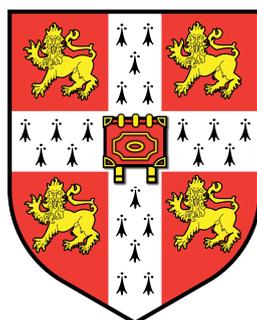


THE UNIVERSITY OF CAMBRIDGE

Taming the Inverse and Forward Problems in Density
Functional Theory



Nicholas David Woods

Clare College

Theory of Condensed Matter, Cavendish Laboratory

*A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy*

in the

Department of Physics

January 2022

Supervisors: Dr. Phil Hasnip
Prof. Mike Payne

To Mum

DECLARATION

This thesis is my own work and conforms to the University of Cambridge's guidelines on plagiarism. Where reference has been made to other research this is acknowledged in the text and bibliography. This thesis does not exceed the word limit for the respective Degree Committee. Chapters three, five, and six are based on work published in,

Computing the Self-Consistent Field in Kohn-Sham Density Functional Theory, N. Woods, P. Hasnip, M. C. Payne, *Journal of Physics: Condensed Matter*, **31**, 453001 (2019),

Insights from Exact Exchange-Correlation Kernels, N. Woods, M. T. Entwistle, R. W. Godby, *Phys. Rev. B*, **103**, 125155 (2021),

Accurate Total Energies from the Adiabatic Connection Fluctuation-Dissipation Theorem, N. Woods, M. T. Entwistle, R. W. Godby, *Phys. Rev. B*, **104**, 125126 (2021).

The following paper was published during the course of the doctoral research, but is not included in the thesis,

Algebraic Classical and Quantum Field Theory on Causal Sets, , E. Dable-Heath, C. J. Fewster, K. Rejzner, N. Woods *Phys. Rev. D*, **101**, 065013, (2020) (Editors' Suggestion).

Software developed during the doctoral research is available at <https://github.com/NickWoods1/>.

Nick Woods

ABSTRACT

Theory of Condensed Matter, Cavendish Laboratory

Taming the Inverse and Forward Problems in Density Functional Theory

by NICHOLAS DAVID WOODS

The ‘forward problem’ of ground-state density functional theory (DFT) constitutes finding the ground-state density $n(x)$ that minimises a Kohn-Sham total energy functional defined using some exchange-correlation (xc) functional $E_{xc}[n]$. Towards this end, the associated Euler-Lagrange equations, i.e. the Kohn-Sham equations, are often solved in practice, which demand a procedure that iterates an initial guess density to a *self-consistent* density (the solution). A new framework is presented for evaluating the performance of self-consistent field methods in Kohn-Sham DFT. The aims of this work are two-fold. First, we explore the properties of Kohn-Sham DFT as it pertains to the convergence of self-consistent field iterations. Sources of inefficiencies and instabilities are identified, and methods to mitigate these difficulties are discussed. Second, we introduce a framework to assess the relative utility of algorithms, comprising a representative benchmark suite of over fifty Kohn-Sham simulation inputs, the SCF- x_n suite. This provides a new tool to develop, evaluate and compare new algorithms in a fair, well-defined and transparent manner.

The ‘inverse problem’ of time-dependent (ground-state) DFT constitutes finding the time-(in)dependent Kohn-Sham potential $v_{KS}(x, t)$ that yields a given reference density $n(x, t)$ upon solution of the time-(in)dependent Kohn-Sham equations. This inverse map can be unstable, particularly in the presence of low-density regions, and thus methods are designed to alleviate numerical difficulties in the present context. On the other hand, linear response time-dependent DFT centres around the first-order response of the xc potential due to perturbing densities – the so-called xc kernel $f_{xc}(x, x', \omega)$. Computing exact xc kernels represents a linearised version of the previous inverse problem: this state of affairs, whilst still challenging, is more manageable. Methods to ensure the robustness of exact numerical f_{xc} computations are set out. In the context of inhomogenous one-dimensional finite systems, these developments permit an improved understanding of f_{xc} in itself, and in relation to various applications, such as the optical spectrum and ground-state correlation energies using the adiabatic connection fluctuation-dissipation theorem. We expect that certain key insights derived from this work will assist in the informed development of improved functional approximations.

ACKNOWLEDGEMENTS

Thank you foremost to my family – Mum, Dad, Liv, Nan & Grandad, Grandma & Grandad – for giving me endless unconditional support and the freedom to pursue my passions from a young age. To Mum & Nan: most of what I have I owe to them, and to Grandad, whose jovial spirit and no-use-crying-over-spilt-milk approach towards life has been an inspiration. To Eleni, the fortuitous meeting of whom has been the highlight of the past years – a truly unique person, and I look forward to many more debates and Souvlakis in times to come.

Thank you to my supervisors Mike and Phil, who have given me generous guidance, the freedom to address a range of topics in my thesis, and continued support. I would also like to thank David Bowler for discussions regarding my work on self-consistency, and for help organising my first publication. Furthermore, I am indebted to Rex Godby, with whom, alongside Mike Entwistle, much of the work in the second half of this thesis was done in collaboration. It was an enjoyable and productive experience, and I learned a great deal from our meetings. Thank you to Mike for helping organise this collaboration, and for our continued friendship.

To Alex, who has been once of my closest friends throughout my time at Cambridge, and to Juraj: while it is not always clear at the time, our countless discussions and debates have had a profound influence on my views. I hope our WTs legacy will live on. Matt – a would-be guardian, no doubt not even this sentence is safe – has been a dear friend for some time, and a source of endless humour spanning the days of the MPhys room & the doctoring station, the days of the Iron Maiden & The Rack Pack curse, and through to present times.

To the innumerable people that have comprised my ‘online’ circles over the past fourteen years: such a sustained friendship across geographical boundaries is a rare gift, which is made possible through our idiosyncratic means of communication, to be understood in more ways than one – *nobody asked for this life*.

While it is impossible to name everyone, a number of others have had a significant positive influence on my life: Iria (fellow member of the Mott-ly crue), Sarah, Alice, Micheal (our final night in Spain before the pandemic was a fitting and memorable send-off), Ben, Joas, Sundeep, Adiran, Callum, Petr, Pavao, Angela, Matt E, Benat, Ollie, Max, Atilla, Gareth, Michael R, and, of course, the remaining of the members of the Centre for Doctoral Training/Theory of Condensed Matter group/Clare College.

Contents

Acknowledgements	iv
Contents	v
1 Introduction	1
1.1 Structure	2
1.2 Units and Conventions	3
2 Theory	4
2.1 Quantum Mechanics and the Many-Body Problem	4
2.2 Electronic Structure Methods	7
2.3 Density Functional Theory	9
2.3.1 Exchange-Correlation Potential	11
2.3.2 Implementation	13
2.4 Time-Dependent Density Functional Theory	14
2.4.1 Time-Dependent Exchange-Correlation Potential	15
2.4.2 Linear Response Function	16
2.4.3 Exchange-Correlation Kernel	18
2.4.4 Optical Properties	19
2.5 The Adiabatic-Connection Fluctuation-Dissipation Theorem	20
2.5.1 Adiabatic Connection	20
2.5.2 Fluctuation-Dissipation Theorem	21
2.5.3 Total Energy Functional	22
2.5.4 Statistical Hole and the Exchange-Correlation Hole	23
3 Computing the Self-Consistent Field in Density Functional Theory	26
3.1 Background	26
3.2 Purpose	28
3.3 Self-Consistency in Kohn-Sham Theory	30
3.3.1 The Kohn-Sham Map	30
3.3.2 Defining Convergence	31
3.3.3 Some Unique Properties of The Kohn-Sham Map	32
3.3.4 Fixed-Point and Damped Iterations	32
3.3.5 The Aufbau Principle and Fractional Occupancy	36
3.3.6 The Initial Guess	39
3.3.7 Ill-Conditioning and Charge Sloshing	41
3.4 Methods and Algorithms	48

3.4.1	The Quasi-Newton Update	49
3.4.1.1	Broyden's methods	50
3.4.1.2	Multisecant Broyden's methods	52
3.4.1.3	Pulay's Method	54
3.4.1.4	Holistic Approaches based on Multisecants	57
3.4.2	Preconditioning	60
3.5	Test Suite	64
3.6	Results and Discussion	67
3.7	Conclusion	70
4	The Exact Time-Dependent Exchange-Correlation Potential	74
4.1	Implementation	74
4.1.1	Ground State	75
4.1.2	Time Dependence	76
4.2	Reverse Engineering	77
4.2.1	Background	77
4.2.2	Stability	79
4.3	Results and Discussion	80
4.3.1	Ground State	80
4.3.2	Time Dependence	83
4.4	Conclusion	84
5	The Exact Exchange-Correlation Kernel and the Optical Spectrum	88
5.1	Methodology	90
5.1.1	Background	90
5.1.2	Challenges in Computing Exact Exchange-Correlation Kernels	91
5.1.3	The Real-Space Truncation Method	93
5.1.4	The Eigenspace Truncation Method	95
5.1.5	Error and Exact Conditions	97
5.1.6	Gauge Freedom	98
5.2	Results and Discussion	98
5.2.1	Atom	99
5.2.2	Infinite Potential Well	104
5.2.3	Model Response Functions	109
5.2.3.1	Scalar Model Response Function	109
5.2.3.2	Two-State Model Response Function	110
5.2.4	Quantum Harmonic Oscillator	113
5.2.5	Gauge Freedom	116
5.2.6	Exchange-Correlation Potential Versus Exchange-Correlation Kernel	118
5.3	Conclusions	121
6	Total Energies from the Adiabatic-Connection Fluctuation-Dissipation Theorem	123
6.1	Methodology	125
6.1.1	Response Functions Along the Adiabatic Connection	125
6.1.2	Integration Schemes	129
6.1.3	One-Shot Calculations	130

6.1.4	Self-Consistent Calculations	131
6.2	Results and Discussion	132
6.2.1	Short-Range Correlations	132
6.2.2	Long-Range Correlations	140
6.2.3	Sources of Error	141
6.3	Conclusions	144
7	Conclusion	147
	References	149

Chapter 1

Introduction

Once engaged, his motivation is of a rather different sort. What then challenges him is the conviction that, if only he is skilful enough, he will succeed in solving a puzzle that no one before has solved [...] a fact that makes [science] no less fascinating to the proper sort of addict.

T. S. Kuhn (1962) [1]

The history of science can be understood as a succession of incommensurate paradigms, wherein each new paradigm is birthed out of an insurmountable tension produced in the one just before it [1, 2]. The scientific revolution, and Newton in particular, brought about a new ‘occult’ view of physics, whereby nature conforms more to abstract mathematics than it does to our mechanical contact-based intuitions [3, 4]. Such reasoning facilitated the development and acceptance of *quantum theory* in the twentieth century. Quantum mechanics identifies discrete microscopic constituents of nature that possess striking non-classical properties, but in doing so, widens the gap between the *fundamental* and the *emergent*. Understanding the emergent quantum mechanical properties of matter from the interactions between its manifold constituents represents a central challenge in modern physics [5]. Nonetheless, harnessing the quantum description of matter has led to a great number of otherwise impossible developments: semiconductors & the modern classical computer, the quantum computer, various medical imaging techniques, superconductors, and so on.

In non-relativistic quantum mechanics, particles are a quantum state $|\Psi\rangle$ that evolves according to the time-dependent Schrödinger equation. The dimension of the Hilbert space \mathcal{H} in which the quantum state lives undergoes exponential growth with increasing particle number, bringing about the so-called curse of dimensionality. This curse can be a blessing in disguise: it provides the requisite room for the richness of nature to manifest itself within the framework. However, in the present context, it lives up to its name as it precludes (numerical) determination of exact many-body quantum states involving more than just a few particles. In order to proceed, it is traditional to follow one of two paths: construct a small-scale analytic model designed to capture and elucidate a certain quantum phenomenon, or develop a *general purpose* framework to approximate the equations of quantum mechanics that scales to many particles when implemented on a computer. This thesis shall be concerned with the latter, and progression in this endeavour enables, amongst other things, more effective materials design, the advantages of which are myriad, e.g. improved batteries [6].

Compute power, which refers here to both the performance and accessibility of computer hardware, has grown significantly over the past half-century. This increase has led to the rise of computational science as a distinct third axis in the practice of science, alongside experimental and theoretical science. We are concerned here with the hierarchy of methods that has emerged for calculating the properties of molecular and solid state systems *in silico* by approximating the time-independent and time-dependent Schrödinger equations [7–15]. In particular, one of the most prominent approaches towards this end has been ground-state and time-dependent density functional theory (DFT) [7–9, 12, 16–19] – a reformulation of quantum mechanics in which the particle density is primary, rather than the quantum state. For a number of reasons, practitioners in both the physics and chemistry communities have deemed this level of theory appropriate to tackle a range of problems at an acceptable computational cost. However, the concept of ‘no free lunch’ is ever present: more demanding yet more accurate methods are often preferred in quantum chemistry, such as coupled cluster theory and many-body perturbation theory. On the other hand, less demanding methods are often preferred in biological applications, such as classical molecular dynamics [20], and, more recently, data-driven approaches [21]. An advantage of Kohn-Sham DFT is that it possesses a great scope for improvement, which comes from developing a better understanding of the so-called *exchange-correlation (xc) functional*.

1.1 Structure

In broad terms, this thesis involves the computational implementation of ground-state DFT in practice, and furthermore progressing towards an improved understanding of various theoretical aspects of ground-state and (linear response) time-dependent DFT. Chapter 3 focuses on the iterative procedure used to converge an initial guess density towards the density that minimises an appropriate Kohn-Sham energy functional, i.e. the *self-consistent field* procedure. These iterations are at the heart of practical DFT calculations, and constitute the most common source of numerical divergence [22, 23]. Such issues will become more pronounced with increasing compute power due to practitioners demanding larger and more exotic systems, thereby accentuating so-called *ill-conditioning*. With the aim of alleviating these difficulties in mind, the plane-wave pseudopotential DFT software CASTEP [24] is used to implement and benchmark a range of modern algorithms and preconditioners. Through understanding the sources of divergences and inefficiencies, this work proposes a benchmarking suite of DFT input systems [25] which is used, alongside an unbiased metric, to compare methodology based on the concept of *Pareto efficiency* [22].

Chapter 4 onward deals with exact quantum mechanics, and exact DFT, in one spatial dimension involving model systems comprising a moderate number of electrons. These simplifications allow us to examine in depth the exact time-dependent xc potential $v_{xc}[n](x, t)$ and xc kernel $f_{xc}[n](x, x', \omega)$ with the aim of improving approximate DFT functionals across a variety of applications. Chapter 4 outlines the software that has been developed to implement these model systems, and in particular focuses on the difficulties arising when computing the *inverse map* in time-dependent Kohn-Sham DFT, that is, finding the unique time-dependent Kohn-Sham potential that generates

a given reference density. We shall provide evidence that certain numerical implementations of this map are *ill-posed*, as dynamical inverse problems often are, and thus to avoid chaotic behaviour a set of algorithms are designed to regularise the iterations as best as possible.

Computing the exact xc kernel $f_{xc}[n](x, x', \omega)$ of linear response time-dependent DFT is the central focus of Chapter 5. Evaluation of $f_{xc}[n]$ is yet again a difficult numerical procedure, although, it is a *linear* inverse problem, and is therefore more manageable than the non-linear inverse map considered in Chapter 4. Nonetheless, these numerical obstacles have prevented or hindered studies of the exact xc kernel, and this work is among the first to provide a careful treatment of exact numerical xc kernels [26]. The properties of the exact xc kernel are examined in isolation and in relation to the optical spectrum. For example, we find that $f_{xc}[n](x, x', \omega)$ contains singularities along the real ω axis, related to non- v -representable density perturbations, and this analytic structure is demonstrated to be critical for recovering certain kinds of optical excitations. Conventional density functionals omit this intricate analytic structure, and are therefore precluded *in principle* from recovering these excitations.

Linear response time-dependent DFT can be used, somewhat counter-intuitively, to reformulate ground-state DFT, and thus calculate ground-state energies of solids and molecules, using the *adiabatic-connection fluctuation-dissipation theorem* (ACFDT) [7, 9, 17, 19, 27]. The developments presented in Chapter 5 allow us to compute the exact coupling-constant dependent exchange-correlation kernel $f_{xc}^\lambda[n](x, x', \omega)$ towards an improved understanding of approximate ground-state total energies derived from the ACFDT [28]. In Chapter 6, it is observed that chemical accuracy is reliably preserved while neglecting entirely the frequency dependence in the exact functional $f_{xc}[n](\omega = 0)$. This performance is understood on the grounds that the exact $f_{xc}[n]$ varies slowly over the most relevant ω range (but not in general), and hence the spatial structure in $f_{xc}[n](\omega = 0)$ is able to largely remedy the principal issue in the present context: self interaction. Computations of the coupling-constant averaged correlation hole facilitate the conclusion that self interaction plagues conventional density-functional approximations, and furthermore provides a perspective on the long-range correlations inherent to the ACFDT formalism.

1.2 Units and Conventions

Hartree atomic units are used throughout unless stated otherwise: $m_e = e = \hbar = 4\pi\epsilon_0 = 1$.

Dirac bra-ket notation is used for linear algebra, e.g. members of a complex Hilbert space $|\psi\rangle, |\phi\rangle \in \mathcal{H}$ have the inner product $\langle\psi|\phi\rangle \in \mathbb{C}$.

Chapter 2

Theory

2.1 Quantum Mechanics and the Many-Body Problem

How many bodies are required before we have a problem? ... In the eighteenth-century, Newtonian mechanics, the three-body problem was insoluble. With the birth of general relativity around 1910 and quantum electrodynamics in 1930, the two- and one-body problems became insoluble. And within modern quantum-field theory, the problem of zero-bodies (vacuum) is insoluble. So, if we are out after exact solutions, no bodies at all is already too many.

R. D. Mattuck (1967) [11]

Particles are quantum states $|\Psi\rangle$, i.e. vectors in a Hilbert space \mathcal{H} with unit norm $\langle\Psi|\Psi\rangle = 1$. A composite system of N particles is a unit-norm state in a much larger tensor product space $\mathcal{H}_1 \otimes \mathcal{H}_2 \dots \otimes \mathcal{H}_N$, where \mathcal{H}_i is the individual Hilbert space of particle i . Consider the spin component of an electron, which exists in a two-dimensional Hilbert space. The spin state of N electrons lives in a composite space of dimension 2^N , in other words, a maximum of 2^N complex numbers are required to represent an arbitrary spin state. Complex numbers are approximated *in silico* using two 64-bit floating-point real numbers, and therefore to simulate the spin state of just $N = 40$ electrons exceeds the typical capacities of a modern computer. Improved hardware does not save us here: the growth in the Hilbert space dimension is exponential. Furthermore, an everyday object, such as the paper on which this thesis is printed, contains $N \approx 10^{24}$ atoms (electrons *and* nucleons) – a less than ideal state of affairs in light of the prior concern¹. These are just two facets of the many-body problem in quantum mechanics, which is a problem that only worsens when the full state of the electrons and nucleons are considered, rather than just the spin state.

In further defining the problem at hand, a great deal of lossless simplification can be made. Namely, we are concerned mostly with many-body quantum states that are

¹Schrödinger in *What is Life?* [29] can assist in intuiting just how small and numerous atoms are: “Suppose that you could mark the molecules in a glass of water; then pour the contents of the glass into the ocean and stir the latter thoroughly so as to distribute the marked molecules uniformly throughout the seven seas; if then you took a glass of water anywhere out of the ocean, you would find in it about a hundred of your marked molecules”.

eigenstates of a self-adjoint operator called the Hamiltonian, H . Eigenstates of the Hamiltonian are quantum states of definite energy, which are relevant in part because nature often finds herself in states of minimum energy, i.e. in the ground state of an appropriate Hamiltonian operator. In analogy with the Hamiltonian formulation of classical mechanics, the N -electron² Hamiltonian operator,

$$H = \sum_{i=1}^N \left(\frac{1}{2} \hat{p}_i^2 + v_{\text{ext}}(\hat{x}_i) \right) + \sum_{i=1}^N \sum_{j>i}^N \frac{1}{|\hat{x}_i - \hat{x}_j|}, \quad (2.1)$$

comprises three terms: a kinetic term, an external potential term, and a Coulomb electrostatic interaction term. All three components are defined using the primitive quantum operators, position \hat{x}_i and momentum \hat{p}_i , where subscript i denotes the operator acting on the Hilbert space \mathcal{H}_i of particle i within the larger composite vector space. Spin states are omitted in the following discussion.

The particular form of the N -electron Hamiltonian restricts the *entanglement structure* of electrons that are in an eigenstate of H . In order to see what is meant by this, consider a system of particles that are identical to electrons, except without Coulomb interaction between the particles. In this case, the N -particle Hamiltonian Eq. (2.1) involves the sum of N one-particle Hamiltonians,

$$H = H_1 + H_2 + \dots + H_N, \quad (2.2)$$

where H_i operates on the i^{th} Hilbert space alongside the identity operator for the remaining subspaces. A simple exercise in linear algebra demonstrates that the ground state of this Hamiltonian is a *product state*, that is, a product of the single-particle ground states of each individual Hamiltonian.

Let us label the eigenstates of H_i as $\{|\phi_n\rangle\}$ for all i , then a general state is expanded in this basis as such,

$$|\Psi\rangle = \sum_{nlm\dots} \alpha_{nlm\dots} |\phi_n\rangle |\phi_l\rangle |\phi_m\rangle \dots, \quad (2.3)$$

where all coefficients $\alpha_{nlm\dots}$ constitute a maximal parameterisation of the general state. In the non-interacting case, most coefficients are zero, and in fact, in the instance that the particles are bosons, all but one coefficient is zero in the ground state: $\alpha_{111\dots} = 1$. Electrons, however, are spin-half particles, from which it can be deduced that all N -electron states are anti-symmetric upon particle exchange, $|\phi_n\rangle |\phi_m\rangle \rightarrow -|\phi_m\rangle |\phi_n\rangle$ (this is the case no matter the basis). This *exchange anti-symmetry* renders a large portion of the coefficients redundant due to symmetry relations, $\alpha_{nm\dots} = -\alpha_{mn\dots}$ for all pairs of indices. (Eigenstates of the real-space Hamiltonian, for example, are either symmetric or anti-symmetric because the Hamiltonian itself is symmetric under interchange $x_i \leftrightarrow x_j$).

²The full atomic quantum state $|\Psi\rangle \in \mathcal{H}^{\text{electron}} \otimes \mathcal{H}^{\text{nucleon}}$ is often omitted in condensed matter calculations because it is possible to approximately decouple the electron and nucleon Hilbert spaces using the adiabatic approximation (owing to the relative mass difference between electrons and nucleons). A general electronic state is therefore the ground state of an electronic Hamiltonian whose external potential is modified to include the nuclear state. Together with the assumption that the position probability distribution of the nucleons is infinitely localised, these observations constitute the *Born-Oppenheimer approximation*. This thesis does not go beyond the Born-Oppenheimer approximation.

Continuing with the non-interacting example, a further consequence of enforcing exchange anti-symmetry is that the ground state must comprise the lowest lying N eigenstates of H_i , rather than N copies of the lowest lying eigenstate. In fact, the ground state has $N!$ non-zero coefficients in the basis set expansion,

$$|\Psi\rangle = \frac{1}{\sqrt{N!}} (|\phi_1\rangle|\phi_2\rangle|\phi_3\rangle\dots - |\phi_2\rangle|\phi_1\rangle|\phi_3\rangle\dots \pm \text{permutations}), \quad (2.4)$$

where $\alpha_{123\dots} = -\alpha_{213\dots} = \dots = 1/\sqrt{N!}$. In other words, it is equiprobable that the ‘first particle’ is in state $|\phi_1\rangle$, or $|\phi_2\rangle$, or $|\phi_3\rangle$, and so on, meaning the ‘first particle’ has lost its identity when brought into the composite system of particles identical to itself. While we have introduced a great number of non-zero coefficients, we are only required to compute and store one number, $1/\sqrt{N!}$, and N single-particle states.

This vast simplification largely circumvents the many-body problem discussed above, but is limited as eigenstates of the interacting Hamiltonian cannot in general be represented in this fashion, i.e. as a so-called Slater determinant of single-particle states, Eq. (2.4). The extent to which an interacting state $|\Psi\rangle$ can’t be represented by a Slater determinant (whose single-particle states are generated using some external potential v_{ext}) is one possible measure of *correlation*. Dissociation of the H_2 molecule yields *static* correlation of the kind just described in the limit of large separation [30].

The least amount of non-zero coefficients $\alpha_{nlm\dots}$ required to represent a general many-body quantum state with Coulomb (two-body) interactions appears to remain an open question [31–33]. It is likely that the effective dimension of such states is far less than the dimension of the full Hilbert space, given that only two-body interactions are involved. In fact, the total energy can be written in terms of the first- and second-order density matrices [12],

$$\gamma_1(x_1, x'_1) = N \int \Psi^*(x'_1, x_2, \dots, x_N) \Psi(x_1, x_2, \dots, x_N) \prod_{i=2}^N dx_i \quad (2.5)$$

$$\gamma_2(x'_1, x'_2, x_1, x_2) = \frac{N(N-1)}{2} \int \Psi^*(x'_1, x'_2, \dots, x_N) \Psi(x_1, x_2, \dots, x_N) \prod_{i=3}^N dx_i \quad (2.6)$$

as such,

$$E[\gamma_1, \gamma_2] = \iint \left(-\frac{1}{2} \nabla_1^2 + v(x_1) \right) \gamma_1(x_1, x'_1) dx_1 dx'_1 + \iint \frac{\gamma_2(x_1, x_2, x_1, x_2)}{|x_1 - x_2|} dx_1 dx_2, \quad (2.7)$$

where the real-space basis is used. These density matrices are of much lower dimension than the many-body wavefunction, and if the density matrices that minimise $E[\gamma_1, \gamma_2]$ were known³, then all observables \hat{O} with at most two-body operators involved could be evaluated using the trace, $o = \text{Tr}(\hat{\gamma}_1 \hat{O}_1) + \text{Tr}(\hat{\gamma}_2 \hat{O}_2)$, where \hat{O}_n is an n -body operator. One could envisage a practical scheme where we minimise over the degrees of freedom in γ_1 and γ_2 (whose dimension does not grow with increasing particle number) in order to obtain ground-state energies, and so on. However, this minimisation must be performed

³The N -representable matrices γ_1 and γ_2 that minimise $E[\gamma_1, \gamma_2]$ are equal to the density matrices constructed from the exact ground-state many-body wavefunction.

over functions γ_1 and γ_2 that can be constructed from many-body wavefunctions as in Eqs. (2.5) and (2.6): necessary and sufficient conditions for this to be the case are not known. This is the N -representability problem, and research along these lines is ongoing [12, 34].

For these reasons, and others (e.g. the long-range and singular nature of the Coulomb interaction), exact computation of many-body quantum states remains a challenge. However, approximations that utilise, and even build on, the single-particle picture have been remarkably successful, largely due to the fact that correlation tends to constitute a modest correction to non-interacting (e.g. mean-field) theories. Frameworks that aim to approximate the quantum state of interacting electrons (or its derived properties) are called electronic structure methods, some of which are now briefly mentioned.

2.2 Electronic Structure Methods

The Hartree and Hartree-Fock methods are paradigmatic non-interacting theories that modify the external potential in order to approximate the interacting state of affairs. The ansatz of Hartree is straightforward: suppose the N -body electronic wavefunction can be written as a product state,

$$\Psi_{\text{H}}(\{x_i\}) = \prod_{i=1}^N \phi_i(x_i). \quad (2.8)$$

The set of N square-integrable functions $\{\phi_i\} \in L^2(\mathbb{R}^3)$ are termed single-particle orbitals. Hartree theory is then defined by

$$E_{\text{H}}[v_{\text{ext}}] = \min_{\phi_i \in L^2} \{ \langle \Psi_{\text{H}} | H[v_{\text{ext}}] | \Psi_{\text{H}} \rangle \mid \langle \phi_i | \phi_j \rangle = \delta_{ij} \}, \quad (2.9)$$

where the minimum is a search over all wavefunctions that are products of orthonormal single-particle orbitals. Note that replacing $\Psi_{\text{H}}(\{x_i\})$ with a general square-integrable anti-symmetric function $\Psi(\{x_i\})$ would yield the exact ground state (lowest energy eigenstate of $H[v_{\text{ext}}]$) – this is often called the variational principle. In other words, the time-independent Schrödinger equation is the Euler-Lagrange equation for this minimisation problem. The set of square-integrable anti-symmetric functions subsumes the set of (wave)functions that can be expressed as a product state, therefore the Hartree energy is a variational maximum bound on the ground-state energy, $E_{\text{H}} \geq E_0$, where the equality is satisfied iff the exact ground-state wavefunction is a product state (i.e. no interaction term in the Hamiltonian, as discussed).

Hartree-Fock theory supposes that the N -body electronic wavefunction can be written as an anti-symmetric combination of product states (a Slater determinant), denoted $\Psi_{\text{HF}}(\{x_i\})$, and subsequently solves

$$E_{\text{HF}}[v_{\text{ext}}] = \min_{\phi_i \in L^2} \{ \langle \Psi_{\text{HF}} | H[v_{\text{ext}}] | \Psi_{\text{HF}} \rangle \mid \langle \phi_i | \phi_j \rangle = \delta_{ij} \}. \quad (2.10)$$

Hartree-Fock theory tends to be much more accurate than Hartree theory as it provides additional well-motivated variational freedom. Furthermore, the Hartree-Fock energy remains an upper bound on the exact ground-state energy. The associated Euler-Lagrange equations that minimise the Hartree-Fock energy functional are

$$\left(-\frac{1}{2}\nabla^2 + v_{\text{ext}}(x) + v_{\text{H}}(x) \right) \phi_i(x) + \int dx' v_{\text{x}}(x, x') \phi_i(x') = \epsilon_i \phi_i(x), \quad (2.11)$$

where,

$$v_{\text{H}}(x) = \int \frac{n(x')}{|x - x'|} dx' \quad (2.12)$$

$$v_{\text{x}}(x, x') = - \sum_{j=1}^N \frac{\phi_j^*(x) \phi_j(x')}{|x - x'|} \quad (2.13)$$

are the Hartree and exchange potentials, and

$$n(x) = N \int |\Psi(x, x_2, \dots, x_N)|^2 \prod_{i=2}^N dx_i, \quad (2.14)$$

is the particle density, i.e. the probability that a particle resides at position x . As expected, the Hartree-Fock equations are free from interaction, and take the form of a single-particle non-linear eigenequation with a non-local potential (the exchange term). These equations reflect the fact that a mean-field approximation has been made: the (non-interacting) particles are influenced by a background external potential generated from the position distribution of those same particles through the Hartree potential, Eq. (2.12). This is a source of non-linearity in the Hartree-Fock equations, meaning a so-called self-consistent solution is required, see Chapter 3. Moreover, an exchange potential is introduced that influences the particles in such a fashion as to mimic what occurs when we ensure the wavefunction is exchange anti-symmetric (in this way, exchange can be seen as a force between the particles).

A systematic route beyond the Hartree-Fock approximation is given by the configuration interaction and coupled cluster methods. Both of these methods utilise the same general observations. First, the set of *all* Slater determinants constructed from (for example) Hartree-Fock orbitals span the entire Hilbert space, see Eq. (2.3). Second, the Hartree-Fock wavefunction (for example) is a reasonable approximation to the many-body wavefunction, and therefore additional Slater determinants constructed from orbitals with higher energy than the lowest lying N orbitals constitute next-to-leading-order corrections to the Hartree-Fock wavefunction. The more Slater determinants included in the basis set expansion, the more accurate the configuration interaction and coupled cluster methods will be. The computational cost involved in including additional Slater determinants is steep, and therefore these methods are often used for moderately-sized molecular systems in quantum chemistry.

The variational and diffusion Monte Carlo methods are accurate stochastic methods that utilise random numbers to converge towards ground-state observables to within some statistical error [10]. In the case of variational Monte Carlo, an ansatz wavefunction is constructed (for example, the Hartree-Fock wavefunction with a so-called Jastrow

factor included) with some free parameters, and then these parameters are adjusted in a certain random walk designed to ensure variational convergence towards the ground-state energy. The free parameters in the Jastrow factor permits the inclusion of correlation. In diffusion Monte Carlo, no ansatz wavefunction is included per se, and stochastic integration techniques are able to project out the exact ground-state wavefunction to within statistical error. Both techniques, but the latter in particular, suffer due to the fermionic sign problem. Namely, the wavefunction must be anti-symmetric, and therefore it must consist of regions of positive and negative sign: the $N - 1$ -dimensional ‘lines’ along which the sign of the wavefunction switches is called the nodal surface. For technical reasons, these stochastic techniques often require the nodal surface to be specified prior to performing the calculation. Approximating the nodal surface *a priori* is a difficult problem, but nonetheless these methods have produced some of the most accurate computations of ground-state energies to date. For example, Monte Carlo calculations of the homogeneous electron gas (HEG) are used to parameterise various density functionals [35], see Section 2.3.

Perturbation theories constitute an additional, distinct course of action. In Møller-Plesset perturbation theory, the Hartree-Fock Hamiltonian is treated as the unperturbed operator, and the Hartree-Fock Slater determinant is the unperturbed wavefunction. Then, the operator that transforms the Hartree-Fock ground state into the interacting ground state is introduced alongside some parameter λ , and the perturbation theory is constructed as usual (the wavefunction becomes a power series in λ). In contrast, one can instead centre the perturbation theory around the Green’s function – an object constructed from the wavefunctions that contains much of the pertinent information about the quantum state, but for various reasons is particularly suited for perturbation theory. One such reason is that perturbation theories in powers of the Coulomb interaction are not efficient per additional term included in the power series. Another reason is that an effective formalism for Green’s function perturbation theory was developed in the context of quantum electrodynamics which makes use of Feynman diagrams. This leads to the popular *GW* approximation, which has proven able to provide an accurate account of both ground- and excited-state phenomena at a reasonable computational cost [10]. For a modern review of electronic structure methods (including more recent machine learning-based methods), see Ref. [14].

2.3 Density Functional Theory

Among the methods just outlined, density functional theory (DFT) and time-dependent DFT stand out in terms of their widespread adoption over the past decades – for a large number of applications, DFT-based methods have become a standard choice [7, 14, 36, 37].

Ground-state DFT [38] refers to a reformulation of ground-state quantum mechanics in which a one-to-one mapping is established between the external potential v_{ext} and particle density n of a ground-state many-body system as described by the Schrödinger equation⁴. Ultimately, the existence of this map renders much of the information in

⁴Note that statistical mechanics and temperature-dependence via ensembles of states is not considered in this thesis.

the $3N$ -dimensional ground-state wavefunction superfluous, as observables are able to be calculated as functionals of a compressed three-dimensional object: the ground-state density n .

Hohenberg and Kohn's influential article [38] established the following: there exists an invertible (one-to-one) map S defined by the Schrödinger equation

$$S : \mathcal{V} \rightarrow \boldsymbol{\psi} \quad (2.15)$$

from the set of all permissible external potentials \mathcal{V} (up to a constant shift, see Ref. [39]) to the set of all possible many-body ground-state wavefunctions $\boldsymbol{\psi}$ (up to a phase), where each member of $\boldsymbol{\psi}$ is generated by solving the Schrödinger equation with some $v_{\text{ext}} \in \mathcal{V}$. Furthermore, there also exists an invertible map

$$N : \boldsymbol{\psi} \rightarrow \mathcal{N} \quad (2.16)$$

from $\boldsymbol{\psi}$ to the set of particle densities \mathcal{N} that can be obtained from the ground state of some Hamiltonian $H[v_{\text{ext}}]$ – this is also known as the set of interacting v -representable densities. Since both S and N are shown to be invertible, then their composition

$$SB : \mathcal{V} \rightarrow \mathcal{N} \quad (2.17)$$

is invertible, and a one-to-one map between potentials and ground-state densities is established.

Given now that a ground-state density uniquely determines the potential from which it came, we can deduce that the ground-state energy (or indeed any ground-state observable) can be evaluated by minimising a *functional* of the density, as such

$$E_0[v_{\text{ext}}] = \min_{\Psi \in \boldsymbol{\psi}} \langle \Psi | \hat{H}[v_{\text{ext}}] | \Psi \rangle, \quad (2.18)$$

$$\implies E_0[v_{\text{ext}}] = \min_{n \in \mathcal{N}} \langle \Psi[n] | \hat{H}[v_{\text{ext}}] | \Psi[n] \rangle. \quad (2.19)$$

$$= \min_{n \in \mathcal{N}} \left\{ F_{\text{HK}}[n] + \int n(x)v_{\text{ext}}(x) dx \right\}, \quad (2.20)$$

where $F_{\text{HK}}[n]$ is the Hohenberg-Kohn functional [38]. While the Hohenberg-Kohn functional is the first of its kind, it is not the only functional that satisfies the conditions for a valid density functional, i.e. the associated total energy functional is minimised by the ground-state density in order to produce the ground-state energy. Indeed, the Hohenberg-Kohn functional has certain issues: 1) the functional needs to allow for degenerate ground-state wavefunctions, 2) the domain of the functional should ideally be a much larger class of densities than just interacting v -representable ones [40–42], 3) the functional should be convex and differentiable [39, 43, 44] (important for setting up a Kohn-Sham scheme, see next section). These concerns are largely remedied by the Levy-Lieb [45, 46] and Lieb [42] functionals.

At this stage, DFT is impractical, because the definition of the functional $F_{\text{HK}}[n]$ requires knowledge of the many-body wavefunction that corresponds to some trial density n . Kohn and Sham [47] developed a practical framework that permits use of the theorems

of ground-state DFT to improve matters beyond the Hartree and Hartree-Fock levels of theory.

2.3.1 Exchange-Correlation Potential

Kohn-Sham theory is a non-interacting single-particle theory defined to reproduce the same ground-state density and energy as the exact many-body theory [47]. Consider the following ‘Kohn-Sham’ objective functional

$$E_{\text{KS}}[n] = E_{\text{H}}[n] + E_{\text{xc}}[n], \quad (2.21)$$

where $E_{\text{xc}}[n]$ is an unknown and thus-far undefined functional of the density, and $E_{\text{H}}[n] = \langle \Psi_{\text{H}}[n] | H[v_{\text{ext}}] | \Psi_{\text{H}}[n] \rangle$ is the Hartree energy functional, see Eqs. (2.8) and (2.9) (though its analytic form in terms of the density is unknown due to the kinetic energy operator). The *exchange-correlation (xc) functional* $E_{\text{xc}}[n]$ is now defined such that, when $E_{\text{KS}}[n]$ is evaluated at the exact ground-state density n_0 , the Kohn-Sham and exact ground-state energies are equal,

$$\begin{aligned} F[n_0] + \int n_0(x)v_{\text{ext}}(x) dx &= E_{\text{KS}}[n_0], \\ &= E_{\text{H}}[n_0] + E_{\text{xc}}[n_0], \end{aligned} \quad (2.22)$$

where $F[n]$ is some valid density functional (e.g. Hohenberg-Kohn or Levy-Lieb). We can see that $E_{\text{xc}}[n]$ is guaranteed to exist *and* be a functional of the density by the Hohenberg-Kohn theorems. Hence, the Kohn-Sham functional has been defined such that it manifestly produces the exact ground-state energy and density as a variational minimum, but does so by requiring exact knowledge of $F[n]$.

A corresponding set of Euler-Lagrange equations can be derived by taking the functional derivative of the Kohn-Sham objective functional with respect to the density, yielding the Kohn-Sham equations:

$$\left(-\frac{1}{2}\nabla^2 + v_{\text{H}} + v_{\text{ext}} + v_{\text{xc}} \right) \phi_i(x) = \epsilon_i \phi_i(x), \quad (2.23)$$

$$n(x) = \sum_{i=1}^N |\phi_i(x)|^2, \quad (2.24)$$

$$v_{\text{xc}}[n](x) := \frac{\delta E_{\text{xc}}[n]}{\delta n}. \quad (2.25)$$

An exact form of the exchange-correlation potential $v_{\text{xc}}[n]$ can be derived as such,

$$v_{\text{xc}}[n] = \frac{\delta F[n]}{\delta n} - v_{\text{H}}[n] - \frac{\delta T_0[n]}{\delta n}, \quad (2.26)$$

where $T_0[n]$ is the single-particle kinetic energy functional (an implicit functional of the density). We are left with a set of equations that look similar to the Hartree equations, but with the inclusion of an exchange-correlation potential, which acts to guide the positions of the non-interacting particles towards the configuration of the interacting electrons. In other words, $v_{\text{xc}}[n]$ is a local, multiplicative potential that must account

for the influence of exchange anti-symmetry *and* Coulomb interactions in the particle density. Notes on solving these equations are given in Chapter 3. Chapter 4 onward involves realising the inverse map: find the Kohn-Sham potential $v_{\text{KS}} = v_{\text{H}} + v_{\text{ext}} + v_{\text{xc}}$ that ensures the non-interacting particles possess a given density $n(x)$ (whether such a potential exists is known as the non-interacting v -representability problem [39]).

The Kohn-Sham scheme as presented remains impractical, as computation of the exchange-correlation energy functional requires knowledge of $F[n]$, which in turn requires the exact many-body wavefunction. However, it is possible to develop some approximate form for the functional $E_{\text{xc}}[n]$, and progress in DFT largely revolves around the development and application of such approximate functional forms.

The most celebrated approximate xc functional is the local density approximation⁵ (LDA). The LDA aims at obtaining an analytic representation of the exact xc functional for the HEG in hopes that this provides a reasonable account of exchange and correlation when used in the context of inhomogeneous systems. The exchange part of the xc functional has a known analytic form [48],

$$E_{\text{x}}[n] = -\frac{3}{4} \left(\frac{3}{\pi} \right)^{\frac{1}{3}} \int n(x)^{\frac{4}{3}} dx, \quad (2.27)$$

whereas the correlation part uses some well-motivated ansatz functional form (obeying certain known limits) which is then parameterised based on Monte Carlo calculations for the HEG [35, 49]. The successes of the LDA [17, 19, 36, 50] can be attributed to a variety of sources. First, the LDA is accurate for slowly-varying densities, and many practical cases of interest fall within this scope. Furthermore, even outside of this domain, the LDA is able to perform well due to it satisfying various exact conditions, such as those relating to the xc hole [50], see Section 2.5.4. Importantly, though, there also exist many cases for which the LDA is not sufficient, for example, materials that involve long-range (e.g. van der Waals) correlations [18, 19, 51], see Chapter 6.

At present time, hundreds of xc functionals exist⁶ with differing domains of application, and differing design philosophies (e.g. functionals can be grouped according to degree of empiricism) [53]. For example, it can be shown that the LDA is the lowest order term in a Taylor expansion of the xc energy functional [39], i.e. the exact xc energy per electron can be expressed as a functional of the density and its first, second, and so on, derivatives. Generalised gradient approximations (GGAs) utilise this fact, and parameterise the xc functional not just with respect to the density, but also with respect to its first derivative [54]. At some additional computational cost, hybrid functionals include the non-local exchange potential from Hartree-Fock theory alongside a local xc potential [55]. To give a final example: as discussed in Chapter 6, advances such as the adiabatic-connection fluctuation-dissipation theorem (ACFDT) (see Section 2.5) have been central in the development of functionals that are able to treat long-range correlations [56, 57].

⁵The term ‘local’ is used because the xc potential at x depends on the density $n(x)$ at x and nowhere else, which is not the case in general.

⁶See also the `libxc` package [52] which provides a modularised implementation of an assortment of functionals, and is used by a number of the most popular DFT codes.

2.3.2 Implementation

A great deal of effort goes into transforming the Kohn-Sham equations with some approximate xc functional into a form suitable to be implemented on a computer. The central approximation involved is called the *finite-basis approximation*, or the *Galerkin approximation* [58]. The orbitals $\phi_i(x)$ are continuous functions of a continuous three-dimensional variable, x . These functions comprise vectors in an infinite-dimensional vector space spanned by a complete (infinite) basis χ_μ . Provided this basis does indeed span the space, the orbitals can be expressed exactly,

$$\phi_i(x) = \sum_{\mu=1}^{\infty} \alpha_{i\mu} \chi_\mu(x). \quad (2.28)$$

Once the basis is specified, the equations can be rearranged and solved for the infinity of coefficients in the basis, $\alpha_{\mu i}$. In practice, one must truncate the basis such that it is no longer complete and instead captures only the most relevant regions of the otherwise infinite Hilbert space⁷,

$$\phi_i(x) \approx \sum_{\mu=1}^{N_b} \alpha_{i\mu} \chi_\mu(x). \quad (2.29)$$

The characteristic size of the basis N_b in a given practical calculation will depend mostly on the choice of basis functions. The choice of basis functions also determines the algorithmic structure of the implementation, the implementations' scaling with respect to its internal variables, and so on.

In Chapter 3, the Kohn-Sham equations are solved in the context of periodic systems (the three-dimensional unit of repetition is called the unit cell) using Bloch's theorem, see Ref. [36] for a review of DFT in the solid state. The basis set used is the collection of plane waves

$$\chi_G(x) = e^{iG \cdot x} \quad (2.30)$$

where $\{G\}$ is the discrete set of plane-wave frequencies with the same periodicity as the unit cell [61, 62]. Plane waves are a simple and general purpose basis set: the basis does not need adaptation to a specific input, and the accuracy of a calculation can be monotonically improved by increasing the number of plane waves. However, the cost of these conveniences is that the size of the basis N_b is large when compared to other kinds of basis set that forego the advantages of the plane-wave basis. Of particular importance in plane-wave codes is the use of pseudopotentials [63], which can account for the effect of core electrons and thus reduce the computational cost of a given N -electron calculation.

In Chapter 4 onward, the real-space basis is used,

$$\chi_{y_n}(x) = \delta(x - y_n), \quad (2.31)$$

⁷Recent articles, e.g. Refs. [59, 60], consider in detail the circumstances under which the eigenvalues of a discretised operator are able to approximate the eigenvalues (or spectrum) of the corresponding continuous operator.

for a set of discrete positions $\{y_n\}$, and the Kohn-Sham equations are solved for finite (non-periodic) systems. This is not an efficient basis set, but it is acceptable in the present context as the goal of the work is to examine model one-dimensional systems. Needless to say, these are two vastly different approaches, and the relevant implementation details for each approach are presented in the corresponding chapters.

2.4 Time-Dependent Density Functional Theory

Until now, we have considered quantum states that are in the lowest energy eigenstate of some Hamiltonian. Provided the external potential is time independent, observables corresponding to such states do not evolve in time, which can be seen from the equation that governs time dependence in quantum mechanics, the time-dependent Schrödinger equation,

$$i\frac{\partial}{\partial t}|\Psi(t)\rangle = H|\Psi(t)\rangle. \quad (2.32)$$

The inclusion of time dependence allows time-dependent phenomena to be studied, such as the transport of electrons through molecules [64], and the response of atoms when subject to strong laser fields [65]. In this work, we are mostly concerned with time dependence as a route towards obtaining certain optical properties [7, 37], see Section 2.4.4.

The practical difficulties with exact time-dependent quantum mechanics are inherited from the difficulties involved in ground-state quantum mechanics. A time-dependent extension of DFT is now given, which ultimately yields a tractable scheme for practical time-dependent calculations.

The time-dependent Schrödinger equation is a first-order differential equation in time, meaning that given some initial state $|\Psi(0)\rangle$ and some time-dependent Hamiltonian $H(t)$, the state will undergo deterministic unitary time evolution $|\Psi(t)\rangle$. In essence, the time-dependent Schrödinger equation is used to define a map,

$$T : \{\mathcal{V}, \psi_0\} \rightarrow \psi \rightarrow \mathcal{N}, \quad (2.33)$$

where $v_{\text{ext}}(x, t) \in \mathcal{V}$ is some permissible⁸ time-dependent external potential, $|\Psi(0)\rangle \in \psi_0$ is some initial state, $|\Psi(t)\rangle \in \psi$ is the time-evolved state, and $n(x, t) \in \mathcal{N}$ is the corresponding time-dependent v -representable density. The Runge-Gross theorem [66] establishes that the inverse map,

$$T^{-1} : \{\mathcal{N}, \psi_0\} \rightarrow \psi \rightarrow \mathcal{V}, \quad (2.34)$$

is one-to-one for systems interacting under the Coulomb potential, alongside other mild constraints. Therefore, given some density $n(x, t)$ and the initial state from which it came $\Psi_0 := |\Psi(0)\rangle$, observables can be calculated as functionals of these quantities,

$$o[n, \Psi_0](t) = \langle \Psi[n, \Psi_0](t) | \hat{O}(t) | \Psi[n, \Psi_0](t) \rangle. \quad (2.35)$$

⁸In the original paper by Runge and Gross [66], for example, the potential was required to be analytic in time, and is yet again defined up to a time-dependent constant shift.

In the case that the initial state Ψ_0 is a ground state, then $\Psi_0 = \Psi_0[n]$, and time-dependent observables⁹ can be calculated from knowledge of the density alone.

The Runge-Gross theorem was generalised by van Leeuwen [67] to include interaction terms of the form $v_{\text{int}}(|x - x'|)$ in the Hamiltonian. Therefore, the special case $v_{\text{int}} = 0$ allows us to establish a connection with single-particle systems as in ground-state DFT, and thereby set up a time-dependent Kohn-Sham scheme.

2.4.1 Time-Dependent Exchange-Correlation Potential

The Runge-Gross and van Leeuwen theorems are used to construct the unique time-dependent Kohn-Sham potential,

$$v_{\text{KS}}(x, t) = v_{\text{KS}}[n, \Psi_0, \Phi_0](x, t), \quad (2.36)$$

which is determined as such: a non-interacting system is defined that begins in some state Φ_0 constructed from N orbitals $\{\phi_i(x, t = 0)\}$, and this state proceeds to evolve in time according to the time-dependent Kohn-Sham equations,

$$\left(-\frac{\nabla^2}{2} + v_{\text{KS}}(x, t)\right) \phi_i(x, t) = i \frac{\partial}{\partial t} \phi_i(x, t). \quad (2.37)$$

The Kohn-Sham potential v_{KS} is defined such that the resultant non-interacting density $n(x, t)$ is equal to the density of a time-dependent interacting system evolving from some initial state Ψ_0 under the influence of an external potential $v_{\text{ext}}(x, t)$. As discussed, provided that the initial many-body state Ψ_0 is a ground state, then the initial non-interacting state will also be a ground state, and therefore the time-dependent Kohn-Sham potential becomes a functional of the density alone, $v_{\text{KS}} = v_{\text{KS}}[n]$. The issue regarding the existence of such a potential is called the (non-)interacting time-dependent v -representability problem [68–71], which is discussed in Chapter 4.

The time-dependent xc potential is defined as one might expect,

$$v_{\text{KS}} = v_{\text{ext}} + v_{\text{H}} + v_{\text{xc}}, \quad (2.38)$$

and satisfies various exact conditions [7, 72], some of which are of importance to this thesis. It follows from the conservation of probability that the xc potential (and Hartree potential) is an *internal* force between the particles, i.e. the exact time-dependent xc potential should not exert a net force on a finite system bounded in some region¹⁰. This

⁹See Ref. [37] for a nice discussion on the different kinds of time-dependent observables that are of practical interest. This includes, in the most simple case, the density, among other observables such as the dipole moment (related to optical properties).

¹⁰The proof of this statement utilises a central object in time-dependent DFT: the current density $j(x, t)$. The current density is a vector quantity that describes the velocities of electrons at x in each Cartesian direction. In this sense, the electron density can be interpreted as a fluid, obeying a continuity equation [7]. Modern reformulations of time-dependent DFT treat the current density as primary [73], which circumvents some peculiar properties of the xc potential as a functional of the density.

zero-force theorem is expressed in mathematical form as

$$\int_{\Sigma} n(x, t) \nabla v_{\text{xc}}(x, t) dx = 0; \quad (2.39)$$

the linear response version of this condition assists in verifying numerical methodologies. Moreover, the harmonic potential theorem [7] states that particles confined to a harmonic potential should undergo a time-dependent rigid translation when a perturbing potential of the form $\delta v_{\text{ext}}(x, t) \propto \pm x$ is applied. This condition also assists in ensuring the correctness of algorithms when attempting to realise the inverse map in Chapter 4.

Propagating the time-dependent Kohn-Sham state as above is often called real-time time-dependent DFT, see Refs. [74, 75] for popular implementations. The time-dependent extension of the LDA xc functional is called the adiabatic LDA (ALDA), which evaluates the LDA xc potential $v_{\text{xc}}^{\text{LDA}}[n](x)$ with some trial $n := n(x, t)$ at each time t to generate the xc potential at time t . Therefore, the xc potential now depends on the density at (x, t) , meaning it is local in space and time. The term adiabatic is used because we have made the assumption that at each time t the time-dependent system can be approximated with the ground-state LDA xc potential evaluated at the density $n(x, t)$ (i.e. the time-dependence follows a succession of ground-state LDA approximations), see Refs. [76, 77] for applications of the ALDA in practice. In general, the xc potential depends on the density at all prior times $0 < t' \leq t$ – this is called ‘memory dependence’ [78–80], and it is closely related to the initial-state dependence in Eq. (2.36). The inclusion of memory dependence can be critical depending on the application at hand, as Chapter 5 demonstrates in the context of linear response time-dependent DFT.

2.4.2 Linear Response Function

The density-density linear response function of a many-body quantum system can be used to extract a great deal of excited-state information about the system, for example, its optical transition probabilities and transition energies when subject to incident light [10]. The interacting density-density response function $\chi(x, x', t - t')$ describes the first-order change in the density due to a perturbation in the external potential [7, 19]:

$$\chi[n](x, x', t - t') = \left. \frac{\delta n(x, t)}{\delta v_{\text{ext}}(x', t')} \right|_{v_0}. \quad (2.40)$$

In other words, suppose a system begins in its ground state with external potential v_{ext} , and then some perturbation is applied: $v_{\text{ext}} \rightarrow v_{\text{ext}} + \delta v_{\text{ext}}$. The predominant (first-order) response in the density is calculated as such,

$$\delta n(x, t) = \iint \chi[n](x, x', t - t') \delta v_{\text{ext}}(x', t') dx' dt', \quad (2.41)$$

where $\chi[n]$ is a functional of the ground-state density due to the Hohenberg-Kohn and Runge-Gross theorems. Since the linear response function is invariant under time translations (see [7]), then it is local in the frequency domain, i.e. a perturbation $\delta v_{\text{ext}}(x, \omega)$ oscillating in time with frequency ω induces a (linear) response in the density $\delta n(x, \omega)$ also oscillating in time with frequency ω .

The response function satisfies certain exact conditions that relate to various concerns in Chapters 5 and 6. First, suppose (although this is not always the case) that we desire the response function to be causal,

$$\chi(x, x', t - t') = 0 \text{ for } t < t', \quad (2.42)$$

then its Fourier transform $\chi(x, x', \omega)$, defined using

$$\chi(x, x', t - t') = \int_{-\infty}^{\infty} e^{-i\omega(t-t')} \chi(x, x', \omega) \frac{d\omega}{2\pi}, \quad (2.43)$$

must possess no poles in the upper half of the complex ω -plane, i.e. $\text{Im}(\omega) \geq 0$ (otherwise integral would be non-zero due to the residue theorem). This ultimately establishes a relationship between $\text{Im}(\omega)$ and $\text{Re}(\omega)$ called the Kramers-Kronig relations, meaning that the response function can be constructed from just the imaginary (or real) component as such,

$$\chi(x, x', \omega) = \int_{-\infty}^{\infty} \frac{\text{Im}(\chi(x, x', \omega'))}{\omega' - \omega - i\eta} \frac{d\omega'}{\pi}. \quad (2.44)$$

The positive infinitesimal η indicates that the integral runs just below the real axis. As Ref. [81] discusses, the imaginary component of the response function is often linked to absorption/dissipation (time-dependent changes of energy) as it is not invariant $t \leftrightarrow -t$. The response function is also even, $\chi(-\omega) = \chi(\omega)$, because it is real in the time domain, and symmetric under interchange of spatial coordinates: $\chi(x, x', \omega) = \chi(x', x, \omega)$.

A convenient expression for the causal response function can be derived utilising the complete set of eigenstates of the Hamiltonian H whose ground state $|\Psi_0\rangle$ is the state in which our system begins,

$$\chi(x, x', \omega) = \lim_{\eta \rightarrow 0} \sum_{n=1}^{\infty} \frac{\langle \Psi_0 | \hat{n}(x) | \Psi_n \rangle \langle \Psi_n | \hat{n}(x') | \Psi_0 \rangle}{\omega - \Omega_n + i\eta} - \frac{\langle \Psi_0 | \hat{n}(x') | \Psi_n \rangle \langle \Psi_n | \hat{n}(x) | \Psi_0 \rangle}{\omega + \Omega_n + i\eta}, \quad (2.45)$$

where $\Omega_n = E_n - E_0$ is the n^{th} excitation energy of the interacting Hamiltonian, and $\hat{n}(x)$ is the density operator,

$$\hat{n}(x) = \sum_{i=1}^N \delta(x - x_i). \quad (2.46)$$

The infinitesimal η indicates that, for $\omega > 0$, the response function has poles just below the real axis at each discrete excitation energy. (Eq. (2.45) is called the Lehmann representation of the response function.) The response function is discussed in more detail in Chapter 5.

The Kohn-Sham (non-interacting) response function $\chi_0(x, x', \omega)$ describes how a system of non-interacting particles in the background potential v_{KS} responds at linear order due to the perturbation $v_{\text{KS}} \rightarrow v_{\text{KS}} + \delta v_{\text{KS}}$,

$$\delta n(x, t) = \iint \chi_0[n](x, x', t - t') \delta v_{\text{KS}}(x', t') dx' dt'. \quad (2.47)$$

The spectral representation of the non-interacting response function is as follows,

$$\chi_0(x, x', \omega) = \lim_{\eta \rightarrow 0} \sum_{i=1}^N \sum_{j=N+1}^{\infty} \frac{\phi_i^*(x) \phi_j(x) \phi_j^*(x') \phi_i(x')}{\omega + \Omega_{ij}^0 + i\eta} - \frac{\phi_i(x) \phi_j^*(x) \phi_j(x') \phi_i^*(x')}{\omega - \Omega_{ij}^0 + i\eta}, \quad (2.48)$$

which has poles at the non-interacting transition energies, $\Omega_{ij}^0 = \epsilon_i - \epsilon_j$. Formally speaking, these energy differences do not relate to interacting transitions (including the band gap, i.e. the smallest energy difference), however, in certain contexts it is possible to assign meaning to the Kohn-Sham energy differences [82], and indeed Chapter 5 provides one class of examples where the non-interacting transitions reproduce the interacting transitions.

A static non-interacting response function $\chi_0(x, x')$ can be derived from perturbation theory [83, 84], which describes the response due to a perturbation in the potential upon solution of the Kohn-Sham equations. The density with which the potential is constructed need not be self-consistent in order to determine this response, knowledge of which can help guide an iterative procedure (discrete iterations replace continuous time here) towards the self-consistent solution, see Chapter 3.

2.4.3 Exchange-Correlation Kernel

The ground-state density of the non-interacting Kohn-Sham system, despite being equal to the ground-state density of a given interacting system by definition, responds differently at linear order to perturbations in the external potential. Linear response time-dependent DFT provides a map between the responses of these two systems using the xc kernel,

$$f_{\text{xc}}[n](x, x', t - t') = \left. \frac{\delta v_{\text{xc}}(x, t)}{\delta n(x', t')} \right|_{n_0}, \quad (2.49)$$

i.e. the first-order change in the xc potential due to a perturbation in the density. This map is often referred to as the Dyson equation of linear response time-dependent DFT,

$$\chi(x, x', \omega) = \chi_0(x, x', \omega) + \iint \chi_0(x, x'', \omega) \{f_{\text{H}}(x'', x''') + f_{\text{xc}}(x'', x''', \omega)\} \chi(x''', x', \omega) dx'' dx''', \quad (2.50)$$

where $f_{\text{H}} = \delta v_{\text{H}} / \delta n$ is the Hartree kernel (the electron-electron interaction), see Ref. [7]. The form of the Dyson equation is fixed with the requirement that the Kohn-Sham and interacting systems are equal in their response to perturbations. The uniqueness of this map is guaranteed by the Runge-Gross and van Leeuwen theorems, and the definition of f_{xc} in Eq. (2.49) demonstrates that, like the ground-state xc potential, f_{xc} is a functional of the density.

The Dyson equation provides an alternate (but equivalent) definition of the xc kernel,

$$f_{\text{xc}}(x, x', \omega) = \chi_0^{-1}(x, x', \omega) - \chi^{-1}(x, x', \omega) - f_{\text{H}}(x, x'). \quad (2.51)$$

The inverse response function (denoted with superscript -1) can be understood as the object that determines the perturbing external potential δv_{ext} that generates a *given* perturbation in the density,

$$\delta v_{\text{ext}}(x, \omega) = \int \chi^{-1}[n](x, x', \omega) \delta n(x', \omega) dx'. \quad (2.52)$$

The existence of the xc kernel, and thus also the existence of the map between the non-interacting and interacting responses, depends on the invertibility of χ and χ_0 , also known as the (non-)interacting linear response v -representability problem. That is, perturbations δn in the non-interacting system that can't be obtained with some perturbing potential δv gives the Dyson equation (xc kernel) an infinitely difficult job in establishing the map to the interacting system (and vice-versa). Indeed, such occurrences are common in both the interacting and non-interacting systems, and can be of practical importance, as Chapter 5 discusses in more detail.

A practical linear response time-dependent DFT calculation thus involves specifying some approximate xc kernel functional $f_{\text{xc}}[n]$ and some xc potential functional $v_{\text{xc}}[n]$ (the former often does not need to be the functional derivative of the latter) which are used to extract properties from the resultant approximate interacting response function χ . Approximate f_{xc} functionals are discussed in more detail in Chapters 5 and 6. In particular, a simple approximation used for comparative purposes in this thesis involves taking the functional derivative of an LDA xc potential to obtain the ALDA xc kernel: $f_{\text{xc}}^{\text{ALDA}}[n] = \delta v_{\text{xc}}^{\text{LDA}}[n] / \delta n$. Moreover, setting the xc kernel equal to zero (the so-called random phase approximation) and otherwise utilising the framework of linear response time-dependent DFT has close ties to many-body perturbation theories centered around the Green's function [85].

2.4.4 Optical Properties

A celebrated aspect of quantum mechanics is the account that it provides for the discrete emission (or absorption) spectrum of solids, molecules, and atoms. That is, a sample in its ground state subject to light of all frequencies absorbs only certain allowed frequencies ω of light, which interacts with the sample as to induce the transition $|\Psi_0\rangle \rightarrow |\Psi_n\rangle$ where $E_n - E_0 \approx \hbar\omega$. The spectral series of atomic hydrogen provides a typical example of this phenomenon, wherein the observed absorptions are those that correspond to the excited states of an appropriate Hamiltonian operator (improved Hamiltonians give rise to the so-called fine, hyperfine, and so on, structure).

In the present context, light can be regarded as a classical electromagnetic wave (a large collection of photons) which couples to the electrons through a vector potential in the Schrödinger equation, see Ref. [7]. Provided also that the wavelength of the light is much larger than the characteristic size of the sample, then the spatial variation of the light wave across the sample takes a simple linear form (the dipole approximation). These considerations lead to electrons in the ground state of some Hamiltonian H_0 coupling to light through the following term,

$$\delta v_{\text{ext}}(x, t) = \mathcal{E}x \cos(\omega t), \quad (2.53)$$

where \mathcal{E} is the strength of the electromagnetic field. In more than one spatial dimension, x contains the polarisation axis of the classical light.

Time-dependent perturbation theory provides an expression for the transitions per unit time that occur when a sample H_0 in its ground state is perturbed using Eq. (2.53),

$$\Gamma(|\Psi_0\rangle \rightarrow |\Psi_n\rangle) = 2\pi\mathcal{E}^2|\langle\Psi_0|\hat{x}|\Psi_n\rangle|^2\delta(E_n - E_0 - \omega), \quad (2.54)$$

otherwise known as Fermi's golden rule. The above expression includes conservation of energy through the delta function, and also demonstrates that the rate of transition is proportional to the dipole matrix elements, $\langle\Psi_0|\hat{x}|\Psi_n\rangle$, i.e. the oscillator strengths. The *optical absorption spectrum* can be derived with reference to Fermi's golden rule, which involves, in large part, rearranging Eq. (2.54) in terms of the response function to give,

$$\sigma(\omega) = -4\pi\omega \iint \text{Im}(\chi(x, x', \omega))xx' dx dx', \quad (2.55)$$

thus relating the linear response function to optical transition energies and transition rates. A host of experimental spectroscopic techniques are available to measure, in some form or another, information contained within the optical absorption spectrum (see, for instance, angle-resolved photoemission spectroscopy [86]). Since linear response time-dependent DFT gives access to an approximate interacting response function χ , the approximations of linear response time-dependent DFT can be used in practice to predict or describe the outcomes of spectroscopic experiments, see Refs. [87, 88]. Chapter 5 uses the optical absorption spectrum to analyse the relationship between various aspects of the exact xc kernel and certain interacting excitations.

2.5 The Adiabatic-Connection Fluctuation-Dissipation Theorem

The *adiabatic-connection fluctuation-dissipation theorem* (ACFDT) [7, 17–19, 89, 90] formalism is a distinctive, powerful approach to calculating ground-state energies of molecular and solid-state systems. The origin of the ACFDT in the context of inhomogenous systems dates back around five decades to the series of articles given in Refs. [91–93]. Since then, a number of resources have covered the derivation of the ACFDT [7, 9, 17, 19, 27, 85, 94, 95], a brief review is given here.

2.5.1 Adiabatic Connection

The adiabatic connection establishes a link between the interacting many-body system and its corresponding non-interacting Kohn-Sham system, ultimately leading to an alternate expression for the xc energy. Towards this end, a one-parameter family of many-body Hamiltonians is defined,

$$H(\lambda) = \hat{T} + \lambda\hat{v}_{ee} + \hat{v}_{\text{ext}} + \hat{v}_{\text{dxc}}(\lambda), \quad (2.56)$$

such that the *deus ex machina* potential [96] $\hat{v}_{\text{dxm}}(\lambda)$ is the unique [38] potential that ensures the ground-state density at all values of $\lambda \in [0, 1]$ is equal to the ground-state electron density at $\lambda = 1$, labeled $n(x)$ – this is the adiabatic connection. Hence, $\hat{v}_{\text{dxm}}(\lambda = 1) = 0$ and $\hat{v}_{\text{dxm}}(\lambda = 0) = \hat{v}_{\text{H}} + \hat{v}_{\text{xc}}$ where \hat{v}_{H} is the Hartree potential. The λ -interacting ground state of $H(\lambda)$ is denoted $|\Psi^\lambda\rangle$, allowing the total energy to be expressed as such,

$$\begin{aligned} E &= \langle \Psi^{\lambda=1} | H(1) | \Psi^{\lambda=1} \rangle \\ &= \langle \Psi^{\lambda=0} | H(0) | \Psi^{\lambda=0} \rangle \\ &= T_0 + E_{\text{H}} + E_{\text{ext}} + E_{\text{xc}}, \end{aligned} \quad (2.57)$$

where the latter two formulae constitute the conventional definition of the Kohn-Sham system [47] ($|\Psi^{\lambda=0}\rangle$ represents the Kohn-Sham Slater determinant ground state).

Rearrangement of the above expressions yields an alternate form for the Hxc energy $E_{\text{Hxc}} = E_{\text{H}} + E_{\text{xc}}$,

$$\begin{aligned} E_{\text{Hxc}} &= \langle \Psi^{\lambda=1} | H(1) | \Psi^{\lambda=1} \rangle - \langle \Psi^{\lambda=0} | H(0) | \Psi^{\lambda=0} \rangle \\ &\quad - \langle \Psi^{\lambda=1} | \hat{v}_{\text{dxm}}(1) | \Psi^{\lambda=1} \rangle + \langle \Psi^{\lambda=0} | \hat{v}_{\text{dxm}}(0) | \Psi^{\lambda=0} \rangle, \end{aligned}$$

which becomes, upon use of the fundamental theorem of calculus and the Hellmann-Feynmann theorem,

$$\begin{aligned} E_{\text{Hxc}} &= \int_0^1 \frac{d}{d\lambda} \left(\langle \Psi^\lambda | H(\lambda) | \Psi^\lambda \rangle - \langle \Psi^\lambda | \hat{v}_{\text{dxm}}(\lambda) | \Psi^\lambda \rangle \right) d\lambda \\ &= \int_0^1 \langle \Psi^\lambda | \hat{v}_{\text{ee}} | \Psi^\lambda \rangle d\lambda. \end{aligned} \quad (2.58)$$

The expression in Eq. (2.58) contrasts with the conventional one Eq. (2.57) as it does not involve the kinetic operator at the seemingly steep price of having to know the xc *potential* energy [17],

$$U_{\text{xc}}(\lambda) = \langle \Psi^\lambda | \lambda \hat{v}_{\text{ee}} | \Psi^\lambda \rangle - \lambda E_{\text{H}}, \quad (2.59)$$

at each value of λ along the adiabatic connection¹¹.

2.5.2 Fluctuation-Dissipation Theorem

Knowledge of the challenging expectation value in Eq. (2.58) is tantamount to knowledge of the static (equal-time) two-point correlator $\langle \Psi^\lambda | \hat{n}(x) \hat{n}(x') | \Psi^\lambda \rangle$, which describes quantum statistical *fluctuations* in the density inherent to the state $|\Psi^\lambda\rangle$ [10]. To see this, observe that the expectation value of the (two-body) Coulomb potential operator can be written as such,

$$\langle \Psi^\lambda | \hat{v}_{\text{ee}} | \Psi^\lambda \rangle = \frac{1}{2} \iint \frac{P^\lambda(x, x')}{|x - x'|} dx dx', \quad (2.60)$$

¹¹The λ -dependent xc energy (and xc kernel) is zero when $\lambda = 0$ because the particles are non-interacting, and so no work is required in mapping this system onto a non-interacting system.

with pair density,

$$P^\lambda(x, x') = \langle \Psi^\lambda | \hat{n}(x) \hat{n}(x') | \Psi^\lambda \rangle \quad (2.61)$$

$$= N(N-1) \int |\Psi^\lambda(x, x', x_3, x_4, \dots, x_N)|^2 \prod_{i=3}^N dx_i. \quad (2.62)$$

The pair density is the probability that, given a collection of λ -interacting electrons in state $|\Psi^\lambda\rangle$, an electron can be found at position x alongside a second electron at x' (i.e. it expresses a two-body correlation between the positions of a given pair of electrons). This observation is important for the definition of the so-called xc hole in DFT, although here it suffices to observe that this object relates the static density-density correlation function to the xc energy through Eq. (2.58).

The fluctuation-dissipation theorem [97] provides a relationship between the response of a system to these spontaneous *internal* changes (fluctuations) in its density, and the response of that same system to *external* perturbations in its density. The latter is described with the density-density linear response function $\chi^\lambda[n](x, x', \omega)$, i.e. the first-order change in the density due to a perturbation in the external potential *within a system of λ -interacting particles* described by $H(\lambda)$ in Eq. (2.56). In the present context, the fluctuation-dissipation theorem takes the form [98]

$$\langle \Psi^\lambda | \hat{n}(x) \hat{n}(x') | \Psi^\lambda \rangle = n(x)n(x') - \frac{2}{\pi} \int_0^\infty \chi^\lambda(x, x', i\omega) d\omega,$$

thus connecting the ground-state xc energy with linear response theory. The peculiar relationship of the ground-state energy to *time-dependent* DFT is now somewhat clearer: the response function is averaged over all frequencies, yielding a static quantity.

2.5.3 Total Energy Functional

The above derivation outlines an exact reformulation of conventional Kohn-Sham DFT [17], meaning the total energy functional

$$E^{\text{ACFD}}[n] := T_0[n] + E_{\text{ext}}[n] + E_{\text{H}}[n] + E_{\text{xc}}^{\text{ACFD}}[n] \quad (2.63)$$

has the correct minimum, i.e. the exact ground-state energy, and this minimum is attained at the interacting ground-state density. Having performed the rearrangements given from the fluctuation-dissipation theorem, the ACFDT xc energy functional $E_{\text{xc}}^{\text{ACFD}}[n] = E_{\text{x}}^{\text{ACFD}}[n] + E_{\text{c}}^{\text{ACFD}}[n]$ becomes

$$E_{\text{x}}^{\text{ACFD}}[n] = E_{\text{x}}[\{|\phi_i[n]\rangle\}], \quad (2.64)$$

$$E_{\text{c}}^{\text{ACFD}}[n] = -\frac{1}{2\pi} \int_0^1 d\lambda \int_0^\infty d\omega \iint dx dx' v_{\text{ee}}(x, x') [\chi^\lambda[n](x, x', i\omega) - \chi_0[n](x, x', i\omega)], \quad (2.65)$$

where $E_{\text{x}}[\{|\phi_i\rangle\}]$ is the exact-exchange functional evaluated at the Kohn-Sham orbitals $\{|\phi_i\rangle\}$, and $\chi_0[n] = \delta n / \delta v_{\text{KS}}$ is the associated response function of the Kohn-Sham system. These are the central expressions around which Chapter 6 is based. Namely,

we evaluate the total energy functional Eq. (2.63) with its exact definition, and then introduce approximations into the functional through $f_{\text{xc}}^\lambda[n]$ (i.e. $\chi^\lambda[n]$).

As discussed, linear response time-dependent DFT establishes a unique map [66, 67] from a tractable non-interacting (Kohn-Sham) response function $\chi_0[n]$ to an otherwise intractable λ -interacting response function $\chi^\lambda[n]$ through the λ -dependent xc kernel $f_{\text{xc}}^\lambda[n](x, x', \omega)$, i.e. the first-order change in the λ -dependent xc potential due to a perturbation in the density oscillating in time with frequency ω . The λ -dependent extension of the Dyson equation is,

$$\chi^\lambda[n] = \chi_0[n] + \chi_0[n] * (\lambda f_{\text{H}} + f_{\text{xc}}^\lambda[n]) * \chi[n], \quad (2.66)$$

where $A * B = \int A(x, x')B(x', x'')dx'$.

The xc kernel $f_{\text{xc}}^\lambda[n]$ is the central subject of approximation in linear response time-dependent DFT [7], and therefore the principal ingredients in an ACFDT total energy calculation are an approximation to the xc kernel functional $f_{\text{xc}}^\lambda[n]$ together with a prescription for determining the density n at which to evaluate the ACFDT total energy functional $E^{\text{ACFDT}}[n]$. In regard to the latter concern: there are two main approaches that are applied in practice to determine this density. The first is the most common, and is often referred to as a one-shot ACFDT calculation, wherein the density n is obtained from a self-consistent solution of the ground-state Kohn-Sham equations with an approximate xc potential $v_{\text{xc}}[n]$ (see Ref. [89] for a recent review). The second, often referred to as a self-consistent ACFDT calculation [99–105], solves

$$E_0 = \min_n E^{\text{ACFDT}}[n], \quad (2.67)$$

where the equations that yield a stationary (presumed to be minimising) density are known [99, 100, 106, 107]. In the instance that both $v_{\text{xc}}[n]$ and $f_{\text{xc}}^\lambda[n]$ are exact, the output of a one-shot and self-consistent ACFDT calculation coincide, however this is not true when approximations are involved, as we shall explore in Chapter 6.

In practice, an approximate xc kernel functional $f_{\text{xc}}^\lambda[n]$ that is also parameterised with respect to λ is *not* required as it is possible to exploit a curious relationship between ground-state wavefunctions along the adiabatic connection and ground-state wavefunctions with scaled spatial coordinates, see Refs. [17, 108]. This relationship permits us to specify a conventional functional $f_{\text{xc}}[n] := f_{\text{xc}}^{\lambda=1}[n]$ and obtain its λ dependence with little-to-no additional expense. Such an observation is central to practical ACFDT calculations, although we are unable to exploit it in this thesis due to using the softened Coulomb interaction.

2.5.4 Statistical Hole and the Exchange-Correlation Hole

In wavefunction-based theories, the statistical hole $n^{\text{hole}}(x, x')$ describes the modification of the distribution $n(x)$ upon measuring a particle at x' ,

$$n(x|x') = n(x) + n^{\text{hole}}(x, x'), \quad (2.68)$$

i.e. $n(x|x')$ is the updated position distribution for the remainder of the particles. Provided the particles are uncorrelated in a statistical sense, as are sequential coin flips, the hole is equal to zero, reflecting an unchanged distribution come the next measurement. However, the mere fact that measuring a particle removes it from the system of remaining particles (similar to picking balls out of a bag¹²) introduces a statistical correlation: $n^{\text{hole}}(x, x') = -\frac{1}{N}n(x)$. That is, the new distribution $n(x|x')$ normalises to $N - 1$, but is otherwise unchanged. Correctly updating the normalisation of the conditional distribution is a general rule for both the statistical hole, and the xc hole, independent of the kinds of correlation present, i.e.

$$\int n^{\text{hole}}(x, x') dx = -1 \text{ for all } x', \quad (2.69)$$

which reflects that fact that we have removed one particle from the system through measurement.

The requirement that states be exchange antisymmetric, alongside the presence of Coulomb interactions, introduces non-trivial correlations into the hole. Exchange alone is enough to preclude two particles being measured at the same position, meaning $n^{\text{hole}}(x, x) = -n(x)$, which constitutes yet another exact condition. Furthermore, since Coulomb interactions are two-body interactions, knowledge of the exact hole is sufficient to determine the ground-state Coulomb energy of the system,

$$\langle \Psi | \hat{v}_{\text{ee}} | \Psi \rangle = \iint \frac{n(x)n(x')}{|x-x'|} dx dx' + \iint \frac{n(x)n^{\text{hole}}(x, x')}{|x-x'|} dx dx', \quad (2.70)$$

see also Eq. (2.60). We can see in the above expression that the Hartree (left-most) term emerges as the uncorrelated component of the Coulomb energy, whereas the hole captures the correction to E_{H} due to the interactions in the Hamiltonian. In this sense, the correction term can be interpreted as the system Coulomb interacting with its hole.

In single-particle theories that depend on the Hohenberg-Kohn theorems, the statistical hole loses its meaning, as it has been defined in terms of the wavefunction, see Eq. (2.62). Namely, an xc potential is introduced to ensure that the densities of the interacting and non-interacting systems are equal, but there is no such guarantee that the Kohn-Sham Slater determinant has the correct pair density $P(x, x')$. In order to restore meaning to the hole in the context of density-based theories, the so-called xc hole $n_{\text{xc}}^{\text{hole}}(x, x')$ is defined. Following Ref. [91], the xc hole is the object with which the system interacts to produce the xc energy,

$$E_{\text{xc}} = \iint \frac{n(x)n_{\text{xc}}^{\text{hole}}(x, x')}{|x-x'|} dx dx', \quad (2.71)$$

which is further specified with the help of the ACFDT expression for the xc energy. That is, comparing the above definition with Eqs. (2.58) and (2.60) yields

$$E_{\text{xc}} = \iiint \frac{n(x)n^{\text{hole}, \lambda}(x, x')}{|x-x'|} d\lambda dx dx', \quad (2.72)$$

¹²In the example of picking balls (point particles) out of a bag, the hole becomes $n^{\text{hole}}(x, x') = -\delta(x-x')$, i.e. the remaining balls cannot be measured at the same position as the original.

where $n^{\text{hole},\lambda}(x, x')$ is the traditional statistical hole of the λ -interacting system, and

$$n_{\text{xc}}^{\text{hole}}(x, x') = \int_0^1 n^{\text{hole},\lambda}(x, x') d\lambda \quad (2.73)$$

is the xc hole, otherwise called the λ -averaged hole.

The ACFDT expression also allows us to split the xc hole into an exchange hole $n_{\text{x}}^{\text{hole}}$ and a correlation hole $n_{\text{c}}^{\text{hole}}$. The exchange hole Coulomb interacts with the system in order to produce the DFT exchange energy, i.e. the exchange functional evaluated at the Kohn-Sham orbitals. The exchange hole can be interpreted as the (strictly) reduced chance of finding a particle at some position x in the vicinity of measured particle due solely to the exchange interaction, where $n_{\text{x}}^{\text{hole}}(x, x) = -n(x)$. The exchange hole integrates to negative one, meaning the correlation hole must integrate to zero. The correlation hole (which would not exist if not for Coulomb interaction in the many-body Hamiltonian) can either increase *or* decrease the chance of finding a remaining particle at some position x after measurement at x' .

The xc hole is an important concept in ground-state DFT. For example, the acclaimed Perdew, Burke, and Ernzerhof (PBE) functional [54] is parameterised such that certain exact conditions relating to the xc hole are satisfied. Furthermore, the LDA approximation is known to be an accurate predictor for the so-called on-top (local) xc hole [91], which contributes towards explaining the various successes of the LDA when used for systems whose densities are not slowly varying. Chapter 6 involves computations of the xc hole, which are used to analyse certain features in the exact xc kernel and their relation to ACFDT total energies.

Chapter 3

Computing the Self-Consistent Field in Density Functional Theory

3.1 Background

The concept of self-consistency has been prevalent across many domains of physics, typically as a characteristic requirement when one invokes a mean-field approximation. For example, Hartree theory replaces the two-body Coulomb interaction between electrically-charged quantum particles with a mean-field, the *Hartree potential*, generated by the distribution of the electric charge in the system. Each particle is influenced by the Hartree potential, which in turn alters the distribution of charge in the system. This charge distribution can then be used to construct a new Hartree potential. The Hartree potential is *self-consistent* when these two fields are the same, i.e. the potential leads to a charge distribution which gives rise to the *same* potential. In fact, this was the context in which self-consistency was first introduced,

“If the final field is the same as the initial field, the field will be called ‘self-consistent’, and the determination of self-consistent fields for various atoms is the main object of this paper.”

– D.R. Hartree (1927) [109].

This chapter concerns the need to achieve self-consistency in the context of DFT simulations of atoms, molecules and materials. Namely, we focus on computing the particle density $n(x)$ for a set of atomic species and positions within the framework of Kohn-Sham DFT. Each of the N particles in the system are influenced by an external potential v_{ext} which is uniquely defined by the species and positions of the atoms, the level of approximation employed, and more. For the purposes of this chapter, finding the ground-state

energy E in Kohn-Sham theory is viewed as a constrained minimisation problem,

$$E[v_{\text{ext}}] = \min_{\{\phi_i\}} \left\{ E_{\text{KS}}[\{\phi_i\}] \mid \phi_i \in \mathcal{H}^1(\mathbb{R}^3), \int \phi_i^*(x)\phi_j(x) dx = \delta_{ij} \text{ for } i, j \in [1, N] \right\}, \quad (3.1)$$

$$E_{\text{KS}}[\{\phi_i\}] = \sum_{i=1}^N \frac{1}{2} \int |\nabla \phi_i(x)|^2 dx + \frac{1}{2} \int \frac{n(x)n(x')}{|x-x'|} dx dx' + \int n(x)v_{\text{ext}}(x) dx + E_{\text{xc}}[n], \quad (3.2)$$

where atomic units are used, and, for now, spin degrees of freedom are omitted. The particle density $n(x)$ is defined in terms of the single-particle orbitals $\{\phi_i\}$ as such,

$$n(x) = \sum_{i=1}^N |\phi_i(x)|^2. \quad (3.3)$$

That is, one must minimise the Kohn-Sham objective functional Eq. (3.1) over a set of N orthogonal, normalisable functions $\{\phi_i\}$ whose first derivative is also normalisable, i.e. they exist in the Sobolev space $\mathcal{H}^1(\mathbb{R}^3)$.

Explicit constrained variation of the orbitals allows one to approach the optimisation problem in Eq. (3.1) directly. This can be done, for example, with a series of line searches in the direction of steepest descent of E_{KS} with respect to the orbitals [110, 111]. Alternatively, assuming differentiability [39], the associated Lagrangian problem can be formulated, and the functional derivative of the Lagrangian set to zero. This yields the Euler-Lagrange equations for the problem, the solution of which is a stationary point of the functional. In the present context, the Euler-Lagrange equations constitute the familiar non-linear eigenvalue problem,

$$H_{\text{KS}}[n]\phi_i(x) = \epsilon_i\phi_i(x), \quad (3.4)$$

i.e. the *Kohn-Sham equations*. The eigenvalues (quasi-particle energies) ϵ_i are the Lagrange multipliers associated with the orbital orthonormality constraint. Solving the Kohn-Sham equations to find a stationary point of the Kohn-Sham functional is a necessary but not sufficient condition for (local) optimality. A sufficient condition would require the second derivative (curvature) about the stationary point to be everywhere positive. Furthermore, in general, the Kohn-Sham functional for some approximate E_{xc} is not a convex functional of the orbitals, meaning that verifying global optimality is a difficult task. In practice, solving the Kohn-Sham equations with certain methods of biasing the solution toward a (possibly local) minimum are often chosen rather than direct minimisation methods [58, 112]. The advantages and drawbacks of each approach will be examined in Section 3.4.

It is now possible to formally define what is meant by self-consistency. In order to construct the Kohn-Sham Hamiltonian, one requires a density as input n^{in} to compute the Hartree and exchange-correlation potentials. An output density n^{out} is then calculated

(non-linearly) from the eigenfunctions of the Kohn-Sham Hamiltonian

$$H_{\text{KS}}[n^{\text{in}}]\phi_i(x) = \epsilon_i\phi_i(x), \quad (3.5)$$

$$n^{\text{out}}(x) = \sum_{i=1}^N |\phi_i(x)|^2. \quad (3.6)$$

In general, the input density is not equal to the output density. For a given external potential and exchange-correlation functional, the density n_* is *self-consistent* when $n_* = n^{\text{in}} = n^{\text{out}}$, and hence the non-linear eigenvalue problem of Eqs. (3.5) and (3.6) is solved. The non-linearity in Eq. (3.6) necessitates an iterative procedure that takes an initial estimate of the density as input and iterates this density toward a self-consistent solution of the Kohn-Sham equations: the *self-consistent field procedure*, Fig. 3.1. As one might expect, an infinity of self-consistent densities exist for a given external potential and exchange-correlation functional [113]. However, we are interested primarily in the subset of these densities that are local minima of the Kohn-Sham objective functional. An alternate route towards convergence is to solve the minimisation problem in Eq. (3.1) directly, rather than resorting to the Euler-Lagrange equations. These *direct minimisation* methods tend to be less efficient, but more robust, than self-consistent field methods.

Modern computational implementations of Kohn-Sham theory can vary significantly due to various factors. The key distinguishing factors are choice of basis set, and whether the boundary conditions of the simulation are periodic or finite. These concerns lead to the related problem of whether one treats all the electrons in the computation explicitly, or treats core electrons with a pseudopotential [114]. Despite these differences, perhaps with the exception of linear scaling methods [115], the self-consistent field techniques to be discussed here are adaptable to most implementations. Indeed, some of the most popular software, such as VASP [61, 62], abinit [116, 117], Quantum Espresso [118], and CASTEP [24], use similar default methods to achieve self-consistency: preconditioned multiseccant methods, which are discussed in Section 3.4. This chapter utilises the CASTEP software, which is a periodic, plane-wave, pseudopotential code.

3.2 Purpose

The overarching goal of this work is to quantify the utility of a given algorithm for reaching self-consistency in Kohn-Sham theory. In turn, this allows us to compare and analyse the performance of a sample of existing algorithms from the literature. Assessing these algorithms requires the creation of a test suite of Kohn-Sham inputs, representative of a range of numerical issues. This test suite generates a standard which can be used to test, improve, and present new algorithms designed by method developers. Furthermore, the test suite allows DFT developers to more effectively assess which algorithms they wish to implement. With these aims in mind, this chapter is structured in two partitions, as follows.

We first collate decades of past literature on self-consistency in order to elucidate conclusions that have become conventional wisdom. Section 3.3 then examines the Kohn-Sham framework abstractly from a mathematical and computational perspective in order to

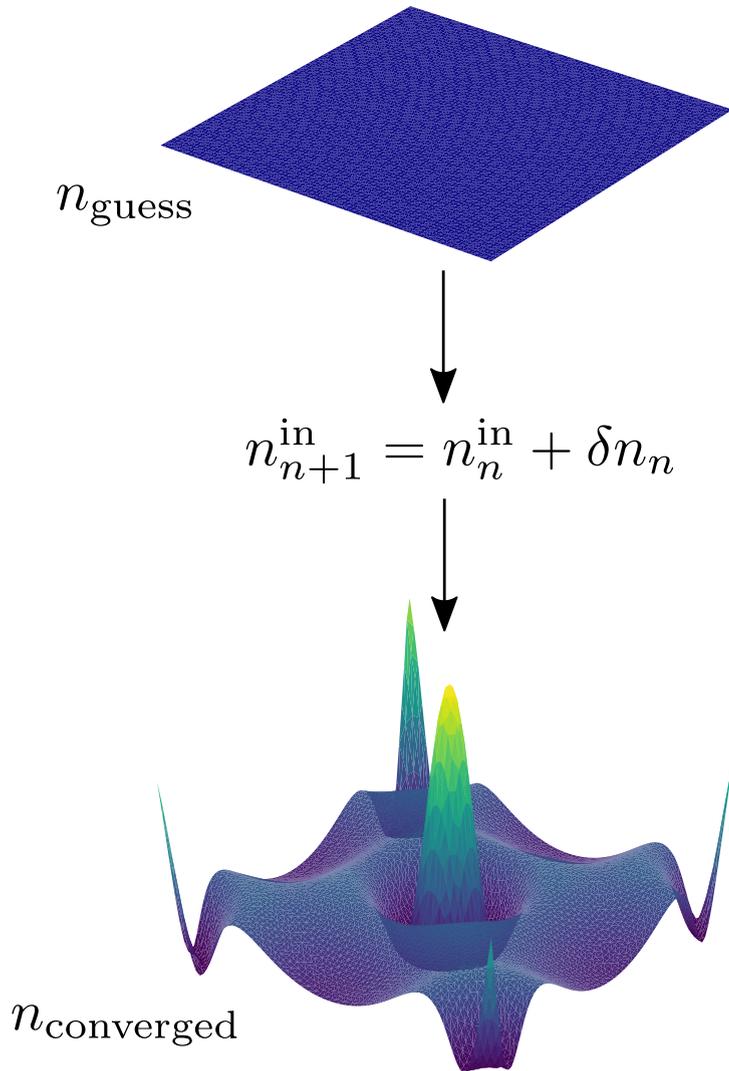


FIGURE 3.1: An iterative algorithm generates a series of perturbations to the density δn_n in order to converge the initial guess density (top) to the fixed-point density n_* (bottom). Example for an *fcc* 4 atom aluminium unit cell.

study *where* and *why* algorithms encounter difficulties. This involves, for example, a discussion on the nature of so-called ‘charge-sloshing’, the initial guess density, sources of ill-conditioning, and more. Section 3.4 then examines and categorises the range of available algorithms in present literature. A focus will be placed on detailing the algorithms which have proven to be particularly successful, and then proceed to detail various novel adjustments made to implement an efficient version of these methods in CASTEP.

The second part then utilises the analysis presented in the prior sections to perform a study akin to recent benchmarking efforts such as *GW100* [119] and the Δ -project [120] for assessing reproducibility in *GW* and DFT codes respectively. Whilst the study presented here will take on a similar structure to these examples, it differs in the following way. The aim of the Δ -project is to assess an ‘error’ for each DFT software in a given

computable property compared to a reference software over a set of test systems. Here, we instead aim to assess the *utility* of an algorithm, rather than an error, which we do with two competing measures: efficiency and robustness, defined in Section 3.5. A test suite of Kohn-Sham inputs is then constructed to target weaknesses in contemporary algorithms and exploit the difficulties discussed in Section 3.3. This test suite is designed to be representative of the range of systems practitioners may encounter and with which they may have difficulties reaching convergence. Each algorithm is then assigned a robustness and efficiency score when tested over the full suite. The methods of Pareto analysis then provide a prescription for the definition of optimal when there exist two or more competing measures of utility. Section 3.6 demonstrates these concepts by using this workflow on a selection of algorithms described in Section 3.4, implemented in the plane-wave, pseudopotential software CASTEP. This study allows conclusions to be drawn about the current state of self-consistency algorithms in Kohn-Sham codes. Finally, we discuss how one might utilise the test suite and workflow demonstrated here to present and assess future methods and algorithms.

3.3 Self-Consistency in Kohn-Sham Theory

In computational implementations of Kohn-Sham theory, when a user has supplied the external potential (e.g. atomic species and positions) and exchange-correlation approximation, the Kohn-Sham energy functional is completely specified. The remaining parameters that are not related to self-consistency, such as Brillouin-zone sampling (*'k-point sampling'*), symmetry tolerances, and so on, tune either the accuracy or efficiency of the calculation. In the context of self-consistency, the user has control over a variety of parameters that can alter the convergence properties of the calculation. Hence, if a calculation is diverging due to the self-consistent field iterations (or converging inefficiently), the user has two options: adjust the parameters of the self-consistency method, or switch to a more reliable fall-back method. This section elucidates the self-consistent field iterations so one can more transparently see why one's iterations may be divergent or inefficient.

3.3.1 The Kohn-Sham Map

As already stated, Kohn-Sham theory is a constrained optimisation problem, Eq. (3.1). The associated Euler-Lagrange equations provide a method for transforming the optimisation problem into a *fixed-point problem*: the Kohn-Sham equations. That is, we seek the density n_* such that it is a *fixed-point* of the discretised *Kohn-Sham map*,

$$K : \mathbb{R}^{N_b} \rightarrow \mathbb{R}^{N_b}, \quad (3.7)$$

$$K[n_*] = n_*. \quad (3.8)$$

where N_b is the number of basis functions (e.g. plane waves) included in the calculation. In general, $K[n^{\text{in}}] = n^{\text{out}}$, where K is defined using Eqs. (3.5) and (3.6). That is, K takes an input density which is used to construct the Hartree and exchange-correlation potentials, then the associated Kohn-Sham Hamiltonian is diagonalised, and an output

density is constructed as the sum of the square of N eigenfunctions. Formally, the Kohn-Sham map is a map from the set of non-interacting v -representable densities onto itself. Here, a non-interacting v -representable density is a density that can be constructed via Eq. (3.6) for a given Kohn-Sham Hamiltonian. The ‘size’ of this set, as a subset of \mathbb{R}^{N_b} , is an open problem [39]. Hence, it is entirely possible that algorithms generate *input* densities that are not non-interacting v -representable; however this appears to not be an issue in practice¹. The aim now is to generate a converging sequence of densities $\{n_0^{\text{in}}, n_1^{\text{in}}, \dots, n_n^{\text{in}}\}$ starting from an *initial guess* density n_0^{in} , where $n_n^{\text{in}} \approx n_*$ to within some desired tolerance. The ease with which this sequence can be generated in practice depends on the functional properties of K , which are examined later in this section.

3.3.2 Defining Convergence

The Kohn-Sham map K , can be used to define a new map R , the *residual*

$$R[n_*] = K[n_*] - n_* = 0, \quad (3.9)$$

which transforms the fixed-point problem into a root-finding problem. An absolute scalar measure of convergence is thus provided by the norm of the residual $\|R[n^{\text{in}}]\|_2$, where $\|\cdot\|_2$ is used to denote the vector L^2 -norm. However, $\|R\|_2$ is a quantity which lacks transparent physical interpretation, making it difficult to assess just *how* converged a calculation is by consideration of $\|R\|_2$ alone. Hence, convergence is conventionally defined in terms of fluctuations in the *total energy*, a more tractable measure. (Indeed, it is possible that relatively large residuals yield small energy fluctuations, and vice versa.) When fluctuations in the total energy are sufficiently low to satisfy the accuracy requirements of the users’ calculation, the iterations are terminated and the calculation is converged. In practice, the total energy is often *not* calculated by evaluating the Kohn-Sham energy functional $E_{\text{KS}}[n_n^{\text{in}}]$. Instead, the *Harris-Foulkes* functional \tilde{E}_{KS} is defined [121],

$$\begin{aligned} \tilde{E}_{\text{KS}}[n^{\text{in}}, n^{\text{out}}] = & \sum_{i=1}^N \frac{1}{2} \int_{\mathbb{R}^3} |\nabla \phi_i|^2 + \frac{1}{2} \int_{\mathbb{R}^3} n^{\text{out}}(x) v_{\text{H}}[n^{\text{in}}] \\ & + \frac{1}{2} \int_{\mathbb{R}^3} (n^{\text{out}}(x) - n^{\text{in}}(x)) v_{\text{H}}[n^{\text{in}}] \\ & + \int_{\mathbb{R}^3} (n^{\text{out}}(x) - n^{\text{in}}(x)) v_{\text{xc}}[n^{\text{in}}] + E_{\text{xc}}[n^{\text{in}}] \\ & + \int_{\mathbb{R}^3} n^{\text{out}}(x) v_{\text{ext}}(x), \end{aligned} \quad (3.10)$$

which can be shown to give the exact ground-state energy correct to quadratic order in the density error about the fixed-point density n_* , i.e. it is correct to $\mathcal{O}((n_* - \delta n)^2)$. Note that it is not the Harris-Foulkes functional that is minimised during the computation, as it possesses incorrect behaviour away from n_* [122, 123]. However, evaluating the energy using this functional *when near* n_* allows one to terminate the iterations at a desired accuracy earlier than if one evaluates the energy using the Kohn-Sham functional,

¹This observation is based on the fact that, in general, one can always find an algorithm that converges to a fixed-point density.

which involves error at linear orders in the density. Finally, recall that $\|R\|_2 \rightarrow 0$ is the criterion for solving the Kohn-Sham equations, not for finding a minimum of the Kohn-Sham functional. Indeed, to verify that a local minimiser of the Kohn-Sham functional is obtained, one would need to ensure all eigenvalues of the Hessian were positive. Such a procedure is not practical in plane-wave codes, and hence the exit criterion for algorithms in Section 3.4 is based solely on fluctuations in the total energy.

3.3.3 Some Unique Properties of The Kohn-Sham Map

Identifying properties unique to the Kohn-Sham map K can help guide and narrow the choice of algorithms in Section 3.4. Firstly, we note that it is computationally expensive to ‘query the oracle’, meaning evaluate K for a given input density to generate the pair $\{n_i^{\text{in}}, n_i^{\text{out}}\}$ on the i^{th} iteration. This is because, when one has specified n^{in} , finding the corresponding n^{out} requires one to construct and diagonalise the Kohn-Sham Hamiltonian. In plane-wave codes, this diagonalisation is done iteratively, and only the relevant N eigenfunctions and eigenvalues are computed. This procedure scales as approximately $\mathcal{O}(N^3)$, and is (in a sense) the bottleneck of the computation [61]. Hence, an algorithm that uses all past iterative data optimally so as to reduce evaluations of K is desirable. Here, the past iterative data constitutes the set of n iterative density pairs $\{(n_i^{\text{in}}, n_i^{\text{out}}) \mid i \in [0, n]\}$. In order to utilise this set to generate the subsequent density n_{n+1}^{in} from some algorithm, one is required to *store* the history of iterative densities in memory. Each density is represented by an array that scales² with N_b , meaning as the iteration number n grows large, so does the memory requirement of storing the entire history. Therefore, a *limited memory algorithm* is also desirable here, meaning no more than m of the most recent density pairs are stored.

The final feature of computational Kohn-Sham theory that we will mention here is the accuracy of the initial guess, n_0^{in} . A discussion on the generation of the initial guess is left to later in this section, but it suffices to note that the initial guess is typically ‘close’ to the converged density n_* . By ‘close’ we mean that a *linear response approximation* can be employed effectively, see Section 3.4. As perhaps would be expected when this is the case, some of the most successful algorithms are able to utilise the past iterations cleverly with limited memory requirements, and employ some form of linearising approximation.

3.3.4 Fixed-Point and Damped Iterations

As mentioned previously, convergence of the self-consistent field iterations depends on the functional properties K , where we recall that each K is specified by the framework of Kohn-Sham theory plus an exchange-correlation approximation and external potential. Despite little being known about the precise functional properties of K [124–126], empirical wisdom allows us to make certain broad statements about it. For the sake of analysis, we now introduce the *fixed-point iteration*,

$$n_{n+1}^{\text{in}} = n_n^{\text{out}} = K[n_n^{\text{in}}]. \quad (3.11)$$

²For example, in plane-wave codes, the density is represented by an array of size $8N_b$, see Ref. [8] for more details.

This is perhaps the most simple iterative scheme one could envisage, yet it remains profoundly important from the point of view of functional analysis [127]. An example algorithm that makes use of the fixed-point iteration scheme is given in Fig. 3.2. This algorithm, on iteration n , constructs and diagonalises the Kohn-Sham Hamiltonian for a given n_n^{in} , and computes the output density n_n^{out} from the N eigenvectors corresponding to the *lowest N eigenvalues*, otherwise known as the *aufbau* principle. The fixed-point iteration is then used as one sets $n_{n+1}^{\text{in}} = n_n^{\text{out}}$, and the procedure is repeated.

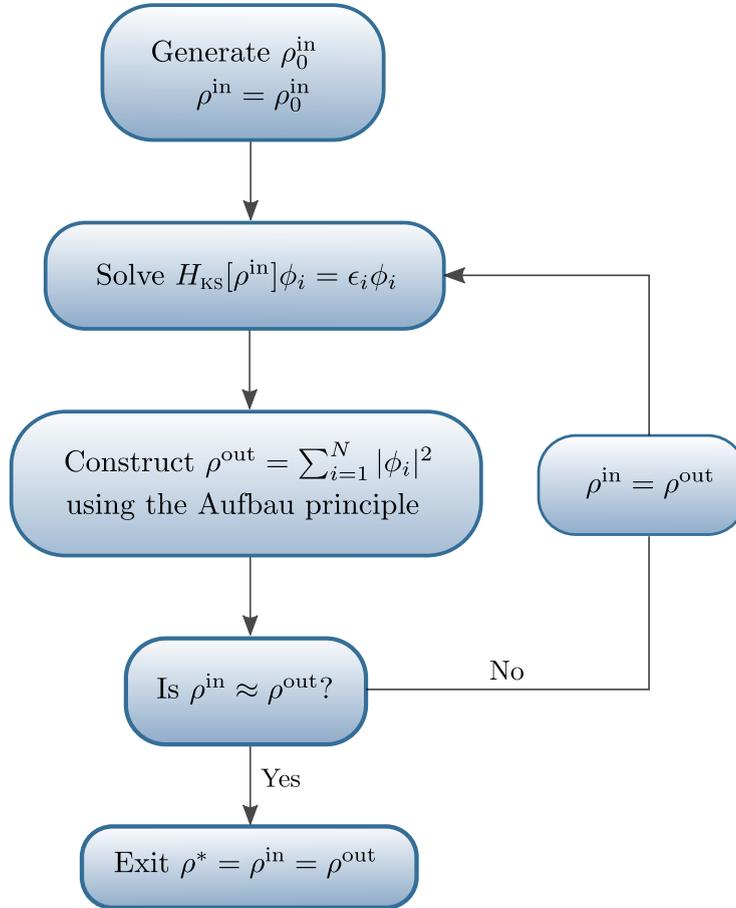


FIGURE 3.2: A flowchart detailing an example algorithm for achieving self-consistency using fixed-point (or Roothaan) iterations.

For the algorithm in Fig. 3.2 to converge, K must be so-called locally k -contractive in the region of the initial guess. For the Kohn-Sham map to be k -contractive under the L^2 -norm, it must satisfy

$$\|K(n_1) - K(n_2)\|_2 \leq k\|n_1 - n_2\|_2 \quad (3.12)$$

for some real number $0 < k < 1$. The intuition here is that, for any two points in the ‘contractive region’, the map K brings these points closer in the L^2 -norm. A contractive

region is simply defined as a domain in density space where Eq. (3.12) is true for all densities within this domain. Successive application of K – the fixed-point iteration scheme – thus continues to bring these points closer toward a locally unique fixed-point, n_* . (See the Banach fixed-point theorem [128] or its generalisations [129] for k -contractive maps.) Unfortunately, as Section 3.6 shows, the Kohn-Sham map is not locally k -contractive for the vast majority of Kohn-Sham inputs. However, perhaps surprisingly, certain calculations do lead to a k -contractive Kohn-Sham map, such as spin-independent *fcc* aluminium at the PBE [54] level of theory, Fig. 3.3. In these cases, sophisticated acceleration algorithms tend to do little-to-nothing to assist convergence. The fixed-point iteration is also referred to as the Roothaan iteration in the physics and quantum chemistry communities [130]. It has been demonstrated that, in the context of Hartree-Fock theory, the Roothaan algorithm either converges linearly toward a solution or oscillates between two densities about the solution [131]. It is expected that this behaviour will carry over to Kohn-Sham theory [132].

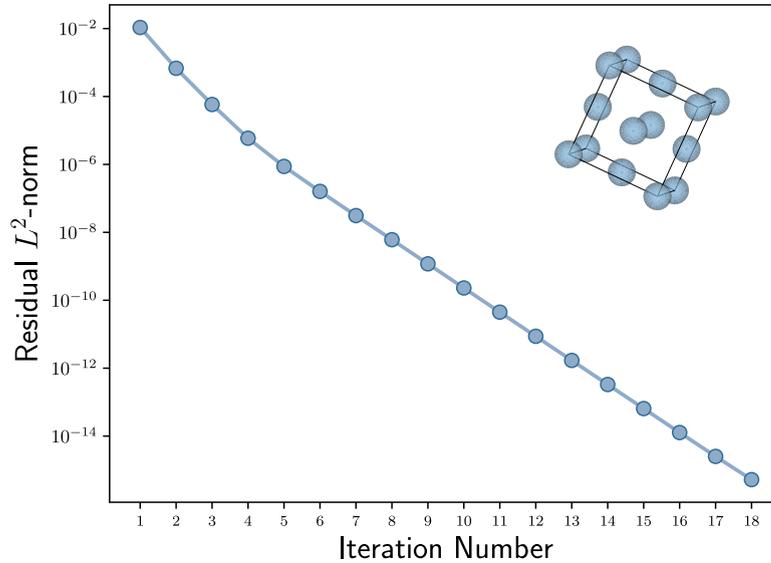


FIGURE 3.3: Iterative convergence in the residual L^2 -norm $\|R\|_2$ toward a fixed-point using fixed-point iterations. Simulation of a four atom *fcc* aluminium unit cell, with k -point spacing of $2\pi \times 0.04\text{\AA}^{-1}$ using the PBE functional.

We now define a new iterative scheme, the damped iteration (or one of its many other aliases, such as Krasnosel'skii-Mann or averaged iteration [133, 134]) such that

$$\begin{aligned} n_{n+1}^{\text{in}} &= n_n^{\text{in}} + \alpha(K[n_n^{\text{in}}] - n_n^{\text{in}}) \\ &= n_n^{\text{in}} + \alpha R[n_n^{\text{in}}]. \end{aligned} \quad (3.13)$$

Hereafter, we refer to a scheme utilising the damped iteration as linear mixing. This scheme constitutes a series of steps in the residual L^2 -norm steepest descent direction R weighted by the parameter $\alpha \in (0, 1)$. It can be shown that provided K is *non-expansive*, there always exists some α such that the damped iteration converges [135, 136]. Here, non-expansive refers to instance whereby $k = 1$ in Eq. (3.12), i.e. densities do not get

further apart upon successive application of K . This property is typically assumed, just as we also assume differentiability of E_{KS} , in theorems relating to convergence features of algorithms discussed in Section 3.4 (e.g. [137]). Indeed, the past few decades of computation using Kohn-Sham theory has led to the wisdom that one can always find some α such that one's calculation converges [132], albeit often impractically slowly. Fortunately, rather large damping parameters of $\alpha \sim 0.5$ are sometimes able to significantly improve convergence, as demonstrated in Fig. 3.4 [138]. In this sense, the Kohn-Sham map is relatively well-behaved, although many problems of physical interest are not so well-behaved. In these cases, sophisticated algorithms are required in order to accelerate and stabilise convergence. However, as Section 3.6 demonstrates, even when recourse to a sophisticated algorithm is required, most inputs excluding those belonging to certain problematic classes are able to converge effectively. This is a testament to the Kohn-Sham map often being dominated by its *linear response* within some relatively large region about the current iterate, a property which is examined further in Section 3.3.7.

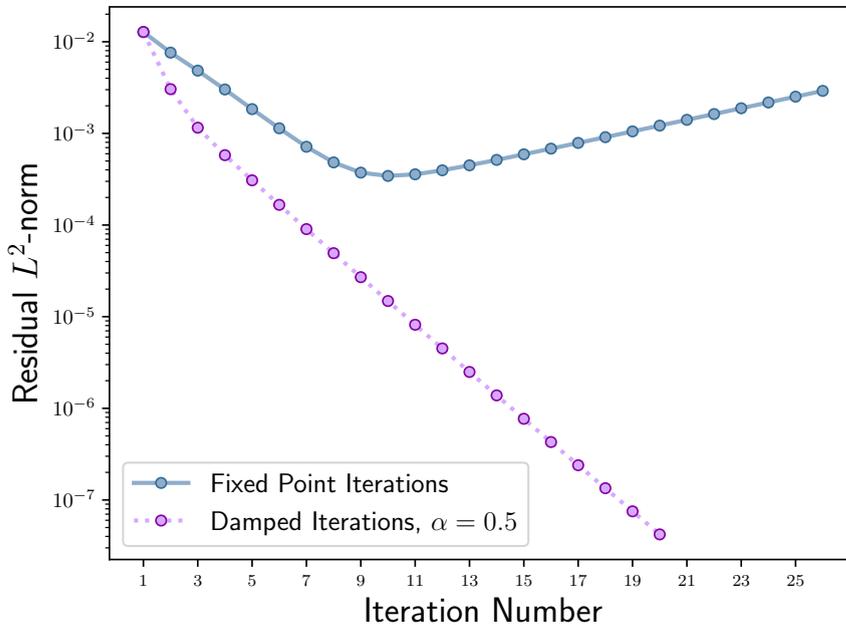


FIGURE 3.4: Iterative convergence in the residual L^2 -norm $\|R\|_2$ using damped iterations with $\alpha = 0.5$, and undamped iterations $\alpha = 1$. Simulation of a four atom *fcc* silicon unit cell, with k -point spacing of $2\pi \times 0.04\text{\AA}^{-1}$ using the PBE functional.

The behaviour of K discussed here could be interpreted as arising due to the lack of convexity of the underlying functional E_{KS} used to generate it. Convexity is defined formally in Section 3.4, but for now it suffices to note that it can be taken to mean E_{KS} has a unique minimum, which is the unique fixed-point of K , and moreover this minimum is global [139]. In other words, solving the Euler-Lagrange equations is a necessary and sufficient condition to verify global optimality. While this is clearly an attractive quality for an energy functional, not least because only the global minimum has direct physical meaning in Kohn-Sham theory, it is not the case here (in general). The lack of convexity of E_{KS} is particularly pronounced in spin-dependent Kohn-Sham theory, where it is not

uncommon for many minima to exist, which are interpreted as representing different meta-stable spin states of the system [140]. In this case, one could, for example, employ some form of global optimisation in an attempt to explore the landscape of local minima with hopes of finding the global minimum.

In summary, while a large class of Kohn-Sham inputs are well-behaved and convergent for relatively high values of the damping parameter, many inputs, especially the increasingly complex ones involved in modern technologies, are not. The remainder of this section explores the precise characteristics of K that lead to ill-behaved convergence.

3.3.5 The Aufbau Principle and Fractional Occupancy

The question remains of how one might go about choosing which N eigenfunctions of the Kohn-Sham Hamiltonian are used to iteratively construct the output densities toward convergence. For $N_b \gg N$, there is of course a large number of permutations of N eigenfunctions from which to choose. While it is perhaps taken for granted, Ref. [113] demonstrates that, in the case of Hartree-Fock theory, the lowest energy solution to the Hartree-Fock equations will necessarily be one which corresponds to the N eigenvectors with the lowest eigenvalues of H_{KS} . This is otherwise known as the *aufbau principle*, and appears in the algorithm presented in Fig. 3.2. These eigenfunctions ϕ_i are termed ‘occupied’ orbitals, with associated quasi-particle energies ϵ_i . However, just because the exact ground state solution satisfies the aufbau principle does not guarantee that doing so *at each iteration* is optimal [126, 131, 141]. Furthermore, iteratively satisfying the aufbau principle does not guarantee a global, or even local, minimum of E_{KS} will be obtained as a solution to the Kohn-Sham equations [58]. Nevertheless, iteratively satisfying the aufbau principle has proven a successful heuristic for finding minima of E_{KS} via the Kohn-Sham equations. Here, the aufbau principle serves to bias our solution of the Kohn-Sham equations toward a minimum of E_{KS} , rather than an inflection point or maximum.

Iterative procedures utilising the aufbau principle are well-defined and work best primarily when the input possesses a Kohn-Sham gap, i.e. when it is *not* a (Kohn-Sham) metal. The Kohn-Sham gap is defined in the limit of large system size as

$$E_{\text{gap}} = \epsilon_{N+1} - \epsilon_N, \quad (3.14)$$

otherwise known as the HOMO-LUMO gap – the difference in energy between the **h**ighest energy **o**ccupied and **l**owest energy **u**noccupied (**m**olecular) **o**rbitals. When this gap disappears, meaning there exists a non-zero density of states at the Fermi energy, convergence becomes increasingly difficult [110, 111]. Here, the Fermi energy μ is defined as the energy of the highest occupied orbital, which is the Lagrange multiplier corresponding to the constraint that the particle number N remain fixed in the minimisation of the Kohn-Sham energy functional. Such cases are prone to the phenomenon of *occupancy sloshing*: iterations become hindered by a continual iterative switching of binary occupation of orbitals whose energies are close to the Fermi energy. In some circumstances, an aufbau solution to the Kohn-Sham equations does not exist for binary occupation of orbitals [39, 142–144].

For example, Ref. [142] demonstrates that, in the case of the C₂ molecule, the Kohn-Sham solution possesses a ‘hole’ below the highest occupied orbital. In the context of self-consistent field iterations, this would mean any algorithm would continue to switch orbital occupancies at each iteration *ad infinitum*. This occurrence is a consequence of degeneracy in the highest occupied Kohn-Sham orbitals, which can occur even in the absence of symmetry and degeneracy in the exact many-body system. Here, and in other cases like this, the density should be constructed from a density matrix

$$D(x_1, \dots, x_N, x'_1, \dots, x'_N) = \sum_{i=1}^q \lambda_i \Phi_i^*(x_1, \dots, x_N) \Phi_i(x'_1, \dots, x'_N) \quad (3.15)$$

via

$$n(x) = \text{Tr} D \hat{n}(x). \quad (3.16)$$

The wavefunctions $\{\Phi_i\}$ are Slater determinants of Kohn-Sham orbitals corresponding to each degenerate solution within some q -fold degenerate subspace. After rearrangement, we find that the density can now be written as

$$n(x) = \sum_{\epsilon_i < \mu} |\phi_i(x)|^2 + \sum_{\epsilon_i = \mu} f_i |\phi_i(x)|^2, \quad (3.17)$$

where the *fractional occupancies* f_i are determined as some combination of the weights λ_i in Eq. (3.15). This form of the density allows one to see more transparently that we have now introduced fractional occupancy of the orbitals whose energy is degenerate at the Fermi energy. In the example of C₂ in Ref. [142], the degenerate subspace is first identified, and then the occupancies f_i are varied smoothly until the energies of the identified orbitals are equal. This procedure, termed *evaporation of the hole*, yielded accurate energy predictions when compared to configuration interaction calculations. In this case, the Kohn-Sham degeneracy is interpreted as being due to the presence of strong electron correlation. These degeneracies lead to densities that are so-called *ensemble non-interacting v -representable*. That is, the exact Kohn-Sham density can no longer be constructed from a *pure state* via the sum of the square of orbitals as in Eq. (3.3), but instead must be constructed from some ensemble of states via Eqs. (3.15) and (3.16). The extension of Kohn-Sham theory to include fractional occupancy is described well in Refs. [145, 146].

This so-called ensemble extension to Kohn-Sham theory is also utilised when constructing a non-interacting theory of Mermin’s finite temperature formulation of DFT [147]. It is this version of Kohn-Sham theory that is usually used in modern Kohn-Sham codes that include fractional occupancy. As we are interested primarily in how this extension mitigates convergence issues, the reader interested in an in-depth discussion of finite temperature Kohn-Sham theory is referred to [110, 145], and references therein. Here,

it suffices to observe that we now seek to minimise the following *free energy* functional

$$\begin{aligned}
 E[\{\phi_i\}, \{f_i\}, T] &= \sum_{i=1}^{\infty} \frac{1}{2} f_i \int |\nabla \phi_i|^2 dx \\
 &\quad + \frac{1}{2} \int \frac{n(x)n(x')}{|x-x'|} dx dx' \\
 &\quad + \int n(x)v_{\text{ext}}(x) dx + E_{\text{xc}}[n] \\
 &\quad - TS[\{f_i\}],
 \end{aligned} \tag{3.18}$$

where the entropy functional and density are defined respectively as

$$S = \sum_{i=1}^{\infty} f_i \ln(f_i) + (1 - f_i) \ln(1 - f_i), \tag{3.19}$$

$$n(x) = \sum_{i=1}^{\infty} f_i |\phi_i(x)|^2. \tag{3.20}$$

The real-valued fractional occupancies $f_i \in [0, 1]$ now constitute additional variational parameters alongside the orbitals. Minimisation of the finite temperature Kohn-Sham functional can be tackled directly as in Ref. [111], which is discussed in Section 3.4 and tested in Section 3.6. Alternatively, the associated fixed-point problem can be formulated, whereby the occupancies are given a fixed functional form dependent on both T and the Kohn-Sham Hamiltonian eigenenergies ϵ_i . This is otherwise known as the *smearing scheme*, an example of which is the Fermi-Dirac function,

$$f_i = \frac{1}{e^{(\epsilon_i - \mu)/T} + 1}. \tag{3.21}$$

The electronic temperature T is now an input parameter which determines the degree of broadening of occupancies about the Fermi energy, Fig. 3.5. At each iteration, the occupancies are updated with new values of ϵ_i , and this process is continued toward convergence. This procedure demonstrably mitigates occupancy sloshing for Kohn-Sham metals with large density of states at the Fermi energy [148, 149]. However, the end result of such a procedure is a ground-state energy at some finite (electronic) temperature T . To recover an estimate for the zero-temperature ground-state energy, the finite-temperature contribution to the energy functional must be subtracted at the end [146].

An additional numerical benefit due to introducing ‘finite temperatures’ is regarding the sampling of the Brillouin zone in periodic Kohn-Sham codes. That is, interpolation techniques for evaluating integrals across the Brillouin zone are inaccurate when many band crossings (discontinuous changes of occupancy) exist, i.e. in Kohn-Sham metals. This necessitates a fine sampling of k -space in order to accurately evaluate the integrals. As discussed, fractional occupancies negate these discontinuities, allowing for a coarser sampling of the Brillouin zone, meaning interpolation techniques become increasingly accurate – see Refs. [148, 150] for more details. In any case, finite electronic temperatures are a valuable numerical tool to assist convergence of the self-consistent field iterations in the event of inputs with large density of states at the Fermi energy. Hence, the test suite in Section 3.5 includes many such systems, and in particular a variety of electronic temperatures are considered.

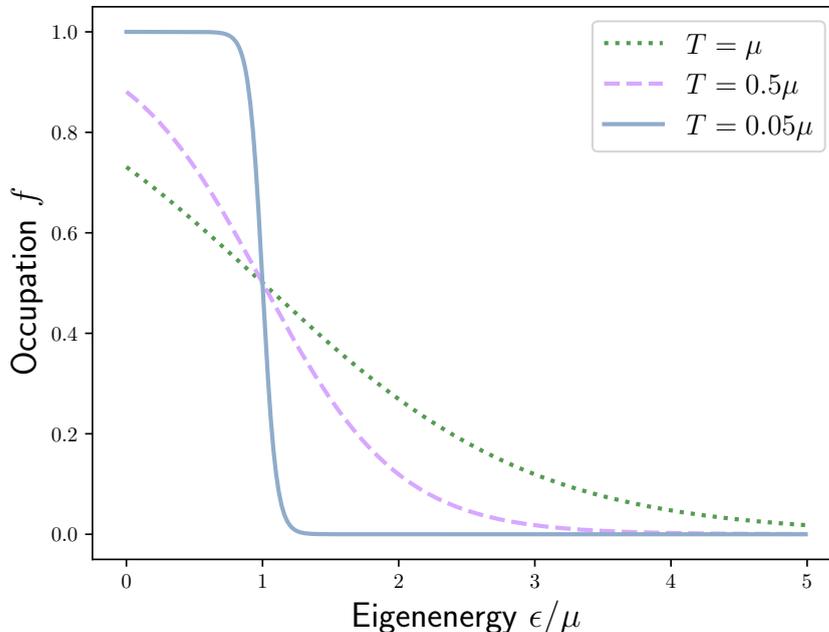


FIGURE 3.5: Functional dependence of the occupancy f of a given eigenenergy ϵ for various temperatures using the Fermi-Dirac smearing scheme.

3.3.6 The Initial Guess

As one might expect, a more accurate initial guess of the variable to be optimised leads monotonically to more efficient and stable convergence rates [151]. In the case of self-consistent field methodology, and for plane-wave and similar codes, the initial guess charge density is often computed as a *sum of pseudoatomic densities* [61, 62]. That is, once the exchange-correlation and pseudopotential for the atomic species in the computation has been specified, the charge density for these atoms *in vacuum* is calculated. Then, each individual density is overlaid at positions centered on the atomic cores in order to construct the initial guess density, Fig. 3.6. This figure demonstrates visually the accuracy of this prescription for generating initial guess charge densities. Note that different considerations are required in order to generate an initial guess for the density matrix or orbitals.

The accuracy of the initial guess is, in part, responsible for the relative success of methods that employ linearising approximations, such as quasi-Newton methods, see Section 3.4. Notable cases in which the initial guess density is relatively poor include polar materials such as magnesium oxide. The initial guess is charge neutral by construction, meaning the charge is required to shift onto the electro-negative species for convergence. Furthermore, inputs whereby the atomic species are subject to large inter-atomic forces can also lead to inaccurate initial guesses. This is partly due to the fact that the initial guess becomes exact in the limit of large atomic separation, and since large inter-atomic forces imply low inter-atomic separation, this can result in potentially inaccurate initial

guess densities. Such inputs are generated routinely during structure searching applications [152]. The test suite includes various examples of these ‘far-from-equilibrium’ systems.

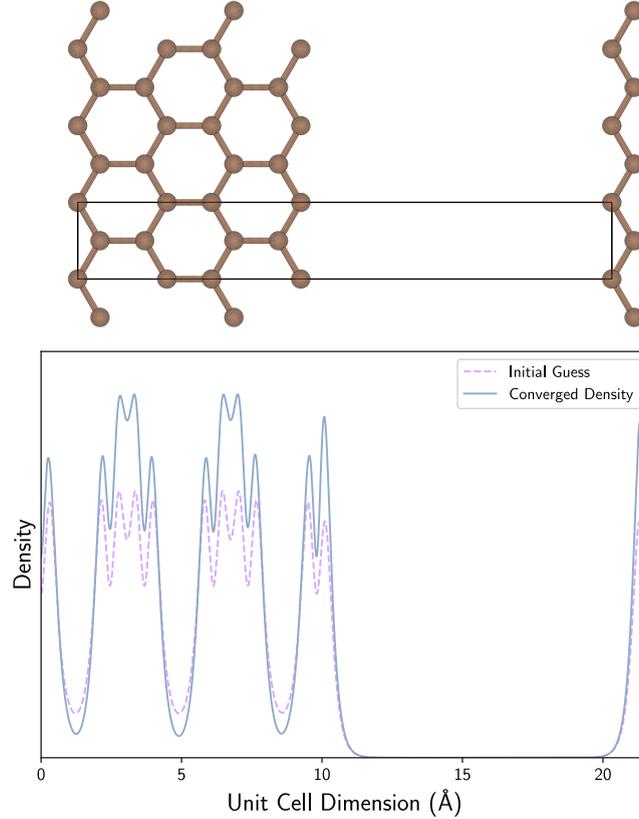


FIGURE 3.6: The difference between the initial guess density, constructed from a sum of isolated pseudoatomic densities, and the converged density, for a graphene nanoribbon along one dimension (above). The density is plotted along a one-dimensional slice across the axis shown in the above image.

Spin polarised Kohn-Sham theory presents more serious issues: there is no widely successful method for generating initial guess spin densities. In the spin polarised or ‘unrestricted’ formalism, the following *spin densities* are introduced (see Ref. [8]),

$$n^\uparrow(x) = \sum_{i=1}^{\infty} f_i^\uparrow |\phi_i^\uparrow(x)|^2 \quad (3.22)$$

$$n^\downarrow(x) = \sum_{i=1}^{\infty} f_i^\downarrow |\phi_i^\downarrow(x)|^2, \quad (3.23)$$

generated from spin up and down particles occupying separate *spin orbitals* $\{\phi_i^\uparrow, \phi_i^\downarrow\}$. This leads now to two coupled non-linear eigenvalue problems, one for each spin. A method for generating the initial guess spin densities is thus required, rather than just the initial guess charge density. As one, in general, has no knowledge of the spin state *a priori*, this initial guess can be relatively far away from the ground state. In practice, one conventionally deals with charge and spin densities, rather than spin up and spin

down densities,

$$n^{\text{charge}}(x) = n^\uparrow(x) + n^\downarrow(x), \quad (3.24)$$

$$n^{\text{spin}}(x) = n^\uparrow(x) - n^\downarrow(x). \quad (3.25)$$

The charge density can be initialised similarly to the spin-independent case, with a sum of independent pseudoatomic charge densities. The spin density can be initialised to zero, or be scaled by specifying some magnetic character on the atoms, e.g. ferromagnetic. Such a prescription typically leads to initial guess densities that are further away in the residual L^2 -norm than spin-independent initial guess charge densities. This observation at least partially accounts for the reason that spin polarised systems tend to be much harder to converge than spin unpolarised systems. For this reason, and others cited in the following section, many spin polarised inputs are included in the test suite. Recently, various schemes have been proposed that aim to better predict self-consistent densities to use as the initial guess [151, 153]. In particular, Ref. [151] considers a data-derived approach to predicting and assessing uncertainty in a guess density away from the ground state.

3.3.7 Ill-Conditioning and Charge Sloshing

The *condition* of a problem, loosely speaking, can be taken as characteristic of the difficulty a black-box algorithm will have in solving the problem. Due to the complexity of the Kohn-Sham map, evaluating its *condition number* directly is impossible in practice. However, within the context of linear response theory, it is possible to explore certain causes of ill-conditioning generic to either all, or certain broad classes, of inputs. Hence, we begin by *linearising* the map K about a fixed-point³,

$$\begin{aligned} n_* + \delta n_n^{\text{out}} &= K[n_* + \delta n_n^{\text{in}}], \\ &\approx K[n_*] + \left. \frac{\delta K[n_n^{\text{in}}]}{\delta n_n^{\text{in}}} \right|_{n_*} \delta n_n^{\text{in}}, \\ \implies \delta n_n^{\text{out}} &= \left. \frac{\delta K[n_n^{\text{in}}]}{\delta n_n^{\text{in}}} \right|_{n_*} \delta n_n^{\text{in}}. \end{aligned} \quad (3.26)$$

This is the definition of linearisation in the present context, i.e. a small change in the input density yields a change in the output density proportional to the initial change, where the constant of proportionality is given by the components of the *Jacobian* of the map K ,

$$\begin{aligned} J_K(x, x') &= \left. \frac{\delta K[n_n^{\text{in}}]}{\delta n_n^{\text{in}}} \right|_{n_*} (x, x') \\ &= \left. \frac{\delta n_n^{\text{out}}}{\delta n_n^{\text{in}}} \right|_{n_*} (x, x'). \end{aligned} \quad (3.27)$$

³Note that one can linearise about any density, not necessarily a fixed-point density. We have chosen the density about which we linearise to be the fixed-point density for the sake of analysis and due to the accuracy of the initial guess.

Within the language of linear response theory, the Jacobian can be identified with the *non-interacting charge dielectric* via

$$\varepsilon_0(x, x') = I - J_K(x, x'), \quad (3.28)$$

which is the linear response function of the *residual map*, rather than the Kohn-Sham map. The dielectric can be expanded as such

$$\varepsilon_0(x, x') = I - \int \frac{\delta v_{\text{hxc}}^{\text{in}}(x'')}{\delta n^{\text{in}}(x')} \frac{\delta n^{\text{out}}(x)}{\delta v_{\text{hxc}}^{\text{in}}(x'')} dx'', \quad (3.29)$$

where $v_{\text{Hxc}} = v_{\text{H}} + v_{\text{xc}}$, which are the only two potentials which have a dependence on the density. Hence, the dielectric is given in terms of the *non-interacting density-density response function* χ_0 as

$$\varepsilon_0(x, x') = I - \int (f_{\text{H}}(x', x'') + f_{\text{xc}}(x', x'')) \chi_0(x, x'') dx'', \quad (3.30)$$

where f_{H} and f_{xc} are the kernels of the Hartree (Coulomb) and exchange-correlation integrals (see Section 2.4.3). Therefore, the linear response of a system to a density perturbation is given by the interplay between the exchange-correlation and Coulomb kernels, and the static response function,

$$\chi_0(x, x') = \frac{\delta n^{\text{out}}(x)}{\delta v_{\text{hxc}}^{\text{in}}(x')}, \quad (3.31)$$

which is highly system dependent [138, 154]. As the non-interacting response function (evaluated at the self-consistent density) plays a central role in the description of many physical phenomena, such as absorption spectra, it is a relatively well-studied object [83, 84, 155–157]. (Note that, during the iterations, the response functions are evaluated at the current iterate). The remainder of this section classifies certain generic behaviours of ε_0 so that causes of divergence in the self-consistency iterations can be studied.

First, in order to see why the linear response function is important for self-consistency iterations, note that one may consider each iteration as a perturbation in the density about the current iterate. Knowledge of the exact response function χ_0 , and subsequently ε_0 , would thus allow one to take a controlled step toward the fixed-point density, depending on how well-behaved⁴ the map is about the current iterate. An iterative scheme utilising the exact response function is given by

$$n_{n+1}^{\text{in}} = n_n^{\text{in}} + \varepsilon_0^{-1} R[n_n^{\text{in}}], \quad (3.32)$$

which one might recognise as Newton's method. While Newton's method does not, in general, guarantee convergence toward a local root, it has many attractive features, see Section 3.4.1. However, one is rarely privileged with knowledge of the exact dielectric as it is vastly expensive to compute and store⁵ [158–161]. In practice, one is left to estimate, or iteratively build, this response function. Cases in which the input is very

⁴By 'well-behaved' here, we mean that the higher order than linear terms can be ignored without much detriment.

⁵In later chapters, these response functions are implemented exactly for finite one-dimensional systems. This affords a great deal of freedom in the context of self-consistency iterations, however, much

sensitive to density perturbations, characterised by large eigenvalues of the discretised ε_0^{-1} , tend to amplify errors in iterates, and thus potentially move one away from the fixed-point.

Consider now completely neglecting higher order terms in the Taylor expansion of the Kohn-Sham map, and let us examine the map as if it were linear. This allows us to borrow results from numerical analysis of linear systems, and apply these results as well-motivated heuristics to convergence in the non-linear case. In particular, assuming linearity, absolute convergence can be identified as

$$\begin{aligned} \delta n_{n+1}^{\text{in}} &\rightarrow 0 \text{ as } n \rightarrow \infty, \\ \implies (\varepsilon_0^{-1})^n \delta n_0^{\text{in}} &\rightarrow 0 \end{aligned} \quad (3.33)$$

using Eq. (3.26), meaning we must have $\lambda_i < 1$ for all i in order to converge, where λ_i are the eigenvalues of the inverse dielectric matrix, which have been shown to be real and positive for some appropriate f_{xc} [138]. Hence, simply multiplying the dielectric by a scalar α such that λ_{max} is below unity can ensure convergence. This comes at the cost of reducing the efficiency of convergence for components of the density corresponding to low eigenvalues of the dielectric matrix. Defining the condition number of the dielectric as

$$\kappa = \frac{\lambda_{\text{max}}(\varepsilon_0^{-1})}{\lambda_{\text{min}}(\varepsilon_0^{-1})}, \quad (3.34)$$

it can be seen that the efficiency of the linear mixing procedure is limited by how close this ratio is to unity (this is an extremely general and foundational conclusion in the numerical analysis of linear systems). One ansatz for the scalar premultiplying the dielectric is

$$\alpha = \frac{2}{\lambda_{\text{max}} + \lambda_{\text{min}}}, \quad (3.35)$$

which ensures, as much as the linear approximation is valid, that components of the density corresponding to the maximal and minimal eigenvalues of ε_0^{-1} converge at the same rate [138, 154, 162]. However, this form of α ignores the distribution (e.g. clustering) of eigenvalues [163], and is not commonly used in conjunction with more sophisticated schemes such as those in Section 3.4. An additional strategy to improve convergence would be to construct a matrix, the *preconditioner*, such that when the preconditioner is applied to ε_0 , the eigenspectrum of the product is compressed toward unity. This is done in practice, see Ref. [61] for example, and is the core idea behind the *Kerker preconditioner* [164, 165], as discussed shortly.

It is clear from Eq. (3.34) that the convergence depends critically on the spectrum of the inverse dielectric matrix. The minimum eigenvalue is unity, and the large eigenvalues are dominated by the contributions from the Coulomb kernel, rather than the exchange-correlation kernel [154, 158, 166]. To see why this is, it is first *asserted* that the discretised Coulomb kernel possesses large (divergent) eigenvalues due to its $|x-x'|^{-1}$

of the ill-conditioning we discuss here is not present in one-dimensional finite systems, thereby restring this line of inquiry.

dependence, which will be demonstrated in the work to follow shortly. These large eigenvalues act as a *generic* amplification factor when multiplied by the eigenvalues of χ_0 in Eq. (3.30), thus resulting in large (divergent) eigenvalues in the dielectric. Conversely, the exchange-correlation kernel does not generically possess such large eigenvalues. In semi-local Kohn-Sham theory, the exchange-correlation kernel has a simple, local and explicit dependence on position,

$$f_{xc}(x, x') \propto \delta(x - x'), \quad (3.36)$$

that does not, in conventional parameterisations, introduce x -dependent ill-conditioning. As such, no generic amplification of the eigenvalues of χ_0 occurs, and hence the exchange-correlation kernel can be ignored relative to the Coulomb kernel from the perspective of ill-conditioning. In other words, the following analysis utilises in the *random phase approximation* (RPA) by setting $f_{xc} = 0$ in Eq. (3.30). As Ref. [158] notes, even in situations whereby the density vanishes in some region, meaning that negative powers of the density are divergent, the linear response function tempers this divergence, and the exchange-correlation contribution remains well-conditioned.

The principle categorisation one can make when analysing generic behaviour of the response function is the distinction between Kohn-Sham metals and insulators. Consider a homogeneous and isotropic system, i.e. the homogeneous electron gas, such that $\chi_0(x, x') \rightarrow \chi_0(|x - x'|)$, which satisfies

$$\delta n^{\text{out}}(x) = \int dx' \chi_0(|x - x'|) \delta v_{\text{h}}^{\text{in}}(x'). \quad (3.37)$$

This is a convolution in real space, and hence a product in reciprocal space

$$\delta \tilde{n}^{\text{out}}(G) = \tilde{\chi}_0(|G|) \delta \tilde{v}_{\text{h}}^{\text{in}}(G), \quad (3.38)$$

where we label the Fourier components G by convention. The susceptibility is local and homogeneous in reciprocal space, and relates perturbations in the input density to a response by the output density (within the RPA) via

$$\begin{aligned} \delta \tilde{n}^{\text{out}}(G) &= f_{\text{H}}(|G|) \tilde{\chi}_0(|G|) \delta \tilde{n}^{\text{in}}(G) \\ &= \frac{4\pi \tilde{\chi}_0(|G|)}{|G|^2} \delta \tilde{n}^{\text{in}}(G). \end{aligned} \quad (3.39)$$

The susceptibility of the homogeneous electron gas, which constitutes a simple metal in the present context, is derived from Thomas-Fermi theory as the Thomas-Fermi wavevector $\tilde{\chi}_0 \sim k_{\text{TF}}^2$, which is constant⁶ [167]. It can therefore be seen that if there is any error in a trial input density, generated by an iterative algorithm, away from the optimal update, then this error is amplified by a factor of $|G|^{-2}$ for $|G| < 1$, where $|G| = 0$ does not contribute.

This sensitivity to error in iterates is identified as the source of *charge sloshing*, and is a somewhat generic feature of Kohn-Sham metals. Whilst the above derivation utilises

⁶A detailed treatment of χ_0 for inhomogeneous inputs is given within the framework of Lindhard theory [155].

Thomas-Fermi theory of the homogeneous electron gas to demonstrate constant susceptibility, it can be shown that all Kohn-Sham metals display this behaviour in the small $|G|$ limit [157, 168]. A demonstration of charge sloshing is illustrated in Fig. 3.7, whereby a linear mixing algorithm purposefully takes slightly too large steps in the density. This leads to vast over-corrections in each iteration, giving the appearance that charge is ‘sloshing’ about the unit cell⁷.

This is not the only source of large eigenvalues of the dielectric in Kohn-Sham metals, as the susceptibility possesses inherently divergent eigenvalues independent from the amplification by the Coulomb kernel. To see this, consider the perturbation theory expression for the linear response function,

$$\chi_0(x, x') = \sum_{n=1}^N \sum_{m=N+1}^{\infty} \frac{\phi_n(x)\phi_m^*(x)\phi_n^*(x')\phi_m(x')}{\epsilon_n - \epsilon_m}, \quad (3.40)$$

otherwise called the Adler-Wiser equation [83, 84]. As Ref. [162] originally noted, the denominator $\epsilon_n - \epsilon_m$ approaches zero when the input is gapless, i.e. it has a large density of states about the Fermi energy. If left untreated, this observation, in conjunction with the amplifying factor from the low $|G|$ components of the Coulomb kernel, leads to significant ill-conditioning. The largest condition numbers arise when $|G|$ is (comparatively) small; since G is a reciprocal lattice vector, this will occur for unit cells that are increasingly large in any (or all) of the three real-space dimensions.

Whilst the dependence of the dielectric eigenvalues on unit cell size is in practice complicated [154], it suffices to note that increased unit cell size is a significant source of ill-conditioning. As compute power continues to grow, larger and larger systems are being tackled using Kohn-Sham theory, and the increase in required number of self-consistency iterations as a result of this instability poses serious issues for Kohn-Sham calculations. Inefficiencies of this kind are best dealt with using preconditioners, as Section 3.4 demonstrates. On the other hand, insulators possess no such divergences in the eigenvalues of the dielectric matrix. It can be shown that in the low $|G|$ limit the behaviour of the susceptibility for gapped materials is [157, 168]

$$\tilde{\chi}_0(G) \propto |G|^2. \quad (3.41)$$

This functional dependence *cancels* the $|G|^{-2}$ dependence from the Coulomb kernel, and thus the eigenvalues of the dielectric become constant for all G . This constant is unknown in general, and guaranteed convergence for simple insulators amounts to finding the damping parameter α such that this constant is below unity. This is in line with the empirical wisdom that insulators are much easier to converge than metals, provided that the insulator does not artificially assume a metallic character during the self-consistency iterations.

In this vein, inputs that are increasingly complex, i.e. deviating from simple metals or insulators, are likely to exhibit problematic behaviour. As discussed in Section 3.4, preconditioners are able to alleviate charge sloshing in simple metals. However, when a metal-insulator interface is used as input, there are regions with starkly different

⁷In the context of numerical analysis, charge sloshing simply amounts to taking steps far outside of the ‘trusted region’ about the current iterate

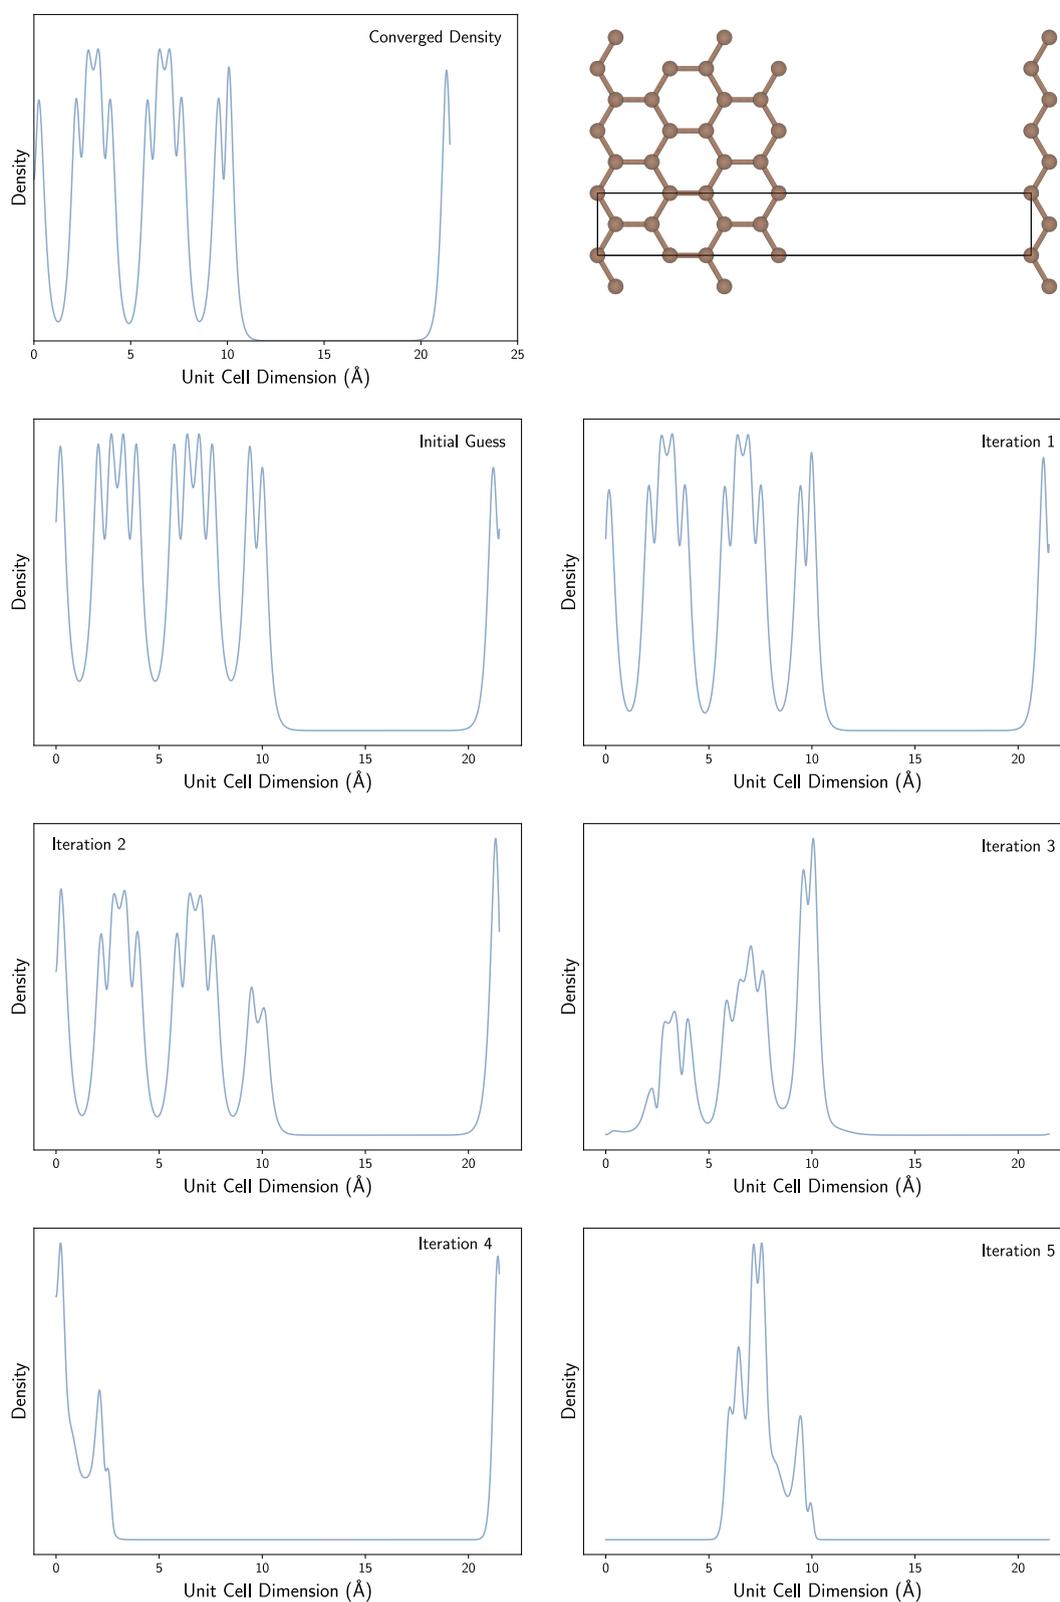


FIGURE 3.7: An illustration of charge sloshing for a graphene nanoribbon unit cell (top right). The linear mixing algorithm is applied with damping parameter $\alpha = 0.8$; this leads to an overcorrection in the density at each iteration, resulting in complete divergence.

behaviour in the response function, which is difficult to capture analytically. Hence, preconditioning techniques may fail to assist, and even hinder, iterations in calculations on interfaces of this kind [154]. Furthermore, it is possible that *artificial phase transitions* between gapped and gapless phases occur during the self-consistency iterations. Many algorithms function by building up an approximation to the dielectric using past iterates. The discontinuous change in behaviour of the dielectric in differing phases causes parts of the iterative history to actively interfere in correctly modelling the dielectric. Hence, iterations become hindered or divergent. An artificial phase change of this kind is demonstrated to occur in Ref. [169] for an isolated iron atom. Various examples of the aforementioned problematic classes of inputs are included in the test suite.

Finally, a brief comment is provided on how the above analysis translates to spin-dependent Kohn-Sham theory. As discussed, in the spin-dependent case one solves two non-linear eigenvalue problems that independently look very similar to Eqs. (3.5) and (3.6), but crucially are coupled through the Hartree and exchange-correlation potential. That is, an algorithm that perturbs the spin up (spin down) density will lead to a response by the spin up (spin down) density given by the prior analysis. However, one must now also consider how a perturbation in the spin up density affects the spin down density, and vice versa, which is again through the Coulomb and exchange-correlation kernels. Hence, all of the above sources of ill-conditioning translate directly to the spin-dependent case, with the added difficulty that the number of optimisation parameters has doubled, and these parameters are coupled in such a way that potentially introduces further ill-conditioning. In other words, the response function now contains four components,

$$\varepsilon_0 = \begin{bmatrix} \varepsilon_0^{\uparrow\uparrow} & \varepsilon_0^{\uparrow\downarrow} \\ \varepsilon_0^{\downarrow\uparrow} & \varepsilon_0^{\downarrow\downarrow} \end{bmatrix}, \quad (3.42)$$

where the diagonal elements describe the response of the spin up (spin down) density to perturbations in the spin up (spin down) density, and the off-diagonal elements describe the response of the spin up (spin down) density to perturbations in the spin down (spin up) density.

Implementations of spin-dependent Kohn-Sham theory often utilise the charge and spin densities, defined in Eqs. (3.24) and (3.25), rather than the spin up and spin down densities. In this representation, the behaviour of the off-diagonal (coupling) elements is determined via the exchange-correlation potential, as the Hartree contribution cancels, see Refs. [8, 138]. To the authors' knowledge, there is less literature on the manifestation of this coupling in the self-consistency iterations than on the spin-independent counterpart. Ref. [138] uses self-consistency in the 'Stoner model' to demonstrate that the condition of the system is indeed worsened in the presence of magnetism due to the coupling. However, it is noted that the charge and spin densities decouple near self-consistency. In any case, for these reasons, and perhaps for reasons yet unexplored, empiricism demonstrates that spin polarised calculations are, in general, more difficult to converge than spin unpolarised calculations. Analysing this coupling in terms of self-consistency and ill-conditioning is an avenue for further research, which could be used to assist self-consistency iterations through the use of improved preconditioners, see Section 3.4.2.

3.4 Methods and Algorithms

Having established a variety of sources of ill-conditioning in the non-linear Kohn-Sham map, we now examine methodology used to find self-consistent densities that are fixed-points of this map. Of course, over the past few decades, a number of differing approaches to the self-consistency problem in Kohn-Sham and Hartree-Fock theory have been reviewed, analysed, and advanced; see, for example, Refs. [23, 58, 131, 154, 169–172] and references therein. The goal of this section is to collate the conclusions from these studies, and more, in order to provide a unifying view of secant-based algorithms in terms of the differing optimisation *sub-problems* underlying each method. These concepts and methods are then adapted (to improve performance) and implemented in CASTEP for the benchmarking effort in Section 3.6.

Consider the general iteration for solving the Kohn-Sham equations,

$$n_{n+1}^{\text{in}} = f_n(\{n_i^{\text{in}}, n_i^{\text{out}}\}), \quad (3.43)$$

where n is the current iteration number, $i \in [1, n]$, and we seek a prescription for generating the update f_n as a function of all past data in the history of iterates. The underlying black-box methodology one uses to generate f_n can be regarded as separate to how one alters f_n by preconditioning. Hence, we first review the black-box methodology, and then review preconditioning strategies in Section 3.4.2. Elementary algorithms for generating f_n were first considered in Section 3.3.4: the fixed-point and linear mixing algorithms,

$$n_{n+1}^{\text{in}} = n_n^{\text{out}}, \quad (3.44)$$

$$n_{n+1}^{\text{in}} = n_n^{\text{in}} + \alpha(K[n_n^{\text{in}}] - n_n^{\text{in}}), \quad (3.45)$$

respectively. As stated, the linear mixing algorithm is a weighted step in the direction of the error and is identically zero at convergence. Hence, assuming K is continuous (in some sense) and non-expansive, this algorithm converges for sufficiently low fixed values of α [138]. It can be shown that this algorithm converges *q-linearly* toward the fixed-point density n_* [173]; where *q-linear* convergence is defined as

$$\|n_{n+1}^{\text{in}} - n_*\|_2 \leq q \|n_n^{\text{in}} - n_*\|_2. \quad (3.46)$$

That is, the error decreases linearly iteration by iteration, and the gradient of this linear decrease is given by the factor $q \in (0, 1)$, which is determined by the initial guess and the *fixed* parameter α . Assuming one chooses an appropriate value for α , the linear mixing algorithm is guaranteed, meaning it converges from any initial guess in the $n \rightarrow \infty$ limit. The price one often pays for guaranteed convergence here is an impractically slow convergence rate, or q factor, for the problematic classes of inputs defined in the prior section. The remainder of this section considers methods for *accelerating* the linear mixing iterations, conventionally referred to as acceleration algorithms. In particular, these algorithms exhibit *q-superlinear* convergence,

$$\|n_{n+1}^{\text{in}} - n_*\|_2 \leq \gamma \|n_n^{\text{in}} - n_*\|^q, \quad (3.47)$$

for some positive real number γ , where $q > 1$ and $q = 2$ defines quadratic convergence. These algorithms tend to sacrifice guaranteed convergence, but can vastly improve the rate of convergence, as demonstrated in Section 3.6.

Before introducing the various acceleration strategies, we remark that the difficulty in solving a constrained functional optimisation problem, or equally the associated Lagrangian fixed-point problem, is not primarily determined based on the linearity of a problem, or lack thereof. Rather, as Ref. [139] asserts and demonstrates, the characteristic difficulty of an optimisation problem depends on whether or not the underlying functional is convex,

$$F[\alpha x + \beta y] \leq \alpha F[x] + \beta F[y]. \quad (3.48)$$

Here, F is a convex functional, x and y are two elements in the domain of the functional, and α and β are two real numbers. Convex functionals have a unique minimum, and minimiser, which can be found, in some sense, in a controlled and efficient manner, see Refs. [139, 173] for more information on convex optimisation. The Kohn-Sham functional is demonstrably not convex in the general case. However, many of the algorithms to follow operate by solving an associated convex problem in order to compute the update f_n . This is typically a *quadratic programming* problem, which is subsequently used to solve the non-convex Kohn-Sham problem. The most popular and successful class of updates in the present context are *quasi-Newton* updates. As we will see, these updates differ chiefly based on the underlying quadratic programming problem one solves to compute f_n .

3.4.1 The Quasi-Newton Update

First, we make some general comments about the Newton update. The Newton update is the optimal *first order* update in the density at the current iteration. In other words, if the current iterate is within the linear response radius of the root, then the exact Newton update would lead to convergence in one iteration by definition. That is, we seek the update δn_n such that

$$R[n_n^{\text{in}} + \delta n_n] \approx R[n_n^{\text{in}}] + J|_{n_n^{\text{in}}} \delta n_n = 0, \quad (3.49)$$

where $J := J_R$ is the Jacobian of the residual, defined similarly to Eq. (3.27), evaluated at the current iterate. Rearranging for δn_n , the update is given as

$$n_{n+1}^{\text{in}} = n_n^{\text{in}} - J|_{n_n^{\text{in}}}^{-1} R[n_n^{\text{in}}]. \quad (3.50)$$

Assuming the Jacobian exists and is Lipschitz continuous⁸, this update is shown to have quadratic convergence in some region about the root [163]. The Jacobian must be computed numerically, which can be done with either the Adler-Wiser equation Eq. (3.40) [83, 84], or with finite-difference numerical differentiation [174]. As Section 3.4.2 will explore in more depth, in the absence of further approximation, both of these techniques are inadequate for modern calculations due to the computational complexity

⁸Lipschitz continuity refers to all real $k \geq 0$ in Eq. (3.12).

and the size of the basis set. The former strategy is an $O(N_b^4)$ process that requires the computation and storage of all eigenvectors of the Kohn-Sham Hamiltonian [158, 159]. The latter strategy requires excessively many evaluations of the Kohn-Sham map K [175]. We now examine the class of methods that can be cast as a Newton step with some iteratively updated approximation to the Jacobian: quasi-Newton methods.

3.4.1.1 Broyden's methods

Consider having knowledge of an approximate Jacobian at the previous iteration, J_{n-1} . We seek a prescription for generating an approximate Jacobian at the current iteration, J_n , such that the following quasi-Newton update can be performed,

$$n_{n+1}^{\text{in}} = n_n^{\text{in}} - J_n^{-1} R_n, \quad (3.51)$$

where $R_n := R[n_n^{\text{in}}]$. First, note that all methods of this kind must begin from some initial estimate of the Jacobian, J_0 . For lack of a better option, this can be taken as a scaled identity, $J_0 = \alpha I$. Although, in the present context, the Kerker matrix is used, which is defined in Section 3.4.2. We begin with a description of Broyden's two methods [176]. These methods, as they are about to be presented, are not commonly used in modern Kohn-Sham software. However, the conceptual foundation of Broyden's methods, that is, *low rank* updates to a Jacobian that satisfies *secant conditions*, remain foundational to contemporary methodology. First, the meaning of a secant condition is defined. For illustrative purposes, a finite-difference approximation for the derivative of a one-dimensional function f at the current iterate x_n is given by

$$f'(x_n) = \frac{f(x_n) - f(x_{n-1})}{x_n - x_{n-1}}, \quad (3.52)$$

which is increasingly accurate as the iterates become closer. Since the Jacobian is the derivative of the residual map, the N_b -dimensional equivalent of this finite-difference equation is

$$J_n(n_n^{\text{in}} - n_{n-1}^{\text{in}}) = R_n - R_{n-1}, \quad (3.53)$$

where hereafter we define $\Delta n_n^{\text{in}} = n_n^{\text{in}} - n_{n-1}^{\text{in}}$ and $\Delta R_n = R_n - R_{n-1}$. If the n^{th} Jacobian satisfies Eq. (3.53), it is said to satisfy the secant condition of the current iteration, and thus belongs to Broyden's family of methods. Since J_n is an $N_b \times N_b$ matrix, and the secant condition only specifies how J_n acts on the vector Δn_n^{in} , there are a remaining $N_b^2 - N_b$ components of the Jacobian that are yet unspecified. Broyden fixes these remaining components by requiring J_n acts on all vectors *orthogonal* to Δn_n^{in} similarly to J_{n-1} . This is equivalent to requiring that the Jacobian of the current iteration solves the following constrained quadratic programming problem,

$$\text{minimise } \|J_n - J_{n-1}\|_f \quad (3.54)$$

$$\text{subject to } J_n \Delta n_n^{\text{in}} = \Delta R_n, \quad (3.55)$$

as demonstrated by Refs. [177, 178], which defines Broyden's *first method*. The Frobenius norm $\|.\|_f$ of a square matrix A is defined as

$$\|A\|_f = \sqrt{\sum_{i,j=1}^N |a_{ij}|^2}. \quad (3.56)$$

In other words, the current Jacobian J_n is required to satisfy the current secant condition, and otherwise minimise the difference between itself and the previous Jacobian J_{n-1} in the sense of the Frobenius norm. Note that the n^{th} Jacobian satisfies *all* of the previous n secant equations provided the past iterates are mutually orthogonal, $(\Delta n_i^{\text{in}})^\dagger \Delta n_j^{\text{in}} = 0$ for $i \neq j$. However, the space of past iterates is often *linearly independent*, but not mutually orthogonal. Therefore, if one requires J_n to satisfy only the most recent secant equation, one *loses* information about past secant equations, i.e. J_n no longer satisfies the past secant equations. Schemes that ensure J_n satisfies multiple previous secant equations are studied in the next section.

The constrained optimisation problem of Eq. (3.54) has a unique analytic solution, which is obtained in Refs. [177, 178] by means of unconstrained optimisation using the method of Lagrange multipliers,

$$J_n = J_{n-1} + \frac{\Delta R_n - J_{n-1} \Delta n_n^{\text{in}}}{\|\Delta n_n^{\text{in}}\|_2^2} (\Delta n_n^{\text{in}})^\dagger. \quad (3.57)$$

The notation uv^\dagger defines the outer product of the vectors u, v (Dirac bra-ket notation is not used in this section). One can now observe from Eq. (3.57) that this prescription has lead transparently to a rank-one update of the Jacobian at each iteration. The full quasi-Newton update for Broyden's first method involves subsequently inverting Eq. (3.57), applying it to residual vector, and performing the quasi-Newton step Eq. (3.51). The apparent excessive cost of inverting Eq. (3.57) is negated as the inverse of a rank-one matrix can be computed analytically using the Sherman-Morrison-Woodbury formula [179]. Furthermore, as matrix-vector multiplication is associative, one can compute the *vector* $J_n^{-1} R_n$ without constructing or storing J_n explicitly, and instead using a series of vector-vector products. This was originally demonstrated in Ref. [180], so that at a given instance Broyden's first method only requires the storage of two N_b -length vectors, and the computation of a few vector-vector products. Broyden's second method optimises the components of the matrix $H_n := J_n^{-1}$ directly via

$$\text{minimise } \|H_n - H_{n-1}\|_f \quad (3.58)$$

$$\text{subject to } \Delta n_n^{\text{in}} = H_n \Delta R_n, \quad (3.59)$$

instead of optimising the Jacobian, then subsequently inverting. Hereafter, methods that optimise the Jacobian are referred to as 'type-I' methods, and methods that optimise the inverse Jacobian are referred to as 'type-II' methods, see Ref. [171]. Note that the constraint in Eq. (3.59) is simply the inverse secant condition. Similarly to Broyden's

first method, this has the analytic solution,

$$H_n = H_{n-1} + \frac{\Delta n_n^{\text{in}} - H_{n-1} \Delta R_n}{\|\Delta R_n\|_2^2} (\Delta R_n)^\dagger, \quad (3.60)$$

which can be substituted directly into the quasi-Newton step⁹. The conventional wisdom has emerged that Broyden's second method tends to provide more robust and efficient convergence than Broyden's first method. However, both methods are shown to be q -superlinearly convergent [163, 177] in the sense that

$$\lim_{n \rightarrow \infty} \frac{\|n_{n+1}^{\text{in}} - n_*\|_2}{\|n_n^{\text{in}} - n_*\|_2} = 0, \quad (3.61)$$

which is a necessary condition for some $q > 1$ in Eq. (3.47). Broyden's second method is implemented and tested in Section 3.6.

3.4.1.2 Multisecant Broyden's methods

A natural extension to Broyden's methods is to consider all prior secant conditions *at each iteration*, rather than just the most recent secant condition. This leads to a so-called *generalised* or *multisecant* version Broyden's methods, which are examined extensively in both optimisation and electronic structure literature [163, 171, 180, 182]. The ensuing summary follows a similar structure to that of Ref. [171]. A multisecant method is defined as a method that generates an iterative Jacobian J_n such that this Jacobian satisfies the most recent m secant conditions. That is, the following $N_b \times m$ matrices are defined

$$\Delta \mathcal{R}_n := (\Delta R_{n-m+1}, \Delta R_{n-m+2}, \dots, \Delta R_n) \quad (3.62)$$

$$\Delta \mathcal{P}_n := (\Delta n_{n-m+1}^{\text{in}}, \Delta n_{n-m+2}^{\text{in}}, \dots, \Delta n_n^{\text{in}}) \quad (3.63)$$

such that a Jacobian satisfying the previous m secant conditions must satisfy the matrix equation

$$J_n \Delta \mathcal{P}_n = \Delta \mathcal{R}_n. \quad (3.64)$$

The parameter m introduced here defines the *history length*, i.e. the number of iterates that are stored and used for secant conditions. If m is less than the full history size n then the method takes on its modified *limited memory* form. If $m = n$, then the method satisfies *all* prior secant conditions. The generalisation of Broyden's two methods is now readily established: alter the constraints in the optimisation problems Eqs. (3.54) and (3.58) to reflect the multisecant condition Eq. (3.64). The multisecant version of

⁹Note that an alternate form of Broyden's updates in terms of the initial estimate J_0 can be determined via recursion. This is omitted here but can be found, for example, in Refs. [61, 62, 181].

Broyden's first and second method respectively are

$$\begin{aligned} & \text{minimise } \|J_n - J_{n-1}\|_f & (3.65) \\ & \text{subject to } J_n \Delta \mathcal{P}_n = \Delta \mathcal{R}_n, \end{aligned}$$

$$\begin{aligned} & \text{minimise } \|H_n - H_{n-1}\|_f & (3.66) \\ & \text{subject to } \Delta \mathcal{P}_n = H_n \Delta \mathcal{R}_n, \end{aligned}$$

which are of type-I and type-II respectively. These both have a unique analytic solution in the form of a rank- m update,

$$\begin{aligned} J_n &= J_{n-1} + (\Delta \mathcal{R}_n - J_{n-1} \Delta \mathcal{P}_n) (\Delta \mathcal{P}_n^\dagger \Delta \mathcal{P}_n)^{-1} \Delta \mathcal{P}_n^\dagger, \\ H_n &= H_{n-1} + (\Delta \mathcal{P}_n - H_{n-1} \Delta \mathcal{R}_n) (\Delta \mathcal{R}_n^\dagger \Delta \mathcal{R}_n)^{-1} \Delta \mathcal{R}_n^\dagger, \end{aligned}$$

which are found by solving the associated Lagrangian problems. The former Jacobian update can be inverted similarly to Broyden's first method with the Sherman-Morrison-Woodbury formula. As Refs. [169, 172] conclude, and Section 3.6 also examines, the type-II variant tends to outperform the type-I variant in the context of multiseant Broyden's methods, in line with the conventional wisdom from Broyden's original methods. As stated previously, if the space of past iterates is mutually orthogonal, this method is equivalent to Broyden's original methods.

Finally, we remark on the connection between the above methods and the method examined by Eyert, Vanderbilt & Louie, and Johnson in Refs. [181–183]. These appear, at first glance, like a natural generalisation of Broyden's methods, but are ultimately equivalent in the optimal case.

First, the following *unconstrained* minimisation problem for variations in H_n is defined,

$$\begin{aligned} & \text{minimise } w_0 \|H_n - H_{n-1}\|_f \\ & \quad + \sum_{i=n-m+1}^n w_i \|\Delta \mathcal{P}_n - H_n \Delta \mathcal{R}_n\|_2^2, \end{aligned} \quad (3.67)$$

where we choose to update the inverse Jacobian H_{n-1} , although a similar method can be formulated in terms of Jacobian updates. The weights $\{w_0, w_i\}$ are introduced as free parameters that act as *penalty coefficients*. That is, the weights are chosen to signify how 'important' it is to satisfy the corresponding constraint. In this sense, inspection of Eq. (3.67) shows that w_0 controls the degree to which the inverse Jacobian can change iteration-to-iteration, and w_i controls the degree to which the i^{th} secant equation should be satisfied by H_n . Therefore, this method also constitutes a multiseant method, but the multiseant conditions are allowed to be weighted according to relative importance. Various common fixed-point methods can be recovered as special cases of these weights. Notably, as Refs. [61, 62] demonstrate, the choice $w_i = 0$ for $i < n$, and $w_0 \ll w_n$, leads to Broyden's second method. This can be intuited from Eq. (3.67): the weights w_i now favour exclusively the most recent secant condition, and in directions orthogonal to that secant condition, the minimum norm condition on H_n is applied. In the original work of Refs. [182, 183], the weights $w_i = (R_i^\dagger R_i)^{-1}$ are considered, which favour secant conditions closer to convergence. This was used in the context of electronic structure

calculations with success in Refs. [181–183]. However, as Ref. [181] demonstrates, the optimal set of weights require $w_0 \rightarrow 0$, and if $\{w_i\}$ are to be non-zero, these weights in fact *cancel* in the update formula. Hence, $w_i = 1$ can be set without loss of generality, and the method can be identified with a standard multiseccant method; see Ref. [181] for additional detail. An interesting aspect of the multiseccant methods discussed here are their relationship Pulay’s or Anderson’s method – a ubiquitous method in electronic structure theory software – which is now examined.

3.4.1.3 Pulay’s Method

Pulay’s method [184, 185], or the discrete inversion in the iterative subspace (DIIS), as it is known in electronic structure literature, or Anderson’s method, as it is known in optimisation literature [186], has proven extremely effective at converging Kohn-Sham calculations. The simplicity of its formulation combined with its impressive efficiency and robustness has lead to Pulay’s method becoming the default algorithm in a range of Kohn-Sham codes [24, 61, 116, 118, 187]. The past few decades of wisdom suggest that Pulay’s method systematically outperforms the unmodified Broyden’s methods in both the single and multiseccant formulation. This conclusion will be tested in Section 3.6. First, a brief review of Pulay’s method as it was originally formulated is given.

Consider constructing a so-called ‘optimum’ residual – a residual whose argument is an optimum density – as a linear combination of past residuals in the m -dimensional iterative subspace,

$$R[n_{\text{opt}}^{\text{in}}] = \sum_{i=n-m+1}^n c_i R[n_i^{\text{in}}]. \quad (3.68)$$

Here, optimum is defined by the method one chooses to fix the coefficients c_i . In Pulay’s method, these coefficients are fixed by requiring that the L^2 -norm of the residual is minimal, i.e. solve

$$\begin{aligned} &\text{minimise } \|R[n_{\text{opt}}^{\text{in}}]\|_2^2 \\ &\text{subject to } \sum_i c_i = 1, \end{aligned} \quad (3.69)$$

where the constraint that the coefficients must sum to unity is an exact requirement at convergence. Substitution of Eq. (3.68) into Eq. (3.69), and use of Lagrange multipliers, allows the optimisation problem to be cast as an $(n+1)$ -dimensional linear system,

$$\begin{pmatrix} R_1^\dagger R_1 & R_1^\dagger R_2 & \dots & R_1^\dagger R_n & 1 \\ R_2^\dagger R_1 & \ddots & & & 1 \\ \vdots & & & & \vdots \\ R_n^\dagger R_1 & & & & c_n \\ 1 & 1 & \dots & \dots & 0 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ c_3 \\ \vdots \\ c_n \\ \lambda \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix},$$

for $n \leq m$, which is readily generalised for $n > m$. Assuming the space of past iterates is of full rank (comprised of linearly independent vectors), solution of this linear system

provides the set of coefficients c_i . Given these coefficients, the density update remains to be defined. Following Refs. [137, 170, 171, 181, 188], the optimum residual can be first be expanded as such,

$$R[n_{\text{opt}}^{\text{in}}] = \sum_{i=n-m+1}^n c_i K[n_i^{\text{in}}] - \sum_{i=n-m+1}^n c_i n_i^{\text{in}}. \quad (3.70)$$

If K is assumed to be linear, the rightmost term in Eq. (3.70) can be interpreted as the optimal input density,

$$n_{\text{opt}}^{\text{in}} = \sum_{i=n-m+1}^n c_i n_i^{\text{in}}. \quad (3.71)$$

Hence, the optimal *update* can take the standard undamped form

$$n_{n+1}^{\text{in}} = n_{\text{opt}}^{\text{in}} + R[n_{\text{opt}}^{\text{in}}] \quad (3.72)$$

$$= \sum_{i=n-m+1}^n c_i K[n_i^{\text{in}}]. \quad (3.73)$$

This update is favoured over $n_{n+1}^{\text{in}} = n_{\text{opt}}^{\text{in}}$ so that the algorithm does not stagnate in the subspace of past input densities. Alternatively, as originally studied in Ref. [186], a damped step can be taken,

$$n_{n+1}^{\text{in}} = n_{\text{opt}}^{\text{in}} + \alpha R[n_{\text{opt}}^{\text{in}}], \quad (3.74)$$

for $\alpha \in (0, 1)$. An example algorithm that implements this formulation of Pulay's method is given in Algorithm. 1.

Algorithm 1 Pulay's Algorithm

- 1: **Input:** tol, n_0^{in} , m , α
 - 2: **for** $n=0,1,2,\dots$
 - 3: Compute $R[n_n^{\text{in}}]$ and store the pair $\{n_n^{\text{in}}, R[n_n^{\text{in}}]\}$
 - 4: Solve Eq. (3.69) for $\{c_i\}$
 - 5: **if** $n \leq m$
 - 6: Set $n_{n+1}^{\text{in}} = \sum_{i=1}^n c_i (n_i^{\text{in}} + \alpha R[n_i^{\text{in}}])$
 - 7: **else**
 - 8: Set $n_{n+1}^{\text{in}} = \sum_{i=n-m+1}^n c_i (n_i^{\text{in}} + \alpha R[n_i^{\text{in}}])$
 - 9: **if** $\|R[n_n^{\text{in}}]\|_2 < \text{tol}$, **exit**
-

At a first glance, Pulay's method bears little resemblance to the secant-based methods discussed in the previous section. However, as described in Refs. [137, 171, 188], a rearrangement of the optimisation problem in Eq. (3.69) reveals a close relationship between Pulay's method and type-II multisection methods. A more detailed treatment of this correspondence is found in Refs. [137, 171, 188]; here, we simply state that the following unconstrained optimisation problem

$$\text{minimise } \|R_n - \Delta \mathcal{R}_n \gamma\|_2, \quad (3.75)$$

for variations in $\gamma = (\gamma_{n-m+1}, \gamma_{n-m+2}, \dots, \gamma_n)$ is equivalent to Pulay's optimisation problem Eq. (3.69). The new coefficients $\{\gamma_i\}$ are related to the old coefficients $\{c_i\}$ such that the update in Eq. (3.74) now takes the form

$$n_{n+1}^{\text{in}} = n_n^{\text{in}} + \alpha R_n - (\Delta \mathcal{P}_n + \alpha \Delta \mathcal{R}_n) \gamma, \quad (3.76)$$

where γ on iteration n is solved by

$$\gamma = (\Delta \mathcal{R}_n^\dagger \Delta \mathcal{R}_n)^{-1} \Delta \mathcal{R}_n^\dagger R_n. \quad (3.77)$$

The parallel between Pulay's method and multiseant methods becomes apparent when these equations are combined to give the final update,

$$n_{n+1}^{\text{in}} = n_n^{\text{in}} + H_n R_n, \quad (3.78)$$

$$H_n = \alpha I - (\Delta \mathcal{P}_n - \alpha \Delta \mathcal{R}_n) (\Delta \mathcal{R}_n^\dagger \Delta \mathcal{R}_n)^{-1} \Delta \mathcal{R}_n^\dagger. \quad (3.79)$$

By comparison with the updated inverse Jacobian in Eq. (3.66), we can observe that Pulay's method is a type-II quasi-Newton step where the iterative inverse Jacobian is updated according to

$$\begin{aligned} & \text{minimise } \|H_n - H_0\|_f & (3.80) \\ & \text{subject to } \Delta \mathcal{P}_n = H_n \Delta \mathcal{R}_n. \end{aligned}$$

In other words, this optimisation problem minimises the difference between the components of the inverse Jacobian H_n and the initial guess inverse Jacobian H_0 , while also requiring the previous m secant conditions to be fulfilled – *there is thus a very close relationship between Pulay's method(s) and Broyden's methods*. Note that $H_0 = \alpha I$ is required in order to recover the update in Eq. (3.79). This reformulation not only connects Pulay's method to the type-II variant of multiseant Broyden's method, but also uncovers another flavour of Pulay's method,

$$\begin{aligned} & \text{minimise } \|J_n - J_0\|_f & (3.81) \\ & \text{subject to } J_n \Delta \mathcal{P}_n = \Delta \mathcal{R}_n, \end{aligned}$$

which is of type-I, i.e. the Jacobian is optimised, rather than the inverse Jacobian. This form of Pulay's method, as originally described in Ref. [171], has seen comparatively less application and testing in the context of Kohn-Sham codes [137]. These methods differ from multiseant Broyden methods precisely when $m > n$, in which case the multiseant Broyden methods retain information from all prior secant equations implicitly, whereas Pulay's method(s) ignore completely secant equations not in the (size m) history.

A few modern modifications to Pulay's method are now examined, although we note that these modifications are adaptable to all the secant-based methods discussed previously. These methods attempt to tackle some of the concerns raised in a recent article by Anderson [189].

Periodic Pulay. The work in Ref. [190], based on Ref. [191], suggests that alternating between Pulay and linear mixing steps can improve the robustness of the iterations over standard Pulay. Each new Pulay step utilises the history from the linear mixing and past Pulay steps to solve the optimisation subproblem Eq. (3.69). This is demonstrated

to have a stabilising effect as the linear mixing history data is used well by the Pulay extrapolation. In Ref. [190], an input parameter k determines the number of linear mixing steps performed between each Pulay step, i.e. k linear steps per Pulay step. As suggested in the original work, the values $k = 2$ are tested in Section 3.6 with a damping parameter of $\alpha = 0.2$ for both the Pulay and linear mixing steps.

Restarted Pulay. The work in Ref. [192] considers occasionally flushing the history every time a certain criterion is met, rather than iteratively overriding the history. This criterion is chosen to be whenever the current iteration number is an integer multiple of the maximum history size, i.e. $n = am$ for some $a \in \mathbb{Z}^+$. For inputs with a considerable degree of non-linearity, either due to a poor initial guess, or inherent to the Kohn-Sham map, the history can actively interfere with modelling an accurate iterative Jacobian at the current iteration. Restarted Pulay thus represents a strategy for dealing with this issue by periodically removing the history. It can also prevent linear dependencies arising in the history.

Guaranteed Reduction Pulay. Before one gets too excited, the ‘guaranteed reduction’ is in the *Pulay-predicted* residual $\|R_{\text{opt}}\|_2$, rather than the *actual* residual [193]. This is achieved by rearranging the stored history of residuals $\{R_i\}$ such that, at a given iteration, the Pulay predicted optimal residual is added to the history, rather than the residual obtained from evaluating the Kohn-Sham map. The subsequent iteration then involves a linear mixing step, which generates a new exact residual that is added to the history. The coefficients $\{c_i\}$ of the now current iteration are determined by solving Pulay’s optimisation problem Eq. (3.69). However, note that the previous Pulay predicted optimal residual is an element of the set of residuals that are used to determine $\{c_i\}$. Hence, the addition of the residual from the linear mixing step can only *lower* the Pulay predicted optimal residual, or at worst leave it the same. This new reduced Pulay predicted optimal residual replaces the exact linear mixing residual in the history, and the process repeats. As expected, this method performs best when the predicted optimal residual accurately models what the residual *would have been* were the optimal density evaluated with the Kohn-Sham map. Pulay’s method predicts the residual increasingly well the closer it is to the linear response regime from the root. Therefore, when the behaviour of Kohn-Sham map is highly non-linear, the guaranteed reductions in the predicted residual tend to stagnate, while the exact residual does not decrease. Variations on a theme of these techniques are benchmarked in Section 3.6.

3.4.1.4 Holistic Approaches based on Multisecants

Here, we highlight a few modern multisecant methods that take a ‘holistic’ approach, i.e. the algorithms have many facets to their design all working in tandem in order to stabilise the iterations. This line of development has, in the opinion of the author, the most favourable future prospects in regards to black-box methodology (alongside more sophisticated ‘direct minimisation’ algorithms [112]).

First, we consider the methods outlined in Refs. [169, 172], which are now default self-consistency methods in WIEN2k [194]. These algorithms are designed with the aim of

converging Kohn-Sham calculations in particular. The range of strategies utilised make Refs. [169, 172] an interesting case¹⁰.

First, the updates considered in Refs. [169, 172] are defined in Eqs. (3.81) and (3.80), which are of the form

$$H_n = \alpha I + (\Delta\mathcal{P}_n - \alpha\Delta\mathcal{R}_n)(\Delta\mathcal{R}_n^\dagger W_n)^{-1} W_n^\dagger. \quad (3.82)$$

The initial guess inverse Jacobian is $H_0 = \alpha I$, and $W_n = \Delta\mathcal{P}_n$, $W_n = \Delta\mathcal{R}_n$ define a type-I and type-II update respectively. It is demonstrated in Ref. [172], and further verified in Section 3.6, that type-II methods are superior for the self-consistency problem than type-I methods. However, if atomic degrees of freedom are included, it is advantageous to consider a linear combination of updates,

$$W_n = Y_n + \beta S_n, \quad (3.83)$$

for $\beta \in \mathbb{R}_{\geq 0}$. The parameter β controls the degree to which the method takes a type-II step, $\beta = 0$, or a type-I step, $\beta \rightarrow \infty$. As noted in Ref. [169], this is similar to a technique used in Ref. [195] whereby a criterion is defined to assess whether a type-I or type-II step will be optimal, and then the corresponding step is taken. In MSR1, the parameter β is determined based on an ansatz that seeks to ensure the eigenvalues of the Jacobian are positive, as they should be in the case that the fixed-point corresponds to a variational minimum. However, as stated, in the context of density mixing type-II methods consistently outperform type-I methods, meaning we hereafter consider $\beta = 0$.

Second, a core novelty of the methods in Refs. [169, 172] involve partitioning of the full update into a *predicted* and *unpredicted* component, which are now defined. Consider the update generated from Eq. (3.82),

$$n_{n+1}^{\text{in}} = n_n^{\text{in}} + H_n R_n, \quad (3.84)$$

which is now split in two,

$$n_{n+1}^{\text{in}} = n_n^{\text{in}} + (H_n^{\text{p}} + H_n^{\text{u}})R_n. \quad (3.85)$$

Since we are now considering the type-II variant, $W_n = \Delta\mathcal{R}_n$, the unpredicted component of the update, $H_n^{\text{u}}R_n$, is defined as the orthogonal projection of the current residual R_n onto the past residual differences,

$$(H_n^{\text{u}}R_n)^\dagger \Delta\mathcal{R}_n = 0. \quad (3.86)$$

In other words, the unpredicted vector is the part of the full update that is not described within the iterative subspace of residuals. In this sense, the remaining update can be considered to be the part of the update that the iterative subspace does describe. Eq. (3.86) is shown in Ref. [172] to have the solution,

$$H_n^{\text{u}}R_n = -(I - \Delta\mathcal{R}_n(\Delta\mathcal{R}_n^\dagger \Delta\mathcal{R}_n)^{-1} \Delta\mathcal{R}_n^\dagger)R_n. \quad (3.87)$$

¹⁰The most recent published form of these algorithms is that given in Ref. [169] titled ‘multisecant rank one’ (MSR1). However, this algorithm is designed to converge both the atomic (geometry optimisation) and electronic degrees of freedom. Hence, we focus on the techniques that are relevant to the self-consistent field iterations, and the reader is referred to Refs. [169, 172] for a more in-depth treatment.

The update now takes the rearranged form

$$H_n^u R_n = (I - \Delta \mathcal{R}_n (\Delta \mathcal{R}_n^\dagger \Delta \mathcal{R}_n)^{-1} \Delta \mathcal{R}_n^\dagger) R_n, \quad (3.88)$$

$$H_n^p R_n = -S_n (\Delta \mathcal{R}_n^\dagger \Delta \mathcal{R}_n)^{-1} W_n^\dagger R_n. \quad (3.89)$$

This partitioning is used to introduce the concept of *algorithmic greed*, which is quantified with an iterative damping parameter α_n that multiplies the *unpredicted update*, rather than the full update,

$$n_{n+1}^{\text{in}} = n_n^{\text{in}} + (H_n^p + \alpha_n H_n^u) R_n. \quad (3.90)$$

In the implementation tested in Section 3.6, the unpredicted direction is also Kerker preconditioned (see Section 3.4.2). Refs. [169, 172] refer to the updating of the parameter α_n as an *implicit trust region*. That is, α_n is allowed to increase, within some bounds, provided the algorithm is performing well, by some definition of ‘well’ which defines the *greed controls*. If the algorithm is performing poorly, the damping parameter is decreased accordingly, thus reducing the need for user intervention. In the CASTEP implementation, a simple set of greed controls is designed that perform well in the context of plane-wave pseudopotential codes. Namely, the parameter α_n is updated as such,

$$\alpha_{n+1} = \alpha_n \left[1 + \beta \left(\frac{C_n}{\gamma} - 1 \right) \right] \quad (3.91)$$

$$0.1 < \alpha_{n+1} < 1.0, \quad (3.92)$$

where C_n is the average over the past k residual ratios,

$$C_n = \left(\sum_{i=n-k-1}^n \frac{|R_{i-1}|}{|R_i|} \right) / k, \quad (3.93)$$

and α_0 , β , and γ are parameters. The parameter α_0 is the initial step length in the unpredicted direction, β scales the strength of the fluctuations in α_{n+1} , and γ determines whether this fluctuation should increase or decrease α_n .

Furthermore, the matrix inverse in Eq. (3.82) is Tikhonov regularised [196] to prevent spurious behaviour due to rank-deficiencies, and the matrices involved are scaled to prevent run-away numbers that are large in magnitude. The method that has been described thus far is similar to ‘multisecant Broyden 2’ (MSB2) of Ref. [172]. The type-I variant, MSB1, can be derived in a similar fashion, and both are tested in Section 3.6.

Greed controls of this kind that have just been discussed are a somewhat simple method of ensuring sure-footed iterative steps are made. Borrowing from optimisation (rather than fixed-point) literature, a further stabilising technique constitutes an explicit trust region. Given R_n as a descent direction on iteration n , the trust region subproblem can take a standard form

$$\text{minimise } \|R_n - H_n^{-1} X\|_2^2 \quad (3.94)$$

$$\text{subject to } \|X\|^2 - \delta^2 < 0 \quad (3.95)$$

for variations in X with some trust region radius δ . The variable X is the new trial step generated from the trust-region subproblem, and the scalar $\|X\|_2$ is the trial total

step length, where the exact step length is $\|H_n R_n\|_2$. If the exact step length exceeds the trust region radius δ , then the trust region problem is required to be solved to generate this new step. Details on how the trust-region subproblem is solved are given in Refs. [163, 169]. Some form of either a trust region or sophisticated greed control will be required for future applications, perhaps also in combination with the ability to ‘back-track’ poor steps.

A variety of methods have been appearing in literature as of late aiming to take a more sophisticated approach beyond those presented in the prior sections. Another notable example is a variant of the type-I Pulay update in Ref. [137]. A number of distinctive concepts ensure convergence toward a local minimum, including weaving in linear mixing steps when convergence is poor, and having a novel criterion to discard linear dependencies in the iterative history.

3.4.2 Preconditioning

Preconditioning refers to the modification of an optimisation problem such that the condition number of the problem is improved. Crucially, the modified problem is required to have the same minimum, and minimiser, as the original problem. Algorithms applied to the modified problem thus have more stable and accelerated convergence. A preconditioner is most transparently defined for linear systems as being the matrix P such that

$$P^{-1}Ax = P^{-1}b, \quad (3.96)$$

where the $P^{-1}A$ has a lower condition number Eq. (3.34) than A . In the case $P = A$ the linear system is solved, and P is the exact preconditioner. The definition of a preconditioner for non-linear systems is less transparent. Consider the optimisation of the Kohn-Sham residual L^2 -norm, which now takes the preconditioned form

$$\|P(R[n_*])\|_2 = 0. \quad (3.97)$$

The preconditioned residual $P(R)$ is required to have the same solution as R , but has, in some sense, improved convergence properties. The perfect preconditioner here would modify R such that only one step of appropriate size in the steepest descent direction is required for convergence. A successfully preconditioned problem therefore represents a problem whose landscape is easier to traverse toward a minimum using, for example, Newton’s algorithm.

It is known that the Jacobian (Hessian) eigenvalue spectrum of the objective function determines the rate of convergence of Newton, and quasi-Newton, methods [139, 163, 177]. Therefore, a preconditioner should accomplish one or multiple of the following: reduce the number of eigenvalue clusters; reduce the width of the eigenvalue clusters; or compress the spectrum as a whole. A discussion on the importance of the clustering of the eigenvalues, rather than just the condition number, can be found in Refs. [163, 169, 197].

To simplify matters, we consider P to be a *linear transformation* (matrix) and furthermore assume it is *constant* with respect to the optimised variable (the density) such

that $P(R) := PR$ where $P \in N_b \times N_b$. Hence, in the present context, preconditioning amounts to finding the matrix P such that the spectrum of the dielectric is more suitable for quasi-Newton algorithms. Note that the preconditioning matrix is permitted to change iteration-to-iteration (but not through some explicit dependence on the density)¹¹. The strategy used in practice is to identify the source of divergent eigenvalues of the dielectric, as examined in Section 3.3, and temper this divergence in such a fashion that is *generally applicable* to all, or large classes, of Kohn-Sham inputs. Depending on the implementation, the preconditioning approach can differ. For example, in augmented plane-wave implementations [194], the unit cell is partitioned into the regions surrounding atomic cores, represented by local basis functions, and an interstitial region, represented by plane-waves. Naturally, due to the differing number of basis functions involved in each region, among other properties, the preconditioning for each region is separate [169, 172]. The following work assumes an entirely plane-wave basis set, although the preconditioners can be adapted for a variety of implementations.

Recall from Section 3.3.7 that the Coulomb kernel, in combination with the susceptibility, is principally responsible for the large eigenvalues of the Kohn-Sham residual linear response function. As discussed, in the general case, the susceptibility is a complicated object about which it is difficult to make sweeping statements. However, when the system is homogeneous and isotropic, the dielectric eigenvalues of gapped and gapless phases are approximately determined by

$$\varepsilon_0^{-1} = \left(1 + \frac{4\pi\gamma}{|G|^2}\right)^{-1}, \quad (3.98)$$

$$\varepsilon_0^{-1} = (1 + 4\pi\gamma)^{-1}, \quad (3.99)$$

for some *a priori* unknown system-dependent constant γ . In the case of the homogeneous electron gas, this constant in Eq. (3.98) is identified with the square of the Thomas-Fermi screening wavevector k_{TF}^2 . Modest departures from homogeneity and isotropy remain accurately modelled by Eqs. (3.98) and (3.99), particularly in the low $|G|$ limit [157, 168]. Therefore, these *model dielectrics* can be used to improve the condition of the residual map by allowing $P = \varepsilon_0^{-1}$ for either Eqs. (3.98) or Eqs. (3.99) depending on whether one suspects the input to be metallic or insulating.

Collecting and relabelling the unknown constants, the preconditioner becomes

$$P = \alpha \frac{|G|^2}{|G|^2 + |G_0|^2}, \quad (3.100)$$

where $|G_0|$ and α are parameters that are determined by the linear response of the input system; e.g. $|G_0| = 0$ for Kohn-Sham insulators, and $|G_0|$ is related to the Wigner-Seitz radius for the homogeneous electron gas. The values of $|G_0|$ and α naturally differ depending on the input, although fixing $|G_0| = 1.5\text{\AA}$ and $\alpha = 0.8$ [61, 62] demonstrably

¹¹This distinction is made because, if P were to have a dependence on n , then the quasi-Newton schemes targeting the optimisation problem in Eq. (3.97) would require derivatives of $P(R)$, rather than just R (and then applying P afterward).

improves convergence, see Section 3.6. The modified Kohn-Sham problem is now solved

$$\left\| \alpha \frac{|G|^2}{|G|^2 + |G_0|^2} R[n_*(G)] \right\|_2 = 0, \quad (3.101)$$

which is referred to as Kerker preconditioned Eq. (3.100) [164, 165]. The Kerker preconditioner suppresses charge sloshing, as defined in Section 3.3.7, by damping eigenvalues of the dielectric corresponding to low $|G|$ components of the density, see Fig. 3.8. It is these components that have a generically amplified response due to the Coulomb kernel. The adaptation of the Kerker preconditioner to real space implementations of Kohn-Sham theory is more difficult. The dielectric response function of the homogeneous electron gas is non-local in real space, meaning the integral in Eq. (3.30) leads to a dense $N_b \times N_b$ matrix that must be computed and stored. The susceptibility takes the Yukawa screening form in real space [155]. An efficient real space implementation of the Kerker preconditioner is given in Ref. [198], and an implementation of the Kerker preconditioner in an augmented plane-wave setting (FLEUR) is given in Ref. [199].

One aspect of the preconditioning presented here is that, for increasingly homogeneous and isotropic inputs, the dependence of the dielectric condition number on unit cell dimension $L \sim G^{-1}$ cancels [154]. This is by construction, and is identified by examining the eigenvalues of the Kerker preconditioned dielectric. For inputs that are not so accommodating, while the Kerker preconditioner does help, the scaling of iterations with unit cell dimension persists [36, 154]. Removing this scaling can be considered one of the primary goals of preconditioners in Kohn-Sham theory as large simulation cells are required for many modern applications. Moderate extensions to the Kerker preconditioner have been proposed [200], which involve more accurately modelling the dielectric response. A number of homogeneous and isotropic dielectric models have been considered in this work, e.g. Refs. [155, 201–203]. However, in general, it is found that these models perform similarly as preconditioners (up to a change of parameters) because the main drawback of the Kerker preconditioner remains: there is no scope for systematically including anisotropy and inhomogeneity. Furthermore, the exchange-correlation kernel is ignored. This is a reasonable practice in non-spin-polarised systems, but in spin polarised systems, the spin density interacts entirely through the exchange-correlation kernel, and is thus not preconditioned.

There have been various efforts to construct a preconditioner that provides an improved description of inhomogeneous and anisotropic inputs. In the most extreme examples, computation of the exact linear response function is considered [158–161, 166, 204]. As discussed, the exact linear response function does not represent a preconditioning scheme, rather it is the exact Newton method. A density-dependent preconditioner necessarily alters the Jacobian, Eq. (3.97), in a non-trivial manner. Hence, it is not obvious the extent to which the condition of the modified problem will improve. Nevertheless, even approximate attempts at computing the exact response function, when treated as a Newton step, are able to improve the iterations over accelerated fixed-point algorithms such as Pulay’s method [158, 204]. The central difficulty in computing the exact susceptibility, and subsequently the exact dielectric, is that it requires a summation over all unoccupied-occupied eigenfunction pairs Eq. (3.40). In plane-wave codes, both the

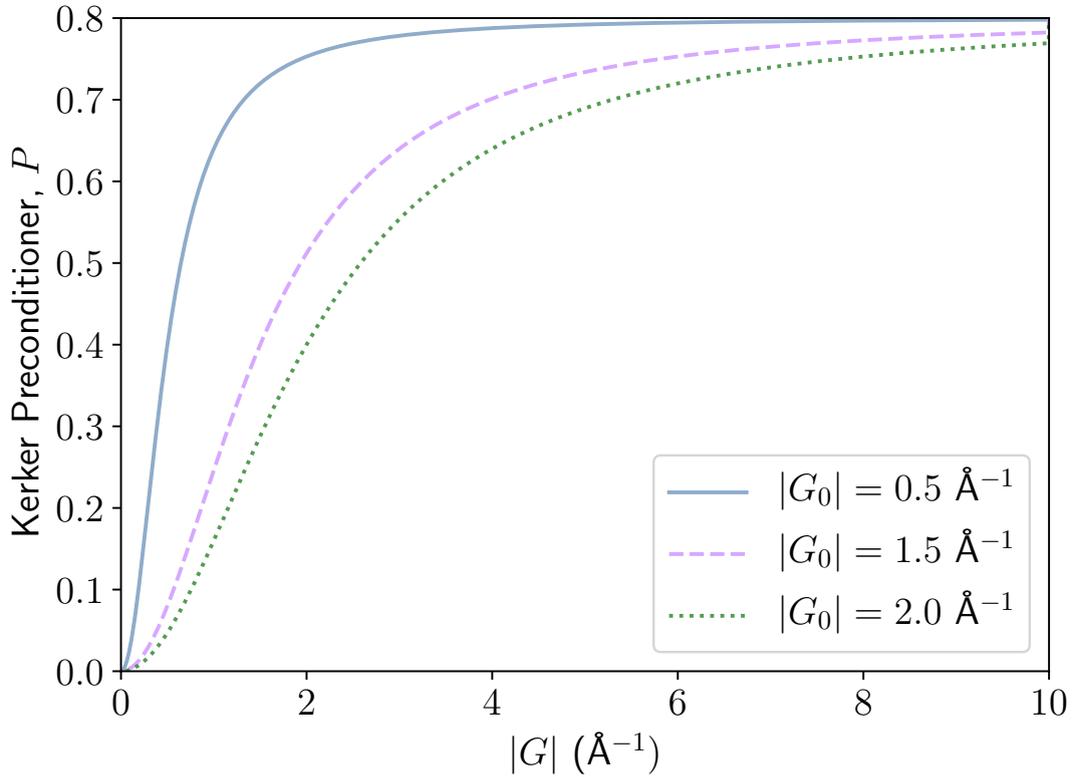


FIGURE 3.8: The Kerker preconditioner for various values of $|G_0|^2$. This figure demonstrates that the low- $|G|$ components of the iterative update are suppressed in order to counteract the low- $|G|$ Coulomb amplification that leads to charge sloshing.

computation and storage of all eigenfunctions of the Kohn-Sham Hamiltonian is infeasible.

Past attempts [159, 160] are able to implement and examine computation of the exact susceptibility with success, albeit with a basis set size now unsuited to modern computation. The problem of having to compute the full set of eigenfunctions, and having too large a basis, can be remedied with a few differing approaches. First, the size of the *effective* basis can be reduced by recalling that the low $|G|$ components of the density are those responsible for divergent eigenvalues of the dielectric. Therefore, the susceptibility need only be computed for a reduced set of plane-waves – those with low frequency G . Working in this reduced space for the purposes of density mixing leads to a significant reduction in compute and memory overhead¹² [61, 62, 158]. Second, we highlight two strategies to remove the infinite summation over unoccupied eigenfunctions in the Adler-Wiser equation Eq. (3.40).

The first strategy utilises density functional perturbation theory, and in particular the *Sternheimer equation* [205]. Solution of the Sternheimer equation allows one to obtain the first-order response in a perturbed quantity – here, the density in the direction of the

¹²In fact, a reduced set of plane-waves is also used for the mixing algorithms of Section 3.4.1 in CASTEP [61, 62]. The components of the density corresponding to G vectors outside this reduced space are treated with the fixed-point algorithm.

residual – without requiring the full eigendecomposition of the Kohn-Sham Hamiltonian. This is utilised in, for example, the implementation of the *GW* approximation [206], and hence is available functionality within many Kohn-Sham codes. The density functional perturbation theory approach is used to construct a Newton step in Ref. [204] and is shown to significantly reduce the number of iterations taken to converge when compared to Pulay’s method.

The second strategy involves exploiting the completeness of eigenfunctions,

$$\delta(x - x') = \sum_{j=0}^{\infty} \phi_j^*(x) \phi_j(x'). \quad (3.102)$$

This identity cannot be applied to the Adler-Wiser equation Eq. (3.40) without further approximation due to the denominator depending on the j^{th} eigenenergy, ϵ_j . However, these eigenenergies can be replaced with some approximate constant value \bar{E} above a certain cut-off number of eigenfunctions $N_{\text{cut-off}}$. This allows Eq. (3.102) to transform the Adler-Wiser equation into a sum over $N_{\text{cut-off}} \ll N_b$ eigenfunctions, which is used [158, 161] to successfully reduce the number of iterations taken to converge. However, note that these methods retain a poor scaling with the number of electrons; $\mathcal{O}(N^4)$ in the most recent examples. Whilst the prefactor of the scaling is much reduced compared to earlier efforts [159], such methods, without further approximation or development, are precluded for larger system sizes. A major advantage of the methods discussed here though is the ability to take into account the spin linear response function. To the authors knowledge, no other methods based around *model* dielectrics attempt to include the spin response function, and thus the spin density is often not preconditioned.

An alternate method that strikes a balance between between model dielectrics, and computations of the response function, is one based on Thomas-Fermi theory [207]. Here, Thomas-Fermi-Von Weizäcker¹³ theory is used, in combination with Pulay’s method, to generate the subsequent density. The method calculates the response that is required to bring the current input density to self-consistency within the framework of Thomas-Fermi-Von Weizäcker theory, and uses this as a model of the exact Kohn-Sham response. Minimisation of the Thomas-Fermi-Von Weizäcker functional, which is done using the conjugate gradient method, is vastly more efficient than minimisation of the Kohn-Sham functional due to that fact it is orbital-free. In certain test cases this method is demonstrated to reduce the time taken to converge by up to a factor of three, and is implemented in the software package *Quantum Espresso* [118].

3.5 Test Suite

The test suite [25], see also Table 3.1, presented here differs from available test suites (e.g. [119, 120]) as it is required to sample the range of sources of ill-conditioning discussed in Section 3.3. Therefore, it is the aim of this test suite that standard algorithms, such as Broyden’s methods, should fail to converge for approximately thirty to forty percent of cases. Furthermore, the standard algorithms should be inefficient,

¹³The Von Weizäcker kinetic energy extension to Thomas-Fermi theory serves to better model inhomogeneities in the density [167].

as defined in Section 3.6, for the majority of the remainder of the systems. The test suite is designed to consume only moderate computational resources. Some of the most taxing inputs, such as large clusters with vacuum, require approximately sixty cores and a few hours. In its current version, which is subject to change, the test suite contains fifty six systems. The geometries and relevant input parameters are given in Ref. [25] in the form of CASTEP input files, but are readily converted using, for example, Ref. [208].

The content of the test suite is compiled from a range of sources, for example, self-consistency articles (e.g. [154, 169, 172, 207, 209–211]), collaboration, and online databases. This content is now briefly motivated in terms of the theory in Section 3.3. First, we recall that within semi-local Kohn-Sham theory, the exchange-correlation approximation can be largely ignored from the point of view of ill-conditioning. Hence, the exchange-correlation approximation is not varied across the test suite, and the PBE [54] level of theory is applied throughout.

A primary source of numerical difficulty, particularly relevant to future applications, was identified as ill-conditioning arising from large unit-cell dimensions. As such, a range of inputs with varying dimensionality is included. For example, we include a relatively large rubidium cluster, a rare earth silicide in the form of a long thin pillar, a slab of gold with a large vacuum, and so on. These additions should allow the test suite to be used to evaluate different preconditioners effectively, rather than just black-box methodology. When necessary, calculations are performed with spin polarisation, where any symmetry is broken by specifying some prior spin state. This initial spin polarisation is applied following Hund’s rules in order to ensure the algorithms converge to the same fixed-point. Approximately fifty percent of the test suite is spin polarised. Moreover, a particular emphases is placed on including inputs that are aligned with contemporary research. For example, superconductivity candidates, perovskites, and phases of matter that are far from their atomic equilibrium such as those generated by structure searching algorithms. The latter in particular can tend to introduce a very high density of states about the Fermi energy, and hence a variety of electronic temperatures is in the test suite for one such out of equilibrium system. Furthermore, isolated atoms in vacuum are conventionally difficult to converge, and in certain cases have been demonstrated to display artificial phase changes during the self-consistency iterations [169]. The test suite includes examples of isolated systems with varying atomic configurations, such as oxygen, nitrogen, iron, titanium, and vanadium. Finally, further to these classes of inputs, we also include examples of interfaces, highly inhomogeneous systems, electronegative systems, supercells of conventional metals, and poorly constructed problems (e.g. undersampling k -space).

TABLE 3.1: The SCF- x_n test suite. The title given to each system corresponds to the names in Ref. [25], in which more information can be found regarding the precise structure of each system, including its parameters. The ‘difficulties’ column motivates the presence of each system in the test suite (note that the ‘empirical’ tag implies that the system was included due to its poor convergence properties).

System	Difficulties
Bulk Al	–
Al slab	Metal, large unit cell in one dimension

Bulk Al supercell	Metal, large unit cell
Au slab	Metal, large unit cell in one dimension
B ₁₂	Large unit cell, empirical
BaCuF ₄	Empirical, transition metal
Bi ₂ Te ₃ with Ge defect	Large unit cell, vacuum, defect
BiSrFeO	Empirical
C adsorption	Large unit cell, vacuum
Caffeine	–
CoFe	Empirical
Fe ₂ Can ₃	Far-from-equilibrium
Fe ₂ CO at $T = 100k$	Far-from-equilibrium, low electronic temperature
Fe ₂ CO at $T = 10k$	Far-from-equilibrium, low electronic temperature
Fe ₂ CO at $T = 1k$	Far-from-equilibrium, low electronic temperature
Fe with H on surface	Large unit cell, transition metal, vacuum
Fe on MgO surface	Large unit cell, polar, transition metal, vacuum
Fe cluster	Large transition metal cluster
Fe cluster (under-converged cut-off energy)	Under-converged parameters
GaAs	Empirical
Graphene	–
IDZ	Empirical
In MOF	Empirical
Isolated Fe	Isolated atom, vacuum
Isolated Ni	Isolated atom, vacuum
Isolated Ni with electric field	Isolated atom, vacuum, electric field
Isolated O	Isolated atom, vacuum
Isolated Ti	Isolated atom, vacuum
Isolated Ti all electron	Isolated atom, vacuum, no pseudopotential
Isolated Ti all electron (over-converged)	Isolated atom, vacuum, over-converged
Isolated V	Isolated atom, vacuum
KPSn	Empirical, far-from-equilibrium
LiCu	Far-from-equilibrium
LiCu (under-converged k -points)	Far-from-equilibrium, under-converged
Mgo	Polar (bad initial guess)
Nb ₄ AlC ₃	Empirical
NdTiO ₃	Empirical
Ni ₃ O ₄	Empirical
Ni surface	Transition metal surface
Ni surface (under-converged k -points)	Transition metal surface, under-converged
Ru cluster	Transition metal cluster
SiC slab	Empirical, large unit cell
SiHER	Empirical
SiHER (no spin)	Empirical
Siderite	Empirical
TiK	Far-from-equilibrium
TiO ₂	Transition metal oxide
TiO ₂ (high cut-off pseudopotential radius)	Transition metal oxide
TiO ₂ (low cut-off pseudopotential radius)	Transition metal oxide
TiSrLaO	Empirical

UF ₄	Empirical, Uranium
Water cluster	Large unit cell, empirical
Water cluster (under-converged cut-off energy)	Large unit cell, empirical, under-converged
YBaCuO	Empirical
Zigzag-edge graphene nanoribbon	–
ZnONb	Empirical
ZnOTa	Empirical

3.6 Results and Discussion

The aim of this section is to combine to the analysis of Section 3.3, a sample of the methods presented in Section 3.4, and the test suite of Section 3.5, in order to arrive at a workflow that can provide insight on the strengths and weaknesses of contemporary self-consistency algorithms. Hence, the following work constitutes a benchmarking effort. However, the conclusions of this benchmark are not intended to be the focus of this section. The reasons for this are two-fold. First, the benchmark is not universal. That is, the benchmark omits a variety of successful methods simply because these methods are unsuitable for the underlying DFT software used, for example, EDIIS [209]. Second, the utility measures are necessarily imperfect as a degree of detail is lost when condensing all convergence data to two scalar measures for each method. We elaborate on this comment after defining the utility measures. This section is instead primarily intended to demonstrate a workflow that can be imitated by both methods developers and DFT software developers. That is, method developers are able to utilise the test suite, and similar measures of efficiency and robustness, to present and analyse new methodology in a more transparent and systematic fashion. Additionally, DFT software developers are able to do the same in order to assess whether they wish to replace old methodology with confidence. Nonetheless, despite the caveats discussed with regards to the benchmark itself, the conclusions of the benchmark remain indicators of the kinds of techniques and principles that are proving successful, and can assist in guiding future method development.

As discussed previously, one must quantify utility precisely in order to compare and contrast differing algorithms. Here, this is done by introducing two separate measures, *robustness* and *efficiency*. Robustness is defined as the percentage of the test suite for which a given algorithm converges in less than a certain cut-off, which depends on how the efficiency measure is defined. Iterations are chosen as the efficiency measure, as discussed shortly, in which case we define two thousand iterations as the cut-off number of iterations after which an algorithm is said to have diverged – this choice depends on the content of the test suite. A robustness measure of $r = 0.6$, for example, would indicate an algorithm converges 60% of the test suite in under two thousand iterations.

Efficiency, in general, is a more complex quantity to measure. Many of the algorithms presented in Section 3.4 require a negligible amount of time to compute the update

in a given iteration, and hence *number of iterations* becomes an effective measure of efficiency. However, there exist many methods that require a significant amount of time per iteration to compute the update. Therefore, these methods demand another measure of efficiency, such as wall-clock time. The use of wall-clock time as a measure of efficiency has transferability issues as it depends heavily on the computer architecture used, number of cores, efficiency and parallel scaling of the implementation, and so on. If one is required to use a measure such as wall-clock time, one must be very careful in assuring all potential sources of interfering causal influence, like changing computer architecture, are held constant. All but one of the methods to be tested here require negligible compute time per iteration, and hence we measure efficiency by number of iterations. For the remaining method that cannot be assessed using number of iterations, namely, ensemble DFT [111], we use wall-clock time to provide an estimated number of iterations, while ensuring all the aforementioned variables are held constant.

The quantity that defines the efficiency of a given algorithm in the present context is given as

$$\eta = \left(\frac{1}{N_{\text{conv}}} \sum_{i \in \text{conv}} n_i \right)^{-1}, \quad (3.103)$$

where N_{conv} is the number of inputs for which the algorithm converged, and n_i is the iterations taken to converge for the i^{th} member of the test suite. The inverse is included such that η is larger for a more efficient method. The normalisation factor N_{conv} is included in order to separate the measures of efficiency and robustness as much as possible. If this normalisation were not included, algorithms that converge a significantly higher percentage of the test suite would spuriously appear more inefficient than they actually are – this effect is still present, but diminished.

A further complication is the following. Inputs which generically take a larger number of iterations to converge (e.g. large systems) will have a much higher method-to-method variance in iteration number than inputs which generically take significantly less iterations to converge. The inputs which lead to a large variance in iteration number can act to drown out the contribution to the total efficiency of the inputs with less absolute variance (but equally high relative variance). For example, a method that performs well on one large system and poor on many small systems may lead to a improved efficiency η when compared to a method which performs well on many small systems and poor on one large system. This issue is not particularly pronounced here because iteration count as a measure of efficiency suffers far less from this pathology than, for example, wall-clock time. Furthermore, one can recover the detail lost due to this issue by comparing any two methods individually, a workflow for which is provided later in this section. Therefore, the efficiency and robustness measures outlined here are able to provide an approximate *guide* toward methodology that performs well. Further inspection of the convergence data can then allow one to make more concrete conclusions. Note that the variance in a given method upon tuning parameters is not significant when close to optimal parameters.

As there exist two separate measures of utility, we must determine a prescription for how optimal can be defined here. This is done using the concept of *Pareto optimality*. Consider an algorithm X with associated efficiency and robustness scores, $\{X, \eta_X, r_X\}$.

If X is Pareto optimal, then there exists no algorithm $\{Y, \eta_Y, r_Y\}$ such that $\eta_Y > \eta_X$ and $r_Y > r_X$. In other words, there is no algorithm that is both more efficient and more robust than X , and hence X has utility. Any algorithm that is not Pareto optimal, or nearly Pareto optimal, has no utility as there exists another algorithm with significantly higher individual utility scores. The set of all Pareto optimal algorithms, which includes differing parameter sets of the same underlying method, define the *Pareto frontier*. Algorithms that lie on, or lie close to, the Pareto frontier can be utilised in the sense that it is up to the developer to make a trade-off between robustness and efficiency. A developer might choose, for example, a particularly robust yet inefficient algorithm as a fall-back, and a slightly less robust yet more efficient algorithm as default, both of which should lie on the Pareto frontier.

As discussed previously, all possible parameters that could influence the convergence behaviour of the algorithm, that are not directly related to the algorithm, must be held constant. In the benchmark presented here, this includes (unless stated otherwise): PBE exchange-correlation functional, Gaussian smearing scheme, electronic temperature $T = 300\text{K}$, history length $m = 20$, k -point spacing $2\pi \times 0.04\text{\AA}^{-1}$, and parallelised over sixty four cores using Intel Xeon Gold 6142 processors at 2.6GHz. Note that the energy tolerance required for convergence, the cut-off energy, and the pseudopotential are varied across the members of the test suite, but *not* across the algorithms. Ultrasoft pseudopotentials are generated following the prescriptions of CASTEP's on-the-fly pseudopotential generator. A summary of these input parameters for each member of the test suite is given in Ref. [25].

The results of the benchmark are given in Table 3.2, and illustrated in Fig. 3.9. The first observation of note is that Pulay's algorithm, Kerker preconditioned using the default parameter set [61, 62], is Pareto optimal. In particular, Pulay's method significantly outperforms Broyden's methods in both the singlesecant and multiseccant form. Despite Pulay's method being Pareto optimal, there exist multiple algorithms that are more stable than Pulay's method while sacrificing little efficiency. The relationship between efficiency and robustness is generally non-linear, meaning it is worth sacrificing more than 10% efficiency for a method that is 10% more robust. Hence, algorithms more robust than Pulay's method, that only incur a relatively small drop in efficiency, can be considered potential upgrades over Pulay's method.

In general, a system that was proven difficult for one method was, more often than not, difficult for the rest. This is most likely due to that system being ill-conditioned, and thus changing the method does not alter the underlying condition of the problem. However, methods that place a particular emphasis on stabilisation were often able to find convergence where certain more efficient methods were not. For example, interleaving linear mixing steps in periodic Pulay was able to converge a large iron cluster, where regular Pulay's method was not.

From the algorithms tested here, these potential upgrades include certain parameterisations of Restarted Pulay [192], Periodic Pulay [190], and Marks & Lukes' MSB2 [172]. The parameters used for these methods, as detailed in Table 3.2, are not necessarily optimal; by this we mean the parameters have not been tailored for performance over the test suite. Rather, these parameter sets are sensible choices that demonstrate improved convergence properties. It is feasible that parameter adjustments could lead to even

more stable and efficient convergence. To this end, we provide a modest demonstration of how the representativeness of how the test suite can be used to determine optimal parameter sets. Fig. 3.10 illustrates the results of calculations using eight different Kerker parameter sets for Pulay’s method over the test suite. As expected, removing the Kerker preconditioner markedly reduces both the efficiency and robustness, as does setting the Kerker parameter $|G_0|$ too high, or too low. In fact, the default parameters $|G_0| = 1.5 \text{ \AA}^{-1}$ and $\alpha = 0.8$ suggested in Refs. [61, 62] are found to be approximately optimal. Reducing the history size to $m = 10$ rather than $m = 20$ had a slight stabilising effect.

As expected, ensemble DFT (EDFT) [110], a variational direct minimisation method of the energy with respect to the orbitals (not discussed in this thesis), is able to converge the vast majority of the test suite – it is guaranteed by design. Note that the method is not 100% robust as two methods took over the maximum allowed time to converge. The cost of guaranteed convergence here is apparent: the efficiency is drastically reduced. Ensemble DFT should be used if and only if one expects divergent iterations with self-consistent field methods. An interesting area of future work is to examine the extent to which self-consistent field methods can match the robustness of EDFT whilst approximately maintaining the efficiency of self-consistent field methods. Recent sophisticated algorithms attempt this [137, 172, 209], see Section 3.4.1.4, using some form of step-length control, i.e. line-searches or trust-regions. Incorporating some of the techniques that demonstrably stabilise iterations, such as adding linear mixing steps to the history or occasionally restarting the history, could be advantageous here.

To conclude, this workflow, namely, assessing an algorithm utilising the test suite and similar measures of performance, can be used to confidently highlight the improvements possible with, for example, guaranteed self-consistent field methods. Note that a one-to-one comparison of algorithms can also be illustrated, Fig. 3.11. Here, we compare the efficiency of Pulay’s method vs. Broyden’s second method, which brings to light the classes of systems for which one method outperforms the other. In this example, Pulay’s method is demonstrated to uniformly outperform Broyden’s second method over the test suite. In the case only two methods are being tested, presenting the results in this form is preferable to Fig. 3.9, as it removes the aforementioned pathologies related to comparing only the efficiency and robustness scores. For example, we are able to observe that Pulay’s method does indeed uniformly outperform Broyden’s second method, as one might expect from a method with a significantly higher efficiency and robustness score.

3.7 Conclusion

Modern research utilising Kohn-Sham theory is progressively demanding self-consistent solutions from inputs that lead to significant ill-conditioning. This ill-conditioning can be a result of increased unit cell sizes, and/or related to the atomic species and positions involved. The core aim of this chapter is to elucidate these issues and provide a clearer path forward for algorithm development. We began in Section 3.3 by examining a variety of properties of the Kohn-Sham map, whose fixed-points define self-consistent densities. The topics covered in this section ranged from, for example, the definition of

TABLE 3.2: A table consisting of each algorithm tested and its corresponding parameter set, efficiency score, and robustness score.

Method	Parameters	Robustness	Efficiency
Pulay II (2)	$\alpha = 0.8, G_0 = 1.5$	0.775	0.0118
Pulay II	$\alpha = 0.8, G_0 = 0.0$	0.637	0.0085
Pulay II	$\alpha = 0.2, G_0 = 1.5$	0.689	0.0088
Pulay II	$\alpha = 0.4, G_0 = 1.0$	0.741	0.0094
Pulay II (3)	$\alpha = 1.0, G_0 = 1.5$	0.741	0.0161
Pulay II	$\alpha = 1.0, G_0 = 1.5,$ $m = 10$	0.827	0.0063
Pulay II	$\alpha = 0.6, G_0 = 2.5$	0.775	0.0097
Pulay II	$\alpha = 0.1, G_0 = 1.5$	0.637	0.0046
Broyden II	$\alpha = 0.8, G_0 = 1.5$	0.706	0.0118
Broyden II	$\alpha = 0.2, G_0 = 1.5$	0.672	0.0090
Multisecant Broyden I	$\alpha = 0.8, G_0 = 1.5$	0.620	0.0056
Multisecant Broyden II	$\alpha = 0.8, G_0 = 1.5$	0.706	0.0179
MSB1	Greed controlled [‡] , $ G_0 = 1.5$	0.689	0.0098
MSB2	Greed controlled, $ G_0 = 1.5$	0.793	0.0097
Two-Step Steepest Descent [†]	N/A	0.689	0.0070
Guar. Red. Pulay	$\alpha = 0.8$	0.448	0.0103
Restarted Pulay	$\alpha = 0.8, G_0 = 1.54,$ $m = 10$	0.819	0.0088
Linear	$\alpha = 0.2$	0.328	0.0154
Linear	$\alpha = 0.05$	0.534	0.0033
Kerker	$\alpha = 0.8, G_0 = 1.5$	0.500	0.0025
Fixed-Point	N/A	0.054	0.0344
Periodic Pulay (1)	$\alpha = 0.2, G_0 = 1.5, k = 2$	0.828	0.0063
Periodic Pulay*	$\alpha = 0.6, G_0 = 1.5, k = 2$	0.705	0.0062
EDFT	$\alpha = 0.8, G_0 = 1.5$	0.948	0.00003

[†] As proposed in Ref. [212].

* Performed with k Pulay steps in between each linear step.

[‡] Greed control parameters: $k = 3, \alpha_0 = 0.3, \beta = 0.2, \gamma = 1.3$

convergence, generation of the initial guess density, and sources of ill-conditioning within the linear response approximation. Following this, an overview of both standard and contemporary methodology was provided. This overview was intended to be fairly brief, and aimed at providing a broad yet digestible introduction for interested practitioners and DFT software developers not actively involved in the development of self-consistency methodology.

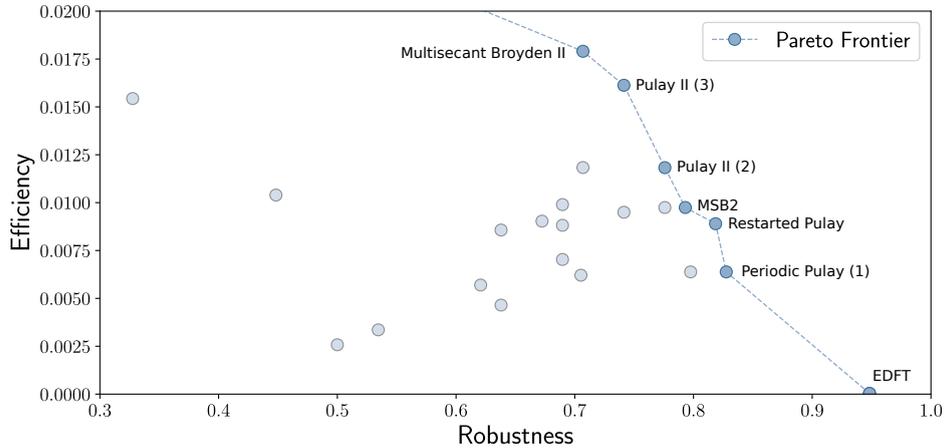


FIGURE 3.9: Results of the benchmark tests of the algorithms in Table 3.2. Each node corresponds to a separate algorithm placed corresponding to its robustness and efficiency across the test suite, with those that are on or close to the Pareto frontier explicitly labelled.

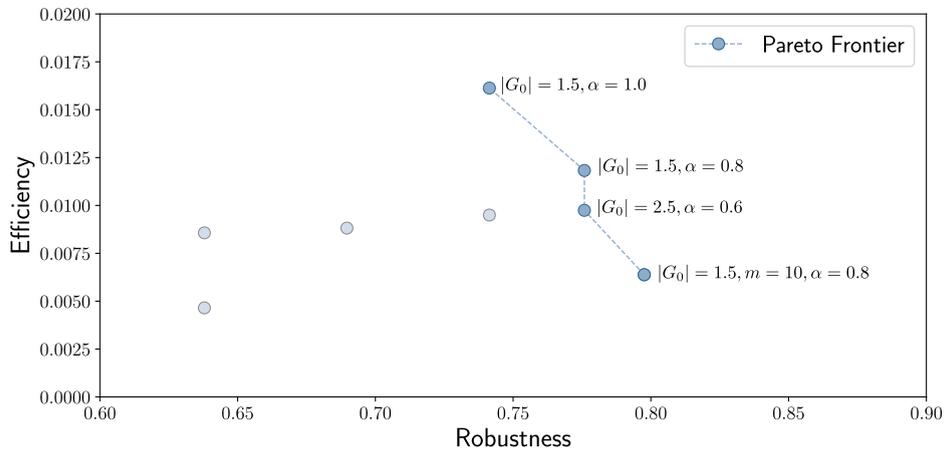


FIGURE 3.10: Results of the tests for differing Kerker parameter sets using Kerker preconditioned Pulay’s method. Each node corresponds to a separate parameter set placed corresponding to its robustness and efficiency across the test suite.

The analysis of Section 3.3 revealed certain classes of inputs that induce difficulty in the self-consistency iterations. These classes include far-from-equilibrium systems, large units cells, highly degenerate systems, complex interfaces with differing electronic behaviour, and others. The insight gained from this analysis led to the creation of a test suite, the SCF- x_n suite, containing over fifty ill-conditioned inputs from a variety of sources. A selection of algorithms suitable to be implemented in CASTEP were then benchmarked using this test suite, and their utility was quantified. The results of this benchmark led to a several observations of note. First, from the standard methods, which include unmodified versions of Pulay and generalised Broyden, the best performing was indeed Pulay’s original method. That being said, relatively simple modifications to these methods were able to demonstrate improved robustness. These modifications involved interweaving linear mixing steps in-between Pulay steps [190, 213, 214], and

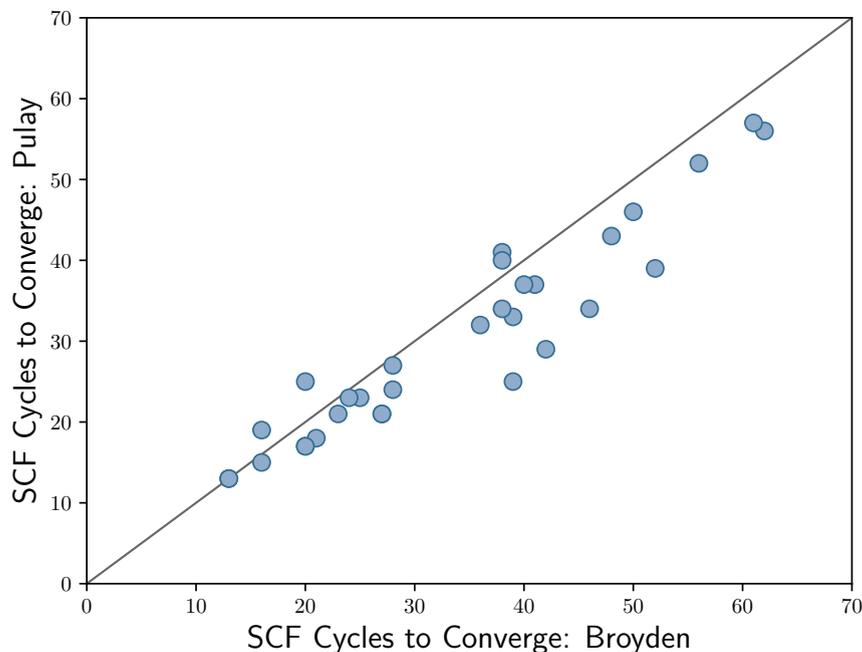


FIGURE 3.11: A direct comparison of the efficiency measure of two algorithms. The plot is restricted to the range of 1-70 iterations for illustrative purposes.

flushing the stored history of iterates after a given number of iterations [192]. Furthermore, considerable promise is shown by more sophisticated methods such as those in Refs. [137, 169, 172, 209]. These methods aim to converge the majority of cases with minimal user intervention. This will become increasingly important in the future as significant adjustments to default parameters may be required to force convergence in difficult cases due to increased ill-conditioning. Finally, the parameter space of Kerker preconditioned Pulay’s method was sampled using the test suite. This confirmed that the default parameters suggested by Kresse [61, 62] are indeed optimal. In particular, lowering the damping parameter too much can negatively impact robustness as well as efficiency, due to the complexities inherent within the Kohn-Sham functional landscape. The damping parameter in Pulay’s method should be kept as close to unity as possible, and reduced if and only if the iterations are divergent.

To conclude, we emphasise that the benchmark itself, while able to reveal certain well-performing methods and parameter sets, is not intended to be the focus of the latter part of this chapter. Rather, the workflow used in Section 3.6 to generate this benchmark is the central development. That is, we present a workflow that comprises a test suite of difficult to converge inputs that are used to compare methodologies with some appropriate measure of efficiency and robustness. Indeed, if one were to replicate this workflow, the test suite need not be exactly the same as the version of $SCF-x_n$ used here. Instead, one can augment the suite with any selection of systems, as long as due care is taken to ensure that the range of sources of ill-conditioning already included here is at least retained. It is hoped that the workflow presented enables and assists the development of self-consistency algorithms that are able to meet the needs of practitioners in modern applications.

Chapter 4

The Exact Time-Dependent Exchange-Correlation Potential

The remainder of this thesis concerns either theoretical or computational developments within the context of exact one-dimensional quantum mechanics. The computational implementation of the relevant theoretical frameworks is first discussed. Following this, an efficient and robust method is presented to obtain exact time-dependent xc potentials. This method constitutes software development, and therefore no particular application of the method is detailed in this chapter. However, the method is central to Chapters 5 and 6, which covers Refs. [26, 28], and has also been used since in papers unrelated to the author, e.g. Ref. [215].

The involvement of the author in this project is through collaboration with Rex Godby (University of York) and Michael Entwisle (University of York, Freie Universität Berlin). For some time, the Godby group’s `iDEA` code has been used to conduct theoretical studies on one-dimensional model systems towards an improved understanding of many-body theory. Central to the `iDEA` code is its *reverse-engineering* algorithm, a name given to the methods involved in obtaining an exact (time-dependent) xc potential given a reference density. Whilst successful at the time [216–218], the limitations of the algorithm had become a hindrance, preventing further studies on certain ‘difficult’ systems, and preventing studies on standard systems beyond some maximum time, after which the algorithms became too inefficient. The present work seeks to expand the scope of the reverse-engineering algorithm through improvements in both robustness and efficiency (in a similar sense to Chapter 3).

4.1 Implementation

This section details the implementation of the author’s code `density2potential`. This code is, in part, a rewrite and restructuring of the relevant modules of `iDEA` that occurred during the developments discussed in the remaining chapters.

The time-independent and time-dependent Schrödinger equations are solved in one spatial dimension in the real space (delta function) basis: $\{\delta(x - y_n) \mid 0 < n \leq N_b\}$, where the discrete positions y_n are the *grid points*. The exact Hamiltonian is a sparse $N_b^N \times N_b^N$ matrix in this basis: the diagonal elements are populated with the local potential, and

the off-diagonal elements are populated with the discretised Laplacian and the Coulomb interaction. The use of a p -point stencil for the Laplacian operator significantly reduces the number of non-zero elements in the Hamiltonian provided $p \ll N_b$. An eleven-point stencil is used for the remainder of this work. In contrast to the previous chapter, we now consider finite systems inside the domain $[-a, a]$ a.u. subject to the Dirichlet boundary conditions, $\Psi(\pm a, \pm a, \dots) = 0$.

The particles are subject to the model interaction

$$v_{ee}(x_i, x_j) = \frac{1}{|x_i - x_j| + \alpha}, \quad (4.1)$$

also known as the softened Coulomb potential; the softening parameter is set to $\alpha = 1$ a.u. This potential is designed to circumvent numerical issues (ill-conditioning due to the singularity at $x = x'$) and provide a reasonable account of certain relevant aspects of the three-dimensional interacting system. For example, a softened electron-nuclear interaction gives rise to the Rydberg series for isolated atoms [219–221]. We consider systems of like-spin fermions as these systems deliver a richer ground state and excitation spectrum for a given number of electrons.

4.1.1 Ground State

The exact Hamiltonian is constructed as a sparse matrix using the `scipy.sparse` package. The ground state of the Hamiltonian is symmetric, and the remainder of the states are either symmetric or anti-symmetric. Since we seek the anti-symmetric solutions, it is possible to reduce the dimension of the Hamiltonian with a transformation of the kind,

$$H \rightarrow A^{-1}HA \quad (4.2)$$

where the unitary projection matrix A is defined as to remove the redundant (symmetry-related) degrees of freedom in the eigenstates,

$$|\Psi^{\text{reduced}}\rangle = A|\Psi\rangle. \quad (4.3)$$

In other words, the matrix A encodes the fact that coefficients to the state in a basis observe the relation $\alpha_{\dots ij \dots} = -\alpha_{\dots ji \dots}$, and so the reduced state includes just one of these two coefficients. The full anti-symmetric state can be obtained by applying the back-transformation A^{-1} on eigenstates of the reduced Hamiltonian.

In cases where the ground state alone is needed, an efficient iterative diagonalisation is performed, whereas, when all excited states are required (e.g. when constructing the linear response function), a full diagonalisation is performed. The code calls LAPACK in the background, which (depending on the routine) is multithreaded across desktop cores. The majority of time spent in a given calculation is in the iterative or full diagonalisation of the Hamiltonian. Given a time-independent external potential v_{ext} , this procedure yields the full set of N -particle eigenstates and eigenenergies of the many-body Hamiltonian, $\{|\Psi_i\rangle, E_i\}$.

4.1.2 Time Dependence

Suppose our electrons exist in some state $|\Psi\rangle$ of the Hamiltonian at $t = 0$, then the perturbation,

$$v_{\text{ext}}(x) \rightarrow v_{\text{ext}}(x) + \delta v_{\text{ext}}(x, t)\Theta(t), \quad (4.4)$$

turned on at $t = 0$ with the step function $\Theta(t)$ will cause time evolution in the state according to

$$|\Psi(t)\rangle = U(t)|\Psi(0)\rangle, \quad (4.5)$$

where

$$U(t) = \mathcal{T}e^{-i \int_0^t H(t')dt'} \quad (4.6)$$

is the time evolution operator, and \mathcal{T} is the time-ordering operator. In discrete time, the time variable takes values in the sequence $0 < \delta t < 2\delta t < \dots < t$ for some increment in time δt with N_t total grid points. Within each interval, the Hamiltonian can be assumed constant for ‘sufficiently’ small δt , in which case the time-evolution operator from t to $t + \delta t$ becomes

$$U(t, t + \delta t) = e^{-iH(t)\delta t}, \quad (4.7)$$

where the Hamiltonian over the interval is approximated as the Hamiltonian at the start of the interval – various time-propagation schemes modify such assumptions. Repeated application of Eq. (4.7) propagates some initial state $|\Psi\rangle$ forward in time.

A standard technique for time propagation (used in `iDEA`) is the so-called Crank-Nicolson scheme [222], which expands the time evolution operator as such,

$$e^{-iH(t)\delta t} = \frac{1 + \frac{\delta t}{2}H(t)}{1 - \frac{\delta t}{2}H(t)}. \quad (4.8)$$

Notably, the Crank-Nicolson scheme retains the fact that the time evolution operator is unitary, whereas a simple Taylor expansion of the exponential operator would not. Ultimately, the Crank-Nicolson scheme became a source of destabilisation in the reverse engineering algorithms, and therefore the more accurate routine `scipy.expm` is used, which is based around a higher-order Padé approximation of the exponential function [223]. These techniques allow us to obtain a time-dependent state $|\Psi(t)\rangle$, and thus also a time-dependent density $n(x, t)$, given an initial state and a perturbing time-dependent potential.

4.2 Reverse Engineering

4.2.1 Background

In order to determine the unique Kohn-Sham potential from which a time-dependent density is generated, we consider the map,

$$T^{-1} : \mathcal{N} \rightarrow \mathcal{V}, \quad (4.9)$$

from time-dependent densities to time-dependent Kohn-Sham potentials, where

$$T[v_{\text{KS}}] = n(x, t) \quad (4.10)$$

for $v_{\text{KS}} \in \mathcal{V}$ and $n \in \mathcal{N}$. This is a fixed-point problem, and in similar fashion to Chapter 3, we shall attempt to solve the associated root-finding problem,

$$R[v_{\text{KS}}] = T[v_{\text{KS}}] - n^{\text{ref}}(x, t) = 0, \quad (4.11)$$

for model one-dimensional systems (although the methodology to follow is not limited to just this context), where n^{ref} is some specified reference density. During the course of this work, various approaches were considered involving the associated optimisation problem,

$$v_{\text{KS}}(x, t) = \arg \min_{v_{\text{KS}}} \|T[v_{\text{KS}}] - n^{\text{ref}}(x, t)\|_2, \quad (4.12)$$

with preconditioning, penalty terms, and so on, however, the root-finding methods were found to be much more effective.

The discrete time-dependent reverse-engineering problem is as follows: assuming we have access to the Kohn-Sham potential $v_{\text{KS}}(x, t - \delta t)$ at time step $t - \delta t$ that generates the density $n^{\text{ref}}(x, t)$ at the subsequent time step, we then proceed to seek the Kohn-Sham potential $v_{\text{KS}}(x, t)$ that generates the reference density $n^{\text{ref}}(x, t + \delta t)$ when the discrete time-evolution operator Eq. (4.7) is applied to the unique Kohn-Sham states $\{\phi_i(x, t)\}$. This takes the form of an iterative procedure that modifies some trial v_{KS} towards convergence, see Fig. 4.1 (right panel).

Of course, we must first begin in some initial state, which in this work is often taken to be the ground state of some Hamiltonian. In other words, the ground-state reverse-engineering problem must first be solved,

$$G[v_{\text{KS}}] - n^{\text{ref}}(x, t = 0) = 0, \quad (4.13)$$

before we can proceed with the time-dependent case, where G is the ground-state DFT map, see Fig. 4.1 (left panel). The same methodology can be applied to both the ground-state and time-dependent reverse-engineering problems, however the former requires iterative solutions of the time-independent Kohn-Sham equations, whereas the latter requires repeated iterative application of the time-evolution operator to the Kohn-Sham states.

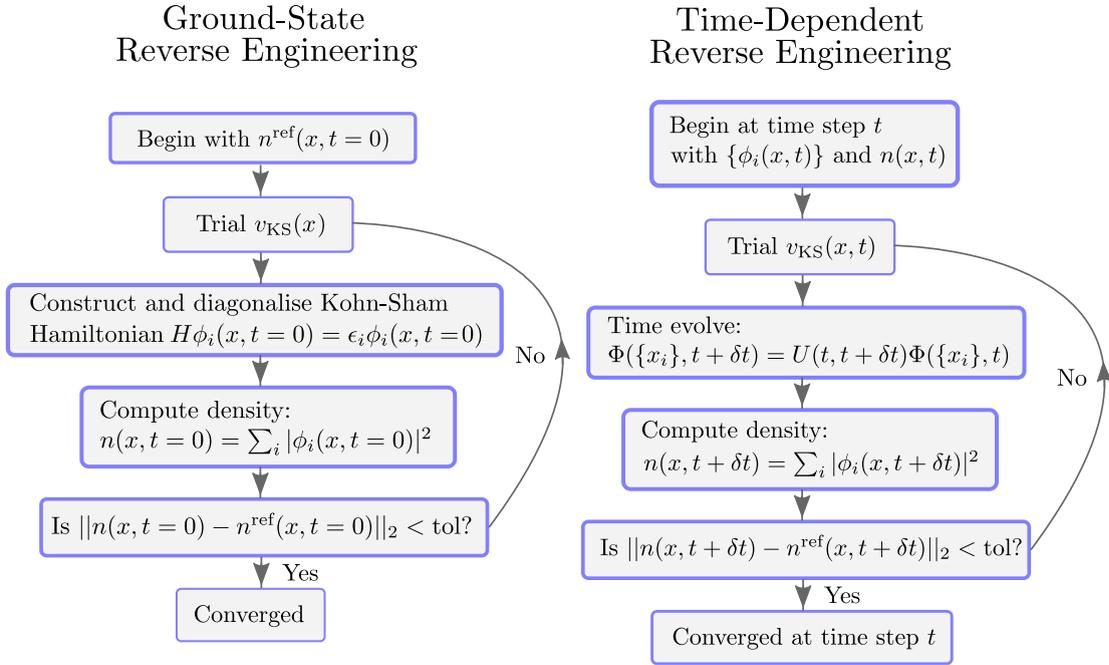


FIGURE 4.1: (Left) An iterative produce to find the time-independent Kohn-Sham potential corresponding to some ground-state reference density. (Right) An iterative produce to find the time-dependent Kohn-Sham potential at a subsequent time step corresponding to some time-dependent reference density.

A number of iterative algorithms have been proposed in recent literature to realise the inverse density-to-potential map in time-dependent and ground-state DFT [224–229]. For example, some of these approaches consider iterative schemes derived from the continuity equation of time-dependent DFT [225, 226], whereas the original time-dependent reverse-engineering algorithm in iDEA utilised iterative updates of the vector potential in the context of time-dependent current DFT (matching the current densities) [218]. The work to follow designs a set of algorithms that perform well when the inverse problem is formulated as in Eq. (4.11).

Numerical studies of the inverse density-to-potential map have been able to elucidate the time-dependent and ground-state v -representability problem [70, 144, 230]. In the ground state, it is possible to show that, given some suitable¹ density n , there exists another (ensemble) non-interacting v -representable density n' that is arbitrarily close to n in the some norm [39]. The time-dependent case is less clear: progress has been made on lattice-based systems [40, 230], but the general problem remains unsolved in a practical sense [42], see Ref. [226] for a recent review. The work to follow provides a range of examples of (non-)interacting v -representable ground-state and time-dependent potentials.

¹Functions that do not integrate to N , for example, are evidently far from being v -representable. Solutions to the v -representability problem seek to find some well-defined subset of densities \mathcal{N} from which the existence of a one-to-one map to some well-defined subset of potentials \mathcal{V} can be established, see Refs. [39, 42] for instructive discussions.

4.2.2 Stability

Standard formulations of DFT and the v -representability problem are mostly concerned with two properties: existence and uniqueness of the inverse map. These are two out of three requirements for a problem to be so-called *well-posed* [231] – the final requirement is stability, which can be understood in the context of some *numerical implementation* as follows. First, recall that numerical implementations deal with a discrete basis in finite precision arithmetic. Thus, given a time-dependent reference density $n^{\text{ref}}(x, t) \in \mathbb{R}^{N_b} \times \mathbb{R}^{N_t}$, a reverse-engineering algorithm is said to be converged when the reverse-engineered density $n(x, t)$ satisfies

$$\|n(x, t) - n^{\text{ref}}(x, t)\|_2 < \varepsilon, \quad (4.14)$$

where ε is some tolerance that, in the most stringent case, is set equal to machine precision. Let us call the potential that generates the density $n(x, t)$ satisfying Eq. (4.14) $v_{\text{KS}}(x, t)$. A lack of stability implies the existence of a second potential, $v'_{\text{KS}}(x, t)$, that satisfies Eq. (4.14) upon numerical solution of the time-dependent Kohn-Sham equations, and furthermore

$$\|v_{\text{KS}}(x, t) - v'_{\text{KS}}(x, t)\|_2 \gg \eta, \quad (4.15)$$

where η is some small parameter related to ε . In other words, there exists two (or more) Kohn-Sham potentials that differ significantly and yet produce the same reference density to within machine precision – a large perturbation to the input did *not* lead to a large perturbation in the output. This is an overlooked aspect of the numerical inverse problem in DFT: if the numerical map were not stable, then certain examples could lead to Kohn-Sham potentials that are not unique to within machine precision.

Baer [71] was the first to pursue this line of reasoning, and since then, it appears there has been little-to-no further comment on the matter. The work in Ref. [71] considers Lyapunov exponents [232], which are used to measure the divergence of neighbouring trajectories in dynamical systems, to argue that certain Kohn-Sham potentials behave chaotically after some characteristic time t_0 , the Lyapunov time². The work to follow provides certain explicit numerical examples of this destabilisation, i.e. it is demonstrated that multiple Kohn-Sham potentials give rise to the same time-dependent density to within machine precision, and therefore certain implementations of the inverse map in time-dependent DFT are ill-posed³. This is problematic because, in the case that the exact unique Kohn-Sham potential is unknown, it can be hard to distinguish which of these potentials corresponds to the platonic exact Kohn-Sham potential. (In the example below, it is clear by eye which potential is correct, and which is incorrect, but argumentation of this kind is dangerous in mathematics.) It is difficult to comment on the extent to which the inverse map is ill-posed independent of implementation.

²A typical example is the N -body problem in Newtonian planetary dynamics: two trajectories with similar initial conditions diverge after some characteristic time t_0 . This effect also occurs with molecular dynamics trajectories, however the ensemble properties are stable.

³There is a sense in which the ground-state map is also ill-posed: high frequency oscillations in the potential between grid points have little-to-no effect on the density.

4.3 Results and Discussion

4.3.1 Ground State

We begin with a ground-state reverse-engineering problem. Consider three interacting particles in an infinite potential well of size $[-15, 15]$ a.u., see Fig. 4.2. The domain is discretised over $N_b = 151$ grid points. The reference density is chosen as the infinite potential well interacting ground-state density.

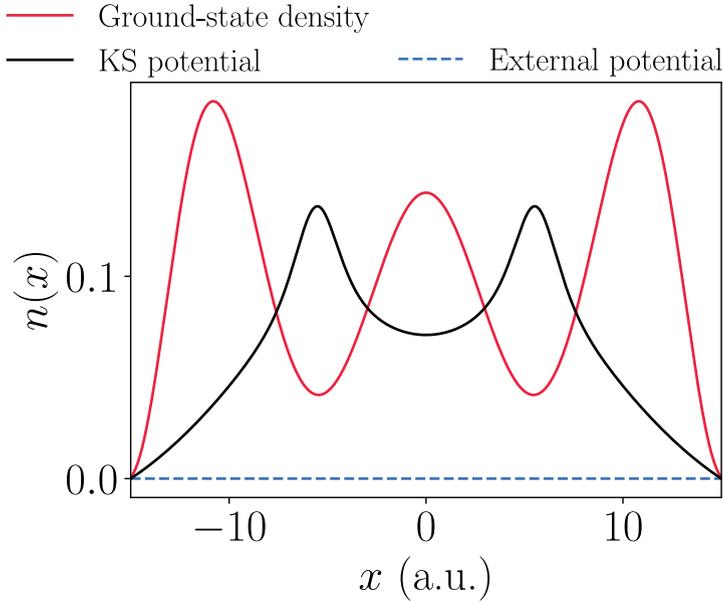


FIGURE 4.2: The interacting ground-state density, external potential, and reverse-engineered Kohn-Sham potential for a three-particle infinite potential well.

The simplest iterative scheme is the damped iteration scheme,

$$v_{\text{KS}}^{\text{new}}(x) = v_{\text{KS}}^{\text{old}}(x) + \alpha(n^{\text{old}}(x) - n^{\text{ref}}(x)), \quad (4.16)$$

where the damping parameter α has appropriate dimensions. Error as a function of iteration count (subscript n) is measured as such,

$$\varepsilon_n = \int |n_n(x) - n^{\text{ref}}(x)| dx. \quad (4.17)$$

The initial guess Kohn-Sham potential is set equal to the external potential. The damped iteration scheme takes $\mathcal{O}(10^4)$ iterations to converge to within 10^{-6} , where the optimal damping parameter for the infinite potential well is $\alpha \approx 1$, see Fig. 4.3. The characteristic number of iterations taken to converge the inverse problem is evidently much increased when compared to that of the forward problem considered in the previous chapter (despite the reduced dimension, $N_b = 151$).

The infinite potential well is a straightforward system to converge, as the main source of ill-conditioning in the present context is regions of low density. Namely, a large perturbation in the Kohn-Sham potential is required to bring density into a region where

it was previously unfavourable to do so. The coexistence of regions with large and low density responses results in ill-conditioning, which calls for a form of preconditioning. Two preconditioning strategies are considered here, leading to the following modified root-finding problems,

$$R'[v_{\text{KS}}] = \alpha \frac{T[v_{\text{KS}}] - n^{\text{ref}}(x)}{n^{\text{ref}}(x)} = 0, \quad (4.18)$$

and,

$$R'[v_{\text{KS}}] = \alpha(T[v_{\text{KS}}]^\gamma - (n^{\text{ref}}(x))^\gamma) = 0, \quad (4.19)$$

where α is again the damping parameter, and $\gamma \in [0, 1]$ is a method parameter. Both of these strategies seek to raise the response of the low density regions towards that of the high density regions⁴.

As Fig. 4.3 shows, the former strategy is much more effective than the latter, where a modest attempt was made in both cases to find a set of optimal parameters for the infinite potential well. We shall hereafter proceed with the preconditioning strategy in Eq. (4.18). These conclusions are transferable across multiple different systems: an atomic system, the quantum harmonic oscillator (QHO), a double well, and more. These systems are important examples in the subsequent chapters, and so are not shown here.

The natural next step is to consider more sophisticated quasi-Newton (or related) algorithms. Three methods additional methods are shown in Fig. 4.3: Pulay’s algorithm [184], the Levenberg–Marquardt algorithm [233], and Powell’s algorithm [234] (`scipy` implementation). The last two methods involve, in part, finite difference computations of the Jacobian at certain iterates. The preconditioned problem Eq. (4.18) is given to these methods. Despite all three methods improving matters significantly over the damped iterations, Pulay’s method is the worst performing of the three, and the difference between Levenberg–Marquardt and Powell is small and depends on the system at hand. For example, Powell’s algorithm with preconditioning is able to reach machine precision $\varepsilon \approx 10^{-15}$ within 700 function evaluations. (This is an acceptable number of iterations to enable exact studies of the adiabatic-connection fluctuation-dissipation theorem in Chapter 6, which requires reverse-engineering the *interacting* system along the adiabatic connection.)

In difficult systems, such as a double well, where the initial guess (external potential) is far from the exact Kohn-Sham potential, the more sophisticated methods can struggle due to beginning outside some radius of convergence. This effect is more pronounced in time-dependent reverse engineering. In such cases, it can be effective to utilise the damped iteration (guaranteed to converge to a local minimum for sufficiently small values of the damping parameter) until the initial guess is brought into a region of phase space suitable for the sophisticated root finder. (This strategy is known as Landweber regularisation [231]). The algorithm that is used for both ground-state and time-dependent reverse engineering is given in Algorithm 2.

⁴Some form of the method presented in Eq. (4.19) in the context of ground-state reverse engineering was implemented in `iDEA` prior to the authors’ involvement. However, all methods presented in this section were studied independently by the author in both the time-dependent and ground-state case.

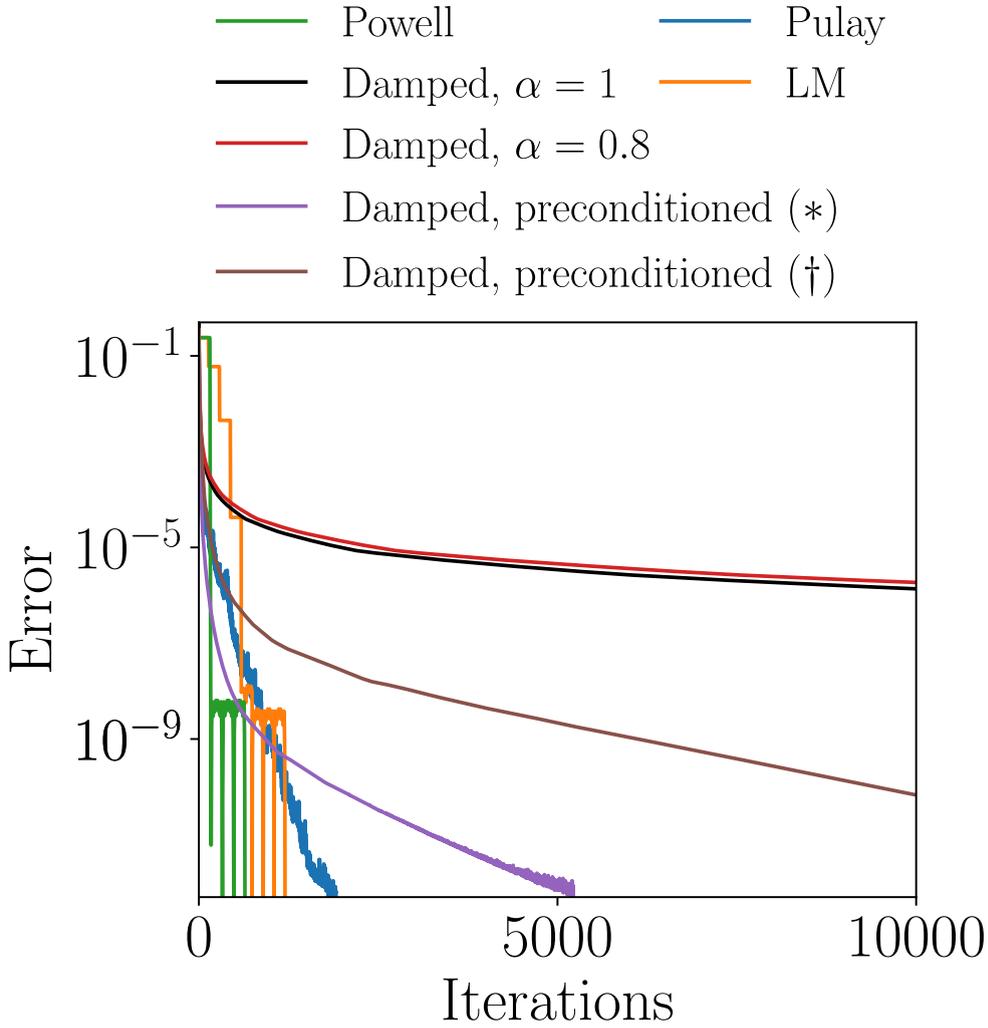


FIGURE 4.3: Reverse engineering the exact three-particle infinite potential well Kohn-Sham potential: error as a function of iteration count across various methods. (†) The preconditioned damped iteration of Eq. (4.18) with $\alpha = 0.05$. (*) The preconditioned damped iteration of Eq. (4.19) with parameters $\alpha = 0.1$, $\gamma = 0.1$. With more sophisticated methods such as the Levenberg–Marquardt (LM) algorithm, iteration count measures function (residual) evaluations.

Algorithm 2 Reverse Engineering

- 1: **Input:** $v_{\text{ext}}, n^{\text{ref}}, \alpha, \text{tol}_1, \text{tol}_2$
 - 2: Set initial guess: $v_{\text{KS}}(x) = v_{\text{ext}}(x)$ or $v_{\text{KS}}(x, t) = v_{\text{KS}}(x, t - \delta t)$
 - 3: **while** error $> \text{tol}_1$:
 - 4: Evaluate map, $n = G[v_{\text{KS}}]$ or $n = F[v_{\text{KS}}]$
 - 5: Update Kohn-Sham potential, $v_{\text{KS}} \leftarrow v_{\text{KS}} + \alpha \frac{n - n^{\text{ref}}}{n^{\text{ref}}}$
 - 6: Update error = $\int |n - n^{\text{ref}}| dx$
 - 7: Update initial guess with output v_{KS}
 - 8: **while** error $> \text{tol}_2$:
 - 9: Powell’s root-finding algorithm (`scipy` implementation)
 - 10: **Output:** v_{KS} that generates n^{ref} to within tol_2
-

4.3.2 Time Dependence

We first consider the two-particle infinite potential well inside the domain $[-8, 8]$ a.u. discretised over $N_b = 121$ grid points. The interacting ground state of this system constitutes the initial state from which our system shall evolve. At $t = 0$, the ground state is subject to the perturbation,

$$\delta v_{\text{ext}}(x, t) = \begin{cases} \frac{1}{20}x & \text{if } t > 0 \\ 0 & \text{if } t = 0, \end{cases} \quad (4.20)$$

which is then allowed to evolve up to $t = 15$ a.u. with time step $\delta t = 0.03$ a.u. The reverse-engineering reference density is thus chosen as the time-dependent interacting density after the perturbation Eq. (4.20) is applied.

Snapshots of the time-evolving interacting *and* reverse-engineered (v -representable) densities are shown in Fig. 4.4, where Algorithm 2 has been used with parameters: $\alpha = 0.05$, $\text{tol}_1 = 10^{-8}$, $\text{tol}_2 = 10^{-15}$. The associated time-dependent Kohn-Sham potential is shown in Fig. 4.5. (Note that the initial guess Kohn-Sham potential at time step t is chosen as the Kohn-Sham potential at the previous time step $t - \delta t$.) The reverse-engineered density matches the reference density to within 10^{-15} at each time step along the evolution, despite signs of minor destabilisation in the Kohn-Sham potential at later times. Notably, the iteration count at each time step does not increase as time proceeds, which was the primary issue with previous methods (thus preventing studies of long-time behaviour). The previous algorithm(s) took $\mathcal{O}(10^4)$ iterations per time step (growing with time) to reach 10^{-10} accuracy, whereas Algorithm 2 now takes $\mathcal{O}(10^2)$ iterations at each time step (constant in time) to reach machine precision. These improvements constitute the central result of this chapter, thus enabling the studies in Chapters 5 and 6.

We now turn to the quantum harmonic oscillator (QHO),

$$v_{\text{ext}}(x) = \frac{1}{2}\omega^2 x^2, \quad (4.21)$$

with $\omega = 0.25$ a.u. to demonstrate a particular kind of destabilisation that can occur. A perturbation of the form,

$$\delta v_{\text{ext}}(x, t) = \begin{cases} \frac{1}{2}x & \text{if } t > 0 \\ 0 & \text{if } t = 0 \end{cases} \quad (4.22)$$

is applied, and the system evolves to $t = 2.5$ a.u. with $\delta t = 0.012$ a.u. In accordance with the harmonic potential theorem [7] (see Section 2.4.1), the QHO ground-state density undergoes a rigid shift to the left, Fig. 4.6. The Kohn-Sham potential that generates this profile is constant (apart from finite-size effects) and given in Fig. 4.7 (using Algorithm 2).

In the instance that ill-conditioning (or noise) is introduced, such as increasing the size of the domain, then the Kohn-Sham potential deteriorates as shown in Fig. 4.8. Notably, *solving the time-dependent Kohn-Sham equations with this ‘spurious’ Kohn-Sham potential yields the correct density profile to within 10^{-15} across all times.* We can

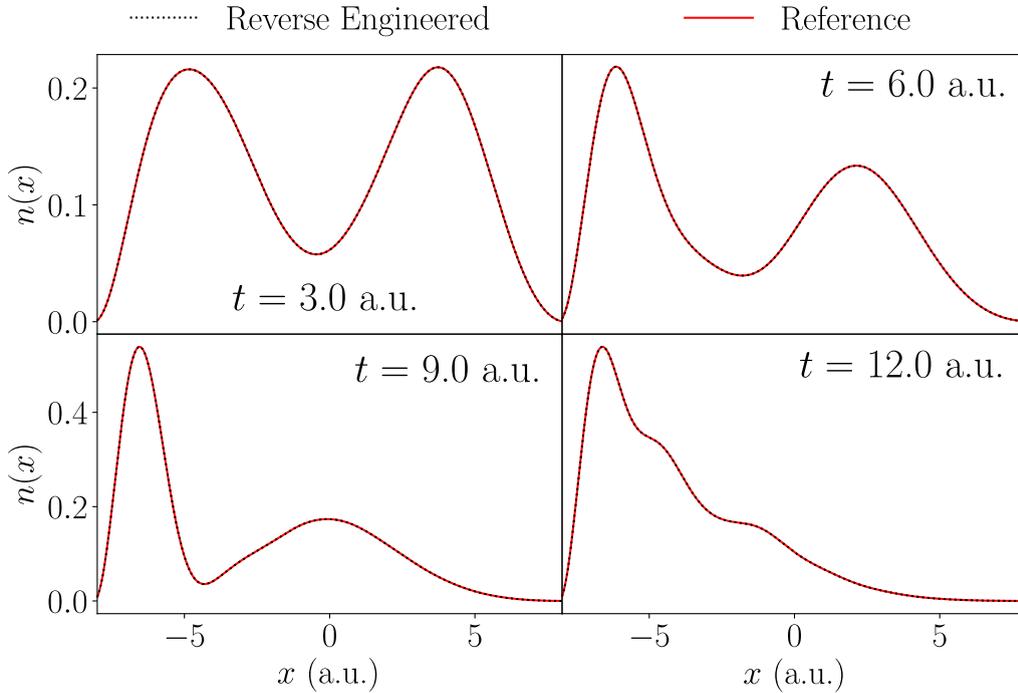


FIGURE 4.4: The time-dependent interacting density of the infinite potential well starting in the ground state, and subject to the perturbation in Eq. (4.20). The interacting density (red) is targeted with a reverse engineering algorithm to produce the non-interacting density (black) – these densities match to within machine precision at all times shown.

see that this destabilisation has a certain character: at one time step, the Kohn-Sham potential produces a large force on the density in one direction, and then, at the next time step, the subsequent Kohn-Sham potential must counteract this with a large but opposite force in the other direction – the net result is a density that shifts gradually to the left as in Fig. 4.6. (Note that this behaviour will manifest itself in the current density. However, the Kohn-Sham current density need not match the interacting current density, thus preventing its use as a regularisation tool without use of time-dependent current DFT.) These errors are evidence that the present (and similar) implementation(s) of the inverse map in time-dependent DFT is ill-posed.

4.4 Conclusion

An algorithm has been presented for the robust and efficient determination of exact ground-state and time-dependent Kohn-Sham potentials that generate a given reference density. This algorithm remedies the shortcomings of previous methods, such as growth in iteration count over time, with the use of preconditioning, regularisation, and advanced root-finding methods. It is hoped that these methods will continue to contribute to the development of an improved understanding of many-body theory, as has been the case in Refs. [26, 28, 215].

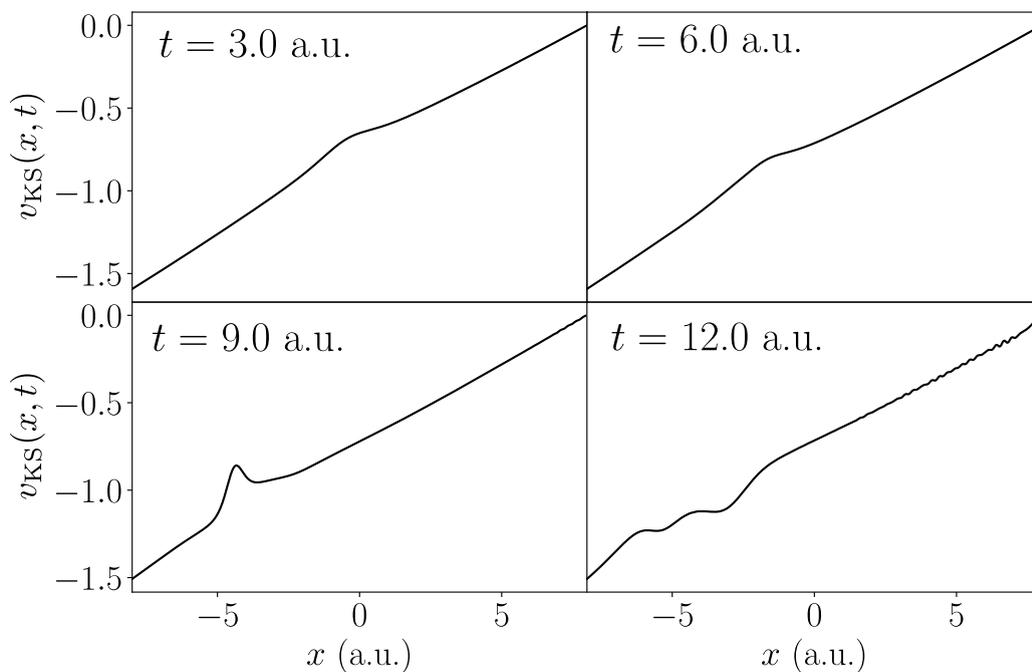


FIGURE 4.5: The Kohn-Sham potential corresponding to the non-interacting density evolution depicted in Fig. 4.4. The final panel contains the onset of mild destabilisation, although this does not (yet) affect the accuracy of the reverse-engineered density. The Kohn-Sham potentials have been shifted (gauge transform) for visual purposes.

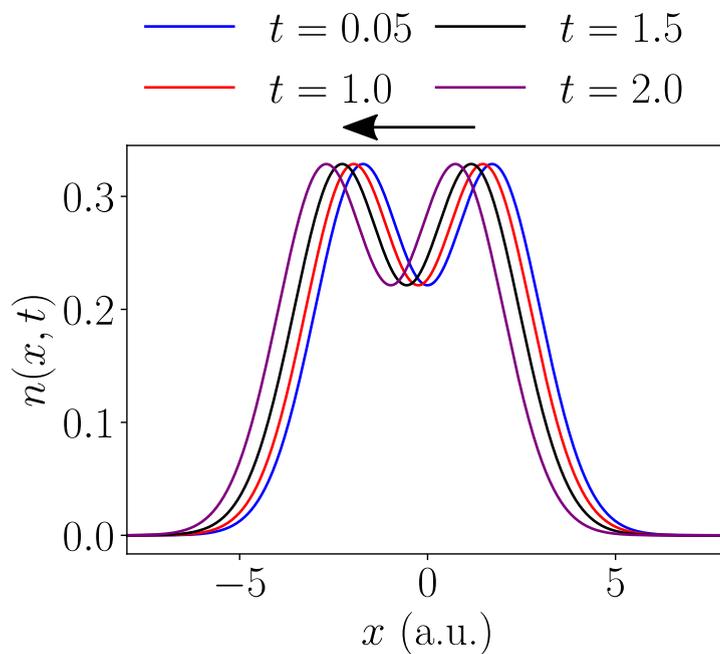


FIGURE 4.6: The time-dependent interacting density of the QHO starting in the ground state, and subject to the perturbation in Eq. (4.22). The density undergoes a rigid shift to the left in accordance with the harmonic potential theorem.

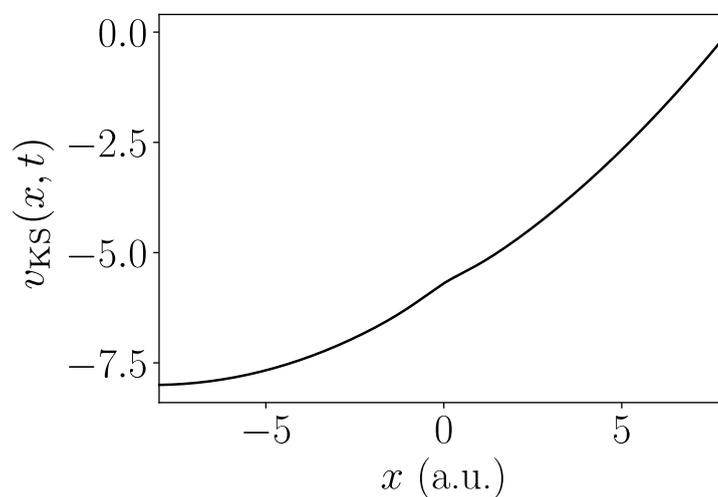


FIGURE 4.7: The Kohn-Sham potential (at all times $t > 0$) that produces the time-dependent density profile in Fig. 4.6.

The particular kind of divergence in the QHO example demonstrates a fundamental obstacle common to numerical implementations of dynamical inverse problems. That is, a large difference in two inputs (Kohn-Sham potentials) produces a negligible difference in the associated outputs (densities). This is an instability, and thus implementations of the sort considered here are ill-posed in certain circumstances. It is possible that other formulations and implementations of the time-dependent DFT inverse map, e.g. Ref. [226], do not suffer from these issues – note that the exact time-dependent DFT map Eq. (4.9) should be continuous or else it would not be differentiable, which is required in order to set up a Kohn-Sham scheme.

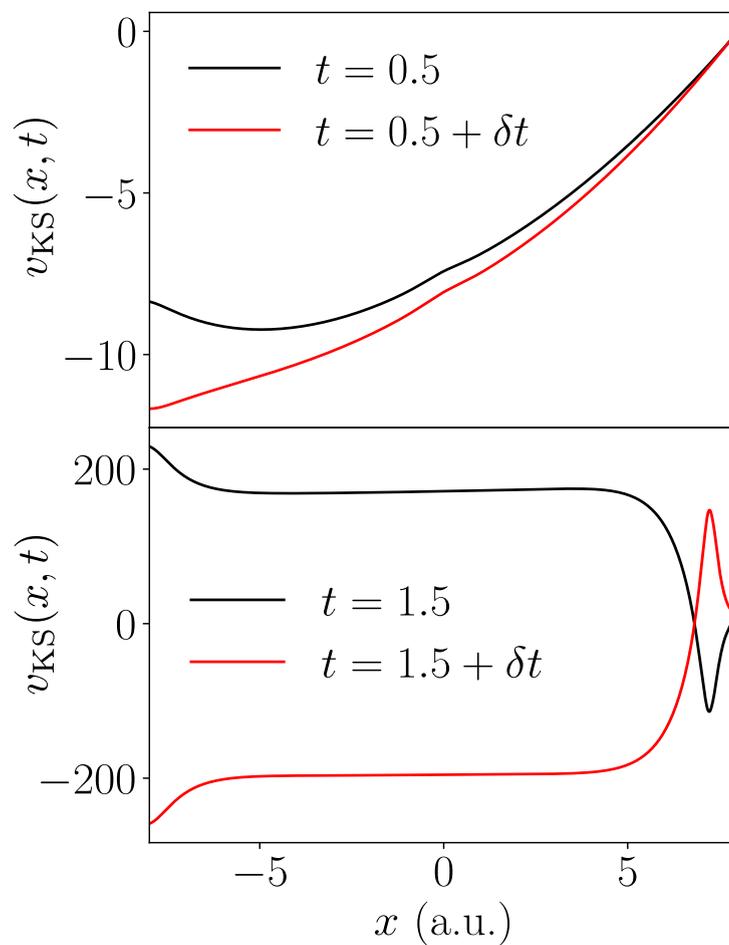


FIGURE 4.8: Snapshots of a Kohn-Sham potential that produces the correct QHO density profile (Fig. 4.6) to within machine precision, and is yet in serious error when compared with the correct Kohn-Sham potential, Fig. 4.7. This Kohn-Sham potential is shown at some time step t and the subsequent time step $t + \delta t$ such that the destabilisation due to successive over-corrections can be observed.

Chapter 5

The Exact Exchange-Correlation Kernel and the Optical Spectrum

An introduction to linear response time-dependent DFT is given in Chapter 2. The present chapter gives more in-depth background and motivation for the project to be detailed herein. The xc kernel, f_{xc} , is an object that fulfils an analogous role to the xc potential, v_{xc} , of ground-state DFT. Namely, f_{xc} establishes a unique map from non-interacting (Kohn-Sham) response functions to interacting response functions through the *Dyson equation*, and it is itself a functional of the density $f_{xc} = f_{xc}[n]$. However, the domain of linear response time-dependent DFT is, in part, *excited-state* properties, and hence the exact f_{xc} must implicitly include *all* requisite many-body effects in order to capture the observables for which it is appropriate to use, e.g. optical properties, correlation energies, and so on.

It should come as no surprise, then, that f_{xc} and v_{xc} are subject to similar scrutiny in the literature, albeit the latter is subject to much more than the former. As Chapter 4 reveals, model studies of the exact time-independent and time-dependent v_{xc} have successfully elaborated a number of its features and idiosyncrasies, such as the static and dynamic step [216], towards improved functional approximations. The aim of this project is to articulate f_{xc} in an analogous fashion, and in particular examine the structure and features of the *exact* $f_{xc}(x, x', \omega)$ for three canonical one-dimensional finite systems, including the full extent of its spatial and frequency dependence. The numerical difficulties inherent to such an endeavour has prevented or hindered studies of f_{xc} along these lines; Refs. [215, 235, 236] are notable prior works on this topic, one of which is authored by the same group with whom this work is done in collaboration.

The development and understanding of approximate xc kernels has been the subject of intense interest over the past decades, see, for example, [7, 237–240] and references therein. In most cases, these approximations can be sorted hierarchically depending on the level of theory involved in the approximation. The lowest orders of this hierarchy contain the random phase approximation (RPA) and the adiabatic local density approximation (LDA). The former ignores exchange and correlation *at the level of the xc kernel* entirely by setting $f_{xc}^{\text{RPA}} = 0$ [241], and the latter includes exchange and correlation within the framework of an LDA,

$$f_{xc}^{\text{ALDA}}(x, x', |t - t'|) = \left. \frac{\delta v_{xc}^{\text{LDA}}(x, t)}{\delta n(x', t')} \right|_{n_0} \propto \delta(t - t') \delta(x - x'), \quad (5.1)$$

which is local in time and local in space [242, 243]. The xc kernel itself, however, is known to possess a range of pathological features that depart significantly from the domain of these approximations, three examples of which are now given.

First, certain circumstances demand a spatial ultra-non-locality¹ in f_{xc} [7], in other words, a perturbing density at x can cause significant change in v_{xc} at some x' *arbitrarily far away from x* . Furthermore, in order to capture excitations of charge transfer character, an f_{xc} that is divergent in proportion to the increasing separation between the subsystems involved is required [245]. A charge transfer excitation is an excitation in which the ground-state density is localised in one region of a system (subsystem A), whereas the excited-state density is localised elsewhere (subsystem B). Excitations of this kind are elusive within present framework [245], one reason for which is the aforementioned divergence in f_{xc} whose form is a step-like structure that can be related to the derivative discontinuity [101].

The most pertinent drawback of the framework is that a non-adiabatic temporal structure is required in principle to capture excitations of a multi-particle character [80, 246–248]. An adiabatic f_{xc} includes the instantaneous response of v_{xc} due to perturbing densities, $f_{xc}(|t - t'|) \propto \delta(t - t')$, but, in true Markovian fashion, does not directly account for the cumulative effects of the past except through its current state at $t = t'$. This definition bears no clear relation to multi-excitations – a multi-excitation is an excitation of the interacting system $|\Psi_0\rangle \rightarrow |\Psi_f\rangle$ where $|\Psi_f\rangle$ has its basis set expansion dominated by a Slater determinant of single-particle states with *more than one particle promoted from its ground-state configuration*. At such an excitation, the Kohn-Sham response function has no pole, whereas the interacting response function *does* have a corresponding pole, see left panel of Fig. 5.1. Inspection of the Dyson equation, or the Casida equation [249], reveals, in fact, that a non-adiabatic f_{xc} cannot introduce poles where there are none in the Kohn-Sham response function, it can merely shift the poles already present. In this sense, a non-adiabatic f_{xc} is required to produce these multi-excitations which are of demonstrable practical importance² [80, 246], and the work to follow elaborates the precise structure of this non-adiabatic behaviour.

The above examples are three manifestations of the fact that the exact f_{xc} contains all correlated many-body effects. More sophisticated approximations to f_{xc} seek to include these effects in some form or another, such as those that utilise the GW approximation and the Bethe-Salpeter equation [239, 251–253], exact exchange kernels [254–258], and long-range corrected kernels [259, 260]. The last is an example of a functional approximation to f_{xc} that is able to capture so-called excitonic effects in the optical spectrum of extended (periodic) systems while remaining relatively inexpensive.

The use of model systems has been effective in developing understanding of f_{xc} . In particular, the frequency dependence of f_{xc} has been the subject of model analytic studies [246, 261–263], numerical studies using model Hamiltonians, e.g. the Hubbard model [264–267], and numerical studies of exact one-dimensional Hamiltonians in a truncated

¹The objects f_{xc} and v_{xc} represent a mathematical mapping, and thus Kohn’s *nearsightedness* principles [244] need not apply.

²Experimental and theoretical studies on the optical properties of certain *polyenes*, a class of organic molecules, constitute landmark examples of double excitations [248, 250]. The double excitations within these polyenes not only serve as confirmation of theoretical developments, but can also account for their presence in natural dyes.

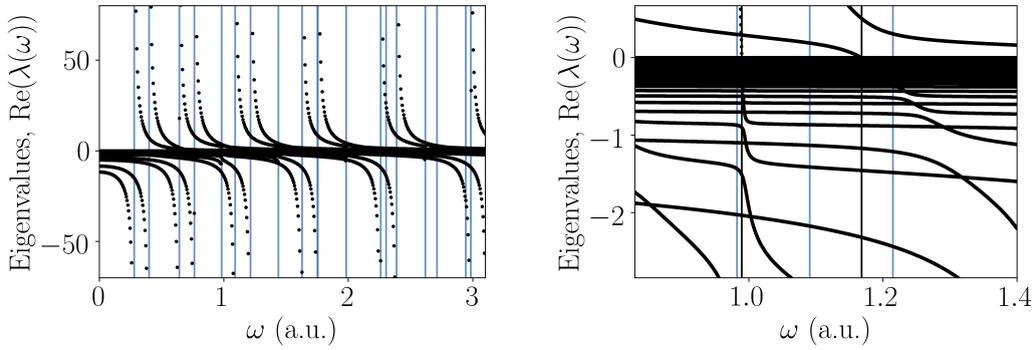


FIGURE 5.1: Eigenvalues of the interacting response function for infinite potential well, see Section 5.2. (Left) A single divergent eigenvalue (simple pole) exists at each interacting excitation energy (blue lines). (Right) Zooming into the range $\omega \in [0.9, 1.4]$, we find two instances of an eigenvalue $\lambda(\omega)$ crossing zero at $\omega = 0.99$ a.u. and $\omega = 1.18$ a.u. (black lines). The eigenvector $|u_i(\omega)\rangle$ associated with this eigenvalue is a *non-v-representable* density perturbation.

Hilbert space [215, 235, 236]. This work continues along the lines of the last approach, and seeks to address f_{xc} at frequencies that far surpass the first few excitation energies of the interacting Hamiltonian. The observed features of f_{xc} are examined in relation to matters of practical interest, such as optical properties. A detailed history of f_{xc} , including the successes and failures of the framework when compared with experiment, is given in [7, 19, 237] – the material therein that is relevant for the work to follow, including certain exact conditions (sum rules), and so on, is introduced where appropriate.

5.1 Methodology

5.1.1 Background

The `density2potential` code is used in order to obtain the interacting and Kohn-Sham response functions, see Chapter 4. For some input external potential $v_{\text{ext}}(x)$, the full set of eigenvectors $\{|\Psi_i\rangle\}$ of the interacting Hamiltonian is found using exact diagonalisation. The corresponding *exact* Kohn-Sham potential $v_{\text{KS}}(x)$ is then reverse-engineered by applying preconditioned root-finding techniques to an appropriate fixed-point map [226]. The full set of Kohn-Sham eigenvectors $\{|\phi_i\rangle\}$ is also obtained using exact diagonalisation. The (causal) response functions are computed in the frequency domain directly using the Lehmann representation,

$$\chi(x, x', \omega) = \lim_{\eta \rightarrow 0} \sum_{n=1}^{\infty} \frac{\langle \Psi_0 | \hat{n}(x) | \Psi_n \rangle \langle \Psi_n | \hat{n}(x') | \Psi_0 \rangle}{\omega - \Omega_n + i\eta} - \frac{\langle \Psi_0 | \hat{n}(x') | \Psi_n \rangle \langle \Psi_n | \hat{n}(x) | \Psi_0 \rangle}{\omega + \Omega_n + i\eta},$$

$$\chi_0(x, x', \omega) = \lim_{\eta \rightarrow 0} \sum_{n=1}^{\infty} \frac{\langle \Phi_0 | \hat{n}(x) | \Phi_n \rangle \langle \Phi_n | \hat{n}(x') | \Phi_0 \rangle}{\omega - \Omega_n^0 + i\eta} - \frac{\langle \Phi_0 | \hat{n}(x') | \Phi_n \rangle \langle \Phi_n | \hat{n}(x) | \Phi_0 \rangle}{\omega + \Omega_n^0 + i\eta},$$

where $\Omega_n = E_n - E_0$ (Ω_n^0) is the n^{th} excitation energy of the interacting (Kohn-Sham) Hamiltonian, and the response functions are zero for times $t > t'$. The state $|\Phi_0\rangle := |\Phi_{(0,1)}\rangle$ denotes the Slater determinant constructed from the zeroth and first Kohn-Sham states, where two-particle systems are considered hereafter. The infinitesimal η is treated as a parameter, and setting $\eta \neq 0$ is tantamount to evaluating the response functions slightly above the real axis, along $\omega + i\eta$, circumventing the poles. The corresponding optical spectrum constitutes distinct peaks with a Lorentzian profile, rather than a series of weighted delta functions – in this work, the *broadening parameter* η has been introduced for visual purposes.

Construction of the interacting response function in the above fashion is an accurate but demanding procedure, whereas the finite-difference methods outlined in [236] to construct the response functions are amenable to larger systems, but more prone to error: either method will suffice. On a finite spatial grid, the response functions at a given ω become *response matrices*, denoted $\chi(\omega)$ and $\chi_0(\omega)$, of dimension $N \times N$. The Dyson equation gives an alternate definition of the xc kernel,

$$f_{\text{xc}}(\omega) = \chi_0^{-1}(\omega) - \chi^{-1}(\omega) - f_{\text{H}}, \quad (5.2)$$

where X^{-1} is to be understood as the matrix inverse of X , and the Hartree kernel $f_{\text{H}} = \delta v_{\text{H}}/\delta n$ becomes softened along with the Coulomb interaction. This expression is used as the definition of f_{xc} in the present context, however, simply inverting the response matrices in order to construct the f_{xc} is not sufficient without further consideration.

5.1.2 Challenges in Computing Exact Exchange-Correlation Kernels

Numerical difficulties arise when attempting to construct the f_{xc} as an object in itself using Eq. (5.2). First, f_{xc} represents the solution to an *inverse problem*, i.e. find the δv_{xc} that produces a *given* δn , and inverse problems are notoriously sensitive to small error [268], such as those introduced by finite-precision arithmetic. As discussed, f_{xc} Eq. (5.2) requires the matrix inverse of the response matrices at a given ω , and hence a naive inversion procedure introduces numerical error at a given ω in proportion to the condition of the response matrices, that is, the ratio of the maximum to minimum eigenvalue. Physical eigenvalues that are close to, or below, machine precision manifest in the response matrices from various sources [269], some of which are now elaborated.

In the case $\eta = 0$, the interacting response function admits the eigendecomposition

$$\chi(\omega) = \sum_{i=0}^N \lambda_i(\omega) |u_i(\omega)\rangle \langle u_i(\omega)|, \quad (5.3)$$

where $\{\lambda_i, |u_i\rangle\}$ are the eigenpairs of the response function. That is to say, a perturbation of the form

$$|v_{\text{ext}}(\omega)\rangle \rightarrow |v_{\text{ext}}(\omega)\rangle + |u_i(\omega)\rangle \quad (5.4)$$

will induce the reaction,

$$|n(\omega)\rangle \rightarrow |n(\omega)\rangle + \lambda_i |u_i(\omega)\rangle. \quad (5.5)$$

A consequence of this view is that a perturbation δv that is oscillating in time with frequency ω yields a density response δn that also oscillates in time with frequency ω . The sources of eigenvalues whose magnitude falls below machine precision, i.e. ill-conditioning, can be analysed on these grounds.

One such source of eigenvalues that approach, and even touch, zero is related to **the linear response v -representability problem**. That is to say, there exists density perturbations $\delta n(\omega)$ that oscillate with some non-resonant frequency ω and cannot be produced by a perturbing potential $\delta v(\omega)$ at linear order³ [254, 269, 270], see the right-hand panel of Fig. 5.1 and its associated caption. The linear response v -representability problem has received much less attention than its analogue in (‘non-linear response’) time-dependent DFT, in part because it has become conventional wisdom that it is not such a problem at all, a wisdom this work refutes. Since perturbing potentials of the ‘switch-on’ form $\delta n(x, t) = \Theta(t - t_0)g(x, t)$ are linear response v -representable [269], where $\Theta(t - t_0)$ is the step function at $t = t_0$, one might expect that non- v -representable density perturbations oscillating with frequency ω *for all time* are a mathematical curiosity, rather than something of practical concern. However, the zero eigenvalue associated with the non- v -representable density manifests as a pole (singularity) in f_{xc} , and the divergent character of f_{xc} around such poles is found to be crucial for recovering optical properties of the interacting system. This behaviour in f_{xc} is missing from most functional approximations.

The most serious cause of ill-conditioning is identified as regions of nearly vanishing *ground-state* density: the response in the density due to a perturbing potential that is localised far outside a high-density region (confining region of the potential) is negligible. The *inverse* response function necessarily diverges for such perturbing potentials in order to produce density in a region where it was previously unfavourable to do so (outside the confined region). Since the aim of this work is, in part, to study the optical response of *confined* systems far beyond the first excitation, an appropriately large spatial domain $[-a, a]$ is required in order to accommodate the more extended excited states without introducing spurious features due to the boundary conditions. In other words, in order to study f_{xc} at frequencies beyond the first excitation, a large spatial domain is unavoidable, and a central novelty of work to follow is a solution to the ensuing ill-conditioning.

Note that these near-machine-precision eigenvalues of the response functions are not necessarily problematic. Rather, the problem occurs when the Kohn-Sham response function differs meaningfully from the interacting response function within the eigenspace that is near or below machine precision. In this case, f_{xc} is divergent and excessively sensitive to small errors – behaviour that is *required* in order to capture the small difference between χ and χ_0 , and in turn this difference is required in order to capture something physical, such as the position and weight of peaks in the optical spectrum. The prospects for constructing such numerical xc kernels are bleak, and indeed impossible, if the meaningful eigenvalues span more than the range afforded by machine precision. A notable instance of this is given in an example above, namely, the case of charge transfer, whereby in order to capture excitations of charge transfer character an f_{xc} that is divergent (ill-conditioned) in proportion to the increasing separation between the subsystems involved is required [245].

³A familiar example: the constant perturbation $\delta v = c(\omega)$ oscillating with frequency ω produces no response in the density δn at *all* orders, including at first order.

Provided the near machine precision difference between the interacting and Kohn-Sham response functions is *not* responsible for some particular physical phenomenon, then an accurate numerical f_{xc} can be obtained, even if the response functions themselves are individually ill-conditioned. That is, the ill-conditioning of the interacting and Kohn-Sham response functions can be assumed to cancel in the definition of f_{xc} , Eq. (5.2), thus producing a regular f_{xc} . This procedure can be viewed as a form of basis set truncation, i.e. assign $\chi = \chi_0$ within some subset of the basis responsible for ill-conditioning and proceed to compute the f_{xc} under this assumption. Two such truncations are considered here: a truncation in real space, and a truncation in eigenspace.

5.1.3 The Real-Space Truncation Method

In finite systems with a confining potential, the response functions tend toward zero outside of the confined region, and this so-called long-range behaviour is known to be relatively unimportant in the present context – this is not the case in periodic systems [168, 237, 239, 271]. Therefore, as we demonstrate in this work, forcing the interacting response function to equal the non-interacting response function within some yet undefined *outer region* does not much alter the derived properties of the interacting response function, such as its optical spectrum. The interacting and Kohn-Sham response functions for the atom, a system borrowed from Section 5.2 for illustrative purposes, are shown in Fig. 5.2. The spatial extent of the confined region, i.e. the atomic core, is apparent, and while the response functions differ inside the confined region, they converge uniformly toward zero outside of the confined region.

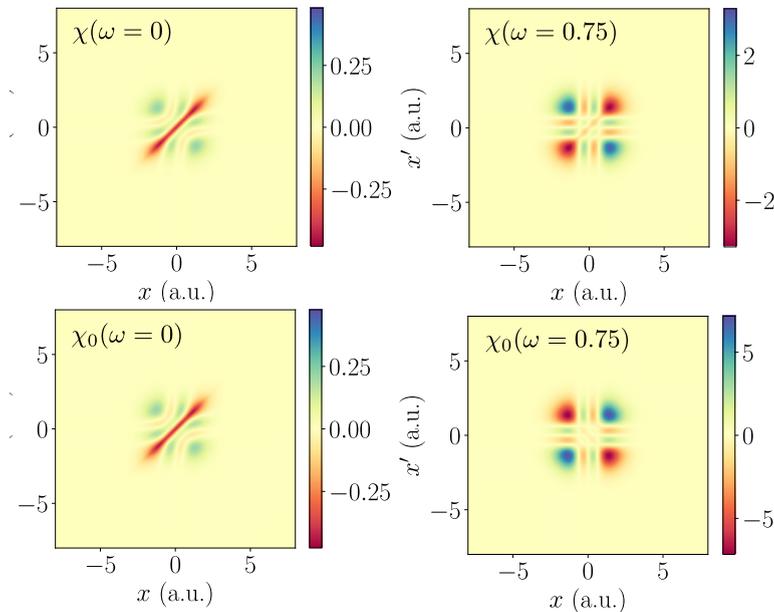


FIGURE 5.2: The atomic response functions of the interacting χ (upper) and Kohn-Sham χ_0 (lower) systems at $\omega = 0$ a.u. and $\omega = 0.75$ a.u. (around the first interacting excitation). The response functions both tend toward zero uniformly outside of the confined region.

To this end, a partition of the spatial domain $[-a, a]$ is made such that an *inner region* is defined where x takes values $-b \leq x \leq b$; the numerical parameter b defines the extent of the truncation. The *outer region* constitutes the remaining space between the inner region and the edges of the domain, $-a$ and a . This partition of the space, as it applies to the response functions, can be seen in Fig. 5.3. The assumption is then made that

$$\begin{aligned}\tilde{\chi}(x, x', \omega) &:= \chi(x, x', \omega) \text{ for } (x, x') \text{ in inner region} \\ \tilde{\chi}(x, x', \omega) &:= \chi_0(x, x', \omega) \text{ for } (x, x') \text{ in outer region}\end{aligned}$$

where $\tilde{\chi}$ is the *truncated* response function. The xc kernel is now defined as the object

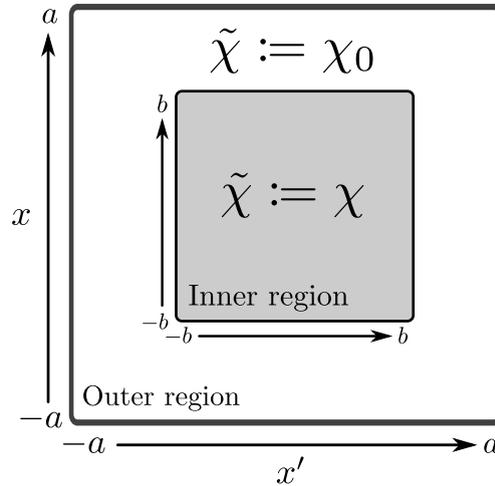


FIGURE 5.3: A schematic depiction of the *real-space truncation* strategy used to regularise computations of f_{xc} , whereby a truncated response function $\tilde{\chi}$ is defined as the interacting response function within some region parameterised by b , the *inner region* (shaded gray), and is otherwise set equal to the Kohn-Sham response function.

that returns the truncated response function $\tilde{\chi}$, rather than the interacting response function χ , upon solution of the Dyson equation. This leads to the following piecewise form for f_{xc} ,

$$f_{\text{xc}} = \begin{cases} \chi_0^{-1} - \chi^{-1} - f_{\text{H}} & \text{for } (x, x') \text{ in inner region} \\ -f_{\text{H}} & \text{for } (x, x') \text{ in outer region.} \end{cases}$$

The regularising effect of the method presented above can be understood by examining the role of the truncation parameter b . In particular, two extremes are considered, first, setting $b = a$ means the inversion of the response matrix at a given ω is performed over the whole domain, and hence is dominated by error due to excessive regions of nearly vanishing ground-state density; this error is hereafter referred to as the *numerical error*. Setting $b = 0$ turns the truncated response function into the Kohn-Sham response function over the entire domain – an evidently unsatisfactory state of affairs – and error of this kind is referred to as *method error*. Whilst this need not be the case in principle, it is the case for the systems studied here that it is possible to choose b such that an acceptable balance is struck between method error and numerical error. In other words, the truncated response function is able to retain all the physical properties of the interacting response function *and* ensure the computation of the resulting piecewise

f_{xc} is well-conditioned. An elaboration on the distinction between method error and numerical error is given in Section 5.1.5 below.

One might take issue that the above piecewise expression for f_{xc} is spuriously discontinuous at the boundary of the inner and outer region, where the extent of this discontinuity depends on the long-range behaviour of f_{xc} . However, we shall be concerned with the behaviour of f_{xc} inside the inner region, i.e. the region where departure of the Hxc kernel $f_{Hxc} = f_{xc} + f_H$ from zero produces meaningful features in the output of the Dyson equation.

5.1.4 The Eigenspace Truncation Method

A second, related, method used in this work in order to regularise the computation of f_{xc} is to truncate the interacting response matrix in the eigenspace of the Kohn-Sham response matrix. This method is much more accurate than the real-space truncation, but is limited to Hermitian response matrices, i.e. response matrices constructed without an artificial broadening η . Consider the eigendecomposition of the interacting and Kohn-Sham response matrices at a given ω , where the eigenpairs are denoted $\{|u_i\rangle, \lambda_i\}$ and $\{|u_i^{KS}\rangle, \lambda_i^{KS}\}$ respectively. Consider further some value $\bar{\lambda}$ such that the *effective null space*, Null_{eff} , is defined as the subspace spanned by eigenvectors whose eigenvalue has modulus below $\bar{\lambda}$, e.g.

$$\text{Null}_{\text{eff}}(\chi_0) = \text{Span}(\{|u_i^{KS}\rangle \mid |\lambda_i^{KS}| < \bar{\lambda}\}). \quad (5.6)$$

The assumption is now made that the *truncated response matrix* $\tilde{\chi}$ operates on vectors that are elements of the effective null space as the Kohn-Sham response matrix,

$$\tilde{\chi}|v\rangle = \chi_0|v\rangle \text{ for } |v\rangle \in \text{Null}_{\text{eff}}(\chi_0), \quad (5.7)$$

see Fig. 5.4. Let us define the projection operator onto the effective null space,

$$P_N = \sum_{i \in \text{Null}_{\text{eff}}} |u_i^{KS}\rangle \langle u_i^{KS}|. \quad (5.8)$$

The expression in Eq. (5.7) is established as follows,

$$\tilde{\chi} = (I - P_N)\chi + P_N\chi_0; \quad (5.9)$$

the first term on the right-hand side removes the effective null space from χ , and the second term ensures $\tilde{\chi}$ operates as intended on elements of the effective null space. Another view of this manipulation is that the truncated and Kohn-Sham response functions are required to share eigenvectors and eigenvalues inside the effective null space,

$$\{|\tilde{u}_i\rangle, \tilde{\lambda}_i\} = \{|u_i^{KS}\rangle, \lambda_i^{KS}\} \text{ for } |\lambda_i^{KS}| < \bar{\lambda}, \quad (5.10)$$

but is otherwise equal to the interacting response function, this is also depicted in Fig. 5.4. The *pseudoinverse* [272] with cutoff $\bar{\lambda}$, i.e. the eigendecomposition with eigenpairs below $\bar{\lambda}$ discarded, is now an exact procedure to obtain f_{xc} that recovers $\tilde{\chi}$ after solution

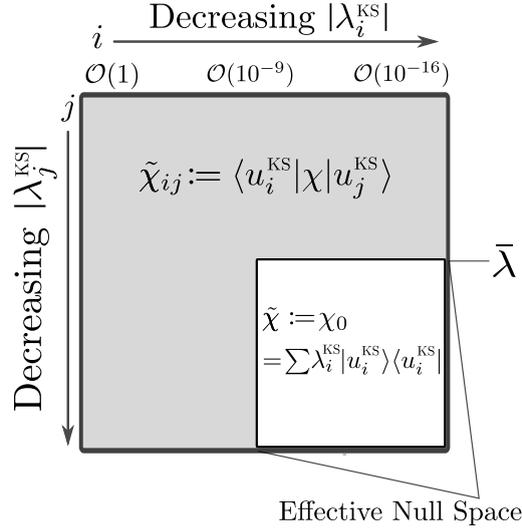


FIGURE 5.4: A schematic depiction of the *eigenspace truncation* strategy used to regularise computations of f_{xc} , whereby the interacting response function is expanded in the basis of eigenvectors of the Kohn-Sham response function, and set equal to the Kohn-Sham response function inside the effective null space, parameterised by $\bar{\lambda}$.

of the Dyson equation,

$$f_{xc} = \chi_0^+ - \tilde{\chi}^+ - f_H, \quad (5.11)$$

where X^+ denotes the pseudoinverse of X . The pseudoinverse shall not be defined in detail here, but note that its use is *not* an approximation (as is often the case), rather, it is an exact procedure to construct the relevant f_{xc} , made possible due to the central assumption that $\tilde{\chi}$ and χ_0 share eigenpairs below $\bar{\lambda}$. In direct analogy with the previous method, the parameter $\bar{\lambda}$ assumes the job of b , namely, it parameterises the extent of the truncation. The difference between the truncated response function Eq. (5.7) and the exact interacting response function is again given the name *method error*. Note that this approach is quite distinct from applying the pseudoinverse to the response functions in Eq. (5.2) – doing so would introduce much more error and the source of this error is not clear. On the contrary, the error inherent in the method presented here is identified as the extent to which the effective null space of the interacting and Kohn-Sham response functions do not overlap, and this error can be tracked without reference to f_{xc} using the method error.

As we shall see shortly, visualisation of the optical spectrum relies on evaluating the response functions slightly above the real axis in the frequency domain, along $\omega + i\eta$. This leads to response matrices that are complex-symmetric, and thus (weakly) non-Hermitian. The eigenspace is no longer orthogonal, and the eigendecomposition used above does not exist. In this case, the real-space truncation of Section 5.1.3 is used, and the eigenspace truncation method is used elsewhere.

5.1.5 Error and Exact Conditions

Having established two separate methods for regularising the computation of the exact f_{xc} , we are able to discuss and quantify error in the present context. The foremost approach used to track error here is the mean absolute error between elements of the interacting response function χ and elements of the output of the Dyson equation χ_{Dyson} over the entire spatial and frequency grid. The quantity χ_{Dyson} is defined as

$$\chi_{\text{Dyson}}(\omega) = \frac{\chi_0(\omega)}{I - \chi_0(\omega)(f_{xc}(\omega) + f_H)}, \quad (5.12)$$

and the mean absolute error between the output of the Dyson equation and the interacting response function is defined as

$$\|\chi - \chi_{\text{Dyson}}\| := \sum_{(x,x',\omega) \text{ grid}} \frac{|\chi(x, x', \omega) - \chi_{\text{Dyson}}(x, x', \omega)|}{N_x^2 N_\omega}. \quad (5.13)$$

The expression for χ_{Dyson} in Eq. (5.12), when evaluated in exact arithmetic with an exact f_{xc} , is equal to the interacting response function. This clarifies the nomenclature used to distinguish the two errors of concern to this work. Namely, the *method error* is understood as the difference between $\tilde{\chi}$ and χ , and the *numerical error* is the difference between $\tilde{\chi}$ and χ_{Dyson} . The truncation parameters b and $\bar{\lambda}$ are chosen such that the mean absolute error between χ and χ_{Dyson} is approximately equal to the method error, meaning little-to-no numerical error is present – a practical example of the process involved in choosing a truncation parameter is given in Section 5.2.1. Corroboration between both methods is also used as a means of extracting signal over noise.

The xc kernel is known to be subject to a multitude of exact conditions [72]; these conditions have been important in both the historical development of f_{xc} approximations [19], and also in other studies on the exact f_{xc} [215, 235, 236]. The xc kernels in this work are complex-symmetric and satisfy the Kramers-Kronig relations by construction. The property that is most useful in the present context is the *zero-force sum rule* [7],

$$\int_{-a}^a n(x') \frac{\partial}{\partial x'} f_{xc}(x, x', \omega) dx' = -\frac{\partial}{\partial x} v_{xc}(x), \quad (5.14)$$

which states that perturbations in the xc potential cannot produce a net momentum over the entire domain at linear order, i.e. total momentum is conserved at zero. The extent to which the xc kernels computed here satisfy the zero-force sum rule can be seen in both a graphical and numerical form, given in Section 5.2. A final test of error is agreement between the optical spectra derived from χ and χ_{Dyson} when using the exact numerical f_{xc} . This agreement would further evidence the fact that the truncation did not change observable properties of interest, while maintaining numerical stability.

5.1.6 Gauge Freedom

The gauge freedom of f_{xc} is seldom discussed in literature, but is of primary importance for this work. As noted in Refs. [101, 254, 255, 264], the following transformation

$$f_{xc}(x, x', \omega) \rightarrow f_{xc}(x, x', \omega) + g(x, \omega) + h(x', \omega) + c(\omega) \quad (5.15)$$

leaves the output of the Dyson equation unchanged, and thus we are, in principle, free to choose the arbitrary complex-valued functions $g(x, \omega)$, $h(x', \omega)$ and $c(\omega)$. All three transformations are a direct manifestation of the invariance of quantum Hamiltonian systems under a constant time-dependent shift of the potential⁴. From the point of view of f_{xc} approximations, two f_{xc} 's are equivalent if they exist within family of functions given in Eq. (5.15)⁵. A preferred gauge is defined by Eq. (5.2), since the objects χ and χ_0 are *themselves* invariant to a shift in the potential. The unique f_{xc} , modulo a constant shift (see below), defined in Eq. (5.2) can be considered the physical f_{xc} , and it is this definition of f_{xc} that is assumed in discussions on its various properties and limits [168, 237, 239, 271, 273]. To modify this f_{xc} using its gauge freedom changes its underlying structure; for example, setting $g \neq 0$ gives f_{xc} spurious long-range behaviour, and setting $g \neq h$ produces an f_{xc} that is not symmetric under interchange of $x \leftrightarrow x'$. In this work, we illustrate f_{xc} as it is defined in Eq. (5.2), which also *defines* $g = h = 0$. The constant shift c has special meaning, since it is itself an eigenvector of χ and χ_0 with eigenvalue zero, i.e. the response functions are non-invertible in this direction. The exact f_{xc} thus has a true arbitrariness corresponding to the constant shift c , meaning there exists no preferred centering of f_{xc} . The convention adopted in this work to centre f_{xc} is such that, in the long-range limit, far outside the confined region, f_{xc} approaches minus the Hartree kernel, $f_{xc} \rightarrow -f_H$.

Having decided upon a preferred gauge, one can consider the possible consequences of this gauge freedom on matters of practical interest. An approximate f_{xc} that differs in relevant structure from the exact f_{xc} largely due to a change of gauge provides at least a partial explanation for the performance of a given approximate f_{xc} . In Section 5.2.5 we consider this line of inquiry, namely, whether or not this structural freedom afforded to f_{xc} is sufficient to provide an explanation for the agreement between two xc kernel approximations, or between an xc kernel approximation and the exact xc kernel.

5.2 Results and Discussion

The generic numerical parameters used to characterise the systems to follow are, unless stated otherwise: a real-space grid of dimension $N_x = 151$, a frequency grid of dimension $N_\omega = 2000$, a Coulomb softening parameter $\alpha = 1$, and a finite broadening (where appropriate) $\eta = 0.0005$.

⁴The derivation of this gauge follows directly from substitution of Eq. (5.15) into the Dyson equation, and use of the identity $\int \chi(x, x', \omega) dx = 0$, which is a statement that the constant function $c(\omega)$ is an eigenvector of χ with eigenvalue 0 (constant shifts produce no response).

⁵Note that the quantities $\langle ij|f_{xc}(\omega)|kl\rangle$, where (i, j, k, l) label indices of single-particle states, are unique [254], and it is these quantities that form the input to the Casida equation [249], for example.

5.2.1 Atom

Two interacting electrons confined in an atom-like (soft core) potential,

$$v_{\text{ext}}(x) = -\frac{2}{|\frac{1}{10}x| + \frac{1}{2}}, \quad (5.16)$$

within the domain $[-8, 8]$ a.u. is considered. An illustration of this system is given in Fig. 5.5, including its corresponding Kohn-Sham potential that is reverse-engineered to within $\mathcal{O}(10^{-15})$ error using the methods of Chapter 4. The purpose of the atom demonstration is two-fold. First, it constitutes a proof-of-concept, and defines a standard of accuracy to which the remainder of the calculations are held unless stated otherwise. Second, the optical spectrum of the atom is calculated in the range $\omega \in [0, 6]$ a.u., which includes many excitations beyond the first, and the efficacy of various approximations to f_{xc} are examined in relation to the optical spectrum.

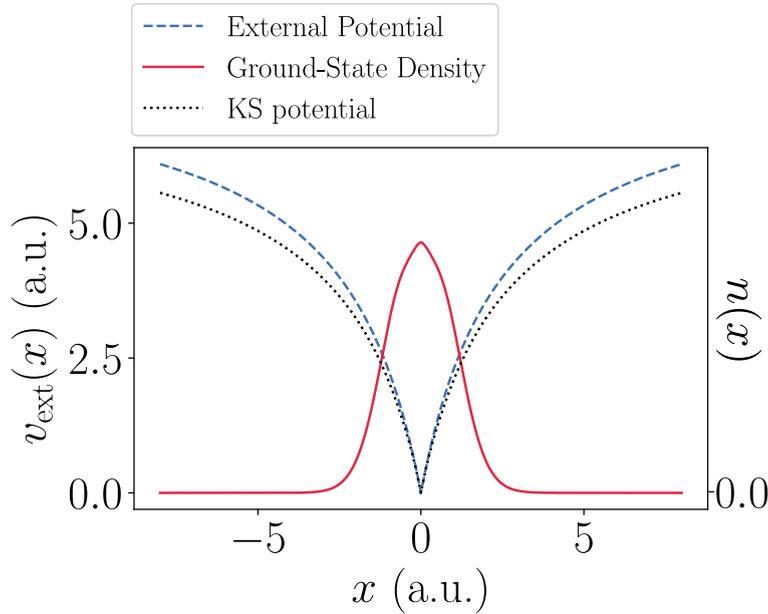


FIGURE 5.5: The ground-state density, external potential, and reverse-engineered Kohn-Sham potential for the atomic system. The external and Kohn-Sham potential have been shifted for illustrative purposes.

The exact f_{xc} , constructed using the real-space truncation method, is shown for the first three visible interacting excitations in the optical spectrum, and at $\omega = 0$ a.u., in Fig. 5.6. The last is sometimes termed the *adiabatic exact* (AE) f_{xc} , and it correctly describes any system in which the response to a perturbation is essentially instantaneous, $f_{\text{xc}}(x, x', \omega = 0)$. The real-space truncation parameter is chosen as $b = 5.8$ a.u., meaning the inner region is defined as $-5.8 < x < 5.8$. It is important to stress that this choice of b is *not* unique, and there exists some feasible range of b within which f_{xc} itself is insensitive to changes. Moreover, within this feasible range, both the method error and numerical error are acceptable and equal approximately to $\|\chi - \chi_{\text{Dyson}}\| = \mathcal{O}(10^{-9})$. As discussed in the previous section, the eigenspace truncation method is *more* accurate than the

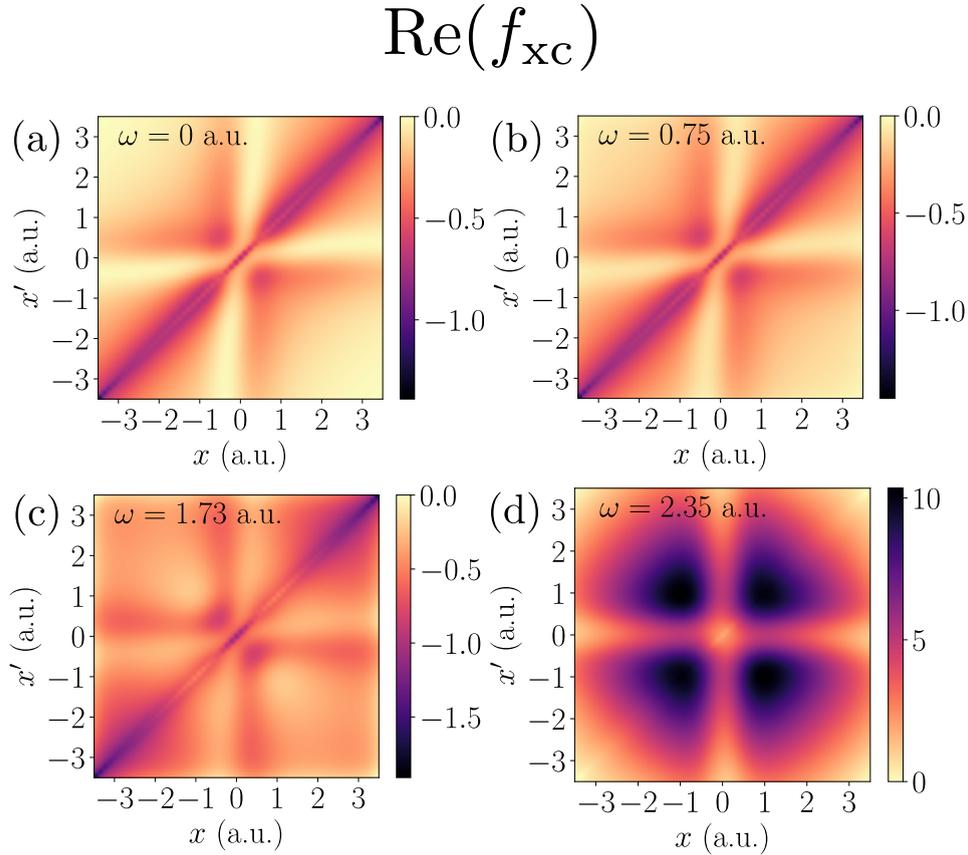


FIGURE 5.6: The real part of the exact numerical xc kernel $f_{xc}(x, x', \omega)$ for the atom at (a) $\omega = 0$ (AE), (b) the first visible interacting excitation, (c) the second visible interacting excitation, and (d) the third visible interacting excitation. For illustrative purposes, f_{xc} is shown for $-3.5 < x < 3.5$, where its essential structure is most visible. The xc kernel across the first two transitions remains approximately equal to the AE $f_{xc}(\omega = 0)$, after which a significant departure from the adiabatic limit is observed. The AE $f_{xc}(\omega = 0)$ displays considerable non-local structure, despite a local dominance along $x = x'$. The centering convention, i.e. the convention used to fix the arbitrary shift c , relates to the long-range limit of f_{xc} , which is discussed in depth in Section 5.2.4.

real-space truncation, and it is now used to demonstrate the relationship between the method error, numerical error, and the truncation parameter(s) in practice. The method error and numerical error, over a range of eigenspace truncation parameters $\bar{\lambda}$, are given in Table 5.1. The first row highlights the essence of the issue at hand: setting the cutoff to machine precision, so that the effective null space no longer exists, gives us a ‘best possible’ error of zero using exact arithmetic, but an excessively large error in finite precision – an error 10^{-2} indicates that the reconstructed response function is poor beyond repair. Moving down the rows of Table 5.1, it can be seen that increasing the cutoff introduces mild method error whose source is the extent to which the effective null space of χ and χ_0 *do not overlap*. The increased cutoff has a significant regularising effect, whereby using $\bar{\lambda} = 10^{-8}$ allowed the truncated response function at $\omega = 0$ to be reconstructed with a numerical error $\mathcal{O}(10^{-13})$. As expected, increasing the cutoff further begins to introduce significant method error that is matched well by the numerical error, indicating $\chi \neq \chi_0$ within this definition of the effective null space. The cutoff is

TABLE 5.1: The *method error* $\|\tilde{\chi} - \chi\|$ is the difference between the truncated response function, and the interacting response function, and the *numerical error* $\|\tilde{\chi} - \chi_{\text{Dyson}}\|$ is the difference between the truncated response function, and the output of the Dyson equation. Both of these are shown here over a range of eigenspace truncation parameters $\bar{\lambda}$. Increasing the truncation parameter introduces error whose source is the approximation inherent to the truncation itself, but doing so has a significant regularising effect, and the numerical error decreases accordingly. It is possible to choose a truncation parameter, e.g. $\bar{\lambda} = 10^{-8}$, such that both the method error *and* numerical error are sufficiently low for the systems studied in this work.

Truncation Parameter $\bar{\lambda}$	Method Error $\ \tilde{\chi} - \chi\ $	Num. Error $\ \tilde{\chi} - \chi_{\text{Dyson}}\ $
1×10^{-17}	0.0000×10^0	3.4852×10^{-2}
1×10^{-14}	7.7170×10^{-18}	1.8161×10^{-9}
1×10^{-11}	1.8947×10^{-17}	2.7201×10^{-11}
1×10^{-8}	6.7188×10^{-14}	2.1082×10^{-13}
1×10^{-3}	1.3512×10^{-8}	2.6115×10^{-8}

thus chosen to represent the best balance between these two errors, which is, in this instance, $\bar{\lambda} = 10^{-9}$ yielding an error $\|\chi - \chi_{\text{Dyson}}\| \approx 10^{-11}$ (two orders of magnitude superior to the real-space truncation method).

The zero-force sum rule is calculated across the the entire ω -range, and gives an average error between the left- and right-hand side in Eq. (5.14) of $\mathcal{O}(10^{-6})$ *within the inner region determined by b* , this agreement is illustrated in Fig. 5.7. Outside of this region, which is not shown, the error in the zero-force condition is significant, but the outer region contributes little to recovering the exact response function and other derived properties.

In order to extract the optical transition energies and transition rates, given a single-particle response function χ_0 and xc kernel f_{xc} , one can construct the entire optical absorption spectrum, see Chapter 2, using the corresponding output of the Dyson equation, denoted $\chi_{\text{Dyson}}(f_{\text{xc}}, v_{\text{xc}})$, where χ_0 is specified with some v_{xc} , see Fig. 5.8. As established in [82, 274, 275], the exact Kohn-Sham single-particle transitions are in excellent agreement with the interacting transitions, but this agreement becomes increasingly poor at higher energies. An approach to understanding this is to consider the overlap between the final states involved in a given interacting $|\Psi_0\rangle \rightarrow |\Psi_f\rangle$ and Kohn-Sham $|\Phi_{(0,1)}\rangle \rightarrow |\Phi_f\rangle$ transition, where $|\Phi_{(i,j)}\rangle$ denotes the Slater determinant constructed from i^{th} and j^{th} Kohn-Sham single-particle states. Such considerations, where possible, permit an identification between the non-interacting and interacting transitions, see Table 5.2, although the existence of this identification is not guaranteed, and depends, amongst other things, on the validity of the single-particle picture.

The overlap of the ground-state is $\langle \Psi_0 | \Phi_{(0,1)} \rangle = 0.99991$, and the overlap of the final states involved in the first transition at $\omega = 0.76$ a.u. is $\langle \Psi_1 | \Phi_{(0,2)} \rangle = 0.9995$. The *static correlation* in the interacting state here is modest, meaning the interacting state has a

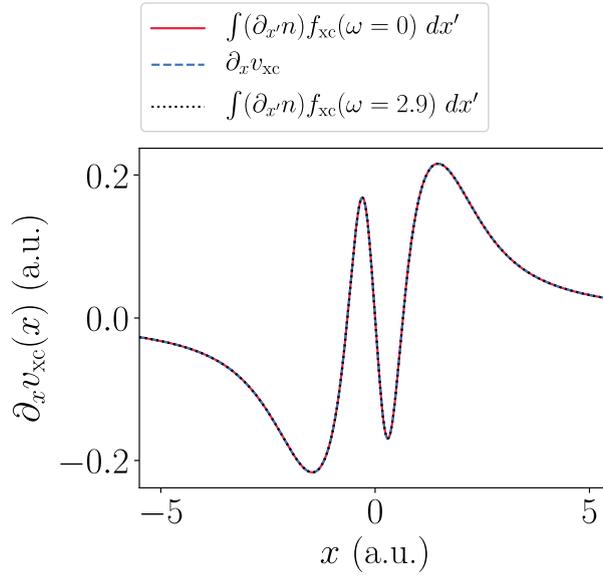


FIGURE 5.7: Agreement between the left-hand and right-hand sides of the zero-force sum rule, Eq. (5.14), for the exact atomic numerical f_{xc} at $\omega = 0$ a.u. and $\omega = 2.9$ a.u.

strong single-particle character, which leads to agreement between the low-energy transitions in the optical spectrum. At higher energies, the overlap decays by multiple orders of magnitude, however this is not the predominant source of error in higher energy transitions. Rather, interacting excitations that correspond to Kohn-Sham single-particle excitations out of the highest occupied (second) state $|\phi_1\rangle$ are much more accurate than single-particle excitations out of the first state, $|\phi_0\rangle$. For example, the interacting excitation $|\Psi_0\rangle \rightarrow |\Psi_{19}\rangle$ at $\omega = 4.41$ a.u. in Fig. 5.8 is captured well with the Kohn-Sham excitation $|\Phi_{(0,1)}\rangle \rightarrow |\Phi_{(0,12)}\rangle$, whereas this is not true of the preceding interacting excitation. This is not surprising, as the highest occupied Kohn-Sham state has energy equal to minus the *exact* electron removal energy [276], and thus at least one energy involved in the transition is correct. This might often be the case, and to comment further would require additional many-body calculations.

The following f_{xc} approximations are now considered: the RPA $f_{xc}^{\text{RPA}} = 0$, an adiabatic LDA $f_{xc}^{\text{ALDA}}[n](x, x', \omega = 0) \propto \delta(x - x')$ parameterised with reference to the homogeneous electron gas in [277], and the AE xc kernel $f_{xc}(x, x', \omega = 0)$, Fig. 5.6. These f_{xc} approximations are used to solve the Dyson equation in conjunction with the *exact Kohn-Sham response function*, for a discussion on the importance of the ground-state xc potential, see Section 5.2.6. The atomic optical spectrum, using the aforementioned series of approximations, is shown in Fig. 5.9 for the first transition and a chosen higher energy transition. The visual structure of the AE f_{xc} , see Fig. 5.6, is observed to be indistinguishable from f_{xc} at the first excitation; this similarity carries over to the optical spectrum, where the peak corresponding to the first excitation is reproduced exactly. This agreement demonstrates that, not only is the adiabatic approximation valid for low-energy transitions, but also the non-local spatial structure in the AE f_{xc} is required in

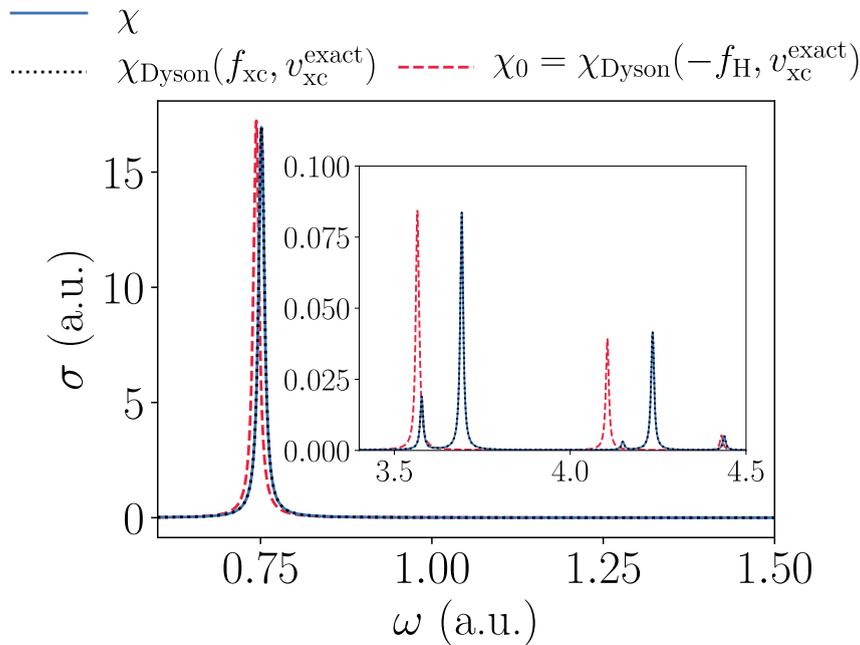


FIGURE 5.8: The optical spectrum for the atom is computed using the interacting response function χ (blue solid), the Kohn-Sham response function χ_0 (red dash), and the output of the Dyson equation χ_{Dyson} with the exact $f_{\text{xc}}(x, x', \omega)$ (black dot). The exact numerical f_{xc} reproduces the interacting peaks perfectly, as expected.

order to reproduce the low-energy peaks in the optical spectrum – in lacking this structure, both the RPA and adiabatic LDA significantly over-correct the underestimation of the non-interacting transition energy.

At higher energies, i.e. beyond the third peak in the optical spectrum (not shown), all three approximations perform similarly, and do not improve matters significantly beyond the corresponding non-interacting peak. This behaviour appears to be generic for all peaks observed up to $\omega = 6$ a.u., namely, the transition energies output from the Dyson equation with these approximate xc kernels are bound to the quality of the non-interacting transition energies. Furthermore, the inset of Fig. 5.9 demonstrates that the AE approximation gives an excitation energy that is worse than the other *adiabatic* approximations. This suggests that, in order to capture higher energy transitions, a frequency dependence is required, and in particular ‘improving’ the spatial structure of adiabatic approximations toward the AE structure does not assist matters – examination of this observation is the basis for the next section.

As alluded to above, the AE approximation ceases to out-perform the adiabatic LDA and RPA beyond the third transition in the optical spectrum, i.e. the same transition for which the corresponding *exact* f_{xc} departs from its adiabatic structure in a serious manner, Fig. 5.6. In fact, as the section to follow demonstrates, this violent departure from adiabaticity has a particular and fairly simple form, and its origin is understood in the context of eigenvalues of $\chi(\omega)$ or $\chi_0(\omega)$ that *cross* zero. The eigenvector corresponding to the eigenvalue that touches zero is a non- v -representable density perturbation at linear order.

TABLE 5.2: Identification of interacting and Kohn-Sham transitions for the atom using the overlap of the interacting and Kohn-Sham final states as a measure. This includes the first double excitation at $\omega = 3.5594$ a.u. that is not visible in the optical spectrum.

Exact Transition	Kohn-Sham transition	Appears in Optical Spectrum	Overlap in Final State	Energy (a.u.)	
				KS	Exact
$ \Psi_0\rangle \rightarrow \Psi_1\rangle$	$ \Phi_{(0,1)}\rangle \rightarrow \Phi_{(0,2)}\rangle$	✓	0.9996	0.7431	0.7513
$ \Psi_0\rangle \rightarrow \Psi_2\rangle$	$ \Phi_{(0,1)}\rangle \rightarrow \Phi_{(0,3)}\rangle$	✗	0.9993	1.3230	1.3291
$ \Psi_0\rangle \rightarrow \Psi_3\rangle$	$ \Phi_{(0,1)}\rangle \rightarrow \Phi_{(0,4)}\rangle$	✓	0.9995	1.7274	1.7325
$ \Psi_0\rangle \rightarrow \Psi_4\rangle$	$ \Phi_{(0,1)}\rangle \rightarrow \Phi_{(0,5)}\rangle$	✗	0.9987	2.0721	2.0762
$ \Psi_0\rangle \rightarrow \Psi_5\rangle$	$ \Phi_{(0,1)}\rangle \rightarrow \Phi_{(1,2)}\rangle$	✓	0.9978	2.2384	2.3118
$ \Psi_0\rangle \rightarrow \Psi_9\rangle$	$ \Phi_{(0,1)}\rangle \rightarrow \Phi_{(1,3)}\rangle$	✓	0.9842	2.8183	2.9398
$ \Psi_0\rangle \rightarrow \Psi_{14}\rangle$	$ \Phi_{(0,1)}\rangle \rightarrow \Phi_{(2,3)}\rangle$	✗	0.9917	3.5594	3.7065
$ \Psi_0\rangle \rightarrow \Psi_{18}\rangle$	$ \Phi_{(0,1)}\rangle \rightarrow \Phi_{(1,7)}\rangle$	✓	0.9791	4.2339	4.1056
$ \Psi_0\rangle \rightarrow \Psi_{19}\rangle$	$ \Phi_{(0,1)}\rangle \rightarrow \Phi_{(0,12)}\rangle$	✓	0.9814	4.4382	4.4301

5.2.2 Infinite Potential Well

The infinite potential well is defined with the external potential $v_{\text{ext}} = 0$ inside the domain $[-5, 5]$ a.u., see Fig. 5.10. This system has no region of near-zero density toward the outer parts of the domain, and therefore the response functions do not succumb to numerical ill-conditioning apart from due to the constant shift c which has eigenvalue zero. This eigenvalue and its eigenvector are discarded⁶ using the singular value decomposition (SVD). Excitations of the infinite potential well are valid up to arbitrary energy, since the external potential coincides with the boundary conditions of the simulation.

The non-adiabatic character of f_{xc} is illustrated up to $\omega \approx 6$ a.u. in Fig. 5.11. Similarly to the atom, f_{xc} across the lowest energy transitions is adiabatic, until a striking frequency dependence is observed that persists and evolves. This behaviour is related to the *poles* that occur in f_{xc} infinitesimally below the real ω -axis [254, 269]. The eigenvalues of f_{xc} as a function of frequency are given in Fig. 5.12, and in particular, three divergences are shown (the singularities are tempered slightly by evaluating the response functions just above the real ω -axis). The lower panels of Fig. 5.12 demonstrate that these divergences coincide with a single eigenvalue of either the interacting or non-interacting response function crossing zero. Thus, the visual character of f_{xc} is dominated by the outer product of the eigenvector whose eigenvalue is either beginning to diverge, or recovering from a divergence.

In other words, the divergent f_{xc} around a singularity has the eigendecomposition

$$f_{\text{xc}}(\omega)|v_i(\omega)\rangle = \lambda_i(\omega)|v_i(\omega)\rangle, \quad (5.17)$$

⁶Since both χ and χ_0 share this eigenpair, a naive removal procedure can be used to circumvent the issues it would otherwise cause.

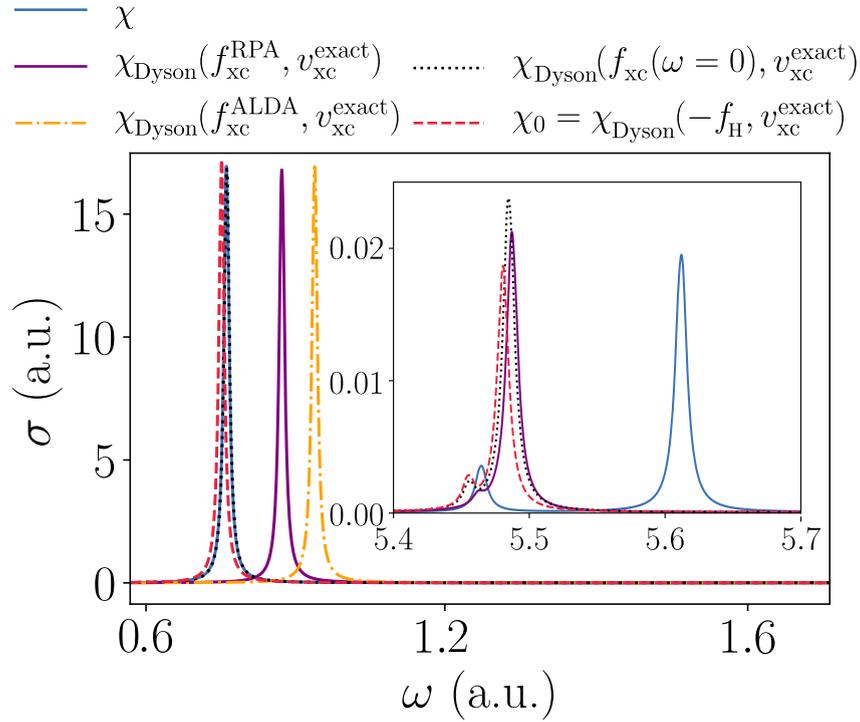


FIGURE 5.9: The optical absorption spectrum for the atom around the first transition, and around two higher energy transitions (inset), calculated at various levels of approximation. The optical spectrum calculated from the interacting and Kohn-Sham response functions are plotted alongside the optical spectrum computed from output of the Dyson equation using the RPA, adiabatic LDA, and AE xc kernels. The adiabatic LDA and RPA fail to accurately reproduce the interacting transitions, whereas the AE approximation successfully describes the low-energy transition, but fails at higher energies.

and, within the relevant frequency range, there is an eigenvalue that is larger in magnitude than the rest. Let us define this as eigenvalue as $|\lambda_0| \gg |\lambda_{i \neq 0}|$, where λ_0 in the upper panel of Fig. 5.12 is apparent. Across the frequency range for which this eigenpair is dominant, f_{xc} has its visual character determined by

$$f_{xc}(\omega) \approx \lambda_0(\omega) |v_0(\omega)\rangle \langle v_0(\omega)|. \quad (5.18)$$

The panels (b), (c), and (d) of Fig. 5.11 all exhibit this outer product structure. The response functions χ and χ_0 at these zeros (infinitesimally below the real ω -axis) are non-invertible, and hence f_{xc} *does not exist*, i.e. f_{xc} has an infinitely difficult job reconstructing χ from χ_0 in the direction corresponding to the eigenvalue that is zero, which can be gathered from inspection of the Dyson equation. The role of these singularities and divergences, i.e. the *analytic structure* of f_{xc} , is the subject of the remainder of this section.

It is argued, for example, in [269], that the singularity itself appears between excitations, and thus is unlikely to be of concern, except in the case where a singularity is condensed between two nearly degenerate excitations. This is contested here on three grounds: first, nothing prevents a zero in χ_0 from occurring at an excitation in χ (and vice-versa),

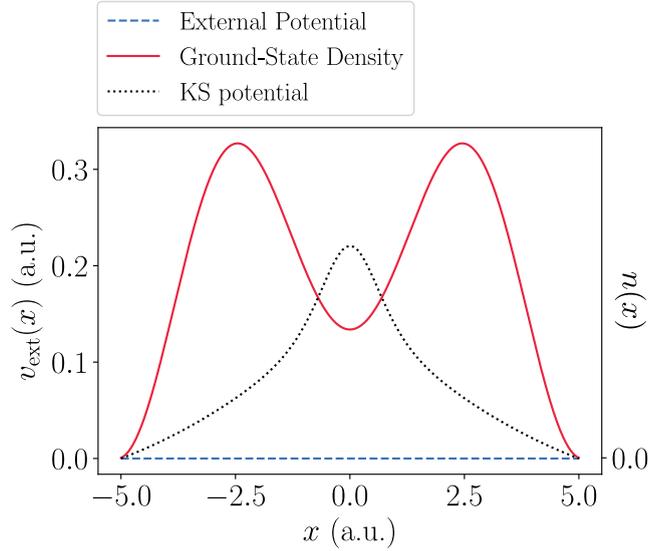


FIGURE 5.10: The ground-state density, external potential, and reverse-engineered Kohn-Sham potential for the infinite potential well. The external and Kohn-Sham potential have been shifted for illustrative purposes.

and while such an occurrence might seem improbable, it almost happens without design in the infinite potential well at $\omega = 1.025$ a.u., see Fig. 5.12. Second, the behaviour *around* the singularity in f_{xc} , whose source is the singularity, can be large in magnitude and is of importance, even if the singularity itself is not. Third, in a finite basis, it is possible *in principle* that a zero of χ can exist on a pole (excitation) in χ , see Section 5.2.3.

Since most f_{xc} approximations lack these divergences, one can question the importance of these divergences in relation to optical properties. In order to determine the impact of the divergences in f_{xc} on the optical spectrum, it is possible to examine how the optical spectrum is affected after projecting out the divergent eigendirection in f_{xc} , and otherwise keeping f_{xc} identical. This is tantamount to setting $\chi|v\rangle := \chi_0|v\rangle$ for some vector $|v\rangle$ within the Hilbert space that is associated with the divergence. Therefore, removal of a divergence is defined by applying the projection operator $P = |v_{div.}\rangle\langle v_{div.}|$ to f_{xc} , where $|v_{div.}\rangle$ is the eigenvector whose eigenvalue is divergent. The output of this procedure is termed the *projected xc kernel* $f_{xc}^{\text{projected}}$, see Fig. 5.13.

The interacting excitation at $\omega = 1.09$ a.u., see Fig. 5.12, is visible in the optical spectrum, and furthermore the character of f_{xc} at this energy, see Fig. 5.11, is dominated by the outer product of an eigenvector whose eigenvalue is much larger in magnitude than the rest, and between the two divergences at $\omega = 1.01, 1.17$ a.u. The removal of these divergences from f_{xc} across the energy range of interest, and subsequently solving the Dyson equation with the projected xc kernel $f_{xc}^{\text{projected}}$, shifts the interacting optical peak back toward the non-interacting peak, Fig. 5.14. Conversely, the much weaker divergence at $\omega = 0.99$ a.u., as seen in Fig. 5.12, has a tail that also yields an eigenvalue much larger in magnitude than the rest at the interacting excitation, and removal of this divergence produces no change in the optical spectrum, i.e. this eigenvector is *not* relevant for capturing the transition in question. In the inset of Fig. 5.14, the visual

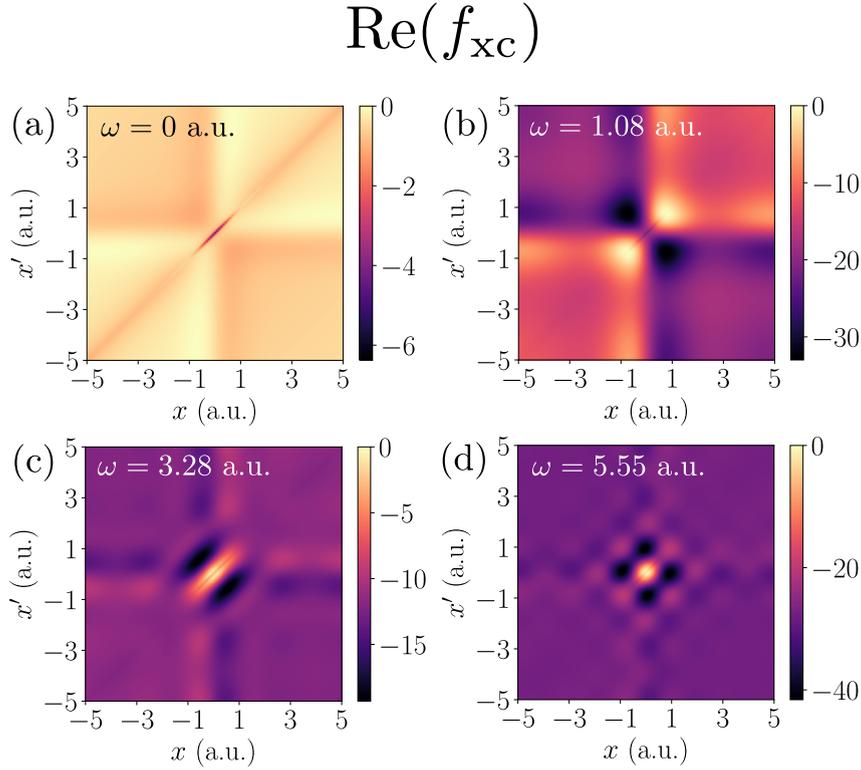


FIGURE 5.11: The real part of the exact numerical xc kernel $f_{\text{xc}}(x, x', \omega)$ for the infinite potential well at (a) $\omega = 0$ (AE), and at three higher frequencies that demonstrate its non-adiabatic character, (b) the sixth interacting excitation, $\omega = 1.09$ a.u., (c) $\omega = 3.28$ a.u., and (d) $\omega = 5.55$ a.u. The xc kernels are shifted such that their maximum value is zero, since there exists no long-range limit, see §5.1.6. A strong frequency dependence, i.e. departure from the adiabatic limit, is observed.

character of f_{xc} is shown at the interacting excitation $\omega = 1.09$ a.u. after the divergences visible in Fig. 5.12 have been removed. Underneath these divergences, f_{xc} is indistinguishable from the adiabatic f_{xc} , meaning the frequency dependence of f_{xc} in this system is largely due to its pole structure. These results suggest that functional approximations that do not capture the non-adiabatic pole structure of f_{xc} will struggle to improve matters beyond the non-interacting peaks for certain transitions.

A further point of note is that there are many more zeros, i.e. non- v -representable densities, in the interacting response function than the non-interacting response function, a simple counting argument to demonstrate this suffices. All N eigenvalues of the response functions begin negative [269], and each excitation brings a negative eigenvalue to a positive eigenvalue across a divergence, which otherwise evolves as a continuous function of ω . For two electrons discretised with a basis set of dimension N , there are $\frac{1}{2}N(N-1) - 1$ interacting excitations, and $2N - 4$ non-interacting excitations, the difference being made up of double (triple, etc. with more than two electrons) excitations that are notoriously not present in the Kohn-Sham response function [246]. Therefore, there are a great deal more eigenvalues that must pass through zero in the interacting response function, and moreover, these eigenvalues that cross zero are connected to the excitations that require them to do so – an account of the precise conditions under

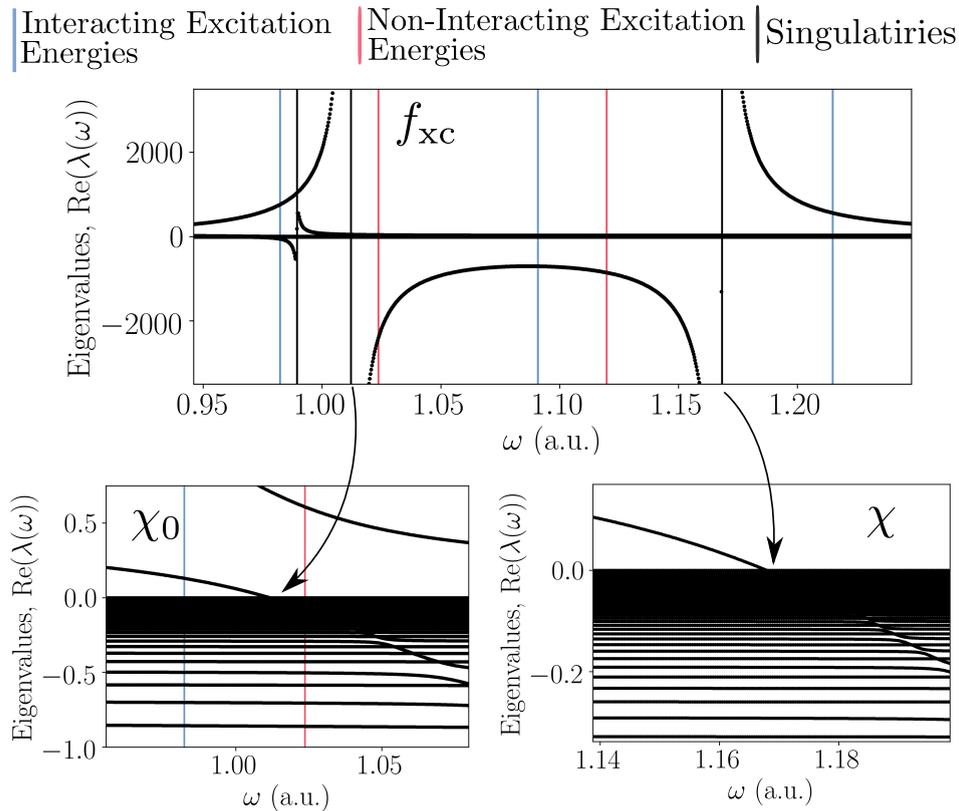


FIGURE 5.12: Eigenvalues as a function of frequency, labelled $\text{Re}(\lambda(\omega))$, of f_{xc} (upper), χ_0 (lower left), and χ (lower right). Within the frequency range shown, the predominant non-adiabatic behaviour of the f_{xc} is a result of three singularities and their surrounding divergences at $\omega = 0.99, 1.01, 1.17$ a.u., the source of the last two is observed to be due to an eigenvalue crossing zero in χ_0 and χ respectively.

which this occurs is given in Section 5.2.3 using a two-state model. In Fig. 5.12, there are three divergences, two of which are *paired*, meaning the zero in the non-interacting response function is a shifted version of the zero in the interacting response function; such behaviour can be related to single excitations.

There also exists an *unpaired* divergence at $\omega = 0.99$ a.u. related to the excitation at $\omega = 0.98$ a.u. which has *double excitation character*. That is, the overlap between the final state involved in this transition $|\Psi_0\rangle \rightarrow |\Psi_5\rangle$ and the Slater determinant $|\Phi_{(2,3)}\rangle$ is $\langle\Phi_{(2,3)}|\Psi_5\rangle = 0.98$. It is perhaps, then, no surprise that removal of eigenvector and eigenvalue whose source is the unpaired divergence did not alter the visible (single excitation) transition in Fig. 5.14. In fact, the pole in the interacting response function relating to the double excitation *disappears* with removal of the unpaired divergence, and so this divergence, and its surrounding character, *is* important in order to capture the transition for which it is relevant. The xc kernel exhibits unpaired divergences after all double excitations up to $\omega = 6$ a.u. at frequencies slightly higher than the interacting double excitation energy. Refs. [246, 248] derive, under certain assumptions, the necessarily divergent character of f_{xc} around a double excitation. The above analysis reveals that divergences in f_{xc} are, in fact, common and associated with the ω neighborhood containing multiple and single excitations alike.

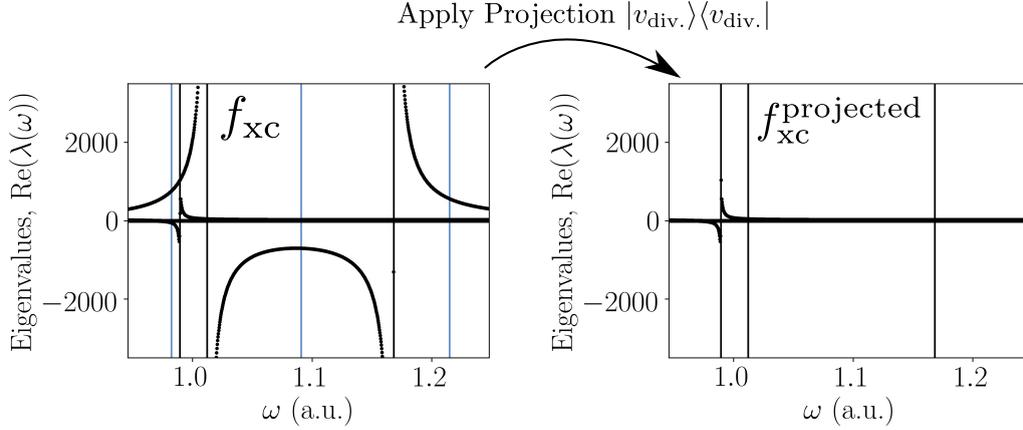


FIGURE 5.13: Eigenvalues of f_{xc} for the infinite potential well as a function of frequency, labelled $\text{Re}(\lambda(\omega))$ (left panel). Within the frequency range shown, the predominant non-adiabatic behaviour of f_{xc} is a result of three singularities (black lines) at $\omega = 0.99, 1.01, 1.17$ a.u. The right panel demonstrates the f_{xc} that has identical eigen-decomposition apart from the divergences related to the singularities at $\omega = 1.01, 1.17$ a.u. in the left panel.

5.2.3 Model Response Functions

A great deal of insight into the structure of χ and χ_0 can be gained with *model response functions*. In particular, the source of zero eigenvalues in χ and χ_0 is elucidated, and in doing so, the singularities in the xc kernel can be related to certain excitations. The conceptual origin of the model to follow is given in Ref. [215].

5.2.3.1 Scalar Model Response Function

Consider a system with two excitations at energies Ω_1 and Ω_2 . In the case $\eta = 0$, the response function χ has the form,

$$\chi(\omega) = \left(\frac{1}{\omega - \Omega_1} - \frac{1}{\omega + \Omega_1} \right) \langle \Psi_0 | \hat{n}(x) | \Psi_1 \rangle \langle \Psi_1 | \hat{n}(x') | \Psi_0 \rangle \quad (5.19)$$

$$+ \left(\frac{1}{\omega - \Omega_2} - \frac{1}{\omega + \Omega_2} \right) \langle \Psi_0 | \hat{n}(x) | \Psi_2 \rangle \langle \Psi_2 | \hat{n}(x') | \Psi_0 \rangle. \quad (5.20)$$

In order to examine the ω -dependence, the gross simplification is made that there is *no* spatial dependence, in which case,

$$\chi(\omega) = a \left(\frac{1}{\omega - \Omega_1} - \frac{1}{\omega + \Omega_1} \right) + a \left(\frac{1}{\omega - \Omega_2} - \frac{1}{\omega + \Omega_2} \right) \quad (5.21)$$

for some constant a . Since this is an algebraic system, the response function crosses zero at $\omega = \sqrt{\Omega_1 \Omega_2}$, which is the geometric mean of the two excitation energies, and thus always between the excitations. As we shall see, this conclusion generalises remarkably well, and similar reasoning allows us to relate zeros of the response function to certain excitations. (In the above example, the zero at $\omega = \sqrt{\Omega_1 \Omega_2}$ is related to excitations one

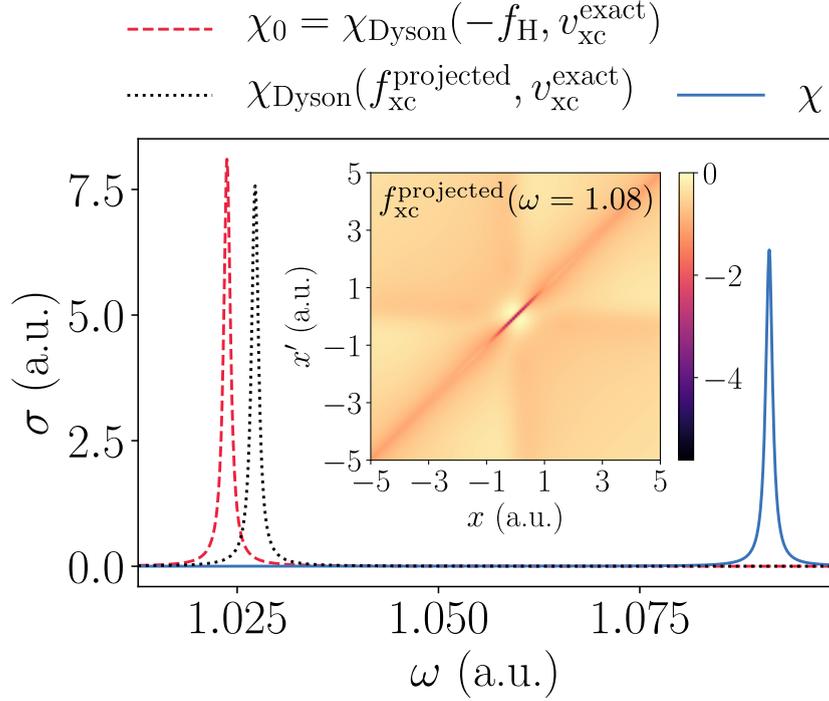


FIGURE 5.14: The optical absorption spectrum for the infinite potential well around a visible interacting excitation at $\omega = 1.09$ a.u., and its corresponding non-interacting excitation at $\omega = 1.02$ a.u. The projected xc kernel (inset), i.e. f_{xc} with its divergences removed, is indistinguishable from the adiabatic xc kernel, see Fig. 5.11(a), and the optical peak associated with $f_{xc}^{\text{projected}}$ is only slightly improved upon the non-interacting peak.

and two).

5.2.3.2 Two-State Model Response Function

The above scalar model is now extended toward more realistic scenarios. Let us first reintroduce spatial dependence, and define the *excitation function*, $f_n(x) = \langle \Psi_0 | \hat{n}(x) | \Psi_n \rangle$, which, in bra-ket notation, is denoted $|f_n\rangle$. The (now exact) response function takes the form

$$\chi(x, x', \omega) = \sum_{i=1}^{\text{num. ex.}} |f_i\rangle \langle f_i| \left(\frac{1}{\omega - \Omega_i} - \frac{1}{\omega + \Omega_i} \right). \quad (5.22)$$

A **two-state model** is established by supposing that there are two excitations with energies Ω_n and Ω_m whose excitation functions are orthogonal to the rest, $\langle f_n | f_i \rangle = 0$ and $\langle f_m | f_i \rangle = 0$ for $i \neq n, m$. The (normalised) excitation functions are given some finite overlap

$$\langle f_n | f_m \rangle = \alpha \quad (5.23)$$

parameterised by $\alpha \in [0, 1]$. Each excitation contributes a term to the response function of the form

$$\chi(\omega) = .. + \frac{|f_n\rangle\langle f_n|}{\omega - \Omega_n} - \frac{|f_n\rangle\langle f_n|}{\omega + \Omega_n} + .. \quad (5.24)$$

The eigenvalues of χ relating to the 2×2 subspace that has been isolated are

$$\lambda_{\pm}(\omega) = x(\omega) + y(\omega)(1 - 2\alpha + 2\alpha^2) \pm \sqrt{(-x - y(1 - 2\alpha + 2\alpha^2))^2 - 4xy(1 - 2\alpha + \alpha^2)}, \quad (5.25)$$

where

$$x(\omega) := \frac{1}{\omega - \Omega_n} - \frac{1}{\omega + \Omega_n} \quad (5.26)$$

$$y(\omega) := \frac{1}{\omega - \Omega_m} - \frac{1}{\omega + \Omega_m}. \quad (5.27)$$

Setting Eq. (5.25) equal to zero gives

$$x(\omega)y(\omega)(1 - \alpha)(1 + \alpha) = 0, \quad (5.28)$$

which reveals that this two-state model crosses zero only when the response functions are parallel, i.e. $\alpha = 1$, since $x(\omega)$ and $y(\omega)$ do not individually cross zero. The frequency at which the zero occurs is equal to the geometric mean of the two excitation energies, $\omega = \sqrt{\Omega_n\Omega_m}$, see left-hand panel of Fig. 5.15. This corresponds to parallel excitation functions, and it turns out the origin and position of the zero is precisely the same as in the scalar model above, i.e. the limit $\alpha = 1$ reduces to the scalar model.

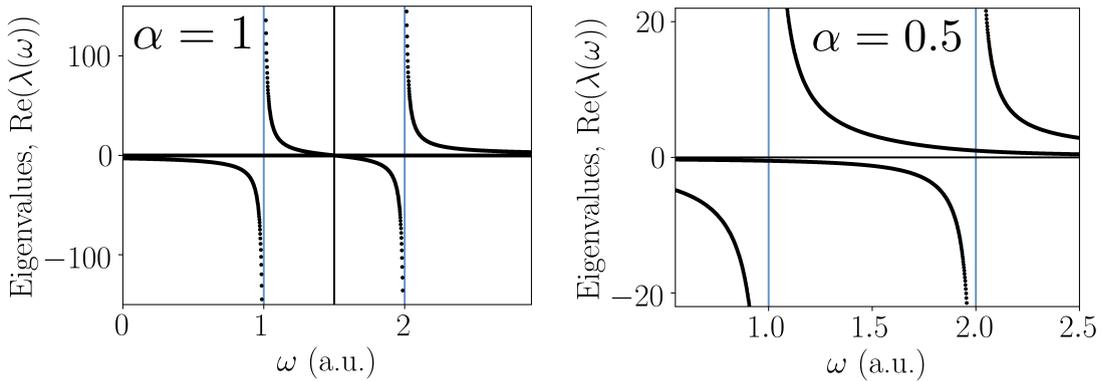


FIGURE 5.15: Eigenvalues of the two-state model χ with excitation energies $\Omega_1 = 1$ a.u. and $\Omega_2 = 2$ a.u. In the left-hand panel, the excitation functions, $|f_1\rangle$ and $|f_2\rangle$, are parallel, in which case, a zero (black line) exists at the geometric mean of the excitation energies (blue lines), $\omega = \sqrt{2}$ a.u. In the right-hand panel, the excitation functions are given overlap $\alpha = 0.5$, which results in an avoided crossing about the ω -axis, and thus an avoided zero.

If the excitation functions have some non-integer overlap $\alpha \neq 1$, which is generally the case in practice, the eigenvalues exhibit an *avoided crossing* about the ω -axis, and thus

never cross zero, see right-hand panel of Fig. 5.15. The eigenvalue spectrum of the response functions exhibit a multitude of avoided crossing-type features, see Fig. 5.1, the source of which is now understood, namely, excitation functions with some non-integer overlap. However, in practice the avoided crossings appear *below* the ω -axis, thus permitting a zero eigenvalue where it is currently forbidden in the model.

To account for this, suppose there exist higher energy excitations $\Omega_i \gg \Omega_{n/m}$ with excitation functions $|f_i\rangle$ that overlap with the excitation functions of interest, $|f_n\rangle$ and $|f_m\rangle$. Since these excitations are much higher in energy, their contribution to the response function χ around the n^{th} and m^{th} excitation is given by subtracting a constant c to all components of the 2×2 subspace (derivation omitted). As Fig. 5.16 demonstrates, this modification of the two-state model, in part, leads to a negative shift in the eigenvalues of Eq. (5.25), which in turn allows χ to cross zero and exhibit an avoided crossing *below* the ω -axis – this is the behaviour observed in the eigenvalues of the exact response function. Therefore, zeros in either the exact or Kohn-Sham response function occur between two isolated excitations whose excitation functions overlap, and this is the sense in which a zero can be related to certain excitations. The distinction made in Section 5.2.2 between *paired* and *unpaired* singularities in f_{xc} , i.e. singularities relating to single- and multi-excitations respectively, should now come as no surprise.

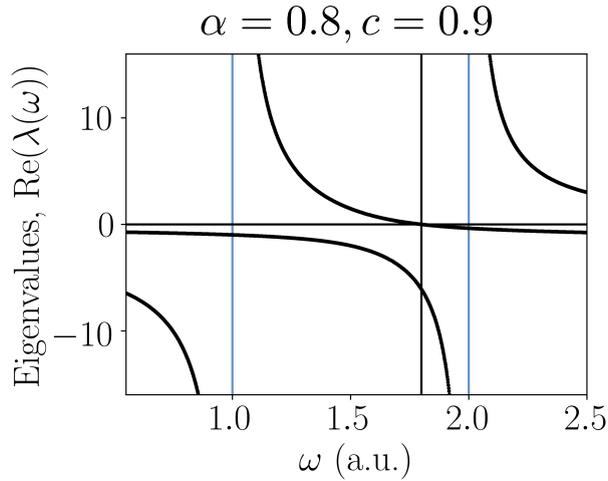


FIGURE 5.16: Eigenvalues of the two-state model χ from Fig. 5.15 with additional excitations at much higher energy, parameterised with $c = 0.9$ and $\alpha = 0.8$. The excitations of higher energy serve to shift the eigenvalues down, allowing a zero (black line) in χ (singularity in f_{xc}) to occur where it was otherwise prohibited in the right-hand panel of Fig. 5.15.

To conclude this section, it is possible to demonstrate that a zero of χ , and thus a singularity in f_{xc} , can occur at the *same* energy as an interacting excitation (pole of χ). This can be realised within the $\alpha = 1$ two-state model by including an additional excitation with two properties: 1) its energy is equal to the geometric mean of the original excitation energies, e.g. those in Fig. 5.15 (left-hand panel), 2) its excitation

function is orthogonal to the original two⁷. Having demonstrated yet another channel through which singularities in f_{xc} can occur *on top of* excitations, the importance of such singularities for optical properties is further established.

5.2.4 Quantum Harmonic Oscillator

The quantum harmonic oscillator is defined in the domain $[-8, 8]$ a.u. with the potential $v_{\text{ext}}(x) = \frac{1}{2}\nu^2 x^2$ where $\nu := 0.45$ a.u., see Fig. 5.17. The eigenspace truncation parameter is $\bar{\lambda} = 10^{-9}$ and the real-space truncation parameter is $b = 6.4$ a.u. First, the spatial

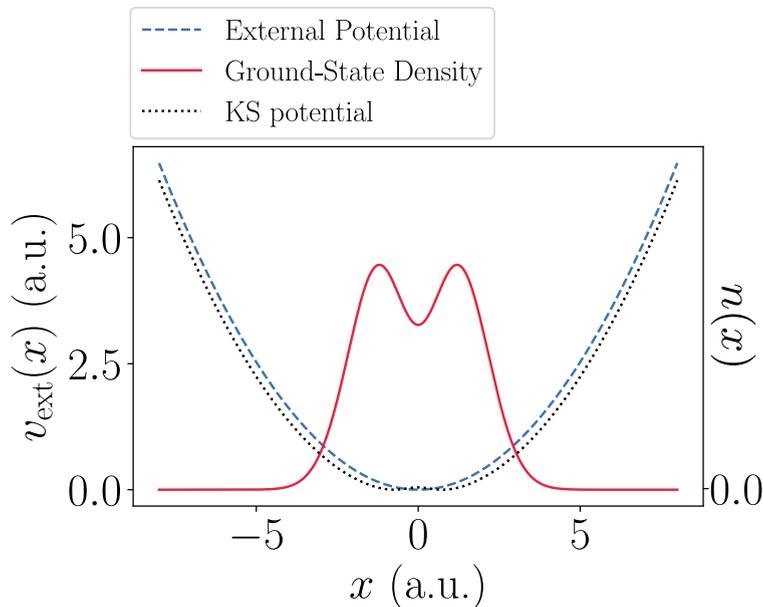


FIGURE 5.17: The ground-state density, external potential, and reverse-engineered Kohn-Sham potential for the quantum harmonic oscillator. The external and Kohn-Sham potential have been shifted for illustrative purposes.

structure, and in particular the *long-range behaviour* [239], of the exact f_{xc} is examined. The delicate nature of the numerics involved in computing the exact f_{xc} is brought to the fore when attempting to capture its long-range limit. The ill-conditioning in the response matrices can be identified with regions of nearly vanishing ground-state density, and it is precisely in these regions that we expect to observe the long-range character of f_{xc} . However, f_{xc} , when evaluated outside some central region where we can be confident there is little-to-no numerical error, diverges in a manner that is not consistent with the known long-range limit of f_{xc} . This spurious divergence in f_{xc} does *not* much affect the accuracy of the output of the Dyson equation, χ_{Dyson} , because the Dyson equation Eq. (5.12) involves the matrix product $\chi_0 f_{xc}$, and the divergent regions of f_{xc} operate on the nearly vanishing regions of χ_0 . If we are to observe the long-range limit of f_{xc} in the present context, we must rely on fortuitous circumstances, namely, that the region where the numerical error is low overlaps with the region where the long-range limit is

⁷The orthogonality of the excitation function implies that the singularity in f_{xc} at the excitation is not relevant for capturing the excitation. However, this example serves to demonstrate a more general feature, the rarity of which in practice is unknown.

observed. This is the case for the AE f_{xc} , which is shown, together with slices of f_{xc} along a particular axis, in Fig. 5.18 and Fig. 5.19.

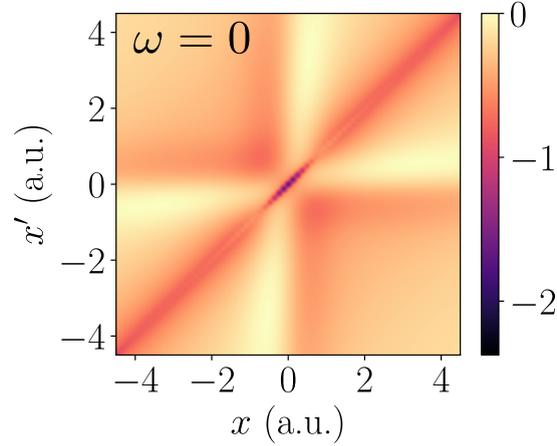


FIGURE 5.18: The AE xc kernel $f_{xc}(x, x', \omega = 0)$ for the quantum harmonic oscillator constructed with the eigenspace truncation method. The precise spatial structure exhibited by the AE kernel, including its long-range limit and non-local character, is discussed in the main text.

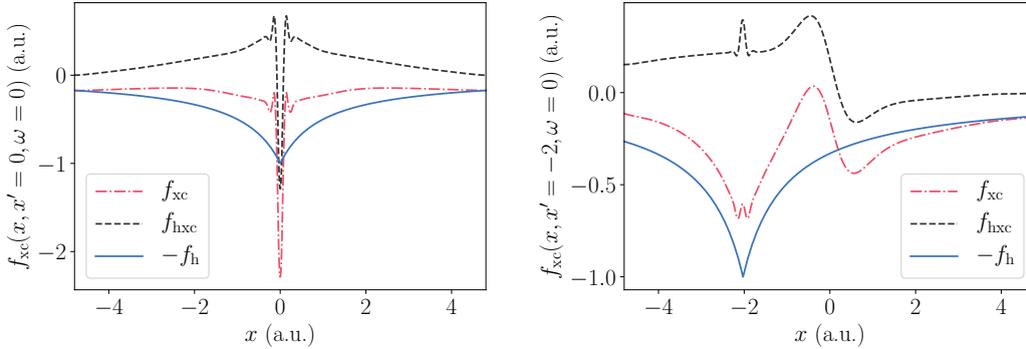


FIGURE 5.19: The AE xc kernel $f_{xc}(x, x', \omega = 0)$ for the quantum harmonic oscillator along $x' = 0$ a.u. (top) and $x' = -2$ a.u. (bottom). The xc kernel, Hxc kernel, and Hartree kernel, are shown, and the long-range limit $f_{xc} \rightarrow -f_H$ is observed.

The centre-most region of the domain is where exchange and correlation effects are most important, and in this region f_{xc} exhibits a strong *local* response that quickly decays as $x \neq x'$. Since this region can be interpreted as the most crucial for recovering accurate observable properties from the Dyson equation, this observation supports, at least in part, local approximations to f_{xc} , such as the adiabatic LDA. The non-local structure of f_{xc} at $\omega = 0$ a.u. can be seen most clearly in the right panel of Fig. 5.19, where perturbations in the density *outside* the centre-most region cause a significant change in the exchange-correlation potential *inside* the centre-most region. The failure of the adiabatic LDA to capture this non-local response leads to fairly poor agreement between

the exact and approximate transition energies, which is seen for the atom in Fig. 5.9, and illustrated for the quantum harmonic oscillator below⁸.

The long-range limit $f_{xc}(x, x', \omega = 0) \rightarrow -f_H(|x - x'|)$ is observed along both $x' = 0$ a.u. and $x' = -2$ a.u. in Fig. 5.19. Due to the rapid decay of the ground-state density in the quantum harmonic oscillator, convergence of the f_{xc} toward $-f_H$ along $x = -2$ a.u. is not directly observed in the $x \rightarrow -\infty$ limit – the atom, whose ground-state density does not decay as quickly, converges toward $-f_H$ in both the positive and negative limits, see Fig. 5.20. It is known that the long-range limit of f_{xc} , in both finite and periodic systems, satisfies $f_{xc}(\omega) \rightarrow -\alpha(\omega)f_H$ [168, 237, 239, 271], where $\alpha(\omega)$ is a frequency-dependent constant that reflects dielectric screening in the system. Therefore, one expects $\alpha = 1$ in finite systems, and at lower frequencies, where the numerical methodology provides a robust long-range limit, the observations are consistent with this, see Fig. 5.19.

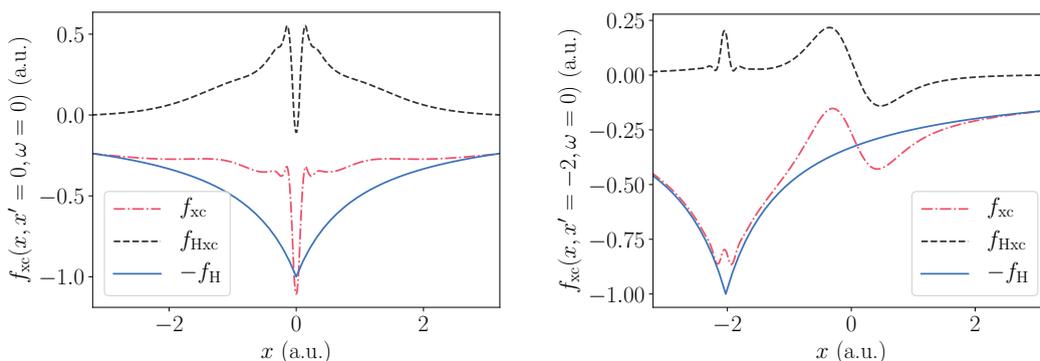


FIGURE 5.20: The AE xc kernel $f_{xc}(x, x', \omega = 0)$ for the atom along $x' = 0$ a.u. (top) and $x' = -2$ a.u. (bottom). The xc kernel (red dot-dash), the Hxc kernel (black dash), and the Hartree kernel (blue solid) are shown, and the long-range limit $f_{xc} \rightarrow -f_H$ is observed. Accordance with this limit defines the procedure used to fix the constant shift arbitrariness in f_{xc} .

To conclude matters for the quantum harmonic oscillator, we examine its optical absorption spectrum, and in particular highlight a failing of the f_{xc} approximations considered in this work when compared to the exact and AE xc kernels. The optical absorption spectrum, using the same range of approximations discussed in the case of the atom, is shown in Fig. 5.21. In the *non-interacting* quantum harmonic oscillator, all transitions but the first are disallowed, otherwise known as its *selection rules*, which can be derived using the harmonic oscillator ladder operators. The inclusion of the Coulomb interaction lifts these special symmetries of the non-interacting quantum harmonic oscillator, but not enough for the previously disallowed transitions to be observed in the optical spectrum – the dipole matrix elements for these *allowed* transitions are $\mathcal{O}(10^{-8})$.

On the other hand, the exact Kohn-Sham potential differs meaningfully from the harmonic oscillator form, which creates a series of visible peaks beyond the first in the optical spectrum; the transition rates for these transitions are vastly overestimated. In

⁸It is possible that the particular form of the non-local f_{xc} functionals considered in [89, 278], in which $f_{xc}(x, x') = f_{xc}[(n(x) + n(x'))/2]$, can capture certain features of the adiabatic non-local response depicted in Fig. 5.18 and Fig. 5.19.

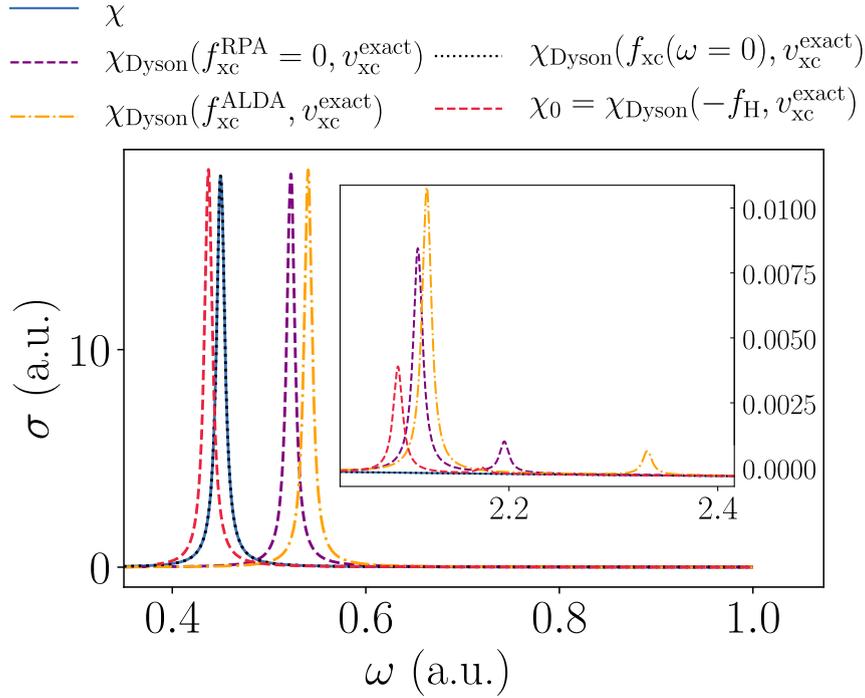


FIGURE 5.21: The optical absorption spectra for the quantum harmonic oscillator around the first transition, and around a chosen higher energy transition (inset), calculated at various levels of approximation. The optical spectrum calculated from the interacting and Kohn-Sham response functions are plotted alongside the optical spectrum computed from output of the Dyson equation using the RPA, ALDA, and AE xc kernels. At higher frequencies (inset), all optical peaks are suppressed for the quantum harmonic oscillator, a state of affairs that the AE f_{xc} reproduces, but the RPA and adiabatic LDA does not.

this situation, the RPA xc kernel and adiabatic LDA xc kernel do not manifest the required suppression of the optical peaks. Interestingly, the AE kernel, as seen in Fig. 5.18, is able to reproduce the exact optical spectrum across the entire frequency range considered $\omega \in [0, 6]$ a.u., perhaps due to the correct excitation functions used in its construction. This suggests that, in cases where the exact system possesses heavily suppressed transitions, perhaps due to symmetries that are not shared by the Kohn-Sham system, the typical f_{xc} approximations are insufficient to recover the exact state of affairs, but improvement of their non-local spatial structure toward the AE kernel can assist matters.

5.2.5 Gauge Freedom

Recall the following transformation to f_{xc} leaves the response function output from the Dyson equation, χ_{Dyson} , unchanged,

$$f_{xc}(x, x', \omega) \rightarrow f_{xc}(x, x', \omega) + g(x, \omega) + h(x', \omega) + c(\omega),$$

which is the familiar definition of the term *gauge*. If two xc kernels (approximate or exact) exist within this family, then the derived properties these xc kernels, such as χ_{Dyson}

and its optical spectrum, are identical. Therefore, two xc kernels that differ approximately by a gauge transformation provides an explanation for the approximate agreement between the derived properties of the two xc kernels in question. The emergence of such a phenomenon in practice is now considered, and in fact we shall demonstrate that the gauge freedom of f_{xc} is *not* sufficient to explain the similarity observed between, for example, the optical properties calculated using the adiabatic LDA and RPA in Section 5.2.1. Moreover, the particular form of the non-local spatial structure within the exact $f_{xc}(x, x', \omega)$ is in general not possible to capture with this gauge freedom, and it is unlikely that this line of reasoning is able to explain the efficacy of approximations of any kind.

Consider a system for which we have defined two xc kernels, $f_{xc}^A(\omega)$ and $f_{xc}^B(\omega)$ – one of these may be the exact xc kernel. The *optimal gauge* is defined as the gauge that brings f_{xc}^A as close as possible to f_{xc}^B in the *Frobenius norm*⁹,

$$\min_{g,h,c} \|M(g, h, c)\|_F := \min_{g,h,c} \|f_{xc}^A - f_{xc}^B - g - h - c\|_F$$

at a given ω . In a finite basis, the matrix M has entries

$$M_{ij} = f_{xc,ij}^A - f_{xc,ij}^B - h_i - g_j - c,$$

which makes it apparent that the the gauge specifies $2N + 1$ degrees of freedom out of a total N^2 . The minimisation problem above admits a *closed-form solution* for these degrees of freedom. In order to derive this solution, first note that the optimal shift c can be absorbed into the definition of g and h . Furthermore, one can observe that, since the xc kernels are always symmetric $f_{xc}(x, x') = f_{xc}(x', x)$, it will never be favourable, from the perspective of the Frobenius norm error, to break this symmetry with the gauge. Thus, we can focus on solving the optimisation problem with the matrix M now having components

$$M_{ij} = f_{xc,ij}^A - f_{xc,ij}^B - \tilde{g}_i - \tilde{g}_j,$$

where \tilde{g} is the function g with the constant c absorbed. The solution of this (convex) optimisation problem is given by

$$\tilde{g}_i = \frac{1}{N} \sum_j (f_{xc,ij}^A - f_{xc,ij}^B) - \frac{1}{2N^2} \sum_{ij} (f_{xc,ij}^A - f_{xc,ij}^B), \quad (5.29)$$

where N is the number of grid points¹⁰.

The atom of Section 5.2.1 is considered, and in particular the optimal gauge is found in order to match the RPA, $f_{xc}^{\text{RPA}} = 0$, with the adiabatic LDA used in this work, f_{xc}^{ALDA} [277], see Fig. 5.22. It can be seen that the vectors g and h introduce an unavoidable degree of non-locality, and the local spatial structure of the adiabatic LDA simply cannot be reproduced with a gauge transform of this kind applied to the RPA.

The story remains much the same when an attempt is made to match the adiabatic LDA xc kernel to the AE xc kernel for the atom, $f_{xc}(x, x', \omega = 0)$, see Fig. 5.23. The

⁹The Frobenius norm of a matrix X defined as $\|X\|_F = \sum_{ij} |X_{ij}|^2$.

¹⁰To derive this, set the derivative of the objective function to zero.

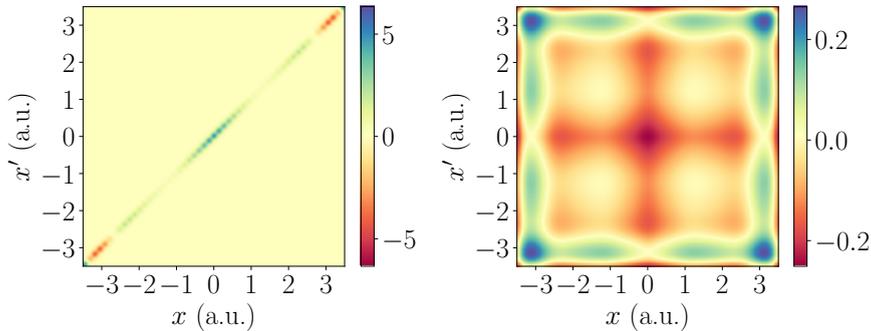


FIGURE 5.22: The adiabatic LDA xc kernel f_{xc}^{ALDA} for the atomic system (left) remains structurally dissimilar to the RPA xc kernel with the optimal gauge transform applied, $f_{xc}^{\text{RPA}} - g - h - c$ (right).

particular form of the spatial non-locality in the AE xc kernel does not lend itself to the fairly inflexible structural freedom afforded by the functions g and h . This conclusion holds true for *almost* all the xc kernels examined in this work, namely, the spatial profile of the exact f_{xc} at any ω is not possible to reproduce with a gauge transform applied to the adiabatic LDA or RPA, despite, at high frequencies, the similarity of the derived optical spectra from these approximations.

An interesting exception to these conclusions is given around the third excitation in the atom. The onset of a divergence occurs before the third visible excitation at $\omega = 2.35$ a.u., and yet the AE approximation captures this excitation well, see Fig. 5.24. The left-hand panel of Fig. 5.25 demonstrates f_{xc} at the third excitation, which has the structure of an outer product $|u\rangle\langle u|$, as expected around a divergence. The xc kernel is therefore mostly composed of N degrees of freedom, a state of affairs that is *a priori* much more amenable to the gauge freedom. The optimal gauge, Eq. (5.29), to transform the AE xc kernel, $f_{xc}^A = f_{xc}(x, x', \omega = 0)$, into the exact xc kernel at the third excitation, $f_{xc}^B = f_{xc}(x, x', \omega = 2.35)$, is shown in the right-hand panel of Fig. 5.25. Indeed, there is a striking resemblance between the two, meaning the AE f_{xc} and the diverging f_{xc} around the third excitation exist within the same family of xc kernels defined by the gauge freedom. This divergence therefore *does not* affect observable properties such as its associated optical peak, however, an occurrence of this kind is not detected again for the remainder of the divergences. Such an observation could account for the performance of certain adiabatic kernels beyond their expected domain of applicability.

5.2.6 Exchange-Correlation Potential Versus Exchange-Correlation Kernel

In finite systems, it is the conventional wisdom that an accurate ground-state xc potential is *more* important for capturing the interacting excitation spectrum than a sophisticated f_{xc} [237, 239, 279–281]. That is, the optical spectrum computed from some non-interacting response function χ_0 benefits much more from improvements in the ground-state wavefunctions and energies that are used to construct χ_0 than improvements in f_{xc} that is used, in conjunction with χ_0 , to solve the Dyson equation. The effect of an improved treatment of exchange and correlation in the ground-state is seen in Fig.

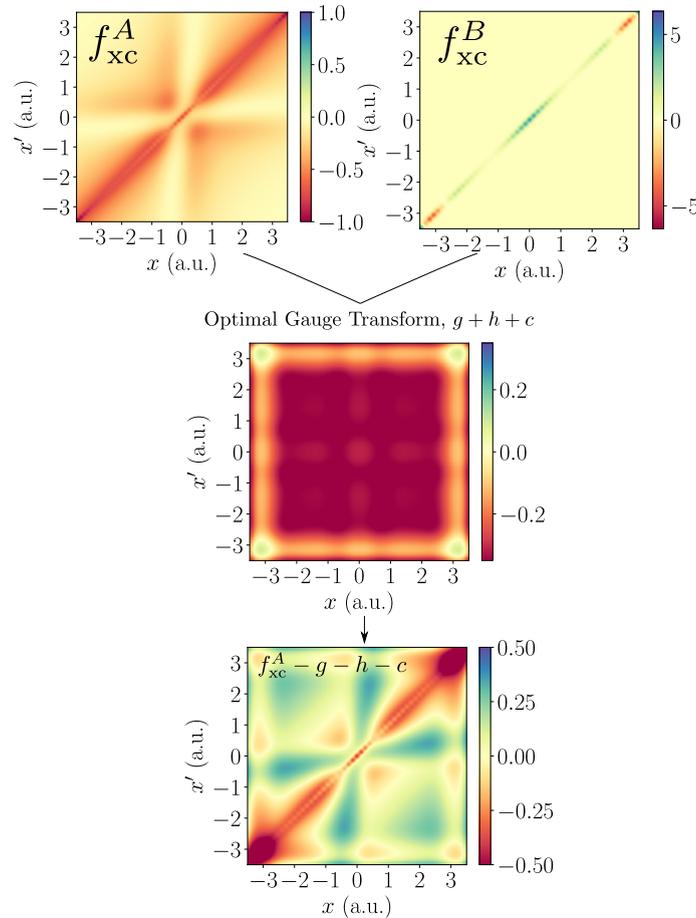


FIGURE 5.23: An attempt is made to transform the AE f_{xc} (upper left) toward the adiabatic LDA xc kernel (upper right) using the optimal gauge (middle). The f_{xc} that is closest in gauge to the adiabatic LDA (bottom) bears no resemblance in structure or magnitude to the adiabatic LDA.

5.26, which demonstrates the optical spectrum for the atomic system calculated from the *non-interacting response function* χ_0 at various levels of approximation. The transition energies and rates calculated from the exact Kohn-Sham system are considerably closer to the exact transitions than, for example, Hartree theory, which is increasingly poor at higher energies.

These non-interacting response functions can be used, in conjunction with their *corresponding* f_{xc} , to solve the Dyson equation, thus shifting the position and weight of the peaks. An xc kernel functional *corresponds to* an xc potential functional if it is the second functional derivative of the xc potential with respect to the density. Thus, the RPA xc kernel corresponds to Hartree theory, and the adiabatic LDA xc kernel corresponds to the ground-state LDA from which it came. To use an xc kernel that does not correspond to the ground-state xc potential can violate various exact conditions [282] – this is evidently the case here for the zero-force sum rule. The aforementioned conventional wisdom is exhibited clearly in Fig. 5.26, namely, use of a corresponding f_{xc} is only able to improve matters slightly beyond the non-interacting peaks, which is most visible at higher frequencies. For this reason, even when using an incompatible f_{xc} , the calculation

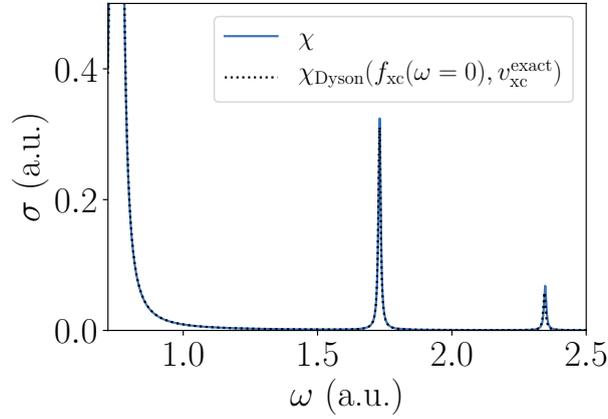


FIGURE 5.24: The optical spectrum for the atom is computed using the interacting response function χ (blue solid), and the output of the Dyson equation χ_{Dyson} with the AE xc kernel $f_{\text{xc}}(x, x', \omega = 0)$ (black dot). The optical spectrum is reproduced accurately for the excitations shown, despite significant non-adiabatic behaviour in $f_{\text{xc}}(\omega)$ beyond $\omega \approx 2$ a.u. The accuracy of the AE approximation is explained with the gauge freedom inherent to f_{xc} .

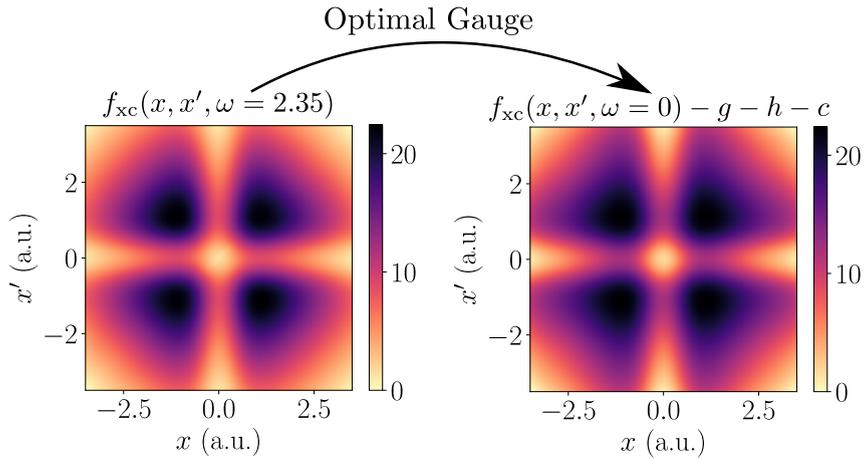


FIGURE 5.25: The exact f_{xc} of the atom around the third excitation, $f_{\text{xc}}(x, x', \omega = 2.35)$ (left panel), is the target of the optimal gauge applied to the AE approximation, $f_{\text{xc}}(x, x', \omega = 0)$ (right panel). The structure of these is similar, meaning the divergence in f_{xc} (left) is captured excellently with the optimal gauge (right).

benefits from an improved treatment of ground-state exchange and correlation. This is evidenced in the case of the atom in §5.2, where using the RPA and adiabatic LDA xc kernels in conjunction with the exact Kohn-Sham ground-state yields a more accurate optical spectrum than if one were to attempt to keep the xc potential and xc kernel compatible.

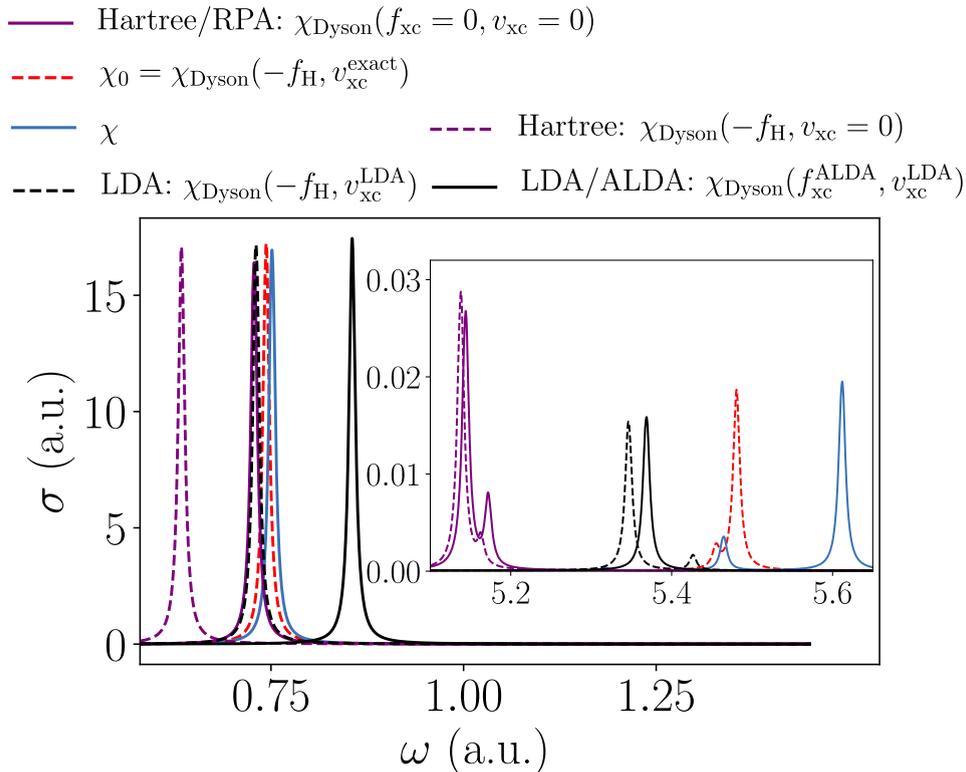


FIGURE 5.26: The exact optical absorption spectrum for the atomic system illustrated alongside the optical absorption spectra from the non-interacting response functions χ_0 calculated with wavefunctions and energies given by Hartree theory, an LDA, and the exact Kohn-Sham potential. The corrected transitions, given upon solution of the Dyson equation with the corresponding f_{xc} , are also shown for Hartree theory/RPA, and LDA/adiabatic LDA. The most accurate optical spectrum is that of the non-interacting exact Kohn-Sham system, and hence improving treatment of ground-state exchange-correlation toward this ideal is found to be most important. In particular, use of f_{xc} to shift the non-interacting peaks toward the interacting peaks is, in general, unable to achieve accuracy improvements comparable to those observed from improving ground-state exchange and correlation – this is seen most clearly at higher frequencies (inset).

5.3 Conclusions

A set of numerical methods is presented to calculate the exact f_{xc} kernel of time-dependent DFT in model finite one-dimensional systems is given, an object whose naive construction succumbs to numerical ill-conditioning. This methodology is then applied to examine the spatial *and* frequency dependence of the exact f_{xc} for the following two-electron systems: the atom, the infinite potential well, and the quantum harmonic oscillator.

The xc kernel exhibits a significant non-local spatial structure at all frequencies, including at $\omega = 0$, i.e. the AE f_{xc} . In lacking this structure, local approximations to f_{xc} are found to be insufficient for recovering the lowest energy excitations, whereas the AE xc kernel performs well. However, beyond the lowest few excitations, all the approximations considered here – the AE xc kernel, the RPA, and the adiabatic LDA – are equally poor, and do not generally improve the optical spectrum obtained directly

from the non-interacting Kohn-Sham response function χ_0 (that is, setting f_{Hxc} to zero everywhere). A notable exception is the quantum harmonic oscillator, whose optical transitions beyond the first are heavily suppressed, a feature that the AE xc kernel is able to capture. In general, improvement of the spatial structure of adiabatic xc kernel approximations toward the AE xc kernel is expected to assist matters for the lowest energy transitions, but beyond these transitions the lack of frequency dependence hinders all adiabatic xc kernels. In addition, the long-range limit of f_{xc} for *finite systems* $f_{\text{xc}} \rightarrow -f_{\text{H}}$ is confirmed, although the long-range character of f_{xc} is demonstrated to be unimportant in the present context, in contrast to its character within and around the region of high density.

Drastic non-adiabatic behavior is observed in f_{xc} for all systems studied in this work, and is, to a considerable extent, attributable to specific aspects of its analytic structure as a function of ω . (Simple) poles in f_{xc} , related to certain interacting or non-interacting transitions that necessitate them, can in practice appear close to interacting excitations, for example, between two nearly degenerate charge-transfer excitations in a double-well system. It is possible that a gauge transform can remove the divergence in specific cases without affecting the optical spectrum, but this is the exception rather than the rule. If f_{xc} is kept identical apart from removal of the diverging eigenvalue and its associated eigenvector, then f_{xc} can be rendered unable to capture the relevant transition. This suggests that an f_{xc} approximation that does not attempt to exhibit the non-adiabatic pole structure of the exact f_{xc} cannot reproduce transitions with energies higher than the first few excitations. This is the case for single, double, triple, and so on, excitations alike. The fact that these divergences can be related to certain excitations that necessitate them provides a new perspective on the divergent character of f_{xc} that is known to exist around double excitations [246, 263].

In general, the subtle spatial structure of the exact f_{xc} cannot be captured by applying a gauge transformation to one of the usual kernel approximations. However, the divergent ω -dependence discussed in the previous paragraph is more amenable to the gauge freedom. Indeed, in the case of the atom, the first divergence (around the third peak in the optical spectrum) turns out to be related to the AE f_{xc} with a gauge transform. Hence, the AE approximation is able to describe the third peak in the optical spectrum, despite the significant frequency dependence of f_{xc} around this peak.

As noted earlier in this section, the simple non-interacting kernel $f_{\text{xc}} = -f_{\text{H}}$ often yields surprisingly good spectra, provided that the exact Kohn-Sham potential, or a good approximation to it, is used to calculate χ_0 . This is in part due to the fact that the exact Kohn-Sham transitions are, in certain circumstances, good approximations to the interacting transitions [281]. By extension, in practical calculations, effort may be usefully devoted to improving approximations used for f_{xc} and v_{xc} individually, without an overriding need to maintain one as the functional derivative of the other. The quality of v_{xc} is of particular importance, as previous authors have observed in specific cases [237, 239, 279–281].

Chapter 6

Total Energies from the Adiabatic-Connection Fluctuation-Dissipation Theorem

The *adiabatic-connection fluctuation-dissipation theorem* (ACFDT) [7, 18, 89, 90, 283, 284] formalism is used to calculate ground-state energies of molecular and solid-state systems. The ACFDT framework is often presented as an expression for the correlation energy to be evaluated *after* performing some ground-state calculation of the practitioners choosing. As alluded to in Chapter 2, we shall instead opt to the view the ACFDT framework as a total energy functional in the Kohn-Sham sense. That is, the functional

$$E^{\text{ACFD}}[n] = T_0[n] + E_{\text{H}}[n] + E_{\text{ext}}[n] + E_{\text{xc}}^{\text{ACFD}}[n], \quad (6.1)$$

where

$$\begin{aligned} E_{\text{x}}^{\text{ACFD}}[n] &= E_{\text{x}}\{|\phi_i[n]\rangle\}, \\ E_{\text{c}}^{\text{ACFD}}[n] &= -\frac{1}{2\pi} \int_0^1 d\lambda \int_0^\infty d\omega \iint dx dx' v_{\text{ee}}(x, x') [\chi^\lambda[n](x, x', i\omega) - \chi_0[n](x, x', i\omega)], \end{aligned} \quad (6.2)$$

constitutes an exact reformulation of conventional ground-state Kohn-Sham DFT, meaning it has the correct minimum, i.e. the exact ground-state energy, and this minimum is attained at the interacting ground-state density. (Recall that $E_{\text{x}}\{|\phi_i[n]\rangle\}$ is the exact-exchange functional evaluated at the Kohn-Sham orbitals.)

The general approach taken in this chapter (which covers the work in Ref. [28]) is to utilise the techniques developed in the previous chapters in order to obtain the exact λ -dependent xc kernel $f_{\text{xc}}^\lambda[n](x, x', \omega)$ and ground-state xc potential $v_{\text{xc}}[n](x)$ with intent to carefully examine the relationship between approximate kernels/potentials and their corresponding inexact ACFDT energies, indicating routes toward more accurate practical versions of the ACFDT scheme.

The ACFDT total energy method occupies the ‘fifth rung’ on ‘Jacob’s ladder of approximate density functionals’ [285, 286] – above local and semi-local approximations to the xc energy $E_{\text{xc}}[n]$, since ACFDT-based calculations involve a full set of Kohn-Sham orbitals and energies $\{|\phi_i\rangle, \varepsilon_i\}$ in order to construct the non-interacting response function χ_0 which is in turn used to solve the equations of linear response time-dependent DFT.

Present-day practical implementations scale as either $\mathcal{O}(N^4)$ or $\mathcal{O}(N^5)$ [100, 287–289] depending on the preferred approximate $f_{xc}[n]$ and whether one performs a so-called *one-shot* or *self-consistent* ACFDT calculation, where N can be taken as the number of constituent particles involved in the calculation.

The most striking successes of published ACFDT calculations have been in describing long-range correlations, e.g. van der Waals correlations between two disjoint subsystems within a larger system [18, 51, 283]. Such correlations evade capture within Kohn-Sham DFT using local or semi-local xc approximations [50, 290, 291], whereas the ACFDT correlation functional includes inherent non-locality even at the lowest level of approximate $f_{xc}[n]$, the random phase approximation (RPA) $f_{xc}^{\text{RPA}} = 0$. For example, ACFDT calculations utilising the RPA are able to accurately describe dissociation curves such as that of N_2 [287] – a notorious challenge for conventional Kohn-Sham DFT at large separations [50, 292]. Furthermore, the ACFDT framework in general is central to the development of systematic van der Waals functionals [56, 57], which have been successful not just in determining dissociation curves, but also in computing van der Waals coefficients, bond lengths, bond energies, and so on.

On the other hand, ACFDT approximations such as the RPA-ACFDT are known to be deficient in regard to *absolute* total energies, owing to weaknesses in their treatment of short-range correlations [100, 293], e.g. in the case of the homogeneous electron gas (HEG) [294, 295]. Indeed, individual RPA-ACFDT energies are often less accurate than those calculated using direct application of local/semi-local approximations to the correlation energy. The leading cause of this is thought to be the effect of spurious self-interaction in the xc kernel [85, 296].

A host of ACFDT approximations that venture beyond the RPA have been considered as possible remedies to this issue. Two such examples include the self-interaction-free exact-exchange kernel f_x [18, 30, 102, 103, 107, 255, 258, 297–301] and approaches that separate treatment of long-range and short-range contributions to the correlation energy, and use the RPA approximation for the former [302–309]. These improvements contribute toward alleviating some of the fundamental issues in the present context, e.g. in the calculation of atomisation energies, and thus advance the on-going effort to further establish the ACFDT approach as a total energy method, an effort with which this work is also concerned.

Although the RPA approximation within the ACFDT framework is amenable to a term-by-term analysis in the context of many-body perturbation theory [85, 310, 311], this connection is lost when moving beyond the RPA approximation, and as such so is a certain degree of transparency. Since f_{xc} implicitly contains all correlated many-body effects, including those required to describe excited-state phenomena such as the optical spectrum, it is imperative to better understand the connection between various aspects of f_{xc} and the ACFDT correlation energy. For example, a major consideration is the extent to which the frequency dependence in f_{xc} is important here [108, 215] – as we saw in the previous chapter, the exact f_{xc} includes a drastic dependence on ω whereby singularities exist along the real ω -axis that are critical for recovering the optical spectrum [26].

In the context of inhomogeneous finite one-dimensional systems, we compute the exact $f_{xc}^\lambda(x, x', \omega)$ in order to elaborate and elucidate the connection between its spatial non-locality/frequency dependence and the ACFDT total energy. This effort extends, in

part, the work of Lein, Gross, and Perdew in Ref. [108]. In particular, the so-called *adiabatic exact* (AE) kernel [235] $f_{xc}^\lambda[n](x, x', \omega = 0)$, i.e. the zero-frequency component of the exact $f_{xc}^\lambda[n]$ functional, is explored in relation to both one-shot and self-consistent ACFDT calculations.

6.1 Methodology

This chapter yet again involves finite systems in one dimension interacting with a softened Coulomb electron-electron interaction, where $\alpha = 1$ a.u. is the chosen softening parameter. We consider four prototype systems, each of which include two like-spin electrons in the external potentials described below.

The central complication when implementing the exact ACFDT total energy functional Eq. (6.1) is evaluation of the ACFDT correlation functional Eq. (6.2), which we now proceed to elaborate, see also Fig. 6.1.

Having chosen some external potential $v_{\text{ext}}(x)$, the exact interacting density $n(x)$ is obtained through solution of the time-independent Schrödinger equation. The corresponding unique Kohn-Sham potential $v_{\text{KS}}(x)$ is then reverse-engineered by applying the preconditioned root-finding techniques to an appropriate fixed-point map as described in Chapter 4. Let us borrow the atomic system from Section 6.2 for illustrative purposes: Fig. 6.2 depicts the ground-state two-electron density of the external potential

$$v_{\text{ext}}(x) = -\frac{1}{|0.05x| + 1}, \quad (6.3)$$

alongside the exact Kohn-Sham potential. The corresponding Kohn-Sham orbitals and energies $\{|\phi_i\rangle, \varepsilon_i\}$ are used to construct the atomic non-interacting response function χ_0 along $i\omega$ in the Lehmann representation [7].

6.1.1 Response Functions Along the Adiabatic Connection

In order to calculate the final ingredient $\chi^\lambda[n]$, we first obtain the λ -dependent wavefunctions $\{|\Psi_i^\lambda\rangle\}$ along the adiabatic connection – this grossly impractical step is for investigative reasons, and constitutes the primary computational expense. In fact, over the course of this project, it is estimated that around $\mathcal{O}(10^7)$ exact solutions of the many-body time-independent Schrödinger equation were performed toward this end. For each value of the coupling constant on some discrete grid, the potential $v_{\text{dxc}}(\lambda, x)$ is obtained by yet again using root-finding techniques to target the $\lambda = 1$ interacting density $n(x)$. That is, *interacting* systems are now reverse-engineered, rather than non-interacting systems as in the previous chapters, which requires solution of the two-particle Schrödinger equation for each iteration of the root finder. An example atomic $v_{\text{dxc}}(\lambda, x)$ for $N_\lambda = 10$ points sampled along the adiabatic connection is given in Fig 6.3.

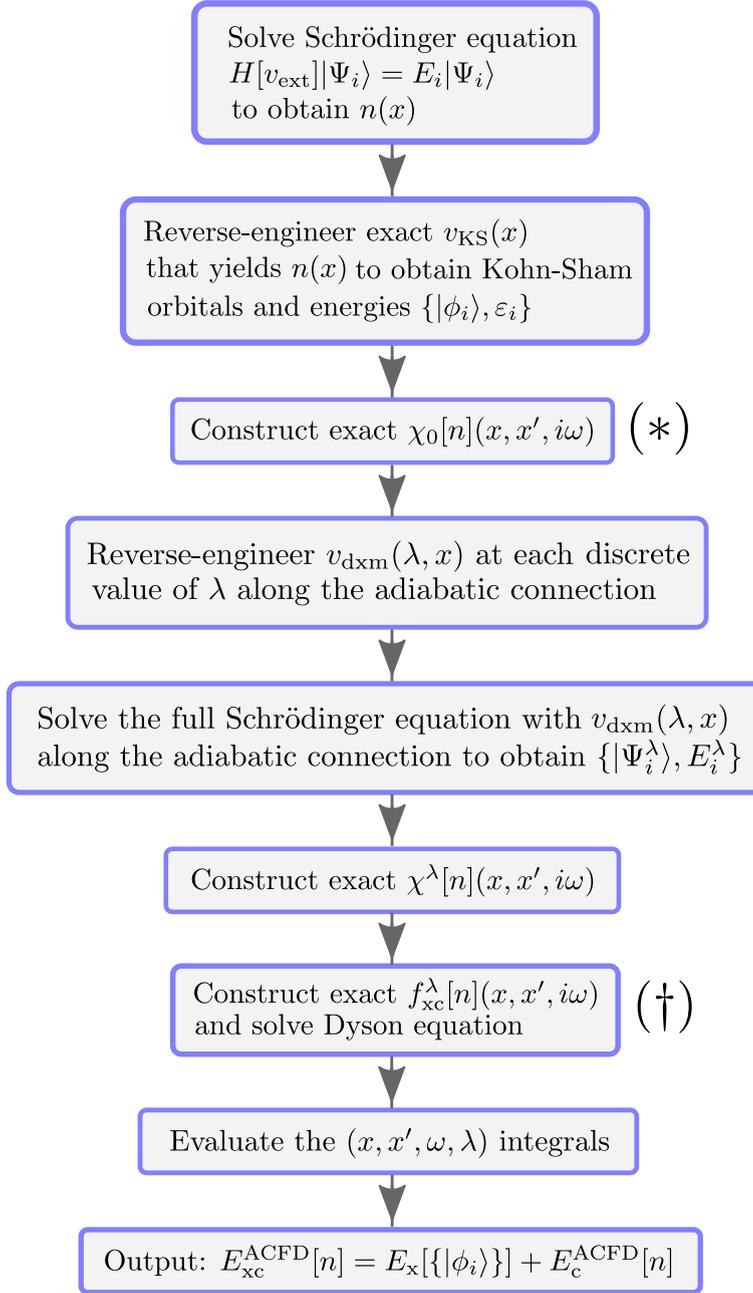


FIGURE 6.1: A flow chart depicting the course of action taken to evaluate the exact ACFDT total energy $E^{\text{ACFD}}[n]$ at the exact interacting ground-state density n . This procedure is modified in order to investigate approximate ACFDT approaches as follows: (*) a non-interacting Kohn-Sham calculation is performed with some approximate $v_{\text{xc}}[n]$, and then the algorithm proceeds with the corresponding approximate non-interacting response function, (\dagger) rather than calculate and utilize the exact $f_{\text{xc}}^\lambda[n]$, an approximate functional $f_{\text{xc}}^\lambda[n]$ is chosen and used alongside $\chi_0[n]$ to solve the Dyson equation, ultimately yielding an approximate many-body response function $\chi^\lambda[n]$.

The full set of λ -dependent wavefunctions and energies $\{|\Psi_i^\lambda\rangle, E_i^\lambda\}$ is then used to calculate the λ -interacting response functions along $i\omega$ in the Lehmann representation,

$$\chi^\lambda(x, x', i\omega) = \sum_{n=1}^{\infty} -2 \frac{\Omega_n^\lambda}{\omega^2 + (\Omega_n^\lambda)^2} \langle \Psi_0^\lambda | \hat{n}(x) | \Psi_n^\lambda \rangle \langle \Psi_n^\lambda | \hat{n}(x') | \Psi_0^\lambda \rangle, \quad (6.4)$$

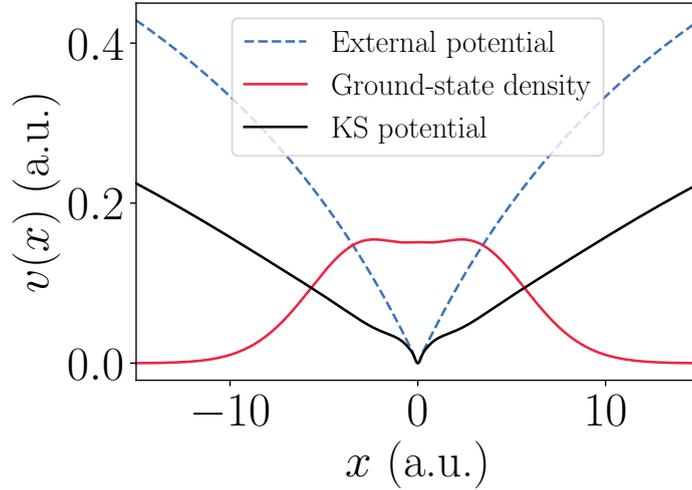


FIGURE 6.2: The ground-state density, external potential, and reverse-engineered Kohn-Sham potential for the atomic system. The external and Kohn-Sham potentials have been shifted for illustrative purposes.

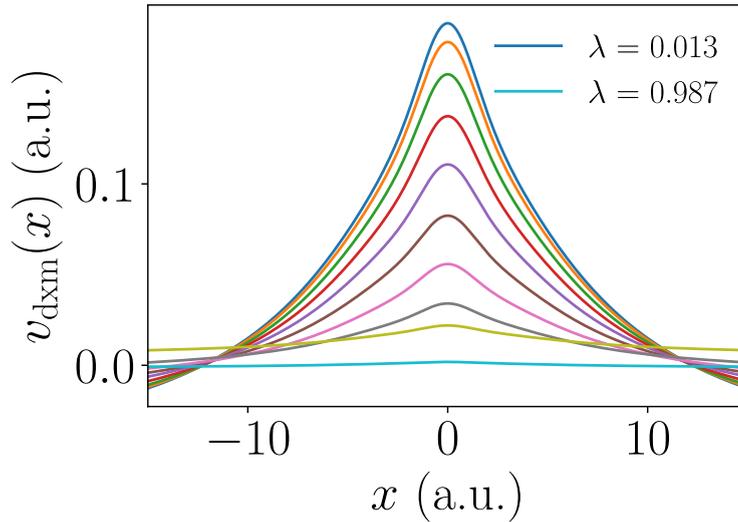


FIGURE 6.3: The potential $v_{\text{dxm}}(x, \lambda)$ that ensures the λ -interacting system shares the same atomic density $n(x)$ as the $\lambda = 1$ interacting system for $\lambda \in [0, 1]$. Note that these potentials interpolate between $v_{\text{dxm}}(x, \lambda = 1) = v_{\text{xc}}(x) + v_{\text{H}}(x)$, and $v_{\text{dxm}}(x, \lambda = 0) = 0$.

where $\Omega_n^\lambda = E_n^\lambda - E_0^\lambda$ is the n^{th} excitation energy of the λ -interacting adiabatic connection Hamiltonian. By evaluating the ω -dependent integral along the $i\omega$ contour, we are able to circumvent the poles lying on the real ω -axis in the response function χ , see Fig. 6.4. In exact arithmetic it does not matter which contour is used, however the new-found regularity along $i\omega$ assists the numerical methods in finite precision arithmetic.

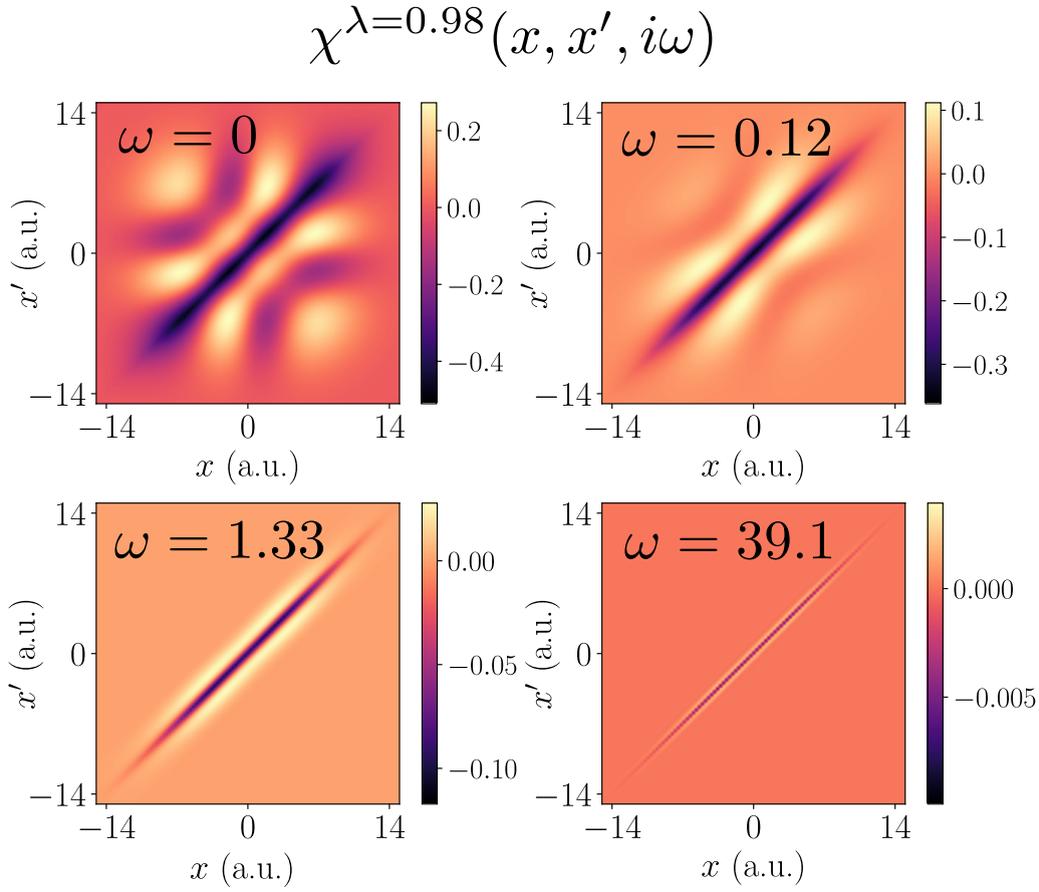


FIGURE 6.4: The atomic interacting response functions along $i\omega$ at $\lambda = 0.98$. The final panel depicts χ at $\omega = 39.1$ a.u. as it approaches its known high- $i\omega$ limit [269] wherein it becomes increasingly localised, and the difference between χ and χ_0 becomes constant.

At this stage it is possible to construct the exact λ -dependent xc kernel using the expression

$$f_{\text{xc}}^{\lambda}(\omega) = \chi_0^{-1}(i\omega) - (\chi^{\lambda})^{-1}(i\omega) - \lambda f_{\text{H}}, \quad (6.5)$$

which comes from inspection of the adiabatic connection Dyson equation (note that superscript -1 again signifies the matrix inverse in a finite spatial basis). In contrast to the previous chapter, we no longer require the systems to possess such large regions of nearly vanishing density toward the edge of the domain. Therefore, computations of f_{xc}^{λ} are much more regular, and in all the cases presented below the exact λ -dependent response function is reconstructed to within machine precision when the exact f_{xc}^{λ} and exact χ_0 are used to solve the Dyson equation. This step is critical as ultimately we shall isolate certain features of the exact f_{xc}^{λ} in order to examine their impact on the correlation energy, for example utilising only the $\omega = 0$ component of f_{xc}^{λ} here defines the AE approximation.

6.1.2 Integration Schemes

The final step toward obtaining E_c^{ACFD} involves evaluating the coupling constant and frequency integrals in Eq. (6.2). Gauss-Legendre quadrature is a numerical integration technique that enables polynomials of degree less than or equal to $2N+1$ to be integrated exactly with just N samples of the underlying polynomial. An arbitrary function $f(x)$ whose integral is computed using Gauss-Legendre quadrature thus involves error to the extent that $f(x)$ is not polynomial. For the systems presented in Section 6.2, the λ -dependent integrand does not deviate much from a linear form, and thus Gauss-Legendre quadrature is able to reach machine precision with $N_\lambda = 10$ grid points.

However, the ω -dependent integrand is not suited to the traditional Gauss-Legendre scheme, and so it is convention to utilise a change of coordinates in order to reduce the number of frequency grid points required to reach a desired accuracy. (The traditional Gauss-Legendre scheme is then applied to the transformed integral.) Namely, the ω -dependent ACFDT correlation energy integrand can be written in the form

$$E_c^{\text{ACFD}} = \int_0^\infty g(\omega) d\omega, \quad (6.6)$$

where we now seek a change of coordinates $\omega \rightarrow \tilde{\omega}$ such that the transformed integral,

$$\int_0^\infty g(\omega) d\omega = \int_{\tilde{\omega}(0)}^{\tilde{\omega}(\infty)} g(\omega(\tilde{\omega})) \frac{d\omega}{d\tilde{\omega}} d\tilde{\omega}, \quad (6.7)$$

is more suited to the traditional Gauss-Legendre integration scheme. Inspection of the ACFDT correlation energy functional demonstrates that the ω -dependence arises through $\chi^\lambda[n](i\omega)$ (and $\chi_0[n](i\omega)$). Therefore, it is sufficient to seek an accurate integration scheme for $\int_0^\infty \chi^\lambda(i\omega) d\omega$, which will ultimately yield an accurate integration scheme for the original integral Eq. (6.6).

Suppose we consider just the $\lambda = 1$ response function, then we find

$$\int_0^\infty \chi(x, x', i\omega) d\omega = -2 \sum_n f_n(x) f_n(x') \int_0^\infty \frac{\Omega_n^2}{\omega^2 + \Omega_n^2} d\omega, \quad (6.8)$$

where $f_n(x) = \langle \Psi_0 | \hat{n}(x) | \Psi_n \rangle$ is the so-called excitation function. This integral has an analytic solution: $\pi/2$ regardless of Ω_n , but to utilise this fact requires us to be in possession of the exact excitation functions, and therefore the exact many-body wavefunctions. In practice, only an approximate form of the curve $g(\omega)$ is attained, the integral of which has no analytic solution, meaning we must construct some accurate and efficient approximate integration scheme.

Given that $g(\omega)$ is demonstrably not a low-order polynomial, we propose a scheme that performs a change of coordinates such that one term in the Lehmann representation of $\chi(i\omega)$ is linear, i.e. integrated exactly with one grid point. In other words, we require

$$\frac{a}{a^2 + \omega^2} \frac{d\omega}{d\tilde{\omega}} = \tilde{\omega}, \quad (6.9)$$

where a is a parameter related to the excitation energies. Solution of this differential equation yields

$$\omega = a \tan\left(\frac{a\tilde{\omega}^2}{2}\right), \quad (6.10)$$

meaning the integration limits are compressed $[0, \infty] \rightarrow [0, \sqrt{\pi/a}]$. Given an even sampling for $\tilde{\omega} \in [0, \sqrt{\pi/a}]$, this scheme heavily biases the low- ω region of $g(\omega)$, the importance of which is discussed in Section 6.2. Applying this coordinate transform to the original integral Eq. (6.6) gives,

$$\int_0^\infty g(\omega) d\omega = \int_0^{\sqrt{\pi/a}} \frac{a^3 \tilde{\omega}}{(\cos(a\tilde{\omega}^2/2))^2} g(\omega(\tilde{\omega})) d\tilde{\omega}. \quad (6.11)$$

We find that the parameter a should be chosen as the average of the excitation energies ‘close to’ the first excitation energy (i.e. the gap), and, in case an estimate for this is not available, it is better to err on the side of larger rather than smaller a . This approach is able to reach more than sufficient accuracies with $N_\omega = 30$ grid points, and outperforms certain other methods from literature [27] for our systems – it would be an interesting exercise to examine how this approach performs in practice.

6.1.3 One-Shot Calculations

The algorithm that has just been outlined captures the exact correlation energy to within $\mathcal{O}(10^{-10})$ a.u. across the forthcoming systems. This procedure can be suitably adapted, see Fig. 6.1, to include an approximate $f_{xc}[n]$ and/or an approximate $v_{xc}[n]$, where the latter is used to determine the density at which $E^{\text{ACFD}}[n]$ is evaluated in a so called *one-shot calculation* (rather than evaluating $E^{\text{ACFD}}[n]$ at the exact interacting density as is done in the first and second panels of Fig. 6.1).

The traditional perspective on one-shot ACFDT calculations is as follows: first, the practitioner performs a ground-state Kohn-Sham calculation with some approximate functional $v_{xc}[n]$ (e.g. LDA), and then treats the ACFDT expression Eq. (6.2) as a *correction* to the approximate ground-state orbitals and energy (in analogy with one-shot GW calculations whose starting point is a Kohn-Sham ground-state). The perspective of this work is, however, to treat the ACFDT framework as an energy functional in its own right, in which case it can be shown that an approximate $v_{xc}[n]$ in a one-shot calculation serves the purpose of generating the density n (upon solution of the Kohn-Sham equations) with which the total energy functional $E^{\text{ACFD}}[n]$ Eq. (6.1) is evaluated.

In order to define the energy functional, we saw in the previous sections that only $f_{xc}[n]$ is required, in which case it is guaranteed that, if $f_{xc}[n]$ is exact, then the resultant energy functional is also exact (i.e. it provides the correct ground-state energy when evaluated at the interacting ground-state density). The question of interest thus becomes: even if we had access to this exact functional, or an accurate representation of it, how would we go about determining the density n at which to evaluate it in practice? The answer is obvious within the above line of reasoning: minimise the functional as in the case of ground-state Kohn-Sham calculations. This defines the much less popular *self-consistent* flavour of ACFDT calculation, which is discussed in the next section.

Instead, the more prominent *one-shot* ACFDT calculation decides to evaluate the functional at a density n that is not the minimum, and in fact does not even attempt to be the minimum, of an *approximate* ACFDT functional. Namely, in the case that $f_{xc}[n]$ is approximate, there is no guarantee that minimising the ACFDT functional brings one closer to the ‘platonic’ exact ground state for a system of N particles. In fact, it is often beneficial to utilise an accurate ground-state $v_{xc}[n]$ functional if one is available, since the majority of the ACFDT functional constitutes the Hartree-Fock energy functional¹. In other words, whether to perform a one-shot calculation or a self-consistent calculation depends on how confident one is in their $v_{xc}[n]$ versus $f_{xc}[n]$ for the kinds of correlations present in one’s system – Section 6.2 elaborates on this assertion.

6.1.4 Self-Consistent Calculations

A *self-consistent* ACFDT total energy calculation is defined as the solution to

$$E_0 = \min_n E^{\text{ACFD}}[n]. \quad (6.12)$$

As we saw in Chapter 3, to solve the minimisation problem in the case of conventional Kohn-Sham calculations, the first-order optimal Euler-Lagrange equations, i.e. the Kohn-Sham equations, for the problem are derived and used in practice. The analogous set of self-consistent equations that minimise the ACFDT energy functional (provided the approximate kernel $f_{xc}[n]$ is expressed in terms of the density and orbitals) are known [99, 100, 106, 107]. Naturally, both deriving and solving these so-called optimised effective potential equations is much more challenging than in a conventional Kohn-Sham setting, see Refs. [99–105] for a series of articles on practical self-consistent ACFDT calculations.

In our case, we are required to minimise the ACFDT energy functional explicitly as the AE kernel $f_{xc}[n](\omega = 0)$ has no analytic representation in terms of the density/orbitals. That is, a self-consistent ACFDT calculation in this work comprises first specifying an initial guess Kohn-Sham potential $v_{KS}(x)$ in place of the first and second panels in Fig. 6.1. Note that since there is a one-to-one correspondence between the density and the Kohn-Sham potential, it is sufficient to minimise over variations in v_{KS} in order to minimise the ACFDT total energy functional as in Eq. (6.12). We are then able to iterate the initial guess toward the minimising Kohn-Sham potential by utilising the BFGS optimisation algorithm – this involves looping over the flow chart in Fig. 6.1. In the event that the ACFDT total energy functional is specified with the exact $f_{xc}[n]$, the minimisation procedure terminates at the exact Kohn-Sham potential $v_{KS}(x)$ /the exact interacting density $n(x)$ *without their explicit inclusion*. In general, iterations are terminated when the Jacobian norm is $\mathcal{O}(10^{-6})$, meaning the BFGS algorithm is making energy variations $\mathcal{O}(10^{-8})$. The calculations are parallelised over λ grid points using `dask` [312].

Minimising the ACFDT total energy functional is able to circumvent the troublesome ‘starting-point dependence’ inherent to a one-shot calculation. Practical self-consistent ACFDT calculations have been performed with the RPA and exact-exchange kernels

¹Theoretical concerns aside, one-shot ACFDT calculations are much simpler and cheaper than self-consistent ones.

[99–105], the latter has been particularly successful. It is unfortunately the case that self-consistent calculations are much more involved and demanding than one-shot calculations. However, the accuracy of self-consistent ACFDT total energies is entirely determined by the approximate $f_{xc}[n]$, which in certain circumstances can be advantageous, as we shall examine in Section 6.2.

6.2 Results and Discussion

Let us begin by first examining one-shot and self-consistent ACFDT total energies in order to explore the significance of the spatial/frequency dependence in $f_{xc}(x, x', i\omega)$ and the influence of the density n at which $E^{\text{ACFD}}[n]$ is evaluated. An atomic system is used to illustrate the results concerning short-ranged correlations, whereas a double well system is used in the later section on long-ranged correlations. A system with a flat slab-like density profile assists in determining various sources of error in Section 6.2.3. Finally, all four systems, including the infinite potential well, enable us to ensure sufficient generality in the conclusions drawn.

This work focuses on three approximations to the xc kernel: the RPA $f_{xc}^{\text{RPA}}[n] = 0$, an adiabatic LDA $f_{xc}^{\text{ALDA}}[n](x, x', \omega = 0) \propto \delta(x - x')$, and the AE xc kernel $f_{xc}[n](x, x', \omega = 0)$. Since this work utilises the softened Coulomb interaction, we are prohibited from using the scaling relationship,

$$f_{xc}^\lambda[n](x, x', \omega) = \lambda^2 f_{xc}[n_{1/\lambda}](x, x', \omega), \quad (6.13)$$

and therefore our LDA energy functional from which $f_{xc}^{\text{ALDA}}[n]$ is constructed is parameterised *at each value of* λ with reference to both the HEG and one-dimensional slab systems² [277].

6.2.1 Short-Range Correlations

We shall now consider two like-spin electrons confined in an atom-like potential $v_{\text{ext}} = -1/(|0.05x| + 1)$ within the domain $[-15, 15]$ a.u. which is discretised over $N_x = 121$ grid points – see Fig. 6.2. The so-called exact *adiabatic connection curve* [284, 313] is given in Fig. 6.5 which provides a geometric interpretation of the λ -dependent ACFDT integrand,

$$E_{xc}^{\text{ACFD}}[n] = E_x[\{|\phi_i\rangle\}] + \int_0^1 \frac{U_{xc}(\lambda)}{\lambda} d\lambda, \quad (6.14)$$

i.e. the λ -dependent xc potential energy, see Eq. (6.2). The slight convex bend in the adiabatic connection curve that is observed here implies a modest static correlation [284]. In this instance, the correlation energy is 1.3% of the xc energy, and the xc energy is 24% of the total energy, $E_{\text{tot}} = -1.510$ a.u.

The relative error in the atomic total energy is illustrated in Fig. 6.6 across the whole range of approximate ground-state xc potentials and xc kernels considered herein. The

²The λ -dependent LDA functional was parameterised by Mike Entwistle (Freie Universität Berlin).

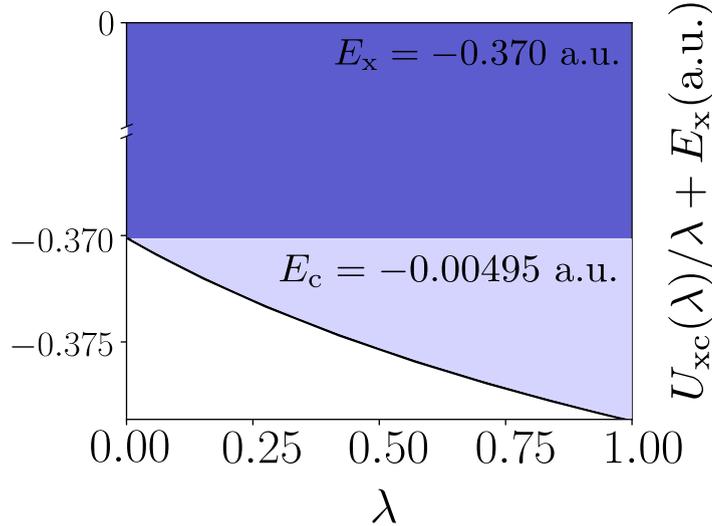


FIGURE 6.5: The exact adiabatic connection curve, i.e. the λ -dependent integrand of the ACFDT formula Eq. (6.14), where the correlation energy and exchange energy are labeled and given geometric context. The white region that compliments the shaded regions has area equal to minus the kinetic correlation energy, T_c .

TABLE 6.1: Absolute error in the atomic total energy (a.u.) across a variety of approximate ACFDT total energies. It is clear that the exact $f_{xc}[n]$ in conjunction with the exact $v_{xc}[n]$ produces a total energy equal to within $\mathcal{O}(10^{-10})$ to the one computed in a conventional Kohn-Sham calculation (labeled ‘No ACFDT’) when the exact $v_{xc}[n]$ is used.

$v_{xc}[n] \backslash f_{xc}[n]$	RPA	ALDA	AE	Exact	No ACFDT
Hartree	0.049	0.037	0.0225	0.0226	0.34
Hartree-Fock	0.024	0.017	0.0024	0.0021	0.0034
LDA	0.052	0.043	0.00064	0.00061	0.0034
Exact KS	0.047	0.039	0.00018	0.0	8.1×10^{-10}
Self-consistent	0.056	0.046	0.00018	2.22×10^{-16}	–

data used to generate Fig. 6.6 is given Table 6.1, where it can be seen that the exact total energy calculated via conventional means and the exact ACFDT total energy calculated using the algorithm presented in Section 6.1 agree to within $\mathcal{O}(10^{-10})$, thus verifying the numerical methods.

With the exception of the energies calculated utilising the notoriously poor Hartree orbitals, the predominant clustering in error appears to be according to the approximate $f_{xc}[n]$, rather than the approximate $v_{xc}[n]$ used to generate the input density. This suggests that there is a fairly general insensitivity to the density at which $E^{\text{ACFD}}[n]$ is evaluated after having been specified with some $f_{xc}[n]$. Arguments have been made that this must be the case in the context of the RPA [18], and we are now able to demonstrate that it is also the case when using the adiabatic LDA and the AE approximations. This conclusion translates to all other systems studied here: an infinite potential well, a slab, and a double well.

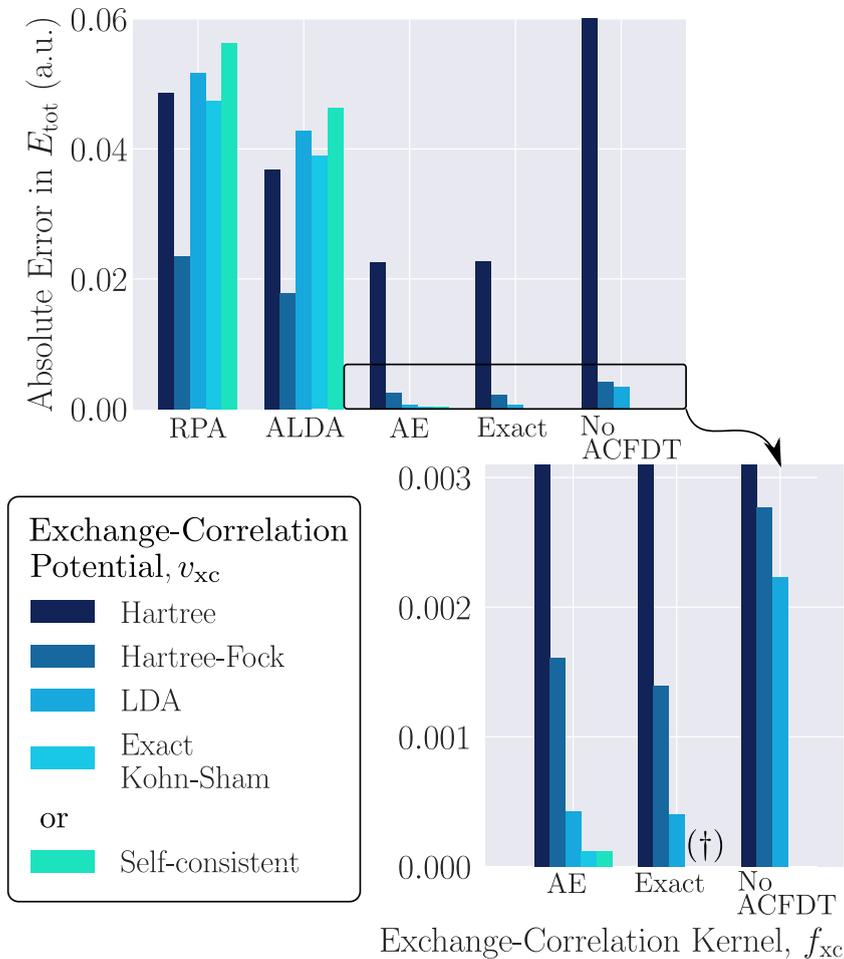


FIGURE 6.6: Absolute error in the atomic total energy across a variety of approximate ACFDT total energies. Each of the first four groups of colored bars is labeled with the approximate $f_{xc}[n]$ used to specify $E^{\text{ACFDT}}[n]$, while the fifth group (‘No ACFDT’) denotes a conventional ground-state Kohn-Sham calculation, for comparison. The first four bars of each group indicate the approximate $v_{xc}[n]$ used to generate the density (and orbitals) at which $E^{\text{ACFDT}}[n]$ is evaluated, while the fifth bar (cyan) indicates that the input density was determined self-consistently (i.e. it minimizes the applicable ACFDT total energy functional). The lower bar chart enlarges the outlined region in the upper bar chart. (†) At this position, there are two bars of zero height, i.e. the exact interacting ground-state density coincides with the density that minimizes the energy functional, both of which return the exact ground-state energy.

Secondly, the AE kernel, i.e. ignoring the frequency dependence in the otherwise exact functional $f_{xc}[n](x, x', \omega = 0)$, reduces the error in the total energy by orders of magnitude when compared to the RPA and adiabatic LDA kernels. For example, the *absolute* error in the AE-ACFDT total energy is 0.0006 a.u. when ground-state LDA density/orbitals are used in a one-shot calculation – significantly better than the usual measure of chemical accuracy, 0.0016 a.u. (1 kcal/mol). Moreover, we find that this level of accuracy is reliably achieved across differing ground-state orbitals (with the exception of the Hartree orbitals) and differing systems.

There is a sense in which this performance can be said to be inherent to the AE approximation $f_{xc}[n](x, x', \omega = 0)$. Namely, in using the exact Kohn-Sham potential $v_{KS}(x)$ to generate the one-shot ground-state orbitals, we are able to isolate error coming solely from the fact that the approximate xc kernel is not exact. Furthermore, minimising the ACFDT total energy functional serves a similar purpose, i.e. error is entirely due to the approximate $f_{xc}[n]$. It can be observed in Fig. 6.6 and Table 6.1 that these two methods produce relative errors in the total energy to within 10^{-6} a.u. of each other. In other words, the implicit (minimising) potential/orbitals that are contained within the self-consistent AE-ACFDT calculation [108] are close to the exact Kohn-Sham potential/orbitals, both of which yield ACFDT total energies well below chemical accuracy.

This similarity between the self-consistent AE-ACFDT orbitals and the exact Kohn-Sham orbitals is manifest when comparing the associated minimising density with the exact interacting density, Fig. 6.7. These observations point toward a central conclusion of this work: the spatial structure in the exact xc kernel at $\omega = 0$, see Fig. 6.8, is sufficient to almost entirely determine the exact total/correlation energy, and in turn the exact interacting density and exact Kohn-Sham potential. Hence, the intricate difference between the non-local spatial structure in Fig. 6.8 versus the local spatial structure in, for example, an adiabatic LDA kernel gives rise to a significant difference in the respective total energies.

A second example is given to illustrate the striking differences (and in some cases similarities) between the exact ground-state density, and densities that minimise an approximate ACFDT total energy functional. The slab has an external potential that has been reverse-engineered in order to produce a slab-like density inside the domain $[-17, 17]$ a.u., see Fig. 6.9. The slab is a delicate system: small changes in the external potential yield large (qualitative) changes in the density. As such, the densities that minimise a given approximate ACFDT total energy functional and the interacting density differ the most in the case of the slab, see Fig. 6.10. The AE minimising density remains able to faithfully reproduce the interacting ground-state density, whereas the RPA and adiabatic LDA minimising densities are woefully poor – it is possible to view their arrangement as being diffused due to self interaction, as we shall now explore.

Turning now toward the RPA and adiabatic LDA kernels, inspection of Fig. 6.6 and Fig. 6.7 leads one to conjecture that these approximate kernels share the same fundamental issues, with the adiabatic LDA suffering slightly less. Both approximations are in serious error when it comes to the correlation energy: an order of magnitude too negative in the case of the atom, and more so in the case of the infinite potential well and double well. Therefore, *minimising* the corresponding ACFDT total energy functionals necessarily makes matters worse – the RPA and adiabatic LDA minimising densities deviate significantly from the interacting ground-state density, and even deviate from most approximations to it, Fig. 6.7 and Fig. 6.10.

The source of this error is understood here in the context of the so-called λ -averaged xc hole. In conventional wavefunction-based theories, the statistical hole $n^{\text{hole}}(x, x')$

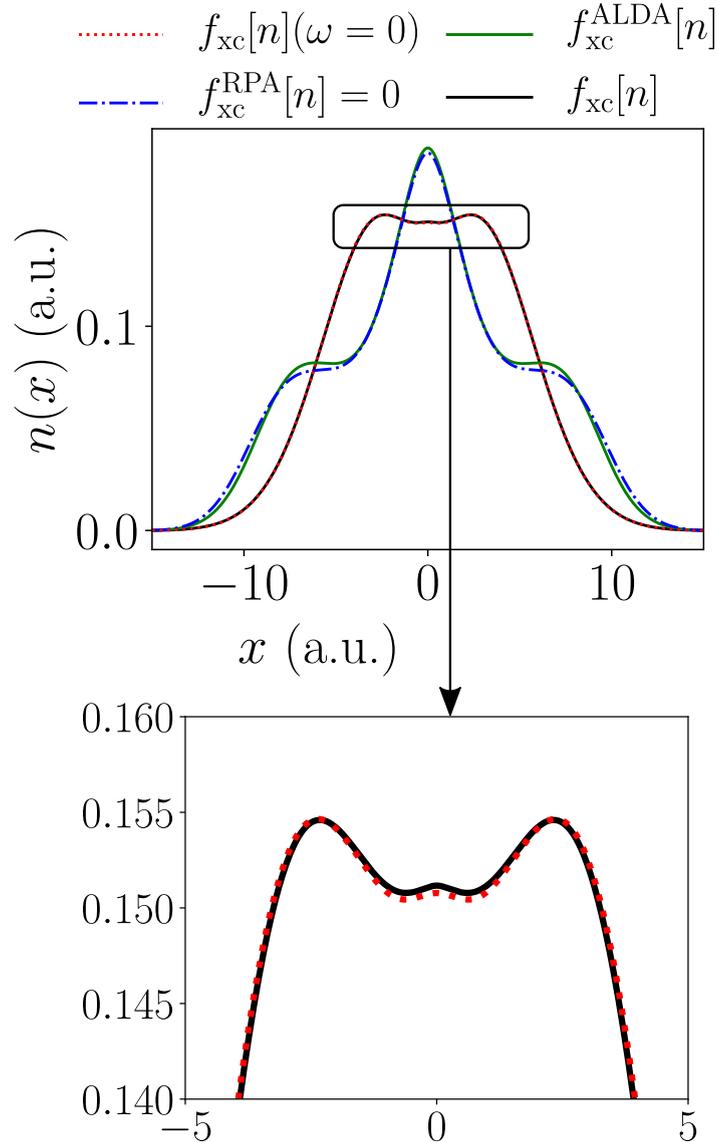


FIGURE 6.7: (Upper panel) The density n that minimises the atomic ACFDT energy functional $E^{\text{ACFDT}}[n]$ specified with some approximate $f_{\text{xc}}[n]$. The minimising density when the *exact* $f_{\text{xc}}[n]$ is used is the interacting ground-state density (black solid). (Lower panel) Zooming in such that the subtle difference between the minimising densities that comes from ignoring the frequency dependence in the exact $f_{\text{xc}}[n]$ can be seen.

describes how the probability distribution of particle positions $n(x)$ changes upon measurement of a particle at x' ³. In density-based theories, however, the xc hole is redefined and is the object with which the density undergoes Coulomb interaction to produce the

³The probability that a particle resides at position x given that a particle has been measured at x' is thus given by the conditional probability $n(x|x') = n(x) - n^{\text{hole}}(x, x')$.

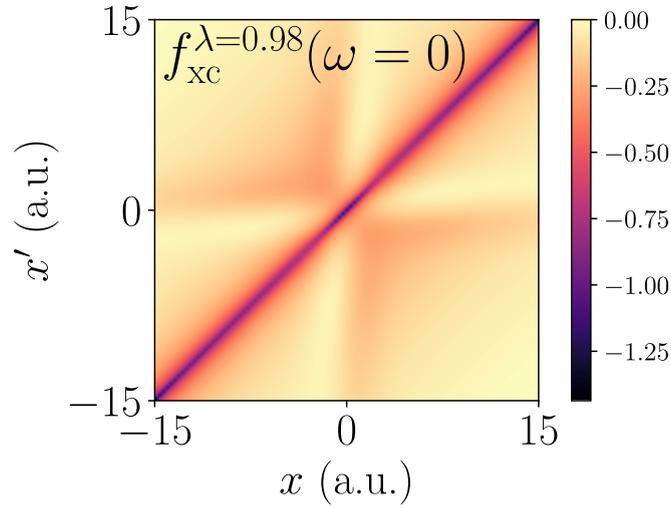


FIGURE 6.8: The AE $f_{xc}^{\lambda}(x, x', \omega = 0)$ for the atomic system at $\lambda = 0.98$, i.e. at a discrete value of λ that is sampled along the adiabatic connection.

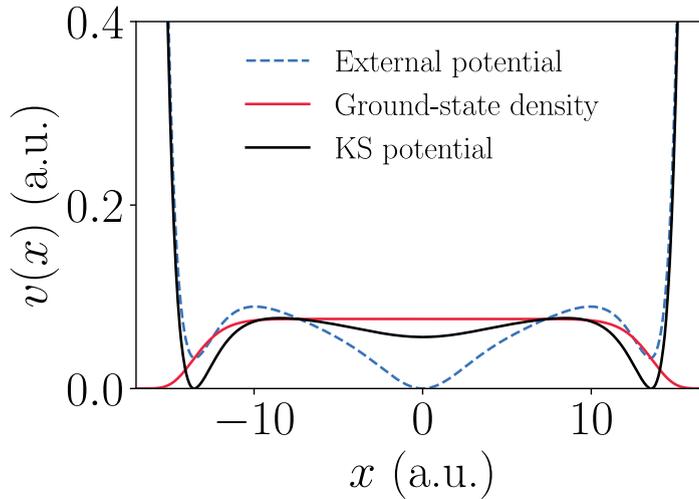


FIGURE 6.9: The ground-state density, reverse-engineered external potential, and reverse-engineered Kohn-Sham potential for the slab system. The external and Kohn-Sham potentials have been shifted for illustrative purposes.

xc energy [50, 91, 284, 290, 314],

$$E_{xc} = \iint \frac{n(x)n_{xc}^{\text{hole}}(x, x')}{|x - x'|} dx dx' \quad (6.15)$$

$$= \iiint \frac{n(x)n^{\text{hole}, \lambda}(x, x')}{|x - x'|} d\lambda dx dx', \quad (6.16)$$

where $n^{\text{hole}, \lambda}$ is the traditional statistical hole of the λ -interacting systems along the adiabatic connection. Comparing Eq. (6.16) with the ACFDT xc energy expression provides a unique definition of the xc hole in the present context, and moreover it

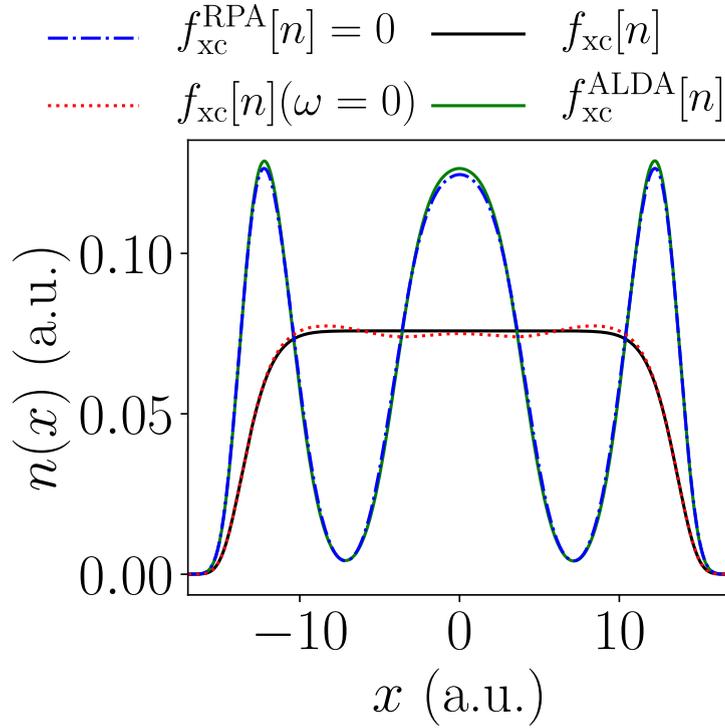


FIGURE 6.10: The density n that Minimises the slab ACFDT energy functional $E^{\text{ACFDT}}[n]$ specified with some approximate $f_{\text{xc}}[n]$. The minimising density when the *exact* $f_{\text{xc}}[n]$ is used is the interacting ground-state density (black solid).

defines an *approximate* xc hole in terms of $f_{\text{xc}}[n]$ [91]. Nevertheless, the two definitions are closely related, and it is instructive to view the λ -averaged xc hole in both contexts, not least because certain important sum rules are shared, e.g. $\int n_{\text{xc}}^{\text{hole}}(x, x') dx = -1$.

As discussed above, in order to isolate deficiencies in the approximate $f_{\text{xc}}[n]$, we shall hereafter consider xc holes that utilise the exact Kohn-Sham orbitals and density in the corresponding one-shot ACFDT calculations⁴. The upper panel in Fig. 6.11 demonstrates that the so-called *on-top* xc hole is far too deep in the case of the RPA and the adiabatic LDA. This can be interpreted as the second particle having *negative* probability to be measured at the position of the first – an artifact due to the original particle interacting with itself at $x' = 2.5$ a.u. This *self interaction* at the level of the xc kernel plagues both the RPA and adiabatic LDA similarly, which can be seen more clearly in the correlation hole [284] (lower panel of Fig. 6.11). In fact, the RPA and adiabatic LDA correlation holes reach a minimum at $x = 2.5$ a.u. where they should be identically zero, that is to say, exchange is entirely responsible for the fact that two fermions cannot be measured at the same position. Minimising the ACFDT energy functional defined with $f_{\text{xc}}^{\text{ALDA}}$ or $f_{\text{xc}}^{\text{RPA}} = 0$ accentuates the self interaction thereby making the on-top correlation hole even deeper.

Note that the traditional on-top LDA xc hole, that is, the on-top hole defined with the exact LDA pair-correlation function [91], is known to be accurate and therefore central

⁴This means that the exchange hole will be exact.

to the success of conventional ground-state LDA calculations [50]. However, the ALDA-ACFDT approximate xc hole, which utilises an adiabatic xc kernel derived from an LDA functional, is distinct from the traditional LDA xc hole, leading to errors of the kind discussed in the previous paragraph.

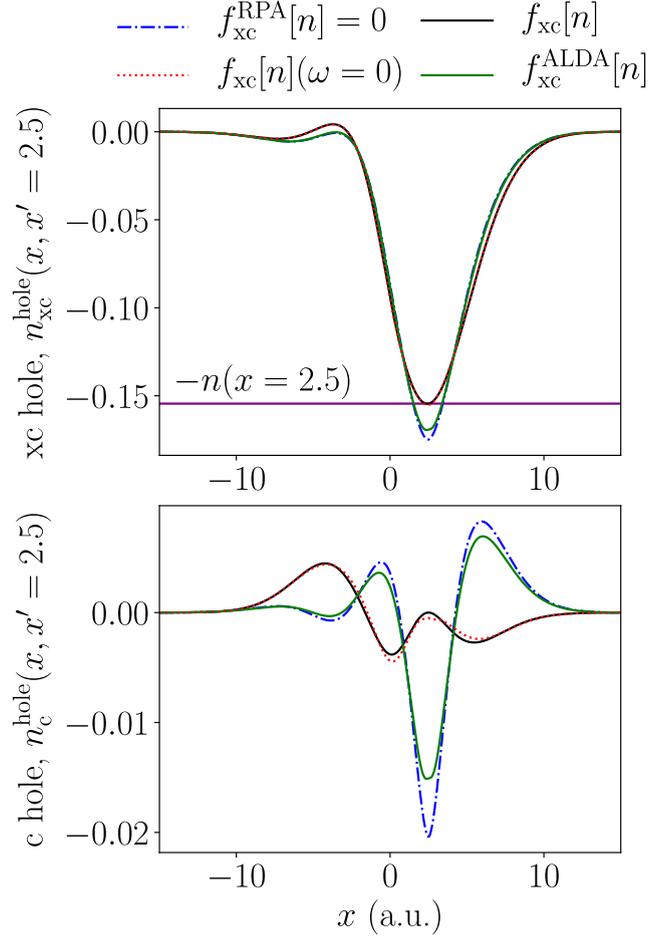


FIGURE 6.11: (Upper panel) The atomic xc hole at $x' = 2.5$ a.u. derived using the approximate xc kernels considered in this work. (Note that in all cases the exchange hole is exact because the exact ground-state Kohn-Sham orbitals were used to evaluate the ACFDT functional.) The density at $x' = 2.5$ a.u. (purple) and the xc hole should satisfy the following relation: $n_{\text{xc}}^{\text{hole}}(x = 2.5, x' = 2.5) = -n(x = 2.5)$, i.e. there is zero chance the second particle is at the same position as the first. (Lower panel) The correlation hole at $x' = 2.5$ a.u., where it can be seen that the adiabatic LDA and RPA on-top correlation holes are much too deep, giving rise to the excessively negative energies seen in Fig. 6.6.

A simple quantitative measure of self interaction is given using one-particle calculations, wherein any correlation present is necessarily due to self interaction. The atomic one-particle adiabatic LDA correlation energy is -0.016 a.u., whereas the corresponding adiabatic LDA two-particle correlation energy is -0.05 a.u. Doubling the spurious one-particle energy reveals that the majority of the two-particle correlation energy constitutes self interaction. The AE approximation has *no* spurious one-particle correlation energy, because the exact one-particle $f_{\text{xc}}[n]$ is itself adiabatic. Whilst this line of reasoning suggests that the AE approximation is self-interaction free, this is not quite the

case: the exact exchange kernel $f_x[n](x, x', \omega)$ provides a more rigorous definition of a self-interaction free kernel, and this includes an ω dependence [255, 298]. Nonetheless, it can be seen in Fig. 6.11 that the non-local spatial structure in the exact $f_{xc}[n]$ at $\omega = 0$ is able to largely remedy the deficiencies due to self interaction – Section 6.2.3 provides more insight on this matter.

6.2.2 Long-Range Correlations

We now examine a system containing long-range correlations, i.e. Van der Waals correlations [98]: a double well, see Fig. 6.12. This system is also known to exhibit a step-like feature in the Kohn-Sham potential [217, 315, 316] in order to localise one electron in each well – without such a step, the non-interacting particles would spuriously collapse into the same well, as is the case in Hartree theory.

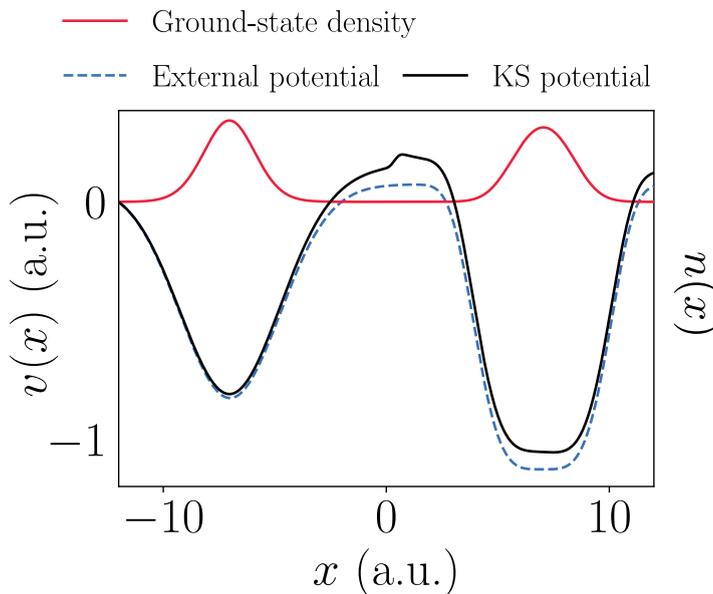


FIGURE 6.12: The ground-state density, external potential, and reverse-engineered Kohn-Sham potential for the double well system. The external and Kohn-Sham potentials have been shifted for illustrative purposes.

The wells are separated at distance $R = 14$ a.u., meaning the correlation energy is low $\mathcal{O}(10^{-6})$, and as such the RPA and adiabatic LDA suffer even more as a result of self interaction: the relative error in E_c is now $\mathcal{O}(10^5)$. However, one might expect such error is alleviated to a degree when computing quantities that involve energy differences, e.g. dissociation curves. In the context of dissociation curves, it is imperative that the total energy method in question provides at least some account of long-range correlations, and it is this phenomenon toward which we now focus our attention.

The correlation hole for the double well is given in Fig. 6.13. As expected, upon measuring an electron in the right-hand well at $x' = 8$ a.u., there is a strong erroneous contribution to RPA and adiabatic LDA correlation holes *in the same well where the*

electron was measured due to the electron Coulomb interacting with itself – the correlation hole should largely reside in the *opposite* well. Inclusion of the non-local spatial structure in the exact xc kernel at $\omega = 0$ is able to almost entirely correct this issue, as discussed in the previous section.

On the other hand, it can also be observed in the lower panel of Fig. 6.13 that the correlation hole in the opposing well, a necessarily long-range feature, is modeled successfully in all three of the approximations considered. This observation represents the systematic non-locality introduced via the Dyson equation, and constitutes the principal advantage of the ACFDT formalism over conventional Kohn-Sham calculations. We expect that these considerations account for the fact that, even in the case of the RPA, accurate dissociation curves can be attained [287]. We further note that the AE approximation $f_{xc}[n](\omega = 0)$ is able to *quantitatively* capture the long-range correlation hole, thus further establishing the notion that the majority of both short-range and long-range correlations reside in the spatial dependence in $f_{xc}[n](x, x', \omega = 0)$ when using the ACFDT framework.

Despite self-consistent RPA-ACFDT and adiabatic LDA-ACFDT calculations accentuating self interaction error by making the on-top correlation hole even deeper, we find that the long-range correlation hole is *improved* by minimising the ACFDT functional. Therefore, even though these approximate self-consistent ACFDT calculations yield poorer absolute energies than their one-shot counterparts, it appears that the long-range properties are improved.

The exact double well f_{xc} computed here contains the step-like features discussed in [317–319] that relate to the derivative discontinuity, Fig. 6.14. Although, these features are *not* necessary to recover the step in the Kohn-Sham potential, which is to say that all three approximations to $f_{xc}[n]$ recover the step when minimising the corresponding ACFDT energy functional. This is due to the fact that the Hartree-Fock optimised effective potential contains the step, and the ACFDT total energy functional comprises in large part the Hartree-Fock functional.

6.2.3 Sources of Error

The central aim of this section is to carefully examine the reasons behind the exceptional performance of the AE approximation $f_{xc}[n](\omega = 0)$ as demonstrated in the previous sections. The slab system, Fig. 6.9, is used to illustrate the forthcoming conclusions. First, it is possible to observe that the exact $f_{xc}[n](x, x', i\omega)$ along $i\omega$ undergoes considerable change away from its adiabatic $\omega = 0$ limit, see Fig. 6.15. (Note, however, that this change is much more regular than along the real ω -axis due to circumventing the singularities described in the previous chapter [26]). Therefore, the performance of the AE approximation *cannot* be attributed to xc kernels generally possessing modest frequency dependence along $i\omega$.

Since we have access to the exact λ -interacting adiabatic connection wavefunctions/energies $\{|\Psi_i^\lambda\rangle, E_i^\lambda\}$ and the exact Kohn-Sham single-particle wavefunctions/energies

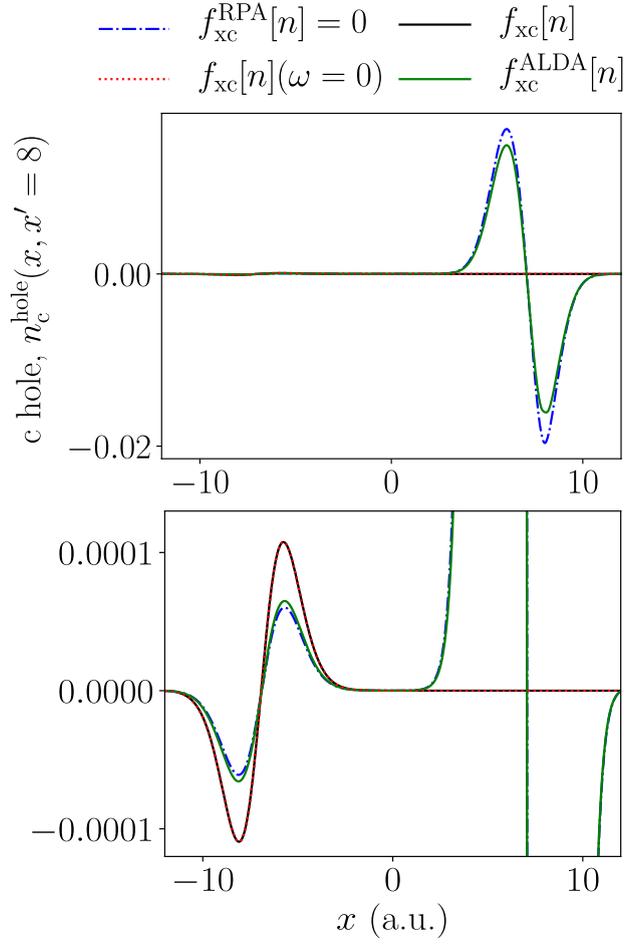


FIGURE 6.13: The double well correlation hole $n_c^{\text{hole}}(x, x' = 8)$ derived using the approximate xc kernels considered in this work. A particle is measured in the right-hand well, see Fig. 6.12, in which a significant spurious correlation hole emerges due to self interaction when using the adiabatic LDA and RPA approximations. The physical correlation is entirely long range here, and the lower panel zooms in on the long-range contribution from the upper panel in order to demonstrate that all three approximations are indeed able to capture this long-range correlation hole.

$\{|\Phi_i\rangle, \varepsilon_i\}$, the ω -dependent integral in $E_c^{\text{ACFD}}[n]$ in Eq. (6.2) can be evaluated analytically up to some arbitrary ω_{max} , which utilises

$$\begin{aligned} \int_0^{\omega_{\text{max}}} \chi(x, x', i\omega) d\omega &= \sum_n f_n(x) f_n(x') \int_0^{\omega_{\text{max}}} \frac{\Omega_n^2}{\omega^2 + \Omega_n^2} d\omega \\ &= \sum_n f_n(x) f_n(x') \arctan\left(\frac{\omega_{\text{max}}}{\Omega_n}\right). \end{aligned} \quad (6.17)$$

It is thus possible to isolate error whose exclusive source is a finite ω_{max} , as depicted in Fig. 6.16 – another perspective on this is that the exact $f_{\text{xc}}[n]$ is used for $\omega \leq \omega_{\text{max}}$, and $f_{\text{xc}}[n] = -\lambda f_{\text{H}}$, i.e. $\chi = \chi_0$, is used for $\omega > \omega_{\text{max}}$.

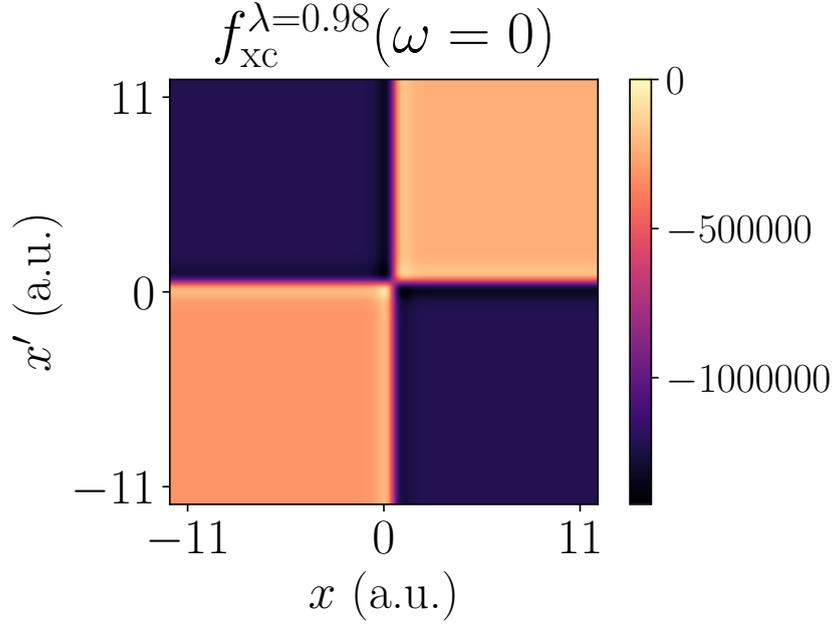


FIGURE 6.14: The exact $f_{xc}^{\lambda}(x, x', \omega = 0)$ for the double well system at $\lambda = 0.98$, i.e. at a discrete value of λ that is sampled along the adiabatic connection.

The absolute error in the correlation energy,

$$\Delta E(\omega_{\max}) = |E_c^{\text{ACFD}}(\infty) - E_c^{\text{ACFD}}(\omega_{\max})|, \quad (6.18)$$

is then plotted in Fig. 6.17, where $E_c^{\text{ACFD}}(\omega_{\max})$ is the ACFDT correlation energy whose analytic ω integration has been terminated at ω_{\max} ($E_c^{\text{ACFD}}(\infty)$ is therefore the exact correlation energy). As Fig. 6.17 demonstrates, chemical accuracy is surpassed prior to $\omega = 1$ a.u., i.e. the capacity for an approximate $f_{xc}[n]$ to yield accurate correlation energies predominantly resides in its structure *below* some characteristic ω_{\max} ($\omega_{\max} \approx 1$ in this case). In fact, beyond this ω_{\max} , the AE approximation $f_{xc}[n](\omega = 0)$ produces approximate interacting response functions upon solution of the Dyson equation that are *as poor as* the RPA, adiabatic LDA, or simply setting $\chi^{\lambda} = \chi_0$.

In all systems studied here, whilst the exact $f_{xc}[n]$ is not adiabatic in general, it *is* adiabatic over the most relevant ω range, and therefore the AE approximation performs accordingly. Such an observation is commensurate with prior findings [26]: in the context of the optical spectrum, $f_{xc}[n](\omega)$ is required *at the ω corresponding to a transition energy*, and thus at frequencies beyond the lowest lying excitations, the AE approximation ceases to perform.

In light of these observations, it is imperative that practical approximate integration schemes target the relevant ω region, whereas the long-range tail of the ω -dependent integrand is less important. Whilst the integration scheme proposed in this work (Section 6.1.2) compresses the domain $[0, \infty]$ to $[0, \sqrt{\pi/a}]$, thus allowing us to capture the tail, its central advantage instead comes from an explicit treatment of the terms responsible for the low ω structure in $g(\omega)$. Namely, the scheme defines an integral change of coordinates such that the most important term(s) in the Lehmann representation of χ

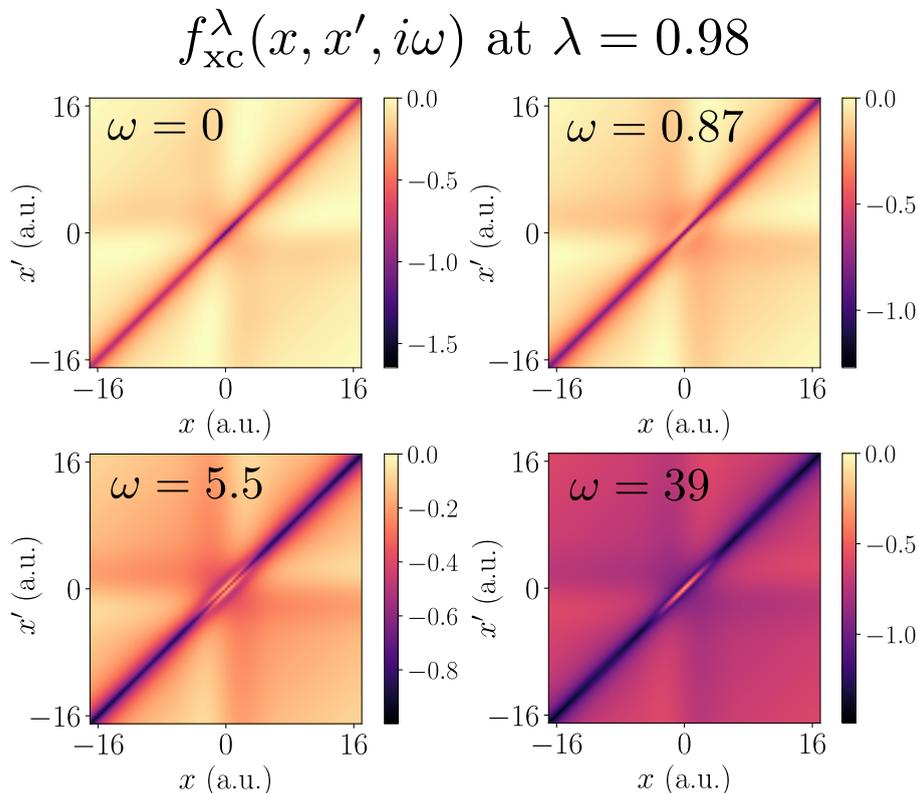


FIGURE 6.15: The exact numerical xc kernel $f_{\text{xc}}^\lambda(x, x', i\omega)$ for the slab system at $\lambda = 0.98$ along the adiabatic connection. The xc kernel is depicted at different stages along the $i\omega$ path, demonstrating that there is little frequency dependence below $\omega = 1$ a.u., whereupon the xc kernel then deviates from its adiabatic form, ultimately approaching its large $i\omega$ limit [269].

(or χ_0) Eq. (6.4) is *linear*, meaning Gauss-Legendre quadrature is exact with just one ω grid point. The characteristic range ω_{max} within which most of the correlation energy resides will depend on a number of factors in practice, such as the size of the interacting and Kohn-Sham gap, and number of excitations clustered around this gap [27].

6.3 Conclusions

The exact coupling constant-dependent xc kernel $f_{\text{xc}}^\lambda[n](x, x', i\omega)$ for four prototype one-dimensional finite systems has been calculated [26] and utilized to better understand the sources of error in practical ACFDT total energy calculations.

Whilst the frequency-dependence in the exact $f_{\text{xc}}[n]$ along the real ω -axis is critical for recovering the optical spectrum [26, 246, 320], this is not the case for ACFDT total energies, where it is demonstrated that chemical accuracy can be consistently attained using the AE kernel $f_{\text{xc}}[n](\omega = 0)$, i.e. neglecting the frequency dependence in the exact xc kernel, but otherwise retaining its non-local spatial structure when solving the Dyson equation of linear response time-dependent DFT. The exact kernel $f_{\text{xc}}[n]$ is shown to exhibit little change within the ω interval that contains the majority of the ground-state

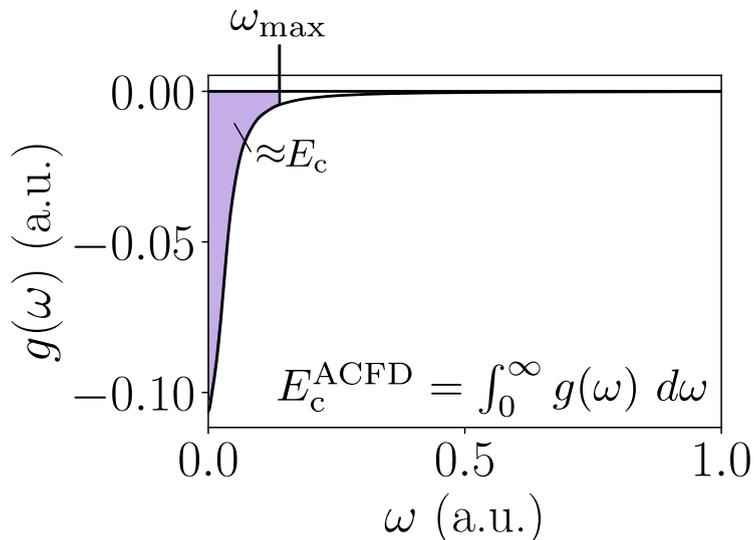


FIGURE 6.16: The exact ω -dependent integrand $g(\omega)$ of the ACFDT correlation energy expression is depicted for the slab system. Analytic integration allows us to terminate the integration at some finite ω_{\max} in order to determine the amount of correlation energy contained in the curve at frequencies $\omega \leq \omega_{\max}$.

correlation energy. Therefore, it is crucial that approximate kernels are accurate within this interval, as is the case for the AE kernel, thereby explaining its success. In light of these findings, a novel change of coordinates is proposed for the ACFDT ω -dependent integral that directly targets the relevant term(s) in the Lehmann representation of the response function in order to reduce the number of Gauss-Legendre grid points required. An interesting course for future work would involve testing this integration scheme in practical settings.

The coupling-constant averaged correlation hole is used, alongside one-particle calculations, to illustrate that strong self-interaction is present in the RPA and adiabatic LDA kernels – in both cases, the on-top correlation hole is far too deep, and the two-particle correlation energy differs little from the twice the spurious one-particle correlation energy. Due to the observations outlined in the previous paragraph, the spatial non-locality in the AE kernel almost entirely remedies the problem of self-interaction, despite the self-interaction-free exact-exchange kernel f_x possessing a frequency dependence.

In the case of a double-well system, we show that all three of the approximate kernels considered in this work – the RPA, adiabatic LDA, and AE kernels – are able to capture the long-range correlation hole. This observation further evidences the central advantage of ACFDT calculations, i.e. in describing long-range van der Waals correlations. Moreover, we find that minimizing the ACFDT total energy functional, while yielding less accurate absolute energies in the case of the RPA and adiabatic LDA, *is* able to improve the description of long-range correlations.

The distinction between one-shot and self-consistent ACFDT total energies is considered in depth, where we recall that the latter minimizes the ACFDT total energy functional, whereas the former evaluates the ACFDT total energy functional at some density n

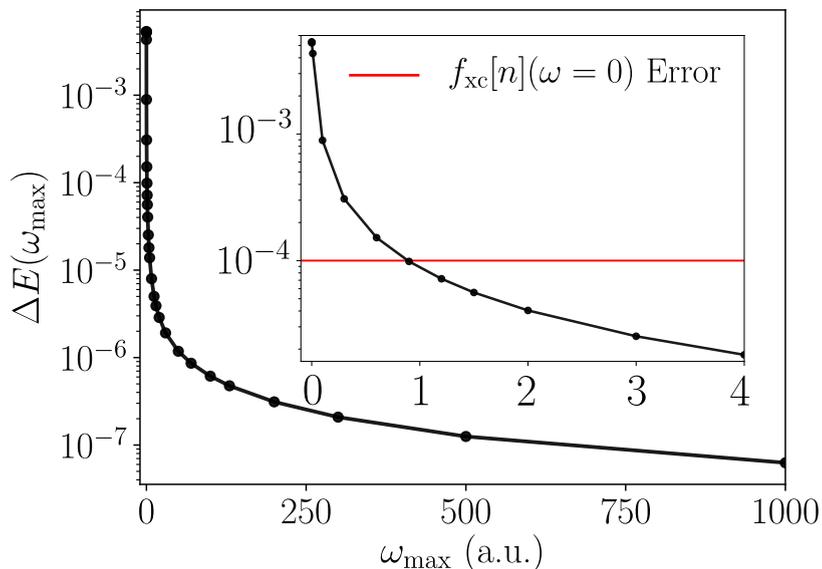


FIGURE 6.17: Absolute error in the slab correlation energy Eq. (6.18) (log scale) with increasing ω_{\max} (upper bound in the ACFDT ω -dependent integral). The curve tends toward zero in the $\omega_{\max} \rightarrow \infty$ limit. The inset enlarges the low- ω region, where it can be seen that the error involved in an AE one-shot ACFDT calculation (red) is contained within $0 < \omega < 1$.

obtained from an approximate ground-state Kohn-Sham calculation. The ACFDT total energy functional is found to be somewhat insensitive to the density at which the functional is evaluated (within reason), meaning the dominant factor that dictates the effectiveness of an ACFDT calculation is the approximate $f_{xc}[n]$ with which $E^{\text{ACFD}}[n]$ is defined. Where self-interaction is present, minimizing the ACFDT functional accentuates issues, thus making the erroneous on-top correlation hole even deeper. However, in the context of self-consistent AE-ACFDT calculations, the exact Kohn-Sham potential is faithfully recovered as the implicit optimized effective potential contained within the calculation, and therefore so are chemically accurate total energies.

The scope for obtaining both improved total energies and improved Kohn-Sham potentials using the ACFDT approach appears to be significant, and depends critically on capturing the spatial non-locality present in the AE kernel. For example, the energies and potentials that would result from combining modern adiabatic non-local kernels [89, 90] with self-consistent ACFDT calculations [99–105, 289] offer an interesting prospect.

Chapter 7

Conclusion

This thesis was largely concerned with theoretical and computational developments regarding ground-state and (linear response) time-dependent DFT. The broad aims of this work were two-fold: to help improve the computational implementation of practical ground-state DFT calculations, and to develop computational tools that can be used to better understand certain theoretical objects, such as the xc potential and xc kernel, that are central to the approximation of many-body quantum mechanics.

We began with a benchmarking effort for self-consistent field methods and preconditioners in popular implementations of ground-state DFT. Sources of numerical instability and inefficiencies were identified, and a host of iterative methods and preconditioners were set out that attempt, in some form or another, to alleviate these difficulties. A subset of these methods were implemented in the plane-wave, pseudopotential DFT code **CASTEP** as part of a benchmarking effort. This benchmark comprised a workflow that involved first assembling a suite of DFT simulation inputs designed to target numerical issues. The techniques of Pareto analysis were then able to reveal certain optimal methods when faced with two competing measures of utility: robustness and efficiency.

To an extent, robustness is to be preferred over efficiency, and it was demonstrated that certain simple modifications to established methods were able to improve robustness at an acceptable sacrifice in efficiency. It is likely that future algorithm design will centre around the use of more sophisticated techniques from optimisation literature in order to ensure convergence in the self-consistent field iterations without having to resort to direct minimisation methods, such as those in Refs. [169, 172]. The workflow presented in Chapter 3 is also able to compare parameter sets for methods and preconditioners. The default parameters proposed in Ref. [61, 62] were found to be Pareto optimal in the context of Kerker preconditioned quasi-Newton methods. Various other homogeneous and isotropic preconditioners (dielectric models) were implemented [155, 201, 202], but these were found to be equivalent to the Kerker preconditioner up to a change of parameters. Improved preconditioners should involve the systematic inclusion of inhomogeneity and anisotropy, such as those in Refs. [158, 161, 203, 207].

Methods were designed that are able to determine exact ground-state and time-dependent Kohn-Sham potentials $v_{\text{KS}}[n](x, t)$ that generate some reference density upon solution of the (time-dependent) Kohn-Sham equations. In conjunction with this, additional numerical methods were set out to calculate the exact xc kernel $f_{\text{xc}}[n](x, x', \omega)$ of linear

response time-dependent DFT for prototype inhomogeneous finite one-dimensional systems. These methods were used in Chapters 5 [26] and 6 [28], and in other works [215]. The regularising techniques used here could assist in future implementations of similar projects on a larger scale.

Computations of exact time-dependent Kohn-Sham potential were able to illustrate that implementations of the inverse map in time-dependent DFT can be ill-posed. This issue could plague the determination of exact time-dependent Kohn-Sham potentials in future studies. On the other hand, computations of the exact xc kernel are more regular, and were used to examine the xc kernel in relation to optical properties and ground-state total energies from the ACFDT. In Chapter 5, the frequency dependence of f_{xc} was found to be largely due to its analytic structure along the real ω -axis, i.e. its singularities at certain frequencies, which are *required* in order to capture particular transitions, including those of double excitation character. However, within the frequency range of the first few interacting excitations, f_{xc} is approximately ω -independent, meaning the *exact adiabatic approximation* $f_{xc}(\omega = 0)$ remedies the failings of the ALDA and RPA for these lowest transitions. The key differences between the exact f_{xc} and its common approximations were analysed, and cannot be eliminated by exploiting the limited gauge freedom in f_{xc} . The optical spectrum benefits from using as accurate as possible an f_{xc} and ground-state xc potential, while maintaining exact compatibility between the two is of less importance.

In the context of the ACFDT, both one-shot and self-consistent calculations were considered, and used to demonstrate that chemical accuracy is reliably preserved when the frequency dependence in the exact functional $f_{xc}[n](\omega = 0)$ is neglected. This performance is understood on the grounds that the exact $f_{xc}[n]$ varies slowly over the most relevant ω range (but not in general), and hence the spatial structure in $f_{xc}[n](\omega = 0)$ is able to largely remedy the principal issue in the present context: self-interaction (examined from the perspective of the exchange-correlation hole). Moreover, we find that the implicit orbitals contained within a self-consistent ACFDT calculation utilizing the adiabatic exact kernel $f_{xc}[n](\omega = 0)$ are remarkably similar to the exact Kohn-Sham orbitals, thus further establishing that the majority of the physics required to capture the ground-state total energy resides in the spatial dependence of $f_{xc}[n]$ at $\omega = 0$. This work helps motivate and direct future efforts in the development of xc kernel functions by providing evidence that much of the total energy is contained in the adiabatic component of the xc kernel.

References

- [1] T. S. Kuhn, *The Structure of Scientific Revolutions* (University of Chicago Press, Chicago, 1962).
- [2] P. K. Feyerabend, *Against Method* (Verso, 1975).
- [3] N. Chomsky, *On Nature and Language* (Cambridge University Press, 2002).
- [4] E. P. Wigner, [Communications on Pure and Applied Mathematics](#) **13**, 1 (1960).
- [5] P. A. M. Dirac and R. H. Fowler, [Proceedings of the Royal Society of London](#). **123**, 714 (1929).
- [6] A. F. Harper, M. L. Evans, J. P. Darby, B. Karasulu, C. P. Koçer, J. R. Nelson, and A. J. Morris, [Johnson Matthey Technology Review](#) **64**, 103 (2020).
- [7] C. A. Ullrich, *Time-dependent density-functional theory: concepts and applications* (Oxford University Press, 2012).
- [8] R. M. Martin, *Electronic Structure* (Cambridge University Press, 2004).
- [9] E. Eberhard and D. M. Reiner, *Density Functional Theory: An Advanced Course* (Springer-Verlag Berlin Heidelberg, 2011).
- [10] R. M. Martin, L. Reining, and D. M. Ceperley, *Interacting Electrons Theory and Computational Approaches* (Cambridge University Press, 2016).
- [11] R. Mattuck, *A Guide to Feynman Diagrams in the Many-body Problem*, Dover Books on Physics Series (Dover Publications, 1992).
- [12] R. G. Parr and W. Yang, *Density Functional Theory of Atoms and Molecules* (Springer, Dordrecht, 1980).
- [13] R. McWeeny, *Methods of molecular quantum mechanics* (Academic Press, 1992).
- [14] N. Marzari, A. Ferretti, and C. Wolverton, [Nature Materials](#) **20**, 736 (2021).
- [15] F. Jensen, *Introduction to Computational Chemistry* (Wiley, 2016).
- [16] K. Burke, [The Journal of Chemical Physics](#) **136**, 150901 (2012).
- [17] K. Burke, ed., *ABC of DFT* (Unpublished, 2007).
- [18] J. Dobson, G. Vignale, and D. Mukunda, eds., *Electronic Density Functional Theory: Recent Progress and New Directions* (Springer Science, 1998).

-
- [19] M. A. Marques, C. A. Ullrich, F. Nogueira, A. Rubio, K. Burke, and E. K. U. Gross, eds., *Time-Dependent Density Functional Theory*, Lecture Notes in Physics, Vol. 706 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2006).
- [20] A. D. Wade and D. J. Huggins, *Journal of Chemical Theory and Computation* **15**, 6504 (2019).
- [21] B. W. J. Irwin, J. R. Levell, T. M. Whitehead, M. D. Segall, and G. J. Conduit, *Journal of Chemical Information and Modeling* **60**, 2848 (2020).
- [22] N. D. Woods, M. C. Payne, and P. J. Hasnip, *Journal of Physics Condensed Matter* **31**, 453001 (2019).
- [23] K. N. Kudin and G. E. Scuseria, *ESAIM: Mathematical Modelling and Numerical Analysis* **41**, 281 (2007).
- [24] S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. I. J. Probert, K. Refson, and M. C. Payne, *Zeitschrift für Kristallographie - Crystalline Materials* **220**, 567 (2005).
- [25] N. Woods, <https://bitbucket.org/NickWoods1/scf-xn-testsuite> (2018).
- [26] N. D. Woods, M. T. Entwistle, and R. W. Godby, *Physical Review B* **103**, 125155 (2021), [arXiv:2101.12207](https://arxiv.org/abs/2101.12207) .
- [27] J. Harl, *The linear response function in density functional theory*, Ph.D. thesis, University of Vienna (2008).
- [28] N. D. Woods, M. T. Entwistle, and R. W. Godby, *Physical Review B* **104**, 125126 (2021).
- [29] E. Schrödinger, *What Is Life?* (Cambridge University Press, 1944).
- [30] M. Fuchs, Y. M. Niquet, X. Gonze, and K. Burke, *Journal of Chemical Physics* **122**, 121104 (2005).
- [31] D. A. Mazziotti, *Physical Review A* **102**, 030802 (2020).
- [32] L. K. Sørensen, [arXiv](https://arxiv.org/abs/1910.06633) (2019), [arXiv:1910.06633](https://arxiv.org/abs/1910.06633) .
- [33] H. Nakatsuji, *The Journal of Chemical Physics* **113**, 2949 (2000).
- [34] M.-E. Pistol, *Chemical Physics Letters* **400**, 548 (2004).
- [35] D. M. Ceperley and B. J. Alder, *Physical Review Letters* **45**, 566 (1980).
- [36] P. J. Hasnip and M. I. J. Probert, (Preprint) (2015).
- [37] C. A. Ullrich and Z. Hui Yang, *Brazilian Journal of Physics* **44**, 154 (2014).
- [38] P. Hohenberg and W. Kohn, *Physical Review* **136**, B864 (1964).
- [39] R. van Leeuwen, *Advances in Quantum Chemistry* **43**, 25 (2003).
- [40] J. T. Chayes, L. Chayes, and M. B. Ruskai, *Journal of Statistical Physics* **38**, 497 (1985).

- [41] K. Capelle, arXiv preprint cond-mat/0211443 (2002).
- [42] E. H. Lieb, *International Journal of Quantum Chemistry* **24**, 243 (1983).
- [43] H. Englisch and R. Englisch, *Physica Status Solidi (b)* **123**, 711 (1984).
- [44] H. Englisch and R. Englisch, *Physica Status Solidi (b)* **124**, 373 (1984).
- [45] M. Levy, *Proceedings of the National Academy of Sciences* **76**, 6062 (1979).
- [46] M. Levy, *Phys. Rev. A* **26**, 1200 (1982).
- [47] W. Kohn and L. J. Sham, *Physical Review* **140**, A1133 (1965).
- [48] P. A. M. Dirac, *Mathematical Proceedings of the Cambridge Philosophical Society* **26**, 376–385 (1930).
- [49] J. P. Perdew and A. Zunger, *Phys. Rev. B* **23**, 5048 (1981).
- [50] K. Burke, J. P. Perdew, and M. Ernzerhof, *Journal of Chemical Physics* **109**, 3760 (1998).
- [51] W. Kohn, Y. Meir, and D. E. Makarov, *Physical Review Letters* **80**, 4153 (1998), [arXiv:9707328 \[cond-mat\]](#) .
- [52] S. Lehtola, C. Steigemann, M. J. Oliveira, and M. A. Marques, *SoftwareX* **7**, 1 (2018).
- [53] N. Mardirossian and M. Head-Gordon, *Molecular Physics* **115**, 2315 (2017).
- [54] J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).
- [55] A. D. Becke, *The Journal of Chemical Physics* **98**, 5648 (1998).
- [56] J. T. Krogel, S. F. Yuk, P. R. C. Kent, and V. R. Cooper, *The Journal of Physical Chemistry A* **124**, 9867 (2020).
- [57] K. Berland, V. R. Cooper, K. Lee, E. Schröder, T. Thonhauser, P. Hyldgaard, and B. I. Lundqvist, *Reports on Progress in Physics* **78**, 066501 (2015).
- [58] C. LeBris, *Computational Chemistry, Volume 10 (Handbook of Numerical Analysis)* (Elsevier Science, 2003).
- [59] M. F. Herbst, A. Levitt, and E. Cancès, *Faraday Discussions* **224**, 227 (2020).
- [60] M. J. Colbrook, B. Roman, and A. C. Hansen, *Physical Review Letters* **122**, 250201 (2019).
- [61] G. Kresse and J. Furthmüller, *Computational Materials Science* **6**, 15 (1996).
- [62] G. Kresse and J. Furthmüller, *Physical Review B* **54**, 11169 (1996).
- [63] D. Vanderbilt, *Physical Review B* **41**, 7892 (1990).
- [64] C. Yam, X. Zheng, G. Chen, Y. Wang, T. Frauenheim, and T. A. Niehaus, *Physical Review B* **83**, 245448 (2011).

- [65] M. A. Marques, C. A. Ullrich, F. Nogueira, A. Rubio, K. Burke, and E. K. U. Gross, eds., *Time-Dependent Density Functional Theory*, Lecture Notes in Physics, Vol. 706 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2006) pp. 358–375.
- [66] E. Runge and E. K. Gross, *Physical Review Letters* **52**, 997 (1984).
- [67] R. van Leeuwen, *Physical Review Letters* **82**, 3863 (1999).
- [68] I. V. Tokatly, *Physical Review B* **83**, 035127 (2011).
- [69] M. Farzanehpour and I. V. Tokatly, *Physical Review B* **86**, 125130 (2012).
- [70] M. Ruggenthaler and R. van Leeuwen, *EPL (Europhysics Letters)* **95**, 13001 (2011).
- [71] R. Baer, *The Journal of Chemical Physics* **128**, 044103 (2008).
- [72] L. O. Wagner, Z.-h. Yang, and K. Burke, in *Fundamentals of Time-Dependent Density Functional Theory* (Springer-Verlag Berlin Heidelberg, 2012) pp. 101–123.
- [73] G. Vignale and W. Kohn, *Physical Review Letters* **77**, 2037 (1996).
- [74] X. Andrade, J. Alberdi-Rodriguez, D. A. Strubbe, M. J. T. Oliveira, F. Nogueira, A. Castro, J. Muguerza, A. Arruabarrena, S. G. Louie, A. Aspuru-Guzik, A. Rubio, and M. A. L. Marques, *Journal of Physics: Condensed Matter* **24**, 233202 (2012).
- [75] R. R. Pela and C. Draxl, *Electronic Structure* **3**, 037001 (2021).
- [76] N. Helbig, J. I. Fuks, M. Casula, M. J. Verstraete, M. A. L. Marques, I. V. Tokatly, and A. Rubio, *Physical Review A* **83**, 032503 (2011).
- [77] F. Calvayrac, P. G. Reinhard, E. Suraud, and C. A. Ullrich, *Physics Reports* **337**, 493 (2000).
- [78] M. Thiele, E. K. U. Gross, and S. Kümmel, *Physical Review Letters* **100**, 153004 (2008).
- [79] N. Maitra, *Lecture Notes in Physics* **706**, 61 (2006).
- [80] N. T. Maitra, K. Burke, and C. Woodward, *Physical Review Letters* **89**, 023002 (2002).
- [81] D. Tong, *Lecture notes (online): Linear response* (2012).
- [82] A. Savin, C. J. Umrigar, and X. Gonze, *Chemical Physics Letters* **288**, 391 (1998).
- [83] S. L. Adler, *Phys. Rev.* **126**, 413 (1962).
- [84] N. Wiser, *Phys. Rev.* **129**, 62 (1963).
- [85] H. Eshuis, J. E. Bates, and F. Furche, *Theoretical Chemistry Accounts* **131**, 1 (2012).
- [86] A. Damascelli, *Physica Scripta* **2004**, 61 (2004).
- [87] W. Zhang, K. H. Sze, C. E. Brion, X. M. Tong, and J. M. Li, *Chemical Physics* **140**, 265 (1990).

-
- [88] W. F. Chan, G. Cooper, X. Guo, and C. E. Brion, *Physical Review A* **45**, 1420 (1992).
- [89] T. Olsen, C. Patrick, J. Bates, A. Ruzsinszky, and K. S. Thygesen, *Computational Materials* **106**, 10.1038/s41524-019-0242-8 (2019).
- [90] A. Görling, *Physical Review B* **99**, 235120 (2019).
- [91] O. Gunnarsson and B. I. Lundqvist, *Physical Review B* **13**, 4274 (1976).
- [92] D. C. Langreth and J. P. Perdew, *Physical Review B* **15**, 2884 (1977).
- [93] D. C. Langreth and J. P. Perdew, *Solid State Communications* **17**, 1425 (1975).
- [94] X. Ren, P. Rinke, C. Joas, and M. Scheffler, *Journal of Materials Science* **47**, 7447 (2012), [arXiv:1203.5536](#) .
- [95] A. Heßelmann and A. Görling, *Molecular Physics* **109**, 2473 (2011).
- [96] D. C. Langreth and J. P. Perdew, *Physical Review B* **21**, 5469 (1980).
- [97] R. Kubo, *Reports on Progress in Physics* **29**, 255 (1966).
- [98] J. Dobson, G. Vignale, and D. Mukunda, eds., *Electronic Density Functional Theory: Recent Progress and New Directions* (Springer Science, 1998) pp. 243–285.
- [99] Y. M. Niquet, M. Fuchs, and X. Gonze, *Physical Review A - Atomic, Molecular, and Optical Physics* **68**, 13 (2003).
- [100] P. Bleiziffer, A. Heßelmann, and A. Görling, *Journal of Chemical Physics* **139**, 84113 (2013).
- [101] M. Hellgren, D. R. Rohr, and E. K. Gross, *Journal of Chemical Physics* **136**, 34106 (2012), [arXiv:1110.6062](#) .
- [102] M. Hellgren and U. Von Barth, *Journal of Chemical Physics* **132**, 34106 (2010).
- [103] P. Bleiziffer, M. Krug, and A. Görling, *Journal of Chemical Physics* **142**, 244108 (2015).
- [104] P. Verma and R. J. Bartlett, *Journal of Chemical Physics* **136**, 44105 (2012).
- [105] A. Thierbach and A. Görling, *Journal of Chemical Physics* **153**, 134113 (2020).
- [106] A. Görling, *Journal of Chemical Physics* **123**, 150901 (2005).
- [107] P. Bleiziffer, A. Heßelmann, and A. Görling, *Journal of Chemical Physics* **136**, 134102 (2012).
- [108] M. Lein and E. Gross, *Physical Review B - Condensed Matter and Materials Physics* **61**, 13431 (2000).
- [109] D. R. Hartree, *Mathematical Proceedings of the Cambridge Philosophical Society* **24**, 111132 (1928).
- [110] N. Marzari, *Ab-initio molecular dynamics for metallic systems*, Ph.D. thesis, University of Cambridge (1996).

- [111] N. Marzari, D. Vanderbilt, and M. C. Payne, *Physical Review Letters* **79**, 1337 (1997).
- [112] É. Cancès, G. Kemlin, and A. Levitt, <https://doi.org/10.1137/20M1332864> **42**, 243 (2021).
- [113] P.-L. Lions, *Communications in Mathematical Physics* **109**, 33 (1987).
- [114] V. Heine, *Solid State Physics* **24**, 1 (1970).
- [115] D. R. Bowler and T. Miyazaki, *Reports on Progress in Physics* **75**, 036503 (2012).
- [116] X. Gonze *et al.*, *Computer Physics Communications* **180**, 2582 (2009).
- [117] X. Gonze *et al.*, *Computational Materials Science* **25**, 478 (2002).
- [118] P. Giannozzi, *Journal of Physics: Condensed Matter* **21**, 395502 (2009).
- [119] M. J. van Setten, F. Caruso, S. Sharifzadeh, X. Ren, M. Scheffler, F. Liu, J. Lischner, L. Lin, J. R. Deslippe, S. G. Louie, C. Yang, F. Weigend, J. B. Neaton, F. Evers, and P. Rinke, *Journal of Chemical Theory and Computation* **11**, 5665 (2015).
- [120] K. Lejaeghere, V. Van Speybroeck, G. Van Oost, and S. Cottenier, *Critical Reviews in Solid State and Materials Sciences* **39**, 1 (2014).
- [121] J. Harris, *Physical Review B* **31**, 1770 (1985).
- [122] B. Farid, V. Heine, G. E. Engel, and I. J. Robertson, *Physical Review B* **48**, 11602 (1993).
- [123] E. Zaremba, *Journal of Physics: Condensed Matter* **2**, 2479 (1990).
- [124] E. Prodan, *Journal of Physics A: Mathematical and General* **38**, 5647 (2005).
- [125] H.-C. Kaiser, H. Neidhardt, and J. Rehberg, *Monatshefte für Mathematik* **158**, 179 (2009).
- [126] E. Cancès, R. Chakir, and Y. Maday, *Journal of Scientific Computing* **45**, 90 (2010).
- [127] E. Zeidler, *Nonlinear Functional Analysis and its Applications : I: Fixed-Point Theorems* (Springer-Verlag New York Inc, 2011).
- [128] Stefan Banach, *Fundamenta Mathematicae* **3**, 133 (1920).
- [129] A. Latif, in *Topics in Fixed Point Theory* (Springer International Publishing, Cham, 2014).
- [130] C. C. J. Roothaan, *Reviews of Modern Physics* **23**, 69 (1951).
- [131] E. Cancès, in *Mathematical Models and Methods for Ab Initio Quantum Chemistry. Lecture Notes in Chemistry*, Vol. 64 (Springer, Berlin, Heidelberg, 2000) pp. 17–43.
- [132] C. Yang, W. Gao, and J. C. Meza, *SIAM Journal on Matrix Analysis and Applications* **30**, 1773 (2009).

- [133] M. A. Krasnosel'skij, *Usp. Mat. Nauk* **10**, 123 (1955).
- [134] W. R. Mann, *Proceedings of the American Mathematical Society* **4**, 506 (1953).
- [135] F. E. Browder, *Archive for Rational Mechanics and Analysis* **24**, 82 (1967).
- [136] N. Mizoguchi and W. Takahashi, *Journal of Mathematical Analysis and Applications* **141**, 177 (1989).
- [137] J. Zhang, B. O'Donoghue, and S. Boyd, <https://doi.org/10.1137/18M1232772> **30**, 3170 (2020).
- [138] P. H. Dederichs and R. Zeller, *Physical Review B* **28** (1983).
- [139] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, 2014).
- [140] J. B. A. Davis, S. L. Horswell, and R. L. Johnston, *The Journal of Physical Chemistry A* **118**, 208 (2014).
- [141] E. Cancès and C. Le Bris, *International Journal of Quantum Chemistry* **79**, 82 (2000).
- [142] P. R. T. Schipper, O. V. Gritsenko, and E. J. Baerends, *Theoretical Chemistry Accounts: Theory, Computation, and Modeling (Theoretica Chimica Acta)* **99**, 329 (1998).
- [143] R. C. Morrison, *The Journal of Chemical Physics* **117**, 10506 (2002).
- [144] R. C. Morrison, *The Journal of Chemical Physics* **117**, 10506 (2002).
- [145] R. K. Nesbet, *Physical Review A* **56** (1997).
- [146] C. A. Ullrich and W. Kohn, *Physical Review Letters* **87** (2001).
- [147] N. D. Mermin, *Physical Review* **137**, A1441 (1965).
- [148] C. L. Fu and K. M. Ho, *Physical Review B* **28**, 5480 (1983).
- [149] M. Verstraete and X. Gonze, *Physical Review B* **65**, 035111 (2001).
- [150] K. Doll, N. M. Harrison, and V. R. Saunders, *Journal of Physics: Condensed Matter* **11**, 5007 (1999).
- [151] A. T. Fowler, C. J. Pickard, and J. A. Elliott, *Journal of Physics Materials* **2**, 34001 (2019).
- [152] C. J. Pickard and R. J. Needs, *Journal of Physics: Condensed Matter* **23**, 053201 (2011).
- [153] M. Lee, K. Leiter, C. Eisner, J. Crone, and J. Knap, *Computational and Theoretical Chemistry* **1062**, 24 (2015).
- [154] L. Lin and C. Yang, *SIAM Journal on Scientific Computing* **35**, S277 (2013).
- [155] N. Ashcroft and N. Mermin, *Solid State Physics* (Saunders College, Philadelphia, 1976).

-
- [156] J. Lindhard, *On the Properties of a Gas of Charged Particles* (E. Munksgaard, 1954).
- [157] N. Wisert, *Physical Review* **129**, 62 (1963).
- [158] P.-M. Anglade and X. Gonze, *Physical Review B* **78**, 045126 (2008).
- [159] K.-M. Ho, J. Ihm, and J. D. Joannopoulos, *Phys. Rev. B* **25**, 4260 (1982).
- [160] A. Sawamura and M. Kohyama, *Materials Transactions* **45**, 1422 (2004).
- [161] J. Auer and E. Krotscheck, *Computer Physics Communications* **151**, 265 (2003).
- [162] J. F. Annett, *Computational Materials Science* **4**, 23 (1995).
- [163] J. Nocedal and S. J. Wright, *Numerical Optimization* (Springer, 1999).
- [164] G. P. Kerker, *Phys. Rev. B* **23**, 3082 (1981).
- [165] M. Manninen, R. Nieminen, P. Hautojärvi, and J. Arponen, *Physical Review B* **12**, 4012 (1975).
- [166] J. Auer and E. Krotscheck, *Computer Physics Communications* **118**, 139 (1999).
- [167] E. H. Lieb, *Rev. Mod. Phys.* **53**, 603 (1981).
- [168] P. Ghosez, X. Gonze, and R. Godby, *Physical Review B - Condensed Matter and Materials Physics* **56**, 12811 (1997).
- [169] L. D. Marks, *Journal of Chemical Theory and Computation* **9**, 2786 (2013).
- [170] T. Rohwedder and R. Schneider, *Journal of Mathematical Chemistry* **49**, 1889 (2011).
- [171] Y. Saad, *Numerical Linear Algebra with Applications* **16**, 197 (2008).
- [172] L. D. Marks and D. R. Luke, *Physical Review B* **78**, 075114 (2008).
- [173] J. M. Borwein, G. Li, and M. K. Tam, *SIAM Journal on Optimization* **27**, 1 (2017).
- [174] L. Lin, Z. Xu, and L. Ying, <http://dx.doi.org/10.1137/16M1077325> **15**, 29 (2017).
- [175] X. Andrade, S. Botti, M. A. L. Marques, and A. Rubio, *The Journal of Chemical Physics* **126**, 184106 (2007).
- [176] C. G. Broyden, *Mathematics of Computation* **19**, 577 (1965).
- [177] J. E. Dennis and J. J. More, *SIAM Review* **19**, 46 (1977).
- [178] J. E. Dennis, Jr. and R. B. Schnabel, *SIAM Review* **21**, 443 (1979).
- [179] J. Sherman and W. Morrison, *Annals of Mathematical Statistics* **20**, 261 (1949).
- [180] G. P. Srivastava, *Journal of Physics A: Mathematical and General* **17**, L317 (1984).
- [181] V. Eyert, *Journal of Computational Physics* **124**, 271 (1996).
- [182] D. D. Johnson, *Physical Review B* **38**, 12807 (1988).

- [183] D. Vanderbilt and S. G. Louie, *Physical Review B* **30**, 6118 (1984).
- [184] P. Pulay, *Chemical Physics Letters* **73**, 393 (1980).
- [185] P. Pulay, *Journal of Computational Chemistry* **3**, 556 (1982).
- [186] D. G. Anderson, *Journal of the ACM* **12**, 4 (1965).
- [187] E. Artacho, E. Anglada, O. Diéguez, J. D. Gale, A. García, J. Junquera, R. M. Martin, P. Ordejón, J. M. Pruneda, D. Sánchez-Portal, and J. M. Soler, *Journal of Physics: Condensed Matter* **20**, 064208 (2008).
- [188] H. F. Walker and P. Ni, *SIAM Journal on Numerical Analysis* **49**, 1715 (2011).
- [189] D. G. M. Anderson, *Numerical Algorithms 2018 80:1* **80**, 135 (2018).
- [190] A. S. Banerjee, P. Suryanarayana, and J. E. Pask, *Chemical Physics Letters* **647**, 31 (2016).
- [191] P. P. Pratapa, P. Suryanarayana, and J. E. Pask, *Journal of Computational Physics* (2016).
- [192] P. P. Pratapa and P. Suryanarayana, *Chemical Physics Letters* **635**, 69 (2015).
- [193] D. R. Bowler and M. J. Gillan, *Chemical Physics Letters* **325**, 473 (2000).
- [194] P. Blaha *et al.*, *Technische Universität Wien, Wien* (2001).
- [195] J. M. Martinez, *Journal of Computational and Applied Mathematics* **124**, 97 (2000).
- [196] A. Tikhonov, *Journal of Physics A: Mathematical and General* **151**, 501 (1963).
- [197] C. T. Kelley, *Iterative Methods for Linear and Nonlinear Equations* (SIAM: Philadelphia, PA, 1995).
- [198] J. Wang, Y. Wang, and S. Yu, *Modelling and Simulation in Materials Science and Engineering* **16**, 035004 (2008).
- [199] M. Winkelmann, E. D. Napoli, D. Wortmann, and S. Blügel, *Physical Review B* **102**, 195138 (2020).
- [200] Y. Zhou, H. Wang, Y. Liu, X. Gao, and H. Song, *Physical Review E* **97**, 33305 (2018).
- [201] Z. H. Levine and S. G. Louie, *Physical Review B* **25**, 6310 (1982).
- [202] G. Cappellini, R. D. Sole, L. Reining, and F. Bechstedt, *Physical Review B* **47**, 9892 (1993).
- [203] N. Woods, (Preprint) (2018), arXiv:1803.01763 .
- [204] E. Krotscheck and M. Liebrecht, *The Journal of Chemical Physics* **138**, 164114 (2013).
- [205] R. M. Sternheimer, *Physical Review* **96**, 951 (1954).

- [206] F. Giustino, M. L. Cohen, and S. G. Louie, *Physical Review B* **81**, 115105 (2009).
- [207] D. Raczkowski, A. Canning, and L. W. Wang, *Physical Review B* **64**, 121101 (2001).
- [208] M. Rutter, *Computer Physics Communications* **225**, 174 (2018).
- [209] K. N. Kudin, G. E. Scuseria, and E. Cancès, *The Journal of Chemical Physics* **116**, 8255 (2002).
- [210] S. Høst, B. Jansík, J. Olsen, P. Jørgensen, S. Reine, and T. Helgaker, *Physical Chemistry Chemical Physics* **10**, 5344 (2008).
- [211] L. Thøgersen, J. Olsen, A. Köhn, P. Jørgensen, P. Salek, and T. Helgaker, *The Journal of Chemical Physics* **123**, 074103 (2005).
- [212] J. Barzilai and J. Borwein, *IMA Journal of Numerical Analysis* **8**, 141 (1988).
- [213] M. Lupo Pasini, *Numerical Linear Algebra with Applications* , 2241 (2019).
- [214] N. C. Henderson and R. Varadhan, <https://doi.org/10.1080/10618600.2019.1594835> **28**, 834 (2019).
- [215] M. T. Entwistle and R. W. Godby, *Physical Review B* **99**, 161102 (2019).
- [216] M. J. Hodgson, J. D. Ramsden, J. B. Chapman, P. Lillystone, and R. W. Godby, *Physical Review B - Condensed Matter and Materials Physics* **88**, 241102 (2013).
- [217] M. J. Hodgson, J. D. Ramsden, and R. W. Godby, *Physical Review B* **93**, 155146 (2016).
- [218] J. D. Ramsden and R. W. Godby, *Phys. Rev. Lett.* **109**, 036402 (2012).
- [219] H. N. Núñez-Yépez, A. L. Salas-Brito, and D. A. Solis, *Physical Review A - Atomic, Molecular, and Optical Physics* **83**, 064101 (2011).
- [220] T. E. Baker, E. M. Stoudenmire, L. O. Wagner, K. Burke, and S. R. White, *Physical Review B - Condensed Matter and Materials Physics* **91**, 235141 (2015), [arXiv:1504.05620](https://arxiv.org/abs/1504.05620) .
- [221] J. H. Lloyd-Williams, R. J. Needs, and G. J. Conduit, *Physical Review B - Condensed Matter and Materials Physics* **92**, 075106 (2015), [arXiv:1508.02004](https://arxiv.org/abs/1508.02004) .
- [222] J. Crank and P. Nicolson, *Mathematical Proceedings of the Cambridge Philosophical Society* **43**, 50–67 (1947).
- [223] A. H. Al-Mohy and N. J. Higham, *SIAM Journal on Matrix Analysis and Applications* **31**, 970 (2010), <https://doi.org/10.1137/09074721X> .
- [224] B. Kanungo, P. M. Zimmerman, and V. Gavini, *Nature Communications* **2019** 10:1 **10**, 1 (2019).
- [225] M. Ruggenthaler, K. J. H. Giesbertz, M. Penz, and R. van Leeuwen, *Physical Review A* **85**, 052504 (2012).

- [226] M. Ruggenthaler, M. Penz, and R. van Leeuwen, *Journal of Physics: Condensed Matter* **27**, 203202 (2015).
- [227] D. S. Jensen and A. Wasserman, *International Journal of Quantum Chemistry* **118**, e25425 (2018).
- [228] J. Brown, J. Yang, and J. D. Whitfield, *Journal of Chemical Theory and Computation* **16**, 6014 (2020).
- [229] D. S. Jensen and A. Wasserman, *Physical Chemistry Chemical Physics* **18**, 21079 (2016).
- [230] P. E. Lammert, *Physical Review A* **82**, 012109 (2010).
- [231] S. I. Kabanikhin, *Journal of Inverse and Ill-Posed Problems* **16**, 317 (2008).
- [232] P. Cvitanović, in *Chaos: Classical and Quantum* (Niels Bohr Inst., Copenhagen, 2016) Chap. 6.
- [233] K. Levenberg, *Quarterly of Applied Mathematics* **2**, 164 (1944).
- [234] M. J. D. Powell, *The Computer Journal* **7**, 155 (1964).
- [235] M. Thiele and S. Kümmel, *Physical Review A - Atomic, Molecular, and Optical Physics* **80**, 012514 (2009).
- [236] M. Thiele and S. Kümmel, *Physical Review Letters* **112**, 083001 (2014).
- [237] S. Botti, A. Schindlmayr, R. Del Sole, and L. Reining, *Reports on Progress in Physics* **70**, 357 (2007).
- [238] M. Marques and E. Gross, *Annual Review of Physical Chemistry* **55**, 427 (2004).
- [239] C. Ullrich and Z.-h. Yang, *Topics in Current Chemistry* **368**, 185 (2016).
- [240] G. Onida, L. Reining, and A. Rubio, *Reviews of Modern Physics* **74**, 601 (2002).
- [241] V. I. Gavrilenko and F. Bechstedt, *Phys. Rev. B* **55**, 4343 (1997).
- [242] M. Corradini, R. Del Sole, G. Onida, and M. Palumbo, *Phys. Rev. B* **57**, 14569 (1998).
- [243] S. Moroni, D. M. Ceperley, and G. Senatore, *Phys. Rev. Lett.* **75**, 689 (1995).
- [244] W. Kohn, *Physical Review Letters* **76**, 3168 (1996).
- [245] N. T. Maitra, *Journal of Physics Condensed Matter* **29**, 423001 (2017).
- [246] N. T. Maitra, F. Zhang, R. J. Cave, and K. Burke, *The Journal of Chemical Physics* **120**, 5932 (2004).
- [247] P. Elliott, S. Goldson, C. Canahui, and N. T. Maitra, *Chemical Physics* **391**, 110 (2011), [arXiv:1101.3379](https://arxiv.org/abs/1101.3379) .
- [248] R. J. Cave, F. Zhang, N. T. Maitra, and K. Burke, *Chemical Physics Letters* **389**, 39 (2004).

- [249] M. E. Casida, *Recent Advances in Density Functional Methods* (World Scientific Publishing, 1995) pp. 155–192.
- [250] B. Hudson and B. Kohler, *Synthetic Metals* **9**, 241 (1984).
- [251] L. Reining, V. Olevano, A. Rubio, A. Rubio, and G. Onida, *Physical Review Letters* **88**, 4 (2002).
- [252] A. Marini, R. Del Sole, and A. Rubio, *Physical Review Letters* **91**, 256402 (2003), [arXiv:0310495 \[cond-mat\]](#) .
- [253] P. Romaniello, D. Sangalli, J. A. Berger, F. Sottile, L. G. Molinari, L. Reining, and G. Onida, *Journal of Chemical Physics* **130**, 044108 (2009).
- [254] M. Hellgren and U. Von Barth, *Journal of Chemical Physics* **131**, 044110 (2009).
- [255] M. Hellgren and U. Von Barth, *Physical Review B - Condensed Matter and Materials Physics* **78**, 115107 (2008).
- [256] Y.-H. Kim and A. Görling, *Phys. Rev. Lett.* **89**, 096402 (2002).
- [257] A. Görling, *Phys. Rev. Lett.* **83**, 5459 (1999).
- [258] A. Görling, *Physical Review A - Atomic, Molecular, and Optical Physics* **57**, 3433 (1998).
- [259] R. Del Sole, G. Adragna, V. Olevano, and L. Reining, *Physical Review B - Condensed Matter and Materials Physics* **67**, 045207 (2003).
- [260] S. Botti, F. Sottile, N. Vast, V. Olevano, L. Reining, H.-C. Weissker, A. Rubio, G. Onida, R. Del Sole, and R. W. Godby, *Phys. Rev. B* **69**, 155112 (2004).
- [261] M. Ruggenthaler, S. E. Nielsen, and R. Van Leeuwen, *Physical Review A - Atomic, Molecular, and Optical Physics* **88**, 022512 (2013).
- [262] N. T. Maitra and D. G. Tempel, *Journal of Chemical Physics* **125**, 184111 (2006).
- [263] O. V. Gritsenko and E. Jan Baerends, *Physical Chemistry Chemical Physics* **11**, 4640 (2009).
- [264] F. Aryasetiawan and O. Gunnarsson, *Phys. Rev. B* **66**, 165119 (2002).
- [265] D. J. Carrascal, J. Ferrer, N. Maitra, and K. Burke, *European Physical Journal B* **91**, 142 (2018), [arXiv:1802.09988](#) .
- [266] V. Turkowski and T. S. Rahman, *Journal of Physics Condensed Matter* **26**, 6 (2014), [arXiv:1210.6397](#) .
- [267] J. I. Fuks and N. T. Maitra, *Phys. Chem. Chem. Phys.* **16**, 14504 (2014).
- [268] A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation* (Society for Industrial and Applied Mathematics, 2005).
- [269] R. Van Leeuwen, *International Journal of Modern Physics B* **15**, 1969 (2001).
- [270] D. Mearns and W. Kohn, *Physical Review A* **35**, 4796 (1987).

- [271] Y.-M. Byun, J. Sun, and C. A. Ullrich, *Electronic Structure* **2**, 023002 (2020).
- [272] G. T. N. E. Ben-Israel A., *Generalized Inverses* (Springer-Verlag, 2003).
- [273] R. W. Godby and R. J. Needs, *Physical Review Letters* **62**, 1169 (1989).
- [274] H. Appel, E. K. Gross, and K. Burke, *Physical Review Letters* **90**, 4 (2003), [arXiv:0203027 \[cond-mat\]](#) .
- [275] A. I. Al-Sharif, R. Resta, and C. J. Umrigar, *Physical Review A - Atomic, Molecular, and Optical Physics* **57**, 2466 (1998).
- [276] J. P. Perdew, R. G. Parr, M. Levy, and J. L. Balduz, *Phys. Rev. Lett.* **49**, 1691 (1982).
- [277] M. T. Entwistle, M. Casula, and R. W. Godby, *Phys. Rev. B* **97**, 235143 (2018).
- [278] T. Olsen and K. S. Thygesen, *Physical Review Letters* **112**, 203001 (2014), [arXiv:1405.0745](#) .
- [279] M. A. L. Marques and E. K. U. Gross, in *A Primer in Density Functional Theory* (Springer, Berlin, Heidelberg, 2003) pp. 144–184.
- [280] M. A. Marques, A. Castro, and A. Rubio, *Journal of Chemical Physics* **115**, 3006 (2001).
- [281] M. Petersilka, E. K. U. Gross, and K. Burke, *International Journal of Quantum Chemistry* **80**, 534 (2000).
- [282] A. Liebsch, *Physical Review B* **32**, 6255 (1985).
- [283] M. A. Marques, C. A. Ullrich, F. Nogueira, A. Rubio, K. Burke, and E. K. U. Gross, eds., *Time-Dependent Density Functional Theory*, Lecture Notes in Physics, Vol. 706 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2006) pp. 443–462.
- [284] K. Burke, ed., *ABC of DFT* (Unpublished, 2007) pp. 93–111.
- [285] L. A. Constantin, *Physical Review B* **99**, 085117 (2019).
- [286] J. P. Perdew, A. Ruzsinszky, J. Tao, V. N. Staroverov, G. E. Scuseria, and G. I. Csonka, *Journal of Chemical Physics* **123**, 6158 (2005).
- [287] F. Furche, *Physical Review B - Condensed Matter and Materials Physics* **64**, 195120 (2001).
- [288] M. Fuchs and X. Gonze, *Physical Review B - Condensed Matter and Materials Physics* **65**, 1 (2002).
- [289] H. V. Nguyen and S. De Gironcoli, *Physical Review B - Condensed Matter and Materials Physics* **79**, 205114 (2009).
- [290] Z. Yan, J. P. Perdew, and S. Kurth, *Physical Review B - Condensed Matter and Materials Physics* **61**, 16430 (2000).
- [291] E. Engel, A. Höck, and R. M. Dreizler, *Physical Review A - Atomic, Molecular, and Optical Physics* **61**, 5 (2000).

- [292] J. P. Perdew, *International Journal of Quantum Chemistry* **48**, 93 (1993).
- [293] J. Harl, L. Schimka, and G. Kresse, *Physical Review B - Condensed Matter and Materials Physics* **81**, 115126 (2010).
- [294] D. L. Freeman, *Physical Review B* **15**, 5512 (1977).
- [295] J. P. Perdew and Y. Wang, *Physical Review B* **45**, 13244 (1992).
- [296] S. Ruan, X. Ren, T. Gould, and A. Ruzsinszky, *Journal of Chemical Theory and Computation* **17**, 2107 (2021).
- [297] A. Görling, *Physical Review A - Atomic, Molecular, and Optical Physics* **57**, 3433 (1998).
- [298] Y. H. Kim and A. Görling, *Physical Review B - Condensed Matter and Materials Physics* **66**, 1 (2002).
- [299] A. Heßelmann and A. Görling, *Molecular Physics* **108**, 359 (2010).
- [300] A. Heßelmann and A. Görling, *Physical Review Letters* **106**, 093001 (2011).
- [301] N. Colonna, M. Hellgren, and S. De Gironcoli, *Physical Review B - Condensed Matter and Materials Physics* **90**, 125150 (2014).
- [302] J. Toulouse, I. C. Gerber, G. Jansen, A. Savin, and J. G. Ángyán, *Physical Review Letters* **102**, 096404 (2009), arXiv:0812.3302 .
- [303] B. G. Janesko, T. M. Henderson, and G. E. Scuseria, *Journal of Chemical Physics* **130**, 94103 (2009).
- [304] W. Zhu, J. Toulouse, A. Savin, and J. G. Ángyán, *Journal of Chemical Physics* **132**, 84119 (2010).
- [305] J. Toulouse, W. Zhu, J. G. Ángyán, and A. Savin, *Physical Review A - Atomic, Molecular, and Optical Physics* **82**, 032502 (2010), arXiv:1006.2061 .
- [306] B. G. Janesko, T. M. Henderson, and G. E. Scuseria, *Journal of Chemical Physics* **131**, 114105 (2009).
- [307] B. G. Janesko and G. E. Scuseria, *Journal of Chemical Physics* **131**, 154106 (2009).
- [308] J. Toulouse, W. Zhu, A. Savin, G. Jansen, and J. G. Ángyán, *Journal of Chemical Physics* **135**, 84119 (2011).
- [309] J. G. Ángyán, R. F. Liu, J. Toulouse, and G. Jansen, *Journal of Chemical Theory and Computation* **7**, 3116 (2011), arXiv:1404.1663 .
- [310] X. Gonze and M. Scheffler, *Physical Review Letters* **82**, 4416 (1999).
- [311] M. Hellgren and U. Von Barth, *Physical Review B - Condensed Matter and Materials Physics* **76**, 075107 (2007), arXiv:0703819 [cond-mat] .
- [312] Dask Development Team, *Dask: Library for dynamic task scheduling* (2016).

-
- [313] A. M. Teale, S. Coriani, and T. Helgaker, *Journal of Chemical Physics* **130**, 104111 (2009).
- [314] J. Dobson, G. Vignale, and D. Mukunda, eds., *Electronic Density Functional Theory: Recent Progress and New Directions* (Springer Science, 1998) pp. 3–27.
- [315] M. J. P. Hodgson, *arXiv preprint arXiv:2104.04588* (2021).
- [316] C. O. Almbladh and U. von Barth, in *Density Functional Methods in Physics*, edited by R. M. Dreizler and J. da Providencia (Plenum Press, New York and London, 1985) p. 209.
- [317] M. Hellgren and E. K. U. Gross, *Physical Review A* **85**, 022514 (2012).
- [318] M. Hellgren and E. K. Gross, *Physical Review A - Atomic, Molecular, and Optical Physics* **88**, 052507 (2013).
- [319] M. Hellgren, *The European Physical Journal B* 2018 91:7 **91**, 1 (2018).
- [320] N. T. Maitra, (Preprint) (2021), *arXiv:2107.05600* .