionally, a recently proposed method simultaneously optimizes excited states by representing the ensemble as a determinant of many-body wave functions, avoiding the need to choose an overlap penalty. [134] In the approaches of Refs. [133], [134], there is a degenerate minimum in which all wave functions are linear combinations of the eigenstates.[148] While in principle one could rediagonalize the Hamiltonian in the subspace, this is inconvenient in variational Monte Carlo (VMC) since the wave function parameterizations are often nonlinear and the off-diagonal matrix elements of the Hamiltonian typically have high variance. Finally, Ref. [135] includes an overlap penalty between states i and j only if $i < j$, effectively implementing the sequential optimization of Refs. [53], [147] in a single calculation by using separate, dynamic cost functionals for each state.

In this paper, we define a single, penalty-based cost functional to optimize an ensemble of states simultaneously. For a finite, given value of penalty (no free parameters), the functional has the eigenstates as its global minimum, with Hessian eigenvalues all greater than 0. Using a weighted average of energies, as proposed in Ref. [149], this cost functional yields the low energy eigenstates, and the Hamiltonian does not need to be rediagonalized in the final low-energy subspace. We derive conditions on the optimization parameters that ensure the ensemble optimizes to the lowest-energy eigenstates. As a practical demonstration, we optimize this cost functional to find excited states of a CO molecule. We believe that strategies based on this functional are likely to be efficient and practical.

6.2 Penalty-Based Cost Functional

Summary and statement We aim to find the lowest N eigenstates of a time-independent Hamiltonian whose eigenstates $|\Phi_i\rangle$ satisfy $H|\Phi_i\rangle = E_i|\Phi_i\rangle$. The minimum of the N state functional of $\{\Psi_i\}$

$$O[\{\Psi_i\}] = \sum_i w_i E[\Psi_i] + \lambda \sum_{i < j} |S_{ij}|^2, \quad (6.1)$$

where $E[\Psi_i] = \langle \Psi_i | \hat{H} | \Psi_i \rangle$ is the energy expectation, $S_{ij} = \langle \Psi_i | \Psi_j \rangle$ is the overlap, $w_i > 0$ are weights, and $\lambda > 0$ penalizes wave function overlap. If $w_i > w_j$ for all i, j where $E_i < E_j$, and

$$\lambda > \max_{i < j} \left[(E_j - E_i) \frac{w_i w_j}{w_i - w_j} \right], \quad (6.2)$$

then O is minimized when $\Psi_i = \Phi_i$ are equal to the lowest N eigenstates of the Hamiltonian H . We prove this condition in sections 6.2.1-6.2.2.

6.2.1 Preliminaries

If the states $\{\Psi_i\}$ are orthogonal, a few properties are known. An orthogonal ensemble satisfies the variational upper bound property[149]

$$\sum_i w_i E[\Psi_i] \geq \sum_i w_i E[\Phi_i]. \quad (6.3)$$

If the weights are all equal, every set of orthogonal states in the low-energy space $\mathcal{LE}_N = \text{span}[\Phi_0, \dots, \Phi_{N-1}]$ has the same cost,[148] resulting in a continuum of critical points. More generally, if the weights in any subset $\{w_j\}$ are equal, then any rotation of the corresponding states $\{\Phi_j\}$ leaves the cost functional unchanged. We therefore require the weights to be strictly decreasing between nondegenerate states, $w_i > w_{i+1}$ if $E_i \neq E_{i+1}$, to ensure that an orthogonal ensemble optimizes to the target ensemble $\{\Phi_i\}$. [149] Of course, within a degenerate subspace, any rotation of states leaves the energy unchanged and is an acceptable solution.

Now we turn to the case of $\{\Psi_i\}$ that are allowed to be nonorthogonal, and the minimization of O . Ψ_i are orthogonal in the limit $\lambda \rightarrow \infty$. At the opposite extreme, if $\lambda = 0$, O is minimized when all states are the ground state. We will show that there is a critical value λ_c , above which O is minimized by the ensemble of eigenstates $\{\Phi_i\}$.

We consider an ensemble $\{\Psi_i\}$ within the low-energy space $\mathcal{LE}_N$. For any choice of w_i and λ , the ensemble that minimizes O is guaranteed to be in $\mathcal{LE}_N$, [149] since any components outside $\mathcal{LE}_N$ can be rotated into $\mathcal{LE}_N$ without increasing the overlap term. We thus can parameterize the ensemble states as linear




Figure 6.1: Overlap of the optimized ensembles in the two-state model system at different values of penalty λ and weight w_0 . The weights are normalized so that $w_0 + w_1 = 1$. The blue line is the calculated critical penalty λ_c , above which the overlap is strictly zero.

combinations of eigenstates in $\mathcal{LE}_N$,

$$\Psi_i = \sum_{j=0}^{N-1} c_{ij} \Phi_j, \quad (6.4)$$

where $c_{ii} = \sqrt{1 - \sum_j |c_{ij}|^2}$ is determined by the normalization. Our target ensemble $\{\Phi_i\}$ is represented by the parameters $c_{ij} = \delta_{ij}$. In general, the cost functional with this parameterization is represented as

$$O = \sum_{ik} w_i |c_{ik}|^2 E_k + \lambda \sum_{i < j} \left| \sum_k c_{ik} c_{jk} \right|^2. \quad (6.5)$$

O is fully parameterized by c_{ij} with $i \neq j$, which are free to vary within $|c_{ij}| \leq 1$.

6.2.2 Determining Critical λ_c

λ_c for Two-State System

We will determine the critical λ_c such that the eigenstates $\{\Phi_0, \Phi_1\}$ minimize O when $\lambda > \lambda_c$. The cost functional for these two states is

$$O = \sum_{i=0}^{N-1} w_i E_i + w_0 |c_{01}|^2 \Delta E - w_1 |c_{10}|^2 \Delta E \quad (6.6)$$

$$+ \lambda |c_{00} c_{10} + c_{01} c_{11}|^2, \quad (6.7)$$

where $\Delta E = E_1 - E_0$. The gradient and Hessian of O give conditions for the critical penalty λ_c . The gradient is

$$\frac{\partial O}{\partial c_{01}} = 2w_0 c_{01} \Delta E + \lambda 2S_{01} \left(c_{11} - \frac{c_{01} c_{10}}{c_{00}} \right), \quad (6.8)$$

using the fact that $\frac{dc_{00}}{dc_{01}} = -c_{01}/c_{00}$. In this form, it is clear that our target ensemble $c_{ij} = \delta_{ij}$ is indeed a critical point.

The second derivatives are

$$\frac{\partial^2 O}{\partial c_{12} \partial c_{21}} = 2\lambda \left[\left(c_{11} - \frac{c_{21} c_{12}}{c_{22}} \right) \left(c_{22} - \frac{c_{12} c_{21}}{c_{11}} \right) - S_{12} \left(\frac{c_{12}}{c_{11}} + \frac{c_{21}}{c_{22}} \right) \right] \quad (6.9)$$

$$\frac{\partial^2 O}{\partial c_{12}^2} = 2w_1 \Delta E + 2\lambda \left(c_{22} - \frac{c_{12} c_{21}}{c_{11}} \right)^2 + 2\lambda S_{12} \left(\frac{c_{12} c_{21}}{c_{11}^2} \frac{c_{12}}{c_{11}} - \frac{c_{21}}{c_{11}} \right). \quad (6.10)$$

The Hessian matrix at the target ensemble $c_{ij} = \delta_{ij}$ is

$$H_O = \begin{pmatrix} 2w_0 \Delta E + 2\lambda & 2\lambda \\ 2\lambda & -2w_1 \Delta E + 2\lambda \end{pmatrix}, \quad (6.11)$$

where the rows and columns are the free parameters c_{01} and c_{10} . It has eigenvalues

$$\alpha = 2\lambda + (w_0 - w_1) \Delta E \pm \sqrt{(w_0 + w_1)^2 \Delta E^2 + 4\lambda^2}. \quad (6.12)$$

This critical point is a minimum when the eigenvalues are both positive, illustrated in Fig. 6.2, which constrains λ to

$$2\lambda + w_0 - w_1 - \sqrt{(w_0 + w_1)^2 + 4\lambda^2} > 0 \quad (6.13)$$

$$2\lambda + w_0 - w_1 > \sqrt{(w_0 + w_1)^2 + 4\lambda^2} \quad (6.14)$$

$$(2\lambda + w_0 - w_1)^2 > (w_0 + w_1)^2 + 4\lambda^2 \quad (6.15)$$

$$4\lambda(w_0 - w_1) - 4w_0 w_1 > 0 \quad (6.16)$$

$$\lambda > \lambda_c = \Delta E \frac{w_0 w_1}{w_0 - w_1}. \quad (6.17)$$

From the denominator, we see that if the weights are equal, $w_0 = w_1$, the critical $\lambda_c \rightarrow \infty$; in this case, no choice of λ will minimize O at the target ensemble. In the other limit $w_1 \rightarrow 0$, $\lambda_c \rightarrow 0$. We expect the optimal choice of weights and λ to be intermediate for efficient optimization.




Figure 6.2: Eigenvalues of the Hessian $\partial^2 O / \partial c_{10} \partial c_{01}$ for the two-state system as a function of λ . The critical value λ_c is the point at which the lower eigenvalue crosses zero from negative to positive, turning the eigenstates from a saddle point to a minimum.

We verify the condition of Eq. 6.17 by optimizing the cost functional in a two-state Hilbert space using different values of λ and weights (w_0, w_1) . The results are shown in Fig. 6.1. The penalty $\lambda/\Delta E$ is plotted against the first weight w_0 , and the weights are normalized, $w_0 + w_1 = 1$. The blue line shows the critical condition $\lambda_c/\Delta E = w_0 w_1 / (w_1 - w_0)$. As expected, the overlap is zero when $\lambda > \lambda_c$, and is nonzero when $\lambda < \lambda_c$. As $\lambda \rightarrow 0$, the overlap goes to one, as both states optimize to the ground state.

λ_c for N -State System

For an ensemble of N states, the cost functional has the form given in Eq. (6.5). We set $c_{ii}^2 = 1 - \sum_j |c_{ij}|^2$. The gradients take the form

$$\frac{\partial O}{\partial c_{ij}} = 2w_i c_{ij} (E_j - E_i) + \sum_{k \neq i} 2\lambda S_{ik} \left(c_{kj} - \frac{c_{ij} c_{ki}}{c_{ii}} \right). \quad (6.18)$$

The derivative is zero when $\Psi_i = \Phi_i$, since $c_{ij} = \delta_{ij}$, and $S_{ij} = \delta_{ij}$.

We now check that the Hessian matrix evaluated at the target ensemble state has positive eigenvalues. The Hessian matrix is made up of the second derivatives $\partial^2 O / \partial c_{ij} \partial c_{kl}$, where $i \neq j$ and $k \neq l$. We start by

identifying the nonzero off-diagonal terms, where $(i, j) \neq (k, l)$,

$$H_O = 2\lambda \sum_{m \neq i} \left[\frac{\partial S_{im}}{\partial c_{kl}} \left(c_{mj} - \frac{c_{ij}c_{mi}}{c_{ii}} \right) \right. \quad (6.19)$$

$$\left. + S_{im} \frac{\partial}{\partial c_{kl}} \left(c_{mj} - \frac{c_{ij}c_{mi}}{c_{ii}} \right) \right] \quad (6.20)$$

$$= 2\lambda \sum_{m \neq i} \frac{\partial S_{im}}{\partial c_{kl}} \delta_{mj} \quad (6.21)$$

$$= 2\lambda \frac{\partial}{\partial c_{kl}} \sum_n c_{in} c_{jn} \quad (6.22)$$

$$= 2\lambda (\delta_{ik} c_{jl} + \delta_{jk} c_{il}), \quad (6.23)$$

where we are at the target ensemble, $c_{ij} = \delta_{ij}$. We see that the only nonzero off-diagonal term is $\partial^2 O / \partial c_{ij} \partial c_{ji}$, where $(k, l) = (j, i)$. These off-diagonal terms are the result of two-state mixing, the situation described in the previous section. Thus, the critical λ_c satisfies the condition

$$\lambda_c = \max_{i < j} \left[(E_j - E_i) \frac{w_i w_j}{w_i - w_j} \right], \quad (6.24)$$

and when $\lambda > \lambda_c$, the target ensemble is a minimum of O .

6.2.3 Choosing Weights

Equal Critical λ

In the sections above, we derived the critical value λ_c that ensures the cost functional is minimized by eigenstates given a set of weights $\{w_i\}$. As defined in Eq. 6.24, λ_c is the maximum of the critical values for each pair of states. One possible choice of weights is to fix all the pairwise critical values to be equal. This strategy has one free parameter, the critical value itself.

Consider the ratios $\tilde{w}_i = w_i/w_0$ and the energy differences $\tilde{E}_i = E_i - E_0$. The critical λ_c^{i0} for states i and 0 are expressed as

$$\frac{\lambda_c^{i0}}{w_0} = \tilde{E}_i \frac{\tilde{w}_i}{1 - \tilde{w}_i}. \quad (6.25)$$

If we set all the λ_c^{i0} equal to a chosen constant λ_c , the weight ratios are given by

$$\tilde{w}_i = \frac{1}{1 + \tilde{E}_i w_0 / \lambda_c}. \quad (6.26)$$

With the weights determined this way, the critical λ_c^{ij} for any pair (i, j) also have the same value λ_c ,

$$\frac{\lambda_c^{ij}}{w_0} = (\tilde{E}_i - \tilde{E}_j) \frac{\tilde{w}_i \tilde{w}_j}{\tilde{w}_j - \tilde{w}_i} \quad (6.27)$$

$$= (\tilde{E}_i - \tilde{E}_j) \frac{1}{1/\tilde{w}_i - 1/\tilde{w}_j} \quad (6.28)$$

$$= (\tilde{E}_i - \tilde{E}_j) \frac{1}{1 + \tilde{E}_i w_0 / \lambda_c - (1 + \tilde{E}_j w_0 / \lambda_c)} \quad (6.29)$$

$$= \frac{\lambda_c}{w_0}. \quad (6.30)$$

The value of w_0 is determined by normalizing the sum of the weights to one,

$$w_0 = \left(\sum_i \tilde{w}_i \right)^{-1}, \quad (6.31)$$

and can be solved for given a desired λ_c and an estimate of the energies E_i . In practice, the energies E_i must be estimated by an approximate method, since the optimized energies are not yet known. The weights affect the efficiency of the optimization, but not the optimal solution, so exact energies are not necessary. Note that it is not necessary for the weights to be normalized; however, only the ratio of the weights to λ impacts the cost function, so we normalize the weights for interpretability.

As an example, we show the weights generated by this method for a range of values of λ_c in Fig. 6.3. At the extreme of small λ_c , w_0 is close to one, and all other weights are near zero. This has the obvious problem that the excited states will not be well distinguished by the cost functional and will be dominated by stochastic noise from the ground state energy. At the other extreme of large λ_c , the weights tend towards being equal. As discussed above, equal weights cause the cost functional to have a degenerate minimum, where any rotation of the first N eigenstates has the same energy as the eigenstates themselves. This effect is expected once noise in the overlap term (magnified by the large λ) dominates the differences in the weighted energies, even when the weights are not exactly equal. In practice, we find that even deterministic optimization on simple models has difficulty reaching the minimum at these extremes of λ_c .

Simultaneous ‘Stack of States’

Another strategy for optimization is to simultaneously optimize all states using the cost functions of Refs. [53], [147], where each state Ψ_j is orthogonalized only to lower states Ψ_i , $i < j$. This strategy, used in Ref. [135], is equivalent to the limit of a certain choice of weights in Eq. 6.1, which we derive in this section.

Consider setting

$$\lambda_{ij} = \alpha_{ij} w_j, \quad (6.32)$$




Figure 6.3: Weights generated by the method in Sec. 6.2.3 at different values of λ_c . The energies in this example are $E = (0.0, 1.0, 1.5)$.

keeping in mind that the overlap sum in Eq. 6.1 is only over $i < j$, and where α_{ij} is a weighting variable chosen so that

$$\alpha_{ij} > (E_j - E_i)/(1 - w_j/w_i). \quad (6.33)$$

Then $\lambda_{ij} > \lambda_c$. The parameter derivative with respect to the parameter set p_j of wave function j is then

$$\nabla_j O = w_j \left(\nabla_j E[\Psi_j] + \sum_{i < j} 2\alpha_{ij} S_{ij} \nabla_j S_{ij} \right) \quad (6.34)$$

$$+ \sum_{k > j} 2w_k \alpha_{jk} S_{jk} \nabla_j S_{jk}. \quad (6.35)$$

Now take the limit as $w_j/w_i \rightarrow 0$ while keeping the sum of weights equal to 1. This satisfies the ensemble criterion $w_i > w_j$ for $E_i < E_j$ and $w_i > 0$ and in this limit, α_{ij} only needs to be greater than $E_j - E_i$. The gradients then become

$$\nabla_j O = w_j \left(\nabla_j E[\Psi_j] + \sum_{i < j} 2\alpha_{ij} S_{ij} \nabla_j S_{ij} \right), \quad (6.36)$$

which is equivalent to the algorithm recently proposed in Ref. [135]. In this version of the algorithm, the

parameter updates can be made completely independent of the weights by scaling them as

$$\Delta p_j = \frac{\tau}{w_j} F^{-1} \nabla_j O, \quad (6.37)$$

where F is a Fisher information matrix[13], [150] and τ is the descent step length that can be optimized using correlated sampling[129] or reduced on a schedule.

We have shown that a practical algorithm proposed in a heuristic manner actually rigorously converges to the eigenstates of the system. One should note that this choice of λ_{ij} hides the upper bound property of O (Eq. 6.3), making comparisons between different wave function ansatzes challenging. However, one can choose to evaluate any objective function on the optimized wave functions with finite weights and λ_{ij} satisfying our λ_c criterion to recover the upper bound property.

6.3 Application to *Ab Initio* System

To demonstrate our penalty-based ensemble optimization on a simple, well-characterized system, we apply it to a CO molecule at the experimentally-determined equilibrium bond length of $r = 2.13$ Bohr [151]. We compute the ground state and degenerate first two excited states using VMC with different values of the overlap penalty λ , verifying the vertical excitation energy of absorption against experiment [151].

We represent our *ab initio* wave functions using a multi-Slater-Jastrow ansatz,

$$\Psi(\mathbf{R}) = e^{J(\mathbf{R};\alpha)} \sum_k c_k D_k(\mathbf{R}; \beta), \quad (6.38)$$

where $\mathbf{R}$ represents all electron coordinates, J is a two-body Jastrow factor with parameters α , D_k are Slater determinants of one-body orbitals parameterized by β , and c_k are determinant expansion coefficients. The determinants and expansion coefficients were generated from a complete active space configuration interaction (CASSCI) calculation with six electrons and six orbitals, giving a basis of 400 determinants. The orbitals were computed using Hartree Fock (HF) with correlation-consistent effective core potentials and corresponding triple- ζ basis functions.[103] The HF and CASSCI calculations were carried out in PySCF, and the wave functions' Jastrow, determinant, and orbital parameters were optimized in PyQMC.

State 0 is initialized as the CASSCI ground state multiplied by the Jastrow from a prior ground state optimization of Jastrow parameters in a multi-Slater-Jastrow wave function. States 1 and 2 are initialized as equal-weight linear combinations of the CASSCI ground state and target excited state multiplied by the same Jastrow as state 0. We start states 1 and 2 at superpositions of CASSCI eigenstates to avoid saddle points and




Figure 6.4: Optimization of the objective functional Eq. 6.5 for the lowest three states of a CO molecule with (a) $\lambda = 0$ Ha or (b) $\lambda = 0.45$ Ha.

ensure the method converges to the correct excited states, even starting far from them. The weights $\{w_i\}$ are set to $(0.41, 0.29, 0.29)$ using Eq. 6.26 with $\lambda_c = 0.27$ Ha, the energy gap from the CASCI calculations. All Jastrow parameters α , orbital parameters β , and determinant coefficients c_k are optimized with 500 walkers. Vertical excitation energies are evaluated from the optimized wave functions using VMC with 10,000 walkers.

Ensemble optimization using our cost functional generates the expected solutions in the limiting values of λ . When $\lambda = 0$ Ha, all three states optimize to the ground state, shown in the upper plot of Fig. 6.4. All three energies decrease over the optimization and approach the same value. Our estimate for the critical penalty based on CASCI energies is $\lambda_c = 0.27$. When $\lambda = 0.45$ Ha, substantially larger than λ_c , state 0 optimizes to the ground state and states 1 and 2 optimize to the lowest two excited states, shown in the right plot of Fig. 6.4. The energies of states 1 and 2, known to be degenerate, approach the same value of $0.220(5)$ Ha above the state 0 energy, even though they were initialized to a state much lower in energy. The results confirm that the optimization works as expected on this system. The VMC-calculated vertical excitation energy is within error bars of the experimentally determined $X^1\Sigma^+ \rightarrow a^3\Pi_r$ value of 0.22 Ha [151].

We also confirm that our ensemble optimization yields the expected results for small and large values




Figure 6.5: Overlap between three wave functions of the CO molecule as a function of the overlap penalty λ . As the penalty approaches the critical value λ_c , the overlaps go to zero. An overlap larger than the critical value is necessary to stabilize the stochastic optimization. A lighted dashed line is shown at zero overlap for reference.

of λ . As λ increases from zero, the optimized states start out perfectly overlapping, and end up mutually orthogonal when λ/λ_c is large, shown in Fig. 6.5. We find that orthogonality is only consistently achieved when $\lambda/\lambda_c > 1.5$, although our derived condition only requires $\lambda/\lambda_c > 1$. We attribute the difference to difficulty optimizing in the region $1 < \lambda/\lambda_c < 1.5$. In this region, the overlaps are zero for some points and not others, a sign of inconsistent optimization. When $\lambda/\lambda_c = 1$, we showed that the Hessian of the cost functional has a zero eigenvalue, meaning the cost functional is flat along at least one direction at the minimum. For λ/λ_c just above 1, this direction has very low curvature, making the minimum difficult to distinguish within stochastic noise. For large enough λ/λ_c , the curvature of the cost functional along this overlap direction is higher, and the minimum is reliably found.

6.4 Conclusion

We have presented a cost functional for simultaneously optimizing an ensemble of states to the low-energy eigenstates of a Hamiltonian in VMC. Our approach is state-specific, allowing all parameters of each ensemble state to vary independently without any restrictions to enforce orthogonality. The cost functional uses a weighted sum of energy expectations, distinguishing it from other recent approaches to ensemble optimization. Whether the cost functional has the correct minimum (at the low energy eigenstates) depends on the choice of λ and the weights. In particular, we showed that there is a critical value of the penalty λ_c , which λ must

exceed; otherwise, the states that minimize the cost functional will not be orthogonal, and therefore not the correct eigenstates. Finally, we demonstrated the proposed cost functional on an *ab initio* VMC calculation of a carbon monoxide molecule. We verified the expected behavior when the penalty λ is below and above the estimated λ_c , and that the method reaches the correct energy eigenstates even when initialized far from them.

We showed that there is a family of numerical methods which can satisfy the constraints noted in Eq. 6.24 and Eq. 6.5, at least one of which has been proposed in the literature without proof in the context of neural network wave functions.[135] It is not clear which in this family are most optimal; an interesting future direction will be to determine the most convenient formulation of excited state optimization. For all these techniques, the total computational cost of the excited state optimization is approximately N times the cost of a single ground state calculation, where N is the number of excited states sought. We anticipate the cost functional presented here to be useful in future excited state quantum Monte Carlo calculations with a variety of wave functions, including traditional as well as neural network and other machine-learning ansatzes.

CHAPTER 7

CONCLUSION

At the first principles level, properties of materials can be substantially affected by interactions between electrons, necessitating many-body computational approaches for complete descriptions of a material's electronic structure. The work in this dissertation focused on extending the capabilities of many-body wave functions in multiple ways.

We analyzed the promise and the flaws of a new many-body formula: Chapter 3 introduced an operator for calculating bulk quadrupole moments of periodic many-body wave functions. The operator successfully reproduced results in non-interacting models with controllable finite size effects, although it was not realized in *ab initio* calculations. Through the many-body lens, we uncovered subtleties in the definition of the quantity due to its intrinsic connection to local dipole moments. Unlike charge, dipole moments are neither conserved nor quantized, making it impossible to disentangle the surface from the bulk. Moreover, the inevitable fluctuations in dipole moments compete with the existence of a higher-order moment. Well-defined bulk quadrupole moments are highly constrained, if they can exist at all, but approximate quadrupole-like quantities may still be found to characterize materials with certain symmetries that exhibit low dipole fluctuations.

We developed software for easier experimentation with many-body wave functions: Chapter 4 introduced PyQMC, implementing quantum Monte Carlo entirely in Python. The use of Python makes the code exceptionally modular and adaptable; sections have already been borrowed by multiple external research projects. Importantly, it greatly simplifies the development and execution of new methods, enabling two other projects in this dissertation as well as several other works from our group, both published[52]–[54], [105] and in progress. Taking advantage of the available optimized Python libraries and the features of PySCF, PyQMC doesn't sacrifice much in performance for its increased usability. As a lightweight, non-compiled software, PyQMC also serves as a minimal-barrier entry point for new users to QMC calculations and will continue to accelerate future developments in the field.

We demonstrated wave function improvements in a new context and assessed new optimization capabilities:

Chapter 5 brought several wave function ideas together in the strongly correlated cuprate CaCuO_2 , leveraging new computational and software capability to capture spin fluctuations in QMC calculations. The computationally demanding optimization of orbital parameters is deceptively subtle; it did not improve our best upper bound of the DMC energy. Despite this, we were successful in optimizing orbitals of multi-determinant wave functions to achieve lower DMC energies. Moreover, we confirmed the presence of spin fluctuations in multi-determinant wave functions in a periodic, strongly correlated system. This successful test on undoped CaCuO_2 establishes the groundwork for a new study on doped cuprates in the superconducting state, where spin fluctuations are not well understood and require accurate characterization.

We proved a new result for assessing quality of many-body excited states: Chapter 6 introduces a variational principle for an ensemble of many-body states. A weighted sum of energies plus an overlap penalty, with certain conditions on the coefficients, provides a single objective function for optimizing the ensemble to the first N eigenstates of a many-body Hamiltonian. We proved that if the overlap penalty exceeds a critical threshold (a function of the weights and exact eigenvalues), the objective function is guaranteed to be minimized at the desired set of eigenstates. Although the threshold requires knowledge of the exact eigenvalues, approximate values are sufficient, since any value above the threshold guarantees the correct minimum. We demonstrated that this objective function reduces to a previously proposed algorithm in a particular limit – this limit obscures the variational principle, but is efficient for optimization. In addition to rigorously validating the existing optimization approach, the objective function we introduced can be used along with any optimization routine to compare ensembles via the variational upper bound. These results improve the search for accurate many-body excited states in strongly correlated systems.

In the problem of identifying the mechanism for superconductivity in the cuprates, the latter three projects provide a foundation for the next steps. The multi-determinant trial functions verified on the undoped CaCuO_2 system can be applied to the more complicated doped system with a clear interpretation. The verification step and the subsequent steps of investigation are accelerated by PyQMC, enabling efficient porting of PySCF determinant expansions and active space analysis, as well as efficient workflow design. After characterizing the ground state of doped CaCuO_2 , my ensemble VMC work provides valuable theoretical foundation and analysis tools for calculation accurate excited states, which are necessary for generating effective models. This progression will be equally effective for studying many other systems and the methods established in my work will support future efforts to understand and model strongly correlated materials.

REFERENCES

- [1] A. Szabo and N. S. Ostlund, *Modern quantum chemistry: introduction to advanced electronic structure theory*. McGraw-Hill, 1989.
- [2] K. Ruedenberg, M. W. Schmidt, M. M. Gilbert, and S. Elbert, “Are atoms intrinsic to molecular electronic wavefunctions? i. the fors model,” *Chemical Physics*, vol. 71, no. 1, pp. 41–49, 1982, ISSN: 0301-0104. DOI: [https://doi.org/10.1016/0301-0104\(82\)87004-3](https://doi.org/10.1016/0301-0104(82)87004-3).
- [3] K. Ruedenberg, M. W. Schmidt, and M. M. Gilbert, “Are atoms sic to molecular electronic wavefunctions? ii. analysis of fors orbitals,” *Chemical Physics*, vol. 71, no. 1, pp. 51–64, 1982, ISSN: 0301-0104. DOI: [https://doi.org/10.1016/0301-0104\(82\)87005-5](https://doi.org/10.1016/0301-0104(82)87005-5).
- [4] K. Ruedenberg, M. W. Schmidt, M. M. Gilbert, and S. Elbert, “Are atoms intrinsic to molecular electronic wavefunctions? iii. analysis of fors configurations,” *Chemical Physics*, vol. 71, no. 1, pp. 65–78, 1982, ISSN: 0301-0104. DOI: [https://doi.org/10.1016/0301-0104\(82\)87006-7](https://doi.org/10.1016/0301-0104(82)87006-7).
- [5] P. Hohenberg and W. Kohn, “Inhomogeneous electron gas,” *Phys. Rev.*, vol. 136, B864–B871, 3B Nov. 1964. DOI: [10.1103/PhysRev.136.B864](https://doi.org/10.1103/PhysRev.136.B864).
- [6] J. P. Perdew, P. Ziesche, and H. Eschrig, *Electronic structure of solids’ 91*, 1991.
- [7] J. P. Perdew, K. Burke, and M. Ernzerhof, “Generalized gradient approximation made simple,” *Phys. Rev. Lett.*, vol. 77, pp. 3865–3868, 18 Oct. 1996. DOI: [10.1103/PhysRevLett.77.3865](https://doi.org/10.1103/PhysRevLett.77.3865).
- [8] D. M. Ceperley and B. J. Alder, “Ground state of the electron gas by a stochastic method,” *Phys. Rev. Lett.*, vol. 45, pp. 566–569, 7 Aug. 1980. DOI: [10.1103/PhysRevLett.45.566](https://doi.org/10.1103/PhysRevLett.45.566).
- [9] R. Jastrow, “Many-body problem with strong forces,” *Phys. Rev.*, vol. 98, pp. 1479–1484, 5 Jun. 1955. DOI: [10.1103/PhysRev.98.1479](https://doi.org/10.1103/PhysRev.98.1479).
- [10] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, “Equation of State Calculations by Fast Computing Machines,” *The Journal of Chemical Physics*, vol. 21, no. 6, pp. 1087–1092, Dec. 2004, ISSN: 0021-9606. DOI: [10.1063/1.1699114](https://doi.org/10.1063/1.1699114).
- [11] R. M. Martin, L. Reining, and D. M. Ceperley, “Interacting Electrons,” en, Jun. 2016.
- [12] P. J. Rossky, J. D. Doll, and H. L. Friedman, “Brownian dynamics as smart Monte Carlo simulation,” *The Journal of Chemical Physics*, vol. 69, no. 10, pp. 4628–4633, Aug. 2008, ISSN: 0021-9606. DOI: [10.1063/1.436415](https://doi.org/10.1063/1.436415).
- [13] S. Sorella, “Green function monte carlo with stochastic reconfiguration,” *Phys. Rev. Lett.*, vol. 80, pp. 4558–4561, 20 May 1998. DOI: [10.1103/PhysRevLett.80.4558](https://doi.org/10.1103/PhysRevLett.80.4558).
- [14] J. Toulouse and C. J. Umrigar, “Optimization of quantum Monte Carlo wave functions by energy minimization,” *The Journal of Chemical Physics*, vol. 126, no. 8, p. 084 102, Feb. 2007, ISSN: 0021-9606. DOI: [10.1063/1.2437215](https://doi.org/10.1063/1.2437215).
- [15] E. Neuscamman, C. J. Umrigar, and G. K.-L. Chan, “Optimizing large parameter sets in variational quantum monte carlo,” *Phys. Rev. B*, vol. 85, p. 045 103, 4 Jan. 2012. DOI: [10.1103/PhysRevB.85.045103](https://doi.org/10.1103/PhysRevB.85.045103).

- [16] D. Ceperley, “The simulation of quantum systems with random walks: A new algorithm for charged systems,” *Journal of Computational Physics*, vol. 51, no. 3, pp. 404–422, 1983, ISSN: 0021-9991. DOI: [https://doi.org/10.1016/0021-9991\(83\)90161-4](https://doi.org/10.1016/0021-9991(83)90161-4).
- [17] H. F. Trotter, “On the product of semi-groups of operators,” *Proceedings of the American Mathematical Society*, vol. 10, no. 4, pp. 545–551, 1959.
- [18] R. Assaraf, M. Caffarel, and A. Khelif, “Diffusion monte carlo methods with a fixed number of walkers,” *Phys. Rev. E*, vol. 61, pp. 4566–4575, 4 Apr. 2000. DOI: [10.1103/PhysRevE.61.4566](https://doi.org/10.1103/PhysRevE.61.4566).
- [19] M. Calandra Buonaurea and S. Sorella, “Numerical study of the two-dimensional heisenberg model using a green function monte carlo technique with a fixed number of walkers,” *Phys. Rev. B*, vol. 57, pp. 11 446–11 456, 18 May 1998. DOI: [10.1103/PhysRevB.57.11446](https://doi.org/10.1103/PhysRevB.57.11446).
- [20] D. H. Davis, “Critical-size calculations for neutron systems by the monte carlo method,” *Lawrence Radiation Lab. Report No. UCRL-6707*, Dec. 1961.
- [21] D. M. Ceperley, “The statistical error of green’s function Monte Carlo,” en, *Journal of Statistical Physics*, vol. 43, no. 5, pp. 815–826, Jun. 1986, ISSN: 1572-9613. DOI: [10.1007/BF02628307](https://doi.org/10.1007/BF02628307).
- [22] P. A. Christiansen, Y. S. Lee, and K. S. Pitzer, “Improved ab initio effective core potentials for molecular calculations,” *The Journal of Chemical Physics*, vol. 71, no. 11, pp. 4445–4450, Jul. 2008, ISSN: 0021-9606. DOI: [10.1063/1.438197](https://doi.org/10.1063/1.438197).
- [23] R. D. King-Smith and D. Vanderbilt, “Theory of polarization of crystalline solids,” *Phys. Rev. B*, vol. 47, pp. 1651–1654, 3 Jan. 1993. DOI: [10.1103/PhysRevB.47.1651](https://doi.org/10.1103/PhysRevB.47.1651).
- [24] D. Vanderbilt and R. King-Smith, “Electric polarization as a bulk quantity and its relation to surface charge,” *Phys. Rev. B*, vol. 48, no. 7, p. 4442, 1993.
- [25] R. Resta, “Quantum-mechanical position operator in extended systems,” *Phys. Rev. Lett.*, vol. 80, pp. 1800–1803, 9 Mar. 1998. DOI: [10.1103/PhysRevLett.80.1800](https://doi.org/10.1103/PhysRevLett.80.1800).
- [26] W. A. Benalcazar, B. A. Bernevig, and T. L. Hughes, “Quantized electric multipole insulators,” *Science*, vol. 357, no. 6346, pp. 61–66, 2017, ISSN: 0036-8075. DOI: [10.1126/science.aah6442](https://doi.org/10.1126/science.aah6442).
- [27] W. A. Benalcazar, B. A. Bernevig, and T. L. Hughes, “Electric multipole moments, topological multipole moment pumping, and chiral hinge states in crystalline insulators,” *Phys. Rev. B*, vol. 96, no. 24, p. 245 115, 2017.
- [28] D. J. Thouless, “Quantization of particle transport,” *Phys. Rev. B*, vol. 27, pp. 6083–6087, 10 May 1983. DOI: [10.1103/PhysRevB.27.6083](https://doi.org/10.1103/PhysRevB.27.6083).
- [29] Q. Niu and D. J. Thouless, “Quantised adiabatic charge transport in the presence of substrate disorder and many-body interaction,” *Journal of Physics A: Mathematical and General*, vol. 17, no. 12, p. 2453, Aug. 1984. DOI: [10.1088/0305-4470/17/12/016](https://doi.org/10.1088/0305-4470/17/12/016).
- [30] I. Souza, T. Wilkens, and R. M. Martin, “Polarization and localization in insulators: Generating function approach,” *Phys. Rev. B*, vol. 62, no. 3, p. 1666, 2000.
- [31] G. Ortiz and R. M. Martin, “Macroscopic polarization as a geometric quantum phase: Many-body formulation,” *Phys. Rev. B*, vol. 49, no. 20, p. 14 202, 1994.
- [32] R. Resta and S. Sorella, “Electron localization in the insulating state,” *Phys. Rev. Lett.*, vol. 82, pp. 370–373, 2 Jan. 1999. DOI: [10.1103/PhysRevLett.82.370](https://doi.org/10.1103/PhysRevLett.82.370).
- [33] B. A. Bernevig, T. L. Hughes, and S.-C. Zhang, “Quantum spin hall effect and topological phase transition in hgte quantum wells,” *Science*, vol. 314, no. 5806, pp. 1757–1761, 2006. DOI: [10.1126/science.1133734](https://doi.org/10.1126/science.1133734).
- [34] T. Senthil, “Symmetry-protected topological phases of quantum matter,” *Annual Review of Condensed Matter Physics*, vol. 6, no. 1, pp. 299–324, 2015. DOI: [10.1146/annurev-conmatphys-031214-014740](https://doi.org/10.1146/annurev-conmatphys-031214-014740).
- [35] J. C. Teo and T. L. Hughes, “Topological defects in symmetry-protected topological phases,” *Annual Review of Condensed Matter Physics*, vol. 8, no. 1, pp. 211–237, 2017. DOI: [10.1146/annurev-conmatphys-031016-025154](https://doi.org/10.1146/annurev-conmatphys-031016-025154).

- [36] X.-G. Wen, “Colloquium: Zoo of quantum-topological phases of matter,” *Reviews of Modern Physics*, vol. 89, no. 4, p. 041004, Dec. 2017. DOI: [10.1103/RevModPhys.89.041004](https://doi.org/10.1103/RevModPhys.89.041004).
- [37] F. Schindler, A. M. Cook, M. G. Vergniory, *et al.*, “Higher-order topological insulators,” *Science Advances*, vol. 4, no. 6, eaat0346, Jun. 2018. DOI: [10.1126/sciadv.aat0346](https://doi.org/10.1126/sciadv.aat0346).
- [38] B. A. Bernevig, C. Felser, and H. Beidenkopf, “Progress and prospects in magnetic topological materials,” en, *Nature*, vol. 603, no. 7899, pp. 41–51, Mar. 2022, ISSN: 1476-4687. DOI: [10.1038/s41586-021-04105-x](https://doi.org/10.1038/s41586-021-04105-x).
- [39] W. A. Wheeler, L. K. Wagner, and T. L. Hughes, “Many-body electric multipole operators in extended systems,” *Phys. Rev. B*, vol. 100, p. 245135, 24 Dec. 2019. DOI: [10.1103/PhysRevB.100.245135](https://doi.org/10.1103/PhysRevB.100.245135).
- [40] R. Resta, “Macroscopic polarization in crystalline dielectrics: The geometric phase approach,” *Rev. Mod. Phys.*, vol. 66, pp. 899–915, 3 Jul. 1994. DOI: [10.1103/RevModPhys.66.899](https://doi.org/10.1103/RevModPhys.66.899).
- [41] R. Resta and D. Vanderbilt, “Theory of polarization: A modern approach,” in *Physics of Ferroelectrics: A Modern Perspective*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 31–68, ISBN: 978-3-540-34591-6. DOI: [10.1007/978-3-540-34591-6_2](https://doi.org/10.1007/978-3-540-34591-6_2).
- [42] D. Vanderbilt, *Berry Phases in Electronic Structure Theory: Electric Polarization, Orbital Magnetization and Topological Insulators*. Cambridge University Press, 2018.
- [43] J. Zak, “Berry’s phase for energy bands in solids,” *Phys. Rev. Lett.*, vol. 62, pp. 2747–2750, 23 Jun. 1989. DOI: [10.1103/PhysRevLett.62.2747](https://doi.org/10.1103/PhysRevLett.62.2747).
- [44] X.-L. Qi, T. L. Hughes, and S.-C. Zhang, “Topological field theory of time-reversal invariant insulators,” *Phys. Rev. B*, vol. 78, no. 19, p. 195424, 2008.
- [45] T. L. Hughes, E. Prodan, and B. A. Bernevig, “Inversion-symmetric topological insulators,” *Phys. Rev. B*, vol. 83, no. 24, p. 245132, 2011.
- [46] A. M. Turner, Y. Zhang, R. S. Mong, and A. Vishwanath, “Quantized response and topology of magnetic insulators with inversion symmetry,” *Phys. Rev. B*, vol. 85, no. 16, p. 165120, 2012.
- [47] A. Aligia and G. Ortiz, “Quantum mechanical position operator and localization in extended systems,” *Physical Rev. Lett.*, vol. 82, no. 12, p. 2560, 1999.
- [48] F. Schindler, A. M. Cook, M. G. Vergniory, *et al.*, “Higher-order topological insulators,” *Science Advances*, vol. 4, no. 6, eaat0346, 2018.
- [49] B. Kang, K. Shiozaki, and G. Y. Cho, “Many-body order parameters for multipoles in solids,” *Phys. Rev. B*, vol. 100, p. 245134, 24 Dec. 2019. DOI: [10.1103/PhysRevB.100.245134](https://doi.org/10.1103/PhysRevB.100.245134).
- [50] S. Ono, L. Trifunovic, and H. Watanabe, “Difficulties in operator-based formulation of the bulk quadrupole moment,” *Phys. Rev. B*, vol. 100, p. 245133, 24 Dec. 2019. DOI: [10.1103/PhysRevB.100.245133](https://doi.org/10.1103/PhysRevB.100.245133).
- [51] W. A. Wheeler, S. Pathak, K. G. Kleiner, *et al.*, “PyQMC: An all-Python real-space quantum Monte Carlo module in PySCF,” *The Journal of Chemical Physics*, vol. 158, no. 11, p. 114801, Mar. 2023, ISSN: 0021-9606. DOI: [10.1063/5.0139024](https://doi.org/10.1063/5.0139024).
- [52] S. Pathak and L. K. Wagner, “A light weight regularization for wave function parameter gradients in quantum monte carlo,” *AIP Advances*, vol. 10, no. 8, p. 085213, 2020. DOI: [10.1063/5.0004008](https://doi.org/10.1063/5.0004008).
- [53] S. Pathak, B. Busemeyer, J. N. B. Rodrigues, and L. K. Wagner, “Excited states in variational Monte Carlo using a penalty method,” en, *The Journal of Chemical Physics*, vol. 154, no. 3, p. 034101, Jan. 2021, ISSN: 0021-9606. DOI: [10.1063/5.0030949](https://doi.org/10.1063/5.0030949).
- [54] S. Yuan, Y. Chang, and L. K. Wagner, “Quantification of electron correlation for approximate quantum calculations,” 2022. DOI: [10.48550/ARXIV.2208.03189](https://doi.org/10.48550/ARXIV.2208.03189).
- [55] X. Li, Z. Li, and J. Chen, “Ab initio calculation of real solids via neural network ansatz,” *arXiv:2203.15472 [cond-mat, physics:physics]*, Mar. 2022.
- [56] C. Filippi, M. Zaccheddu, and F. Buda, “Absorption spectrum of the green fluorescent protein chromophore: A difficult case for ab initio methods?” *Journal of Chemical Theory and Computation*, vol. 5, no. 8, pp. 2074–2087, 2009, PMID: 26613149. DOI: [10.1021/ct900227j](https://doi.org/10.1021/ct900227j).

- [57] L. K. Wagner, “Ground state of doped cuprates from first-principles quantum Monte Carlo calculations,” *Physical Review B*, vol. 92, no. 16, p. 161116, Oct. 2015. DOI: [10.1103/PhysRevB.92.161116](https://doi.org/10.1103/PhysRevB.92.161116).
- [58] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, “Quantum Monte Carlo simulations of solids,” *Reviews of Modern Physics*, vol. 73, no. 1, pp. 33–83, Jan. 2001. DOI: [10.1103/RevModPhys.73.33](https://doi.org/10.1103/RevModPhys.73.33).
- [59] L. K. Wagner and D. M. Ceperley, “Discovering correlated fermions using quantum Monte Carlo,” en, *Reports on Progress in Physics*, vol. 79, no. 9, p. 094501, Aug. 2016, ISSN: 0034-4885. DOI: [10.1088/0034-4885/79/9/094501](https://doi.org/10.1088/0034-4885/79/9/094501).
- [60] R. J. Needs, M. D. Towler, N. D. Drummond, P. López Ríos, and J. R. Trail, “Variational and diffusion quantum Monte Carlo calculations with the CASINO code,” *The Journal of Chemical Physics*, vol. 152, no. 15, p. 154106, Apr. 2020, ISSN: 0021-9606. DOI: [10.1063/1.5144288](https://doi.org/10.1063/1.5144288).
- [61] S. Pilati, E. M. Inack, and P. Pieri, “Self-learning projective quantum monte carlo simulations guided by restricted boltzmann machines,” *Phys. Rev. E*, vol. 100, p. 043301, 4 Oct. 2019. DOI: [10.1103/PhysRevE.100.043301](https://doi.org/10.1103/PhysRevE.100.043301).
- [62] D. Pfau, J. S. Spencer, A. G. D. G. Matthews, and W. M. C. Foulkes, “Ab initio solution of the many-electron schrödinger equation with deep neural networks,” *Phys. Rev. Research*, vol. 2, p. 033429, 3 Sep. 2020. DOI: [10.1103/PhysRevResearch.2.033429](https://doi.org/10.1103/PhysRevResearch.2.033429).
- [63] A. Acevedo, M. Curry, S. H. Joshi, B. Leroux, and N. Malaya, “Vandermonde wave function ansatz for improved variational monte carlo,” in *2020 IEEE/ACM Fourth Workshop on Deep Learning on Supercomputers (DLS)*, 2020, pp. 40–47. DOI: [10.1109/DLS51937.2020.00010](https://doi.org/10.1109/DLS51937.2020.00010).
- [64] J. Hermann, Z. Schätzle, and F. Noé, “Deep-neural-network solution of the electronic Schrödinger equation,” en, *Nature Chemistry*, vol. 12, no. 10, pp. 891–897, Oct. 2020, ISSN: 1755-4349. DOI: [10.1038/s41557-020-0544-y](https://doi.org/10.1038/s41557-020-0544-y).
- [65] M. Wilson, N. Gao, F. Wudarski, E. Rieffel, and N. M. Tubman, “Simulations of state-of-the-art fermionic neural network wave functions with diffusion Monte Carlo,” *arXiv:2103.12570 [physics, physics:quant-ph]*, Mar. 2021.
- [66] J. A. R. Shea and E. Neuscamman, “Size Consistent Excited States via Algorithmic Transformations between Variational Principles,” *Journal of Chemical Theory and Computation*, vol. 13, no. 12, pp. 6078–6088, Dec. 2017, ISSN: 1549-9618. DOI: [10.1021/acs.jctc.7b00923](https://doi.org/10.1021/acs.jctc.7b00923).
- [67] M. Dash, J. Feldt, S. Moroni, A. Scemama, and C. Filippi, “Excited States with Selected Configuration Interaction-Quantum Monte Carlo: Chemically Accurate Excitation Energies and Geometries,” *Journal of Chemical Theory and Computation*, vol. 15, no. 9, pp. 4896–4906, Sep. 2019, ISSN: 1549-9618. DOI: [10.1021/acs.jctc.9b00476](https://doi.org/10.1021/acs.jctc.9b00476).
- [68] L. Otis, I. M. Craig, and E. Neuscamman, “A hybrid approach to excited-state-specific variational Monte Carlo and doubly excited states,” *The Journal of Chemical Physics*, vol. 153, no. 23, p. 234105, Dec. 2020, ISSN: 0021-9606. DOI: [10.1063/5.0024572](https://doi.org/10.1063/5.0024572).
- [69] L. N. Tran and E. Neuscamman, “Improving Excited-State Potential Energy Surfaces via Optimal Orbital Shapes,” *The Journal of Physical Chemistry A*, vol. 124, no. 40, pp. 8273–8279, Oct. 2020, ISSN: 1089-5639. DOI: [10.1021/acs.jpca.0c07593](https://doi.org/10.1021/acs.jpca.0c07593).
- [70] J. Feldt and C. Filippi, “Excited-state calculations with quantum Monte Carlo,” en, Feb. 2020.
- [71] M. Dash, S. Moroni, C. Filippi, and A. Scemama, “Tailoring CIPSI Expansions for QMC Calculations of Electronic Excitations: The Case Study of Thiophene,” *Journal of Chemical Theory and Computation*, vol. 17, no. 6, pp. 3426–3434, Jun. 2021, ISSN: 1549-9618. DOI: [10.1021/acs.jctc.1c00212](https://doi.org/10.1021/acs.jctc.1c00212).
- [72] J. T. Krogel, M. Yu, J. Kim, and D. M. Ceperley, “Quantum energy density: Improved efficiency for quantum Monte Carlo calculations,” *Physical Review B*, vol. 88, no. 3, p. 035137, Jul. 2013. DOI: [10.1103/PhysRevB.88.035137](https://doi.org/10.1103/PhysRevB.88.035137).
- [73] K. Ryczko, J. T. Krogel, and I. Tamblyn, “Machine Learning Diffusion Monte Carlo Energy Densities,” *arXiv:2205.04547 [cond-mat]*, May 2022.

- [74] L. K. Wagner, “Types of single particle symmetry breaking in transition metal oxides due to electron correlation,” *The Journal of Chemical Physics*, vol. 138, no. 9, p. 094106, 2013. DOI: [10.1063/1.4793531](https://doi.org/10.1063/1.4793531).
- [75] H. J. Changlani, H. Zheng, and L. K. Wagner, “Density-matrix based determination of low-energy model Hamiltonians from ab initio wavefunctions,” *The Journal of Chemical Physics*, vol. 143, no. 10, p. 102814, 2015. DOI: [10.1063/1.4927664](https://doi.org/10.1063/1.4927664).
- [76] H. Zheng, H. J. Changlani, K. T. Williams, B. Busemeyer, and L. K. Wagner, “From Real Materials to Model Hamiltonians With Density Matrix Downfolding,” *Frontiers in Physics*, vol. 6, p. 43, 2018, ISSN: 2296-424X. DOI: [10.3389/fphy.2018.00043](https://doi.org/10.3389/fphy.2018.00043).
- [77] Y. Chang and L. K. Wagner, “Effective spin-orbit models using correlated first-principles wave functions,” *Physical Review Research*, vol. 2, no. 1, p. 013195, Feb. 2020, ISSN: 2643-1564. DOI: [10.1103/PhysRevResearch.2.013195](https://doi.org/10.1103/PhysRevResearch.2.013195).
- [78] A. Zen, S. Sorella, M. J. Gillan, A. Michaelides, and D. Alfè, “Boosting the accuracy and speed of quantum monte carlo: Size consistency and time step,” *Phys. Rev. B*, vol. 93, p. 241118, 24 Jun. 2016. DOI: [10.1103/PhysRevB.93.241118](https://doi.org/10.1103/PhysRevB.93.241118).
- [79] T. A. Anderson and C. J. Umrigar, “Nonlocal pseudopotentials and time-step errors in diffusion Monte Carlo,” *The Journal of Chemical Physics*, vol. 154, no. 21, p. 214110, Jun. 2021, ISSN: 0021-9606. DOI: [10.1063/5.0052838](https://doi.org/10.1063/5.0052838).
- [80] C. J. Umrigar, K. G. Wilson, and J. W. Wilkins, “Optimized trial wave functions for quantum monte carlo calculations,” *Phys. Rev. Lett.*, vol. 60, pp. 1719–1722, 17 Apr. 1988. DOI: [10.1103/PhysRevLett.60.1719](https://doi.org/10.1103/PhysRevLett.60.1719).
- [81] S. Sorella, M. Casula, and D. Rocca, “Weak binding between two aromatic rings: Feeling the van der waals attraction by quantum monte carlo methods,” *The Journal of Chemical Physics*, vol. 127, no. 1, p. 014105, 2007. DOI: [10.1063/1.2746035](https://doi.org/10.1063/1.2746035).
- [82] L. K. Wagner, M. Bajdich, and L. Mitas, “QWalk: A quantum Monte Carlo program for electronic structure,” in *Journal of Computational Physics*, vol. 228, no. 9, pp. 3390–3404, May 2009, ISSN: 0021-9991. DOI: [10.1016/j.jcp.2009.01.017](https://doi.org/10.1016/j.jcp.2009.01.017).
- [83] L. Mitáš, E. L. Shirley, and D. M. Ceperley, “Nonlocal pseudopotentials and diffusion monte carlo,” *The Journal of Chemical Physics*, vol. 95, no. 5, pp. 3467–3475, 1991. DOI: [10.1063/1.460849](https://doi.org/10.1063/1.460849).
- [84] W. L. McMillan, “Ground state of liquid He^4 ,” *Phys. Rev.*, vol. 138, A442–A451, 2A Apr. 1965. DOI: [10.1103/PhysRev.138.A442](https://doi.org/10.1103/PhysRev.138.A442).
- [85] C. Lin, F. H. Zong, and D. M. Ceperley, “Twist-averaged boundary conditions in continuum quantum monte carlo algorithms,” *Phys. Rev. E*, vol. 64, p. 016702, 1 Jun. 2001. DOI: [10.1103/PhysRevE.64.016702](https://doi.org/10.1103/PhysRevE.64.016702).
- [86] S. Chiesa, D. Ceperley, R. Martin, and M. Holzmann, “Finite-Size Error in Many-Body Simulations with Long-Range Interactions,” *Physical Review Letters*, vol. 97, no. 7, Aug. 2006, ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.97.076404](https://doi.org/10.1103/PhysRevLett.97.076404).
- [87] M. Casula and S. Sorella, “Geminal wave functions with jastrow correlation: A first application to atoms,” *The Journal of Chemical Physics*, vol. 119, no. 13, pp. 6500–6511, 2003. DOI: [10.1063/1.1604379](https://doi.org/10.1063/1.1604379).
- [88] A. J. Williamson, R. Q. Hood, R. J. Needs, and G. Rajagopal, “Diffusion quantum Monte Carlo calculations of the excited states of silicon,” *Physical Review B*, vol. 57, no. 19, pp. 12140–12144, May 1998. DOI: [10.1103/PhysRevB.57.12140](https://doi.org/10.1103/PhysRevB.57.12140).
- [89] K. K. Docken and J. Hinze, “Lih potential curves and wavefunctions for $x\ 1\Sigma^+$, $a\ 1\Sigma^+$, $b\ 1\Pi$, $3\Sigma^+$, and 3Π ,” *The Journal of Chemical Physics*, vol. 57, no. 11, pp. 4928–4936, 1972. DOI: [10.1063/1.1678164](https://doi.org/10.1063/1.1678164).
- [90] H.-J. Werner and P. J. Knowles, “A second order multiconfiguration scf procedure with optimum convergence,” *The Journal of Chemical Physics*, vol. 82, no. 11, pp. 5053–5063, 1985. DOI: [10.1063/1.448627](https://doi.org/10.1063/1.448627).

- [91] F. Schautz and C. Filippi, “Optimized jastrow-slater wave functions for ground and excited states: Application to the lowest states of ethene,” *The Journal of Chemical Physics*, vol. 120, no. 23, pp. 10931–10941, 2004. DOI: [10.1063/1.1752881](https://doi.org/10.1063/1.1752881).
- [92] A. Cuzzocrea, A. Scemama, W. J. Briels, S. Moroni, and C. Filippi, “Variational Principles in Quantum Monte Carlo: The Troubled Story of Variance Minimization,” *Journal of Chemical Theory and Computation*, vol. 16, no. 7, pp. 4203–4212, Jul. 2020, ISSN: 1549-9618. DOI: [10.1021/acs.jctc.0c00147](https://doi.org/10.1021/acs.jctc.0c00147).
- [93] I. Tobias, R. J. Fallon, and J. T. Vanderslice, “Potential energy curves for co,” *The Journal of Chemical Physics*, vol. 33, no. 6, pp. 1638–1640, 1960. DOI: [10.1063/1.1731475](https://doi.org/10.1063/1.1731475).
- [94] G. Ortiz, D. M. Ceperley, and R. M. Martin, “New stochastic method for systems with broken time-reversal symmetry: 2d fermions in a magnetic field,” *Phys. Rev. Lett.*, vol. 71, pp. 2777–2780, 17 Oct. 1993. DOI: [10.1103/PhysRevLett.71.2777](https://doi.org/10.1103/PhysRevLett.71.2777).
- [95] M. Casula, “Beyond the locality approximation in the standard diffusion monte carlo method,” *Phys. Rev. B*, vol. 74, p. 161102, 16 Oct. 2006. DOI: [10.1103/PhysRevB.74.161102](https://doi.org/10.1103/PhysRevB.74.161102).
- [96] L. Dalcín, R. Paz, and M. Storti, “MPI for Python,” en, *Journal of Parallel and Distributed Computing*, vol. 65, no. 9, pp. 1108–1115, Sep. 2005, ISSN: 0743-7315. DOI: [10.1016/j.jpdc.2005.03.010](https://doi.org/10.1016/j.jpdc.2005.03.010).
- [97] L. Dalcín, R. Paz, M. Storti, and J. D’Elía, “MPI for Python: Performance improvements and MPI-2 extensions,” en, *Journal of Parallel and Distributed Computing*, vol. 68, no. 5, pp. 655–662, May 2008, ISSN: 0743-7315. DOI: [10.1016/j.jpdc.2007.09.005](https://doi.org/10.1016/j.jpdc.2007.09.005).
- [98] L. D. Dalcin, R. R. Paz, P. A. Kler, and A. Cosimo, “Parallel distributed computing using Python,” en, *Advances in Water Resources*, New Computational Methods and Software Tools, vol. 34, no. 9, pp. 1124–1139, Sep. 2011, ISSN: 0309-1708. DOI: [10.1016/j.advwatres.2011.04.013](https://doi.org/10.1016/j.advwatres.2011.04.013).
- [99] L. Dalcin and Y.-L. L. Fang, “Mpi4py: Status Update After 12 Years of Development,” *Computing in Science & Engineering*, vol. 23, no. 4, pp. 47–54, Jul. 2021, ISSN: 1558-366X. DOI: [10.1109/MCSE.2021.3083216](https://doi.org/10.1109/MCSE.2021.3083216).
- [100] Message Passing Interface Forum, *MPI: A message-passing interface standard version 4.0*, Jun. 2021. [Online]. Available: <https://www.mpi-forum.org/docs/mpi-4.0/mpi40-report.pdf>.
- [101] Dask Development Team, *Dask: Library for dynamic task scheduling*, 2016.
- [102] R. Okuta, Y. Unno, D. Nishino, S. Hido, and C. Loomis, “CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations,” *31st Conference on Neural Information Processing Systems (NIPS 2017)*, Long Beach, CA, USA., p. 7, 2017.
- [103] M. C. Bennett, C. A. Melton, A. Annaberdiyev, G. Wang, L. Shulenburger, and L. Mitas, “A new generation of effective core potentials for correlated calculations,” *The Journal of Chemical Physics*, vol. 147, no. 22, p. 224106, Dec. 2017, ISSN: 0021-9606. DOI: [10.1063/1.4995643](https://doi.org/10.1063/1.4995643).
- [104] A. Annaberdiyev, G. Wang, C. A. Melton, M. C. Bennett, L. Shulenburger, and L. Mitas, “A new generation of effective core potentials from correlated calculations: 3d transition metal series,” *The Journal of Chemical Physics*, vol. 149, no. 13, p. 134108, 2018, ISSN: 0021-9606. DOI: [10.1063/1.5040472](https://doi.org/10.1063/1.5040472).
- [105] Y. Chang, S. Joshi, and L. K. Wagner, “Learning emergent models from *ab initio* many-body calculations,” 2023. arXiv: [2302.02899](https://arxiv.org/abs/2302.02899) [cond-mat.str-el].
- [106] Opensource.org, <https://opensource.org/licenses/MIT>. (visited on 11/04/2022).
- [107] SPDX Workgroup a Linux Foundation Project, <https://spdx.org/licenses/MIT.html>, 2018. (visited on 11/04/2022).
- [108] J. H. Saltzer, “The origin of the “MIT License”,” *IEEE Annals of the History of Computing*, vol. 42, no. 4, pp. 94–98, 2020. DOI: [10.1109/MAHC.2020.3020234](https://doi.org/10.1109/MAHC.2020.3020234).
- [109] C. R. Harris, K. J. Millman, S. J. van der Walt, *et al.*, “Array programming with NumPy,” *Nature*, vol. 585, no. 7825, pp. 357–362, Sep. 2020. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2).

- [110] P. Virtanen, R. Gommers, T. E. Oliphant, *et al.*, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020. DOI: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2).
- [111] A. Collette, *Hdf5 for python*, 2008. [Online]. Available: <http://h5py.alfven.org>.
- [112] J. G. Bednorz and K. A. Müller, “Possible highT_c superconductivity in the Ba-La-Cu-O system,” en, *Zeitschrift für Physik B Condensed Matter*, vol. 64, no. 2, pp. 189–193, Jun. 1986, ISSN: 1431-584X. DOI: [10.1007/BF01303701](https://doi.org/10.1007/BF01303701).
- [113] V. J. Emery, “Theory of high-T_c superconductivity in oxides,” *Phys. Rev. Lett.*, vol. 58, pp. 2794–2797, 26 Jun. 1987. DOI: [10.1103/PhysRevLett.58.2794](https://doi.org/10.1103/PhysRevLett.58.2794).
- [114] F. C. Zhang and T. M. Rice, “Effective hamiltonian for the superconducting cu oxides,” *Phys. Rev. B*, vol. 37, pp. 3759–3761, 7 Mar. 1988. DOI: [10.1103/PhysRevB.37.3759](https://doi.org/10.1103/PhysRevB.37.3759).
- [115] V. J. Emery and G. Reiter, “Mechanism for high-temperature superconductivity,” *Phys. Rev. B*, vol. 38, pp. 4547–4556, 7 Sep. 1988. DOI: [10.1103/PhysRevB.38.4547](https://doi.org/10.1103/PhysRevB.38.4547).
- [116] F. C. Zhang, “Validity of the *t*-*J* model,” *Physical Review B*, vol. 41, no. 10, pp. 7243–7246, 1990. DOI: [10.1103/PhysRevB.41.7243](https://doi.org/10.1103/PhysRevB.41.7243).
- [117] E. Pavarini, I. Dasgupta, T. Saha-Dasgupta, O. Jepsen, and O. K. Andersen, “Band-Structure Trend in Hole-Doped Cuprates and Correlation with T_{cmax} ,” *Physical Review Letters*, vol. 87, no. 4, p. 047003, Jul. 2001. DOI: [10.1103/PhysRevLett.87.047003](https://doi.org/10.1103/PhysRevLett.87.047003).
- [118] A. Macridin, M. Jarrell, T. Maier, and G. A. Sawatzky, “Physics of cuprates with the two-band hubbard model: The validity of the one-band hubbard model,” *Phys. Rev. B*, vol. 71, p. 134527, 13 Apr. 2005. DOI: [10.1103/PhysRevB.71.134527](https://doi.org/10.1103/PhysRevB.71.134527).
- [119] D. C. Peets, D. G. Hawthorn, K. M. Shen, *et al.*, “X-ray absorption spectra reveal the inapplicability of the single-band hubbard model to overdoped cuprate superconductors,” *Phys. Rev. Lett.*, vol. 103, p. 087402, 8 Aug. 2009. DOI: [10.1103/PhysRevLett.103.087402](https://doi.org/10.1103/PhysRevLett.103.087402).
- [120] N. Mott, “The spin-polaron theory of high-t c superconductivity,” *Advances in Physics*, vol. 39, no. 1, pp. 55–81, 1990. DOI: [10.1080/00018739000101471](https://doi.org/10.1080/00018739000101471).
- [121] L. Hozoi, S. Nishimoto, and C. de Graaf, “Renormalization of quasiparticle hopping integrals by spin interactions in layered copper oxides,” *Phys. Rev. B*, vol. 75, p. 174505, 17 May 2007. DOI: [10.1103/PhysRevB.75.174505](https://doi.org/10.1103/PhysRevB.75.174505).
- [122] C. H. Patterson, “Small polarons and magnetic antiphase boundaries in $\text{Ca}_{2-x}\text{Na}_x\text{CuO}_2\text{Cl}_2$ ($x = 0.06, 0.12$): Origin of striped phases in cuprates,” *Phys. Rev. B*, vol. 77, p. 094523, 9 Mar. 2008. DOI: [10.1103/PhysRevB.77.094523](https://doi.org/10.1103/PhysRevB.77.094523).
- [123] G. Yu, Y. Li, E. M. Motoyama, *et al.*, “Magnetic resonance in the model high-temperature superconductor $\text{HgBa}_2\text{CuO}_{4+\delta}$,” *Phys. Rev. B*, vol. 81, p. 064518, 6 Feb. 2010. DOI: [10.1103/PhysRevB.81.064518](https://doi.org/10.1103/PhysRevB.81.064518).
- [124] D. J. Scalapino, “A common thread: The pairing interaction for unconventional superconductors,” *Rev. Mod. Phys.*, vol. 84, pp. 1383–1417, 4 Oct. 2012. DOI: [10.1103/RevModPhys.84.1383](https://doi.org/10.1103/RevModPhys.84.1383).
- [125] Q. Sun, T. C. Berkelbach, N. S. Blunt, *et al.*, “PySCF: The Python-based simulations of chemistry framework,” en, *Wiley Interdisciplinary Reviews: Computational Molecular Science*, vol. 8, no. 1, e1340, Jan. 2018, ISSN: 1759-0884. DOI: [10.1002/wcms.1340](https://doi.org/10.1002/wcms.1340).
- [126] Q. Sun, X. Zhang, S. Banerjee, *et al.*, “Recent developments in the PySCF program package,” *The Journal of Chemical Physics*, vol. 153, no. 2, p. 024109, Jul. 2020, ISSN: 0021-9606. DOI: [10.1063/5.0006074](https://doi.org/10.1063/5.0006074).
- [127] J. P. Perdew, M. Ernzerhof, and K. Burke, “Rationale for mixing exact exchange with density functional approximations,” *The Journal of Chemical Physics*, vol. 105, no. 22, pp. 9982–9985, Dec. 1996, ISSN: 0021-9606. DOI: [10.1063/1.472933](https://doi.org/10.1063/1.472933).
- [128] G. Wang, A. Annaberdiyev, C. A. Melton, M. C. Bennett, L. Shulenburger, and L. Mitas, “A new generation of effective core potentials from correlated calculations: 4s and 4p main group elements and first row additions,” en, *The Journal of Chemical Physics*, vol. 151, no. 14, p. 144110, Oct. 2019, ISSN: 0021-9606. DOI: [10.1063/1.5121006](https://doi.org/10.1063/1.5121006).

- [129] W. A. Wheeler, S. Pathak, K. G. Kleiner, *et al.*, “PyQMC: An all-Python real-space quantum Monte Carlo module in PySCF,” *The Journal of Chemical Physics*, vol. 158, no. 11, p. 114801, Mar. 2023, ISSN: 0021-9606. DOI: [10.1063/5.0139024](https://doi.org/10.1063/5.0139024).
- [130] F. A. Reboredo, R. Q. Hood, and P. R. C. Kent, “Self-healing diffusion quantum monte carlo algorithms: Direct reduction of the fermion sign error in electronic structure calculations,” *Phys. Rev. B*, vol. 79, p. 195117, 19 May 2009. DOI: [10.1103/PhysRevB.79.195117](https://doi.org/10.1103/PhysRevB.79.195117).
- [131] F. A. Reboredo, “Many-body calculations of low-energy eigenstates in magnetic and periodic systems with self-healing diffusion Monte Carlo: Steps beyond the fixed phase,” *The Journal of Chemical Physics*, vol. 136, no. 20, p. 204101, May 2012, ISSN: 0021-9606. DOI: [10.1063/1.4711023](https://doi.org/10.1063/1.4711023).
- [132] W. A. Wheeler, K. G. Kleiner, and L. K. Wagner, “Ensemble variational monte carlo for optimization of correlated excited state wave functions,” *Electronic Structure*, vol. 6, no. 2, p. 025001, Apr. 2024. DOI: [10.1088/2516-1075/ad38f8](https://doi.org/10.1088/2516-1075/ad38f8).
- [133] S. Shepard, R. L. Panadés-Barrueta, S. Moroni, A. Scemama, and C. Filippi, “Double Excitation Energies from Quantum Monte Carlo Using State-Specific Energy Optimization,” *Journal of Chemical Theory and Computation*, vol. 18, no. 11, pp. 6722–6731, Nov. 2022, ISSN: 1549-9618. DOI: [10.1021/acs.jctc.2c00769](https://doi.org/10.1021/acs.jctc.2c00769).
- [134] D. Pfau, S. Axelrod, H. Sutterud, I. von Glehn, and J. S. Spencer, “Natural Quantum Monte Carlo Computation of Excited States,” Aug. 2023. arXiv: [2302.02899](https://arxiv.org/abs/2302.02899).
- [135] M. T. Entwistle, Z. Schätzle, P. A. Erdman, J. Hermann, and F. Noé, “Electronic excited states in deep variational Monte Carlo,” en, *Nature Communications*, vol. 14, no. 1, p. 274, Jan. 2023, ISSN: 2041-1723. DOI: [10.1038/s41467-022-35534-5](https://doi.org/10.1038/s41467-022-35534-5).
- [136] A. Lüchow, “Quantum Monte Carlo methods,” en, *WIREs Computational Molecular Science*, vol. 1, no. 3, pp. 388–402, 2011, ISSN: 1759-0884. DOI: [10.1002/wcms.40](https://doi.org/10.1002/wcms.40).
- [137] L. Shulenburger, “Quantum Monte Carlo applied to solids,” *Physical Review B*, vol. 88, no. 24, 2013. DOI: [10.1103/PhysRevB.88.245117](https://doi.org/10.1103/PhysRevB.88.245117).
- [138] K. T. Williams, Y. Yao, J. Li, *et al.*, “Direct Comparison of Many-Body Methods for Realistic Electronic Hamiltonians,” en, *Physical Review X*, vol. 10, no. 1, p. 011041, Feb. 2020, ISSN: 2160-3308. DOI: [10.1103/PhysRevX.10.011041](https://doi.org/10.1103/PhysRevX.10.011041).
- [139] N. S. Blunt, A. Alavi, and G. H. Booth, “Krylov-projected quantum monte carlo method,” *Phys. Rev. Lett.*, vol. 115, p. 050603, 5 Jul. 2015. DOI: [10.1103/PhysRevLett.115.050603](https://doi.org/10.1103/PhysRevLett.115.050603).
- [140] N. S. Blunt, G. H. Booth, and A. Alavi, “Density matrices in full configuration interaction quantum Monte Carlo: Excited states, transition dipole moments, and parallel distribution,” *The Journal of Chemical Physics*, vol. 146, no. 24, p. 244105, Jun. 2017, ISSN: 0021-9606. DOI: [10.1063/1.4986963](https://doi.org/10.1063/1.4986963).
- [141] K. Choo, G. Carleo, N. Regnault, and T. Neupert, “Symmetries and many-body excitations with neural-network quantum states,” *Phys. Rev. Lett.*, vol. 121, p. 167204, 16 Oct. 2018. DOI: [10.1103/PhysRevLett.121.167204](https://doi.org/10.1103/PhysRevLett.121.167204).
- [142] C. L. Benavides-Riveros, L. Chen, C. Schilling, S. Mantilla, and S. Pittalis, “Excitations of quantum many-body systems via purified ensembles: A unitary-coupled-cluster-based approach,” *Phys. Rev. Lett.*, vol. 129, p. 066401, 6 Aug. 2022. DOI: [10.1103/PhysRevLett.129.066401](https://doi.org/10.1103/PhysRevLett.129.066401).
- [143] L. Otis and E. Neuscamman, “A promising intersection of excited-state-specific methods from quantum chemistry and quantum monte carlo,” *WIREs Computational Molecular Science*, vol. 13, no. 5, e1659, 2023. DOI: <https://doi.org/10.1002/wcms.1659>.
- [144] L. Zhao and E. Neuscamman, “An Efficient Variational Principle for the Direct Optimization of Excited States,” *Journal of Chemical Theory and Computation*, vol. 12, no. 8, pp. 3436–3440, Aug. 2016, ISSN: 1549-9618. DOI: [10.1021/acs.jctc.6b00508](https://doi.org/10.1021/acs.jctc.6b00508).
- [145] S. D. Pineda Flores and E. Neuscamman, “Excited State Specific Multi-Slater Jastrow Wave Functions,” *The Journal of Physical Chemistry A*, vol. 123, no. 8, pp. 1487–1497, Feb. 2019, ISSN: 1089-5639. DOI: [10.1021/acs.jpca.8b10671](https://doi.org/10.1021/acs.jpca.8b10671).

- [146] L. Otis and E. Neuscamman, “Optimization Stability in Excited-State-Specific Variational Monte Carlo,” *Journal of Chemical Theory and Computation*, vol. 19, no. 3, pp. 767–782, Feb. 2023, ISSN: 1549-9618. DOI: [10.1021/acs.jctc.2c00642](https://doi.org/10.1021/acs.jctc.2c00642).
- [147] O. Higgott, D. Wang, and S. Brierley, “Variational Quantum Computation of Excited States,” *Quantum*, vol. 3, p. 156, Jul. 2019, ISSN: 2521-327X. DOI: [10.22331/q-2019-07-01-156](https://doi.org/10.22331/q-2019-07-01-156).
- [148] A. K. Theophilou, “The energy density functional formalism for excited states,” *Journal of Physics C: Solid State Physics*, vol. 12, no. 24, p. 5419, Dec. 1979. DOI: [10.1088/0022-3719/12/24/013](https://doi.org/10.1088/0022-3719/12/24/013).
- [149] E. K. U. Gross, L. N. Oliveira, and W. Kohn, “Rayleigh-ritz variational principle for ensembles of fractionally occupied states,” *Phys. Rev. A*, vol. 37, pp. 2805–2808, 8 Apr. 1988. DOI: [10.1103/PhysRevA.37.2805](https://doi.org/10.1103/PhysRevA.37.2805).
- [150] S.-i. Amari, “Natural gradient works efficiently in learning,” *Neural Computation*, vol. 10, no. 2, pp. 251–276, 1998. DOI: [10.1162/089976698300017746](https://doi.org/10.1162/089976698300017746).
- [151] I. Tobias, R. J. Fallon, and J. T. Vanderslice, “Potential Energy Curves for CO,” *The Journal of Chemical Physics*, vol. 33, no. 6, pp. 1638–1640, Dec. 1960, ISSN: 0021-9606, 1089-7690. DOI: [10.1063/1.1731475](https://doi.org/10.1063/1.1731475).