

This is a repository copy of *Computing the self-consistent field in Kohn-Sham density functional theory*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/185647/>

Version: Published Version

Article:

Woods, N. D., Payne, M. C. and Hasnip, P. J. orcid.org/0000-0002-4314-4093 (2019) Computing the self-consistent field in Kohn-Sham density functional theory. *Journal of Physics Condensed Matter*. 453001. ISSN: 1361-648X

<https://doi.org/10.1088/1361-648X/ab31c0>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

TOPICAL REVIEW • OPEN ACCESS

Computing the self-consistent field in Kohn–Sham density functional theory

To cite this article: N D Woods *et al* 2019 *J. Phys.: Condens. Matter* **31** 453001

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

You may also like

- [Accurate effective potential for density amplitude and the corresponding Kohn–Sham exchange–correlation potential calculated from approximate wavefunctions](#)
Ashish Kumar, Rabeet Singh and Manoj K Harbola
- [Time-dependent density-functional theory for periodic solids: assessment of excitonic exchange–correlation kernels](#)
Young-Moo Byun, Jiuyu Sun and Carsten A Ullrich
- [Electronic structure of CuWO₄: dielectric-dependent, self-consistent hybrid functional study of a Mott–Hubbard type insulator](#)
Ho Viet Thang, Elisa Albanese and Gianfranco Pacchioni



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Topical Review

Computing the self-consistent field in Kohn–Sham density functional theory

N D Woods^{1,3}, M C Payne¹ and P J Hasnip²

¹ Theory of Condensed Matter, Cavendish Laboratory, University of Cambridge, Cambridge, CB3 0HE, United Kingdom

² Department of Physics, University of York, Heslington, York YO10 5DD, United Kingdom

E-mail: nw361@cam.ac.uk

Received 25 May 2018, revised 5 July 2019

Accepted for publication 12 July 2019

Published 14 August 2019



Abstract

A new framework is presented for evaluating the performance of self-consistent field methods in Kohn–Sham density functional theory (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 in the present context, 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.

Keywords: self-consistent field methods, density mixing, density functional theory, non-linear optimisation, numerical analysis

(Some figures may appear in colour only in the online journal)

Contents

1. Preface	2	3.6. The Aufbau principle and fractional occupancy	8
2. Introduction	2	3.7. The initial guess	10
2.1. Background	2	3.8. Ill-conditioning and charge sloshing	11
2.2. Review purpose and structure	4	4. Methods and algorithms	14
3. Self-consistency in Kohn–Sham theory	4	4.1. The quasi-Newton update	15
3.1. Computational implementation	4	4.1.1. Broyden’s methods.	15
3.2. The Kohn–Sham map	5	4.1.2. Multisecant Broyden’s methods.	16
3.3. Defining convergence	5	4.1.3. Pulay’s method.	17
3.4. Some unique properties of K	6	4.1.4. Modern multisecant-based algorithms.	19
3.5. Fixed-point and damped iterations	6	4.2. Density matrix optimisation	21
		4.3. Preconditioning	22
		4.4. Direct minimisation	25
		5. Test suite	26
		6. Results and discussion	26
		7. Conclusion	29
		Acknowledgments	30
		References	30

³ Author to whom any correspondence should be addressed.



Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1. Preface

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 discipline. In the present context, we are concerned with the hierarchy of methods that has emerged for calculating the properties of molecular and solid state systems by approximating the Schrödinger equation [1–3]. In particular, the most prominent method from this hierarchy over the past few decades has proven to be density functional theory (DFT) within the Kohn–Sham framework [4, 5]. For a variety 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 [6–8]. It is, therefore, of paramount importance that implementations of Kohn–Sham DFT optimally utilise the available computational resources.

Many distinct implementations to Kohn–Sham theory exist, differing according to the choice of basis set, whether to use a density matrix or explicit wavefunction formulation etc, each with advantages and disadvantages in the computational domain [2, 9–18]. When one has decided on such an approach, its effectiveness is limited by the efficiency and reliability of the available numerical algorithms. This work reviews an aspect of Kohn–Sham theory that is more-or-less universal across many of these approaches; that is, how one iterates a density towards so-called *self-consistency*. This is conventionally referred to as the self-consistent field procedure, and is the most common source of numerical divergence when solving the equations of Kohn–Sham theory *in silico* [19]. This work examines the effectiveness of the methods and algorithms used in the self-consistent field procedure, reviewing a wide range of available methods drawn from the literature, studying the causes of divergences and inefficiencies and exploring how the available algorithms mitigate these potential issues. In order to assess the performance of the algorithms, a test suite is presented comprising a wide range of representative simulations. This test suite allows the algorithms to be judged according to both their robustness (ability to find a solution to the Kohn–Sham equations) and efficiency (speed with which a given solution is found) in a transparent and unbiased manner. The test suite and the associated workflow constitute a powerful new framework for the development, testing and assessment of new methods and algorithms. Throughout this work care has been taken to present the wide range of different methods in a consistent way, such that the similarities and differences of the methods are readily apparent.

2. Introduction

2.1. 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) [20].

Later refined by Fock [21] and Slater [22], Hartree–Fock theory became widely adopted in computational quantum chemistry to compute ground state properties of molecules [3]. Whilst Hartree and Hartree–Fock theory are mean-field approximations, Hohenberg *et al* [5, 23] showed that a mean-field exists which reproduces the ground-state energy and particle density *exactly*. This ‘density functional theory’ allows, in principle, the computation of the exact electronic structure of any quantum system; however the exact density functional is not known, and must be approximated in any practical application of DFT. For a more detailed examination of the origins and physical foundations of Kohn–Sham theory, the reader is directed to the following resources [1, 4, 24], and references therein.

This work 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 $\rho(x)$ 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 article, finding the ground state energy E in Kohn–Sham theory is viewed as a constrained minimisation problem,

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

$$E_{\text{KS}}[\{\phi_i\}] = \sum_{i=1}^N \frac{1}{2} \int_{\mathbb{R}^3} |\nabla \phi_i(x)|^2 + \frac{1}{2} \int_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho(x) \rho(x')}{|x - x'|} \\ + \int_{\mathbb{R}^3} \rho(x) v_{\text{ext}}(x) + E_{\text{xc}}[\rho], \quad (2)$$

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

$$\rho(x) = \sum_{i=1}^N |\phi_i(x)|^2. \quad (3)$$

That is, one must minimise the Kohn–Sham objective functional equation (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 $\mathcal{H}^1(\mathbb{R}^3)$. The *exchange-correlation functional* E_{xc} is a yet undetermined functional of the density designed to capture the effects of exchange and correlation missing from the remainder of the functional. In principle, the Hohenberg–Kohn theorems guarantee that the Kohn–Sham objective functional is a functional of the density alone [23]. However, in the case of Kohn–Sham theory, recourse to an orbital-dependent functional is necessitated by the definition of the single-particle kinetic energy.

Explicit constrained variation of the orbitals allows one to approach the optimisation problem in equation (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 [25, 26]. Alternatively, assuming differentiability [27], 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 a *nonlinear eigenvalue problem*,

$$H_{KS}[\rho]\phi_i(x) = \epsilon_i\phi_i(x), \quad (4)$$

where the Hamiltonian operator H_{KS} depends on its eigenvectors via

$$H_{KS}[\rho] = -\frac{1}{2}\nabla^2 + v_{ext} + v_h[\rho] + v_{xc}[\rho], \quad (5)$$

$$v_h[\rho](x) = \int_{\mathbb{R}^3} \frac{\rho(x')}{|x - x'|}, \quad (6)$$

$$v_{xc}[\rho] = \frac{\delta E_{xc}}{\delta \rho}. \quad (7)$$

These are 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 [28]. The advantages and drawbacks of each approach will be examined in section 4.




Figure 1. An iterative algorithm generates a series of perturbations to the density $\delta\rho_n$ in order to converge the initial guess density (top) to the fixed-point density ρ_* (bottom). Example for an *fcc* four atom aluminium unit cell.

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 ρ^{in} to compute the Hartree and exchange-correlation potentials. An output density ρ^{out} is then calculated (non-linearly) from the eigenfunctions of the Kohn–Sham Hamiltonian

$$H_{KS}[\rho^{in}]\phi_i(x) = \epsilon_i\phi_i(x), \quad (8)$$

$$\rho^{out}(x) = \sum_{i=1}^N |\phi_i(x)|^2. \quad (9)$$

In general, the input density is not equal to the output density. For a given external potential and exchange-correlation functional, the density ρ_* is *self-consistent* when $\rho_* = \rho^{in} = \rho^{out}$, and hence the non-linear eigenvalue problem of equations (8) and (9) is solved. The non-linearity in equation (9) 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*, figure 1. As one might expect, an infinity of self-consistent densities exist for a given external potential and exchange-correlation functional [29]. However, we are interested primarily in the subset of these densities that are local minima of the Kohn–Sham objective functional.

Modern computational implementations of Kohn–Sham theory can vary significantly due to various factors. The key distinguishing factor is the choice of basis set, which leads to the related problem of whether one treats all the electrons in the computation explicitly, or treats core electrons with a

pseudopotential [30]. Despite these differences, perhaps with the exception of linear scaling methods [14], the self-consistent field techniques to be discussed here are adaptable to most implementations. Indeed, some of the most popular software, such as VASP [15, 31], ABINIT [10, 32], QUANTUM ESPRESSO [33], and CASTEP [34], use similar default methods to achieve self-consistency: preconditioned multiseant methods, which are discussed in section 4.

2.2. Review purpose and structure

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 article is structured in two partitions, as follows.

The first part constitutes a review of self-consistency in Kohn–Sham theory. As such, the relevant sections are ideal for an interested party who is not actively involved in development to gain a more in-depth understanding of self-consistency from an algorithmic perspective. In particular, this review collates decades of past literature on self-consistency in Kohn–Sham theory, thus elucidating conclusions that have become conventional wisdom. Section 3 examines the Kohn–Sham framework abstractly from a mathematical and computational perspective in order to 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 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.

The second part then utilises the analysis presented in the prior sections to perform a study akin to recent benchmarking efforts such as GW100 [35] and the Δ -project [36] 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 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. 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 6 demonstrates these concepts by using this workflow on a selection of algorithms described in section 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. 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. No claim is made for providing a strictly detailed and rigorous treatment of the mathematical problem at hand. Instead, literature is cited throughout such that the interested reader can venture further in detail than this article provides.

3.1. Computational implementation

The central approximation involved in converting the framework of Kohn–Sham theory into a form suitable for computation is called the *finite-basis approximation*, or the *Galerkin approximation* [28]. The orbitals ϕ_i are continuous functions of a continuous three dimensional variable, x . These functions are equivalent to vectors existing in an infinite dimensional vector space, spanned by a complete basis χ_μ . Provided this basis does indeed span the space, the orbitals can be expressed exactly as

$$\phi_i(x) = \sum_{\mu=1}^{\infty} \alpha_{i\mu} \chi_\mu(x). \quad (10)$$

Note that the basis set is required to be complete, but not necessarily orthogonal,

$$\int_{\mathbb{R}^3} \chi_\mu(x) \chi_\nu(x) = S_{\mu\nu} = \delta_{\mu\nu}. \quad (11)$$

If the basis is non-orthogonal, the finite-basis Kohn–Sham equations generalise to include the overlap matrix, $S_{\mu\nu} \neq \delta_{\mu\nu}$, see [28] for more details. This distinction does not affect the ensuing analysis, and so the basis set is hereafter taken to be

orthogonal for clarity. Once the basis is specified, the equations can be rearranged and solved for the infinity of coefficients to 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 formally infinite Hilbert space,

$$\phi_i(x) \approx \sum_{\mu=1}^{N_b} \alpha_{i\mu} \chi_{\mu}(x). \quad (12)$$

The characteristic size of the basis N_b will depend primarily on the choice of basis functions. Within the finite-basis approximation, the Kohn–Sham Hamiltonian becomes an $N_b \times N_b$ matrix, of which a subset of the eigenvalues and eigenvectors is required to progress toward a solution of the non-linear eigenvalue problem, equations (8) and (9). Basis functions which are *localised* about the atomic cores [37] are a popular choice. These tend to be particularly accurate per basis function, meaning N_b is typically the same order of magnitude as the number of electrons, N . Methods utilising local basis functions are often able to form and diagonalise the Kohn–Sham Hamiltonian matrix explicitly. In such implementations, the Kohn–Sham energy functional can be expressed in terms of the *density matrix*,

$$D(x, x') = \sum_{i=1}^N \phi_i^*(x) \phi_i(x') \\ = \sum_{\mu\nu=0}^{N_b} D_{\mu\nu} \chi_{\mu}^*(x) \chi_{\nu}(x'), \quad (13)$$

$$\rho(x) = D(x, x), \quad (14)$$

rather than the orbitals, where the density matrix is also of dimension $N_b \times N_b$. The Kohn–Sham energy functional has a closed-form expression in terms of the density matrix (see [28] and section 4.2), and therefore the constrained optimisation in equation (1) becomes an optimisation over allowed variations in the density matrix. From the point of view of the work to follow, the ability to construct, store, and optimise the density matrix directly is the distinguishing characteristic of localised basis sets with respect to the basis set considered in the following work: namely, the set of N_b plane-waves,

$$\chi_G(x) = e^{iGx}, \quad (15)$$

with the same periodicity as the unit cell [15, 31], labelled by the frequency of the plane-wave G . This basis set is *delocalised*, meaning the functions χ_G are non-zero across the whole unit cell. The introduction of a delocalised basis results in a reduction in accuracy per basis function, which in turn necessitates a much larger value of N_b to reproduce the same accuracy as a computation using localised basis sets. The advantage of a plane-wave, or similar, basis set lies elsewhere [15, 31]. This will become relevant in section 4, as certain algorithms exploit the ability to construct $D(x, x')$ explicitly. Nevertheless, much of the analysis to follow in this section will remain largely independent of basis set. The discussion will, however, be framed in the language of an entirely plane-wave basis set.

3.2. The Kohn–Sham map

As already stated, Kohn–Sham theory is a constrained optimisation problem, equation (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 ρ_* such that it is a *fixed-point* of the discretised Kohn–Sham map,

$$K : \mathbb{R}^{N_b} \rightarrow \mathbb{R}^{N_b}, \quad (16)$$

$$K[\rho_*] = \rho_*. \quad (17)$$

In general, $K[\rho^{\text{in}}] = \rho^{\text{out}}$, where K is defined using equations (8) and (9). 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 equation (9) for a given Kohn–Sham Hamiltonian. The ‘size’ of this set, as a subset of \mathbb{R}^{N_b} , is an open problem [27]. 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 practice⁴. The aim now is to generate a converging sequence of densities $\{\rho_0^{\text{in}}, \rho_1^{\text{in}}, \dots, \rho_n^{\text{in}}\}$ starting from an *initial guess* density ρ_0^{in} , where $\rho_n^{\text{in}} \approx \rho_*$ 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. Defining convergence

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

$$R[\rho_*] = K[\rho_*] - \rho_* = 0, \quad (18)$$

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 $\|R[\rho^{\text{in}}]\|_2$, 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. 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 $E_{\text{KS}}[\rho_n^{\text{in}}]$. Instead, the *Harris–Foulkes* functional E_{KS} is defined [38],

⁴ This observation is based on the fact that, in general, one can always find an algorithm that converges to a fixed-point density.

$$\begin{aligned}
\tilde{E}_{\text{KS}} = & \sum_{i=1}^N \frac{1}{2} \int_{\mathbb{R}^3} |\nabla \phi_i|^2 + \frac{1}{2} \int_{\mathbb{R}^3} \rho^{\text{out}}(x) v_h[\rho^{\text{in}}] \\
& + \frac{1}{2} \int_{\mathbb{R}^3} (\rho^{\text{out}}(x) - \rho^{\text{in}}(x)) v_h[\rho^{\text{in}}] \\
& + \int_{\mathbb{R}^3} (\rho^{\text{out}}(x) - \rho^{\text{in}}(x)) v_{\text{xc}}[\rho^{\text{in}}] + E_{\text{xc}}[\rho^{\text{in}}] \\
& + \int_{\mathbb{R}^3} \rho^{\text{out}}(x) v_{\text{ext}}(x), \quad (19)
\end{aligned}$$

which can be shown to give the exact ground state energy correct to quadratic order in the density error about the fixed-point density ρ_* —i.e. it is correct to $\mathcal{O}((\rho_* - \delta\rho)^2)$. Note that it is not the Harris-Foulkes functional that is minimised during the computation, as it possesses incorrect behaviour away from ρ_* [39, 40]. However, evaluating the energy using this functional *when near* ρ_* allows one to terminate the iterations at a desired accuracy earlier than if one evaluates the energy using the Kohn–Sham functional, which is correct to linear order 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 4 is based solely on fluctuations in the total energy.

3.4. Some unique properties of K

Identifying properties unique to K can help guide and narrow the choice of algorithms in section 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 $\{\rho_i^{\text{in}}, \rho_i^{\text{out}}\}$ on the i th iteration. This is because, when one has specified ρ^{in} , finding the corresponding ρ^{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 [15]. 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 $\{(\rho_i^{\text{in}}, \rho_i^{\text{out}}) \mid i \in [0, n]\}$. In order to utilise this set to generate the subsequent density ρ_{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 scales with N_b^5 , 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, ρ_0^{in} . A discussion on the generation of the initial guess is left to later in this

⁵ For example, in plane-wave codes, the density is represented by an array of size $8N_b$, see [2] for more details.

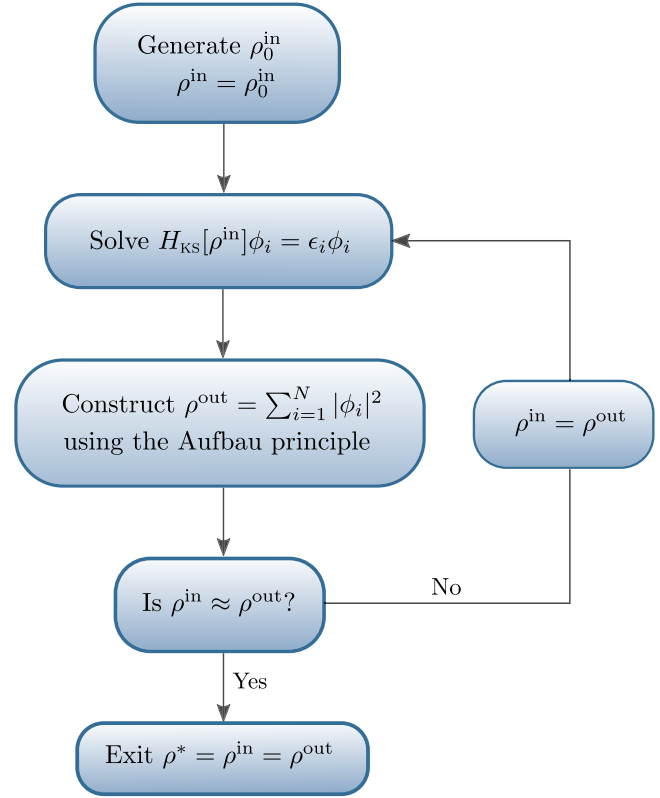


Figure 2. A flowchart detailing an example algorithm for achieving self-consistency using fixed-point (or Roothaan) iterations.

section, but it suffices to note that the initial guess is typically ‘close’ to the converged density ρ_* . By ‘close’ we mean that a *linear response approximation* can be employed effectively, see section 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.

3.5. 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 [41–43], empirical wisdom allows us to make certain broad statements about it. For the sake of analysis, we now introduce the *fixed-point iteration*,

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{out}} = K[\rho_n^{\text{in}}]. \quad (20)$$

This is perhaps the most simple iterative scheme one could envisage, yet it remains profoundly important from the point of view of functional analysis [44]. An example algorithm that makes use of the fixed-point iteration scheme is given in figure 2. This algorithm, on iteration n , constructs and diagonalises the Kohn–Sham Hamiltonian for a given ρ_n^{in} , and computes the output density ρ_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




Figure 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.

as one sets $\rho_{n+1}^{\text{in}} = \rho_n^{\text{out}}$, and the procedure is repeated. For the algorithm in figure 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

$$\|K(\rho_1) - K(\rho_2)\|_2 \leq k \|\rho_1 - \rho_2\|_2 \quad (21)$$

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 equation (21) 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, ρ_* . (See the Banach fixed-point theorem [45] or its generalisations [46] for k -contractive maps). Unfortunately, as section 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 [47] level of theory, figure 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 [48]. 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 [49]. It is expected that this behaviour will carry over to Kohn–Sham theory [50].

We now define a new iterative scheme, the damped iteration (or one of its many other aliases, such as Krasnosel’skiĭ–Mann or averaged iteration [51, 52]) such that

$$\begin{aligned} \rho_{n+1}^{\text{in}} &= \rho_n^{\text{in}} + \alpha(K[\rho_n^{\text{in}}] - \rho_n^{\text{in}}) \\ &= \rho_n^{\text{in}} + \alpha R[\rho_n^{\text{in}}]. \end{aligned} \quad (22)$$

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 [53–55]. Here, non-expansive refers to an instance whereby $k = 1$ in equation (21), 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 4 (e.g. [56]). 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 [50], albeit often impractically slowly. Fortunately, rather large damping parameters of $\alpha \sim 0.5$ are sometimes able to significantly improve convergence, as demonstrated in figure 4 [57]. 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 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 3.7.

The behaviour of K discussed here could be interpreted as arising due to the lack of convexity of the underlying




Figure 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.

functional E_{KS} used to generate it. Convexity is defined formally in section 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 [58]. 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 [59]. 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.

3.6. 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, [29] 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 figure 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 [43, 49, 60]. 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 [28]. 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

$$E_{\text{gap}} = \epsilon_{N+1} - \epsilon_N, \quad (23)$$

otherwise known as the HOMO–LUMO gap—the difference in energy between the highest energy occupied and lowest energy unoccupied (molecular) orbitals. When this gap disappears, meaning there exists a non-zero density of states at the Fermi energy, convergence becomes increasingly difficult [25, 26]. 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 [27, 61–63]. For example, [61] 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

$$D = \sum_{i=1}^q \lambda_i \Psi_i^*(x_1, \dots, x_N) \Psi_i(x'_1, \dots, x'_N) \quad (24)$$

via

$$\rho(x) = \text{Tr} D \hat{\rho}(x). \quad (25)$$

The wavefunctions $\{\Psi_i\}$ 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

$$\rho(x) = \sum_{\epsilon_i < \mu} |\phi_i(x)|^2 + \sum_{\epsilon_i = \mu} f_i |\phi_i(x)|^2, \quad (26)$$

where the *fractional occupancies* f_i are determined as some combination of the weights λ_i in equation (24). 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 [61], 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 equation (3), but instead must be constructed from some ensemble of states via equations (24) and (25). The extension of Kohn–Sham theory to include fractional occupancy is described well in [64, 65].

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 [66]. 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 [25, 64], and references therein. Here, it suffices

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

$$E[\{\phi_i\}, \{f_i\}, T] = \sum_{i=1}^{\infty} \frac{1}{2} f_i \int_{\mathbb{R}^3} |\nabla \phi_i|^2 + \frac{1}{2} \int_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{\rho(x)\rho(x')}{|x-x'|} + \int_{\mathbb{R}^3} \rho(x) v_{\text{ext}}(x) + E_{\text{xc}}[\rho] - TS[\{f_i\}], \quad (27)$$

where the entropy functional and density are defined respectively as

$$S = \sum_{i=1}^{\infty} f_i \ln(f_i) + (1-f_i) \ln(1-f_i), \quad (28)$$

$$\rho(x) = \sum_{i=1}^{\infty} f_i |\phi_i(x)|^2. \quad (29)$$

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 [26], which is discussed in section 4 and tested in section 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,

$$f_i = \frac{1}{e^{(\epsilon_i - \mu)/T} + 1}. \quad (30)$$

The electronic temperature T is now an input parameter which determines the degree of broadening of occupancies about the Fermi energy, figure 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 [67, 68]⁶. Furthermore, introducing finite temperature also assists with 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 [67, 69] 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 5 includes many such

⁶ Note that it is possible to approximately recover the zero-temperature solution [65].




Figure 5. Functional dependence of the occupancy f of a given eigenenergy ϵ for various temperatures using the Fermi-Dirac smearing scheme.

systems, and in particular a variety of electronic temperatures are considered.

3.7. 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 [70]. 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* [15, 31]. 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, figure 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 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 [71]. The test suite includes various examples of these ‘far-from-equilibrium’ systems.

Spin polarised Kohn–Sham theory presents more serious issues: there is no widely successful method for generating




Figure 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.

initial guess spin densities. In the spin polarised or ‘unrestricted’ formalism, the following *spin densities* are introduced (see [2]),

$$\rho^\uparrow(x) = \sum_{i=1}^{\infty} f_i^\uparrow |\phi_i^\uparrow(x)|^2 \quad (31)$$

$$\rho^\downarrow(x) = \sum_{i=1}^{\infty} f_i^\downarrow |\phi_i^\downarrow(x)|^2, \quad (32)$$

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,

$$\rho^{\text{charge}}(x) = \rho^\uparrow(x) + \rho^\downarrow(x), \quad (33)$$

$$\rho^{\text{spin}}(x) = \rho^\uparrow(x) - \rho^\downarrow(x). \quad (34)$$

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 [70, 72]. In particular, [70] considers a data-derived approach to predicting and assessing uncertainty in a guess density away from the ground state.

3.8. 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-point⁷,

$$\begin{aligned}\rho_* + \delta\rho_n^{\text{out}} &= K[\rho_* + \delta\rho_n^{\text{in}}], \\ &\approx K[\rho_*] + \left. \frac{\delta K[\rho_n^{\text{in}}]}{\delta\rho_n^{\text{in}}} \right|_{\rho_*} \delta\rho_n^{\text{in}}, \\ \Rightarrow \delta\rho_n^{\text{out}} &= \left. \frac{\delta K[\rho_n^{\text{in}}]}{\delta\rho_n^{\text{in}}} \right|_{\rho_*} \delta\rho_n^{\text{in}}.\end{aligned}\quad (35)$$

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 shown by the components of the *Jacobian* of the map K ,

$$\begin{aligned}J_K(x, x') &= \left. \frac{\delta K[\rho_n^{\text{in}}]}{\delta\rho_n^{\text{in}}} \right|_{\rho_*} (x, x') \\ &= \left. \frac{\delta\rho_n^{\text{out}}}{\delta\rho_n^{\text{in}}} \right|_{\rho_*} (x, x').\end{aligned}\quad (36)$$

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

$$\varepsilon_0(x, x') = I - J_K(x, x'), \quad (37)$$

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

$$\varepsilon_0(x, x') = I - \int_{\mathbb{R}^3} dx'' \frac{\delta v_{\text{hxc}}^{\text{in}}(x'')}{\delta\rho^{\text{in}}(x')} \frac{\delta\rho^{\text{out}}(x)}{\delta v_{\text{hxc}}^{\text{in}}(x'')}, \quad (38)$$

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 susceptibility* χ_0 as

$$\varepsilon_0(x, x') = I - \int dx'' (f_{\text{h}}(x', x'') + f_{\text{xc}}(x', x'')) \chi_0(x, x''), \quad (39)$$

where f_{h} and f_{xc} are the kernels of the Hartree (Coulomb) and exchange–correlation integrals. 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 susceptibility

$$\chi_0 = \frac{\delta\rho^{\text{out}}}{\delta v_{\text{hxc}}^{\text{in}}}, \quad (40)$$

which is highly system dependent [57, 73]. As the non-interacting susceptibility plays a central role in the description of many physical phenomena, such as absorption spectra, it is a relatively well-studied object [74, 75]. 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-behaved⁸ the map is about the current iterate. An iterative scheme utilising the exact response function is given by

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{in}} + \varepsilon_0^{-1} R[\rho_n^{\text{in}}], \quad (41)$$

which one may recognise as Newton’s method. While Newton’s method is not global, it has many attractive features, see section 4.1. However, one is rarely privileged with knowledge of the exact dielectric as it is vastly expensive to compute and store [78–81]. In practice, one is left to estimate, or iteratively build, this response function. Cases in which the input is very 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

$$\begin{aligned}\delta\rho_{n+1}^{\text{in}} &\rightarrow 0 \text{ as } n \rightarrow \infty, \\ \Rightarrow (\varepsilon_0^{-1})^n \delta\rho_0^{\text{in}} &\rightarrow 0\end{aligned}\quad (42)$$

⁷ Note 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.

⁸ By ‘well-behaved’ here, we mean that the higher order than linear terms can be ignored without much detriment.

using equation (35), meaning $\lambda_i < 1$ for all i , 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} [57]. 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

$$\kappa = \frac{\lambda_{\max}(\varepsilon_0^{-1})}{\lambda_{\min}(\varepsilon_0^{-1})}, \quad (43)$$

it can be seen that the efficiency of the linear mixing procedure is limited by how close this ratio is to unity. One ansatz for the scalar premultiplying the dielectric is

$$\alpha = \frac{2}{\lambda_{\max} + \lambda_{\min}}, \quad (44)$$

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 [57, 73, 82]. However, this form of α ignores the distribution (e.g. clustering) of eigenvalues [83], and is not commonly used in conjunction with more sophisticated schemes such as those in section 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 [15] for example, and is the core idea behind the *Kerker preconditioner* [84, 85], as discussed shortly.

It is clear from equation (43) 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 [73, 78, 86]. To see why this is, it is first *asserted* that the discretised Coulomb kernel possesses large (divergent) eigenvalues due to its $|x - x'|^{-1}$ 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 equation (39), 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 is a polynomial of the density and potentially its higher order derivatives. Crucially, it has a simple, local (explicit) dependence on position given by $\delta(x - x')$,

$$f_{xc}(x, x') = \delta(x - x') \frac{\delta v_{xc}}{\delta \rho}, \quad (45)$$

that cannot, by definition, 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 equation (39). As [78] 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

$$\delta \rho^{\text{out}}(x) = \int dx' \chi_0(|x - x'|) \delta v_h^{\text{in}}(x'). \quad (46)$$

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

$$\delta \tilde{\rho}^{\text{out}}(G) = \tilde{\chi}_0(|G|) \delta \tilde{v}_h^{\text{in}}(G), \quad (47)$$

where we label the Fourier components G by convention. This 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

$$\begin{aligned} \delta \tilde{\rho}^{\text{out}}(G) &= f_h(|G|) \tilde{\chi}_0(|G|) \delta \tilde{\rho}^{\text{in}}(G) \\ &= \frac{4\pi \tilde{\chi}_0(|G|)}{|G|^2} \delta \tilde{\rho}^{\text{in}}(G). \end{aligned} \quad (48)$$

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 constant⁹ [87]. 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 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 [75, 88]. A demonstration of charge sloshing is illustrated in figure 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 cell. 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 Adler–Wiser equation which is defined as

$$\chi_0(x, x') = \sum_{n=1}^N \sum_{m=N+1}^{\infty} \frac{\phi_n(x) \phi_m^*(x) \phi_n^*(x') \phi_m(x')}{\epsilon_n - \epsilon_m}, \quad (49)$$

which is an expression from perturbation theory for the exact Kohn–Sham susceptibility [74, 75]. As [82] 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

⁹ A detailed treatment of χ_0 for inhomogeneous inputs is given within the framework of Lindhard theory [76].

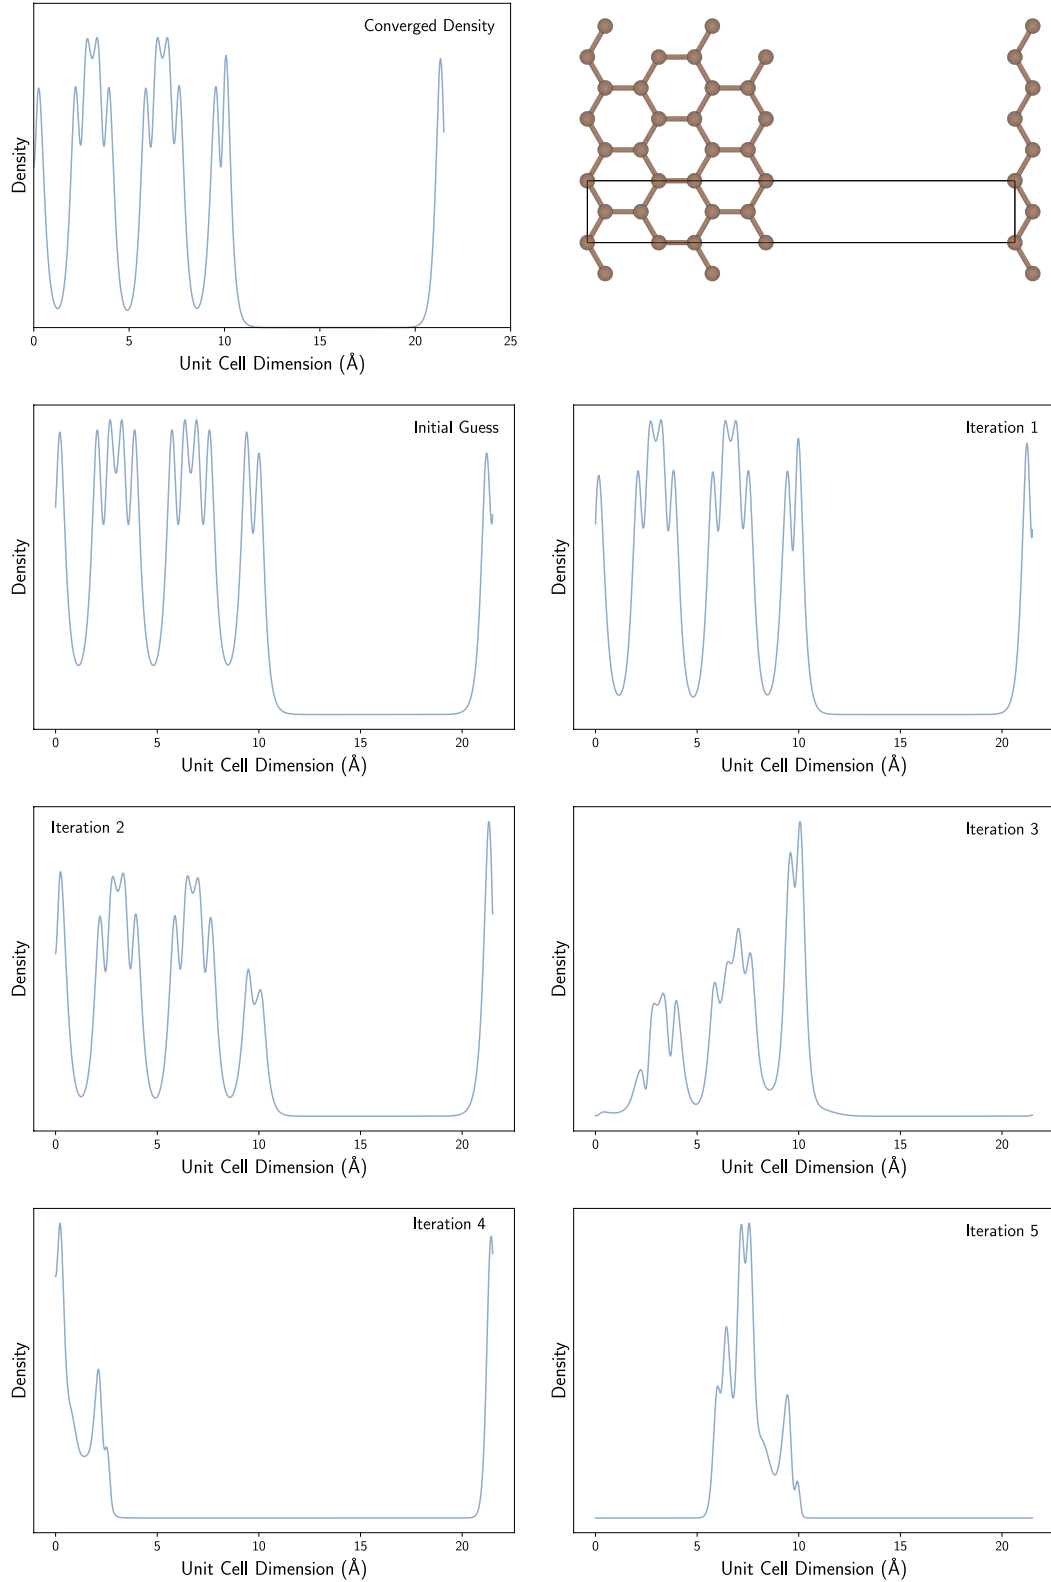


Figure 7. An illustration of charge sloshing for a graphene nanoribbon unit cell (top right). The linear mixing algorithm is applied with damping parameter $\alpha = 0.8$; this leads to an overcorrection in the density at each iteration, resulting in complete divergence.

energy. If left untreated, this observation, in conjunction with the amplifying factor from the low $|G|$ components of the Coulomb kernel, lead to significant ill-conditioning. The largest condition numbers arise when $|G|$ is extremely small; since

G is a reciprocal lattice vector, this will occur for unit cells that are large in any (or all) of the three real-space dimensions. Whilst the dependence of the eigenvalues of the dielectric on unit cell size is in practice complicated [73], 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 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 [75, 88]

$$\tilde{\chi}_0 \propto |G|^2. \quad (50)$$

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 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 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 [73]. 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 [89] 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 equations (8) and (9), 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,

$$\varepsilon_0 = \begin{bmatrix} \varepsilon_0^{\uparrow\uparrow} & \varepsilon_0^{\uparrow\downarrow} \\ \varepsilon_0^{\downarrow\uparrow} & \varepsilon_0^{\downarrow\downarrow} \end{bmatrix}, \quad (51)$$

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 equations (33) and (34), 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 [2, 57]. 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. Dederichs and Zeller [57] 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 4.3.

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, [19, 28, 49, 73, 89–92] and references therein. The aim of this section is to collate conclusions from these studies, and many others, in order to provide a contemporary survey of self-consistency methodology in a pedagogical manner. This survey includes methodology suitable for software utilising either a localised or delocalised basis set. However, only the subset of algorithms suitable for a delocalised basis set are implemented in CASTEP for the benchmarking effort in section 6.

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

$$\rho_{n+1}^{\text{in}} = f_n(\{\rho_i^{\text{in}}, \rho_i^{\text{out}}\}), \quad (52)$$

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 4.3. Elementary algorithms for generating f_n were first considered in section 3.5: the fixed-point and linear mixing algorithms,

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{out}}, \quad (53)$$

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{in}} + \alpha(K[\rho_n^{\text{in}}] - \rho_n^{\text{in}}), \quad (54)$$

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 [57]. It can be shown that this algorithm converges q -linearly toward the fixed-point density ρ_* [93]; where q -linear convergence is defined as

$$\|\rho_{n+1}^{\text{in}} - \rho_*\|_2 \leq q \|\rho_n^{\text{in}} - \rho_*\|_2. \quad (55)$$

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 global, meaning it converges from any initial guess in the $n \rightarrow \infty$ limit. The price one often pays for global 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,

$$\|\rho_{n+1}^{\text{in}} - \rho_*\|_2 \leq \gamma \|\rho_n^{\text{in}} - \rho_*\|^q, \quad (56)$$

for some positive real number γ , where $q > 1$ and $q = 2$ defines quadratic convergence. These algorithms tend to sacrifice guaranteed global convergence, but can vastly improve the rate of convergence, as demonstrated in section 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 [58] asserts and demonstrates, the characteristic difficulty of an optimisation problem depends on whether or not the underlying functional is convex,

$$F[\alpha x + \beta y] \leq \alpha F[x] + \beta F[y]. \quad (57)$$

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 [53, 58, 93] 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 .

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 $\delta\rho_n$ such that

$$R[\rho_n^{\text{in}} + \delta\rho_n] \approx R[\rho_n^{\text{in}}] + J|_{\rho_n^{\text{in}}} \delta\rho_n = 0, \quad (58)$$

where $J := J_R$ is the Jacobian of the residual, defined similarly to equation (36), evaluated at the current iterate. Rearranging for $\delta\rho_n$, the update is given as

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{in}} - J|_{\rho_n^{\text{in}}}^{-1} R[\rho_n^{\text{in}}]. \quad (59)$$

Assuming the Jacobian exists and is Lipschitz continuous¹⁰, this update is shown to have quadratic convergence in some region about the root [83]. The Jacobian must be computed numerically, which can be done with either the Adler–Wiser equation equation (49) [74, 75], or with finite-difference numerical differentiation [94]. As section 4.3 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 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 [78, 79]. The latter strategy requires excessively many evaluations of K [95]. 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.

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,

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{in}} - J_n^{-1} R_n, \quad (60)$$

where $R_n := R[\rho_n^{\text{in}}]$. 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 4.3. We begin with a description of Broyden’s two methods [96]. These methods, as they are about to be presented, are not commonly used in modern Kohn–Sham software. However, the conceptual foundation of

¹⁰ Lipschitz continuity refers to all real $k \geq 0$ in equation (21).

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

$$f'(x_n) = \frac{f(x_n) - f(x_{n-1})}{x_n - x_{n-1}}, \quad (61)$$

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

$$J_n(\rho_n^{\text{in}} - \rho_{n-1}^{\text{in}}) = R_n - R_{n-1}, \quad (62)$$

where hereafter we define $\Delta\rho_n^{\text{in}} = \rho_n^{\text{in}} - \rho_{n-1}^{\text{in}}$ and $\Delta R_n = R_n - R_{n-1}$. If the n th Jacobian satisfies equation (62), 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 $\Delta\rho_n^{\text{in}}$, 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 $\Delta\rho_n^{\text{in}}$ similarly to J_{n-1} . This is equivalent to requiring that the Jacobian of the current iteration solves the following constrained quadratic programming problem,

$$\text{minimise } \|J_n - J_{n-1}\|_f \quad (63)$$

$$\text{subject to } J_n \Delta\rho_n^{\text{in}} = \Delta R_n, \quad (64)$$

as demonstrated by [97, 98], which defines Broyden's *first method*. The Frobenius norm $\|A\|_f$ of a square matrix A is defined as

$$\|A\|_f = \sqrt{\sum_{i,j=1}^N |a_{ij}|^2}. \quad (65)$$

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, $(\Delta\rho_i^{\text{in}})^\dagger \Delta\rho_j^{\text{in}} = 0$ 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 equation (63) has a unique analytic solution, which is obtained in [97, 98] by means of unconstrained optimisation using the method of Lagrange multipliers,

$$J_n = J_{n-1} + \frac{\Delta R_n - J_{n-1} \Delta\rho_n^{\text{in}}}{\|\Delta\rho_n^{\text{in}}\|_2^2} (\Delta\rho_n^{\text{in}})^\dagger. \quad (66)$$

The notation uv^\dagger defines the outer product of the vectors u, v . One can now observe from equation (66) 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 equation (66), applying it to residual vector, and performing the quasi-Newton step equation (60). The apparent excessive cost of inverting equation (66) is negated as the inverse of a rank-one matrix can be computed analytically using the Sherman–Morrison–Woodbury formula [99]. 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 [100], 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

$$\text{minimise } \|H_n - H_{n-1}\|_f \quad (67)$$

$$\text{subject to } \Delta\rho_n^{\text{in}} = H_n \Delta R_n, \quad (68)$$

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 [91]. Note that the constraint in equation (68) is simply the inverse secant condition. Similarly to Broyden's first method, this has the analytic solution,

$$H_n = H_{n-1} + \frac{\Delta\rho_n^{\text{in}} - H_{n-1} \Delta R_n}{\|\Delta R_n\|_2^2} (\Delta R_n)^\dagger, \quad (69)$$

which can be substituted directly into the quasi-Newton step¹¹. 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 [83, 97] in the sense that

$$\lim_{n \rightarrow \infty} \frac{\|\rho_{n+1}^{\text{in}} - \rho_*\|_2}{\|\rho_n^{\text{in}} - \rho_*\|_2} = 0, \quad (70)$$

which is a necessary condition for some $q > 1$ in equation (56). Broyden's second method is implemented and tested in section 6.

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 of Broyden's methods, which are examined extensively in both optimisation and electronic structure literature [83, 91, 100, 102]. The ensuing summary follows a similar structure to that of [91]. A multisecant method is defined as a method that

¹¹ Note 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 [15, 31, 101].

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

$$\Delta \mathcal{R}_n := (\Delta R_{n-m+1}, \Delta R_{n-m+2}, \dots, \Delta R_n) \quad (71)$$

$$\Delta \mathcal{P}_n := (\Delta \rho_{n-m+1}^{\text{in}}, \Delta \rho_{n-m+2}^{\text{in}}, \dots, \Delta \rho_n^{\text{in}}) \quad (72)$$

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

$$J_n \Delta \mathcal{P}_n = \Delta \mathcal{R}_n. \quad (73)$$

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 equations (63) and (67) to reflect the multisecant condition equation (73). The multisecant version of Broyden's first and second method respectively are

$$\begin{aligned} &\text{minimise } \|J_n - J_{n-1}\|_f \\ &\text{subject to } J_n \Delta \mathcal{P}_n = \Delta \mathcal{R}_n, \end{aligned} \quad (74)$$

$$\begin{aligned} &\text{minimise } \|H_n - H_{n-1}\|_f \\ &\text{subject to } \Delta \mathcal{P}_n = H_n \Delta \mathcal{R}_n, \end{aligned} \quad (75)$$

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 [89, 92] conclude, and section 6 also examines, the type-II variant tends to outperform the type-I variant in the context of multisecant 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 [101–103]. First, the following *unconstrained* minimisation problem for variations in H_n is defined,

$$\begin{aligned} &\text{minimise } w_0 \|H_n - H_{n-1}\|_f \\ &\quad + \sum_{i=n-m+1}^n w_i \|\Delta \mathcal{P}_n - H_n \Delta \mathcal{R}_i\|_2^2, \end{aligned} \quad (76)$$

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 equation (76) 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 multisecant method, but the multisecant 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 [15, 31] 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 equation (76): 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 [102, 103], 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 [101–103]. However, as [101] 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 multisecant method; see [101] for additional detail. An interesting aspect of the multisecant methods discussed here are their relationship Pulay's or Anderson's method—a ubiquitous method in electronic structure theory software—which is now examined.

4.1.3. Pulay's method. Pulay's method [104, 105], 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 [106], 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 [10, 15, 33, 34, 107, 108]. The past few decades of wisdom suggest that Pulay's method systematically outperforms the unmodified Broyden's methods in both the single and multisecant formulation. This conclusion will be tested in section 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,

$$R[\rho_{\text{opt}}^{\text{in}}] = \sum_{i=n-m+1}^n c_i R[\rho_i^{\text{in}}]. \quad (77)$$

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

$$\begin{aligned} &\text{minimise } \|R[\rho_{\text{opt}}^{\text{in}}]\|_2^2 \\ &\text{subject to } \sum_i c_i = 1, \end{aligned} \quad (78)$$

Table 1. 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 ^a	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 ^b	$\alpha = 0.6, G_0 = 1.5, k = 2$	0.705	0.0062
EDFT	$\alpha = 0.8, G_0 = 1.5$	0.948	0.00003

^a As proposed in [149].^b Performed with k Pulay steps in between each linear step.

where the constraint that the coefficients must sum to unity is an exact requirement at convergence. Substitution of equations (77) into (78), and use of Lagrange multipliers, allows the optimisation problem to be cast as an $(n + 1)$ -dimensional linear system,

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

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 [56, 90, 91, 101, 109], the optimum residual can be first be expanded as such,

$$R[\rho_{\text{opt}}^{\text{in}}] = \sum_{i=n-m+1}^n c_i K[\rho_i^{\text{in}}] - \sum_{i=n-m+1}^n c_i \rho_i^{\text{in}}. \quad (79)$$

If K is assumed to be linear, the rightmost term in equation (79) can be interpreted as the optimal input density,

$$\rho_{\text{opt}}^{\text{in}} = \sum_{i=n-m+1}^n c_i \rho_i^{\text{in}}. \quad (80)$$

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

$$\rho_{n+1}^{\text{in}} = \rho_{\text{opt}}^{\text{in}} + R[\rho_{\text{opt}}^{\text{in}}] \quad (81)$$

$$= \sum_{i=n-m+1}^n c_i K[\rho_i^{\text{in}}]. \quad (82)$$

This update is favoured over $\rho_{n+1}^{\text{in}} = \rho_{\text{opt}}^{\text{in}}$ so that the algorithm does not stagnate in the subspace of past input densities. Alternatively, as originally studied in [106], a damped step can be taken,

$$\rho_{n+1}^{\text{in}} = \rho_{\text{opt}}^{\text{in}} + \alpha R[\rho_{\text{opt}}^{\text{in}}], \quad (83)$$

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,  $\rho_0^{\text{in}}, m, \alpha$ 
2: for  $n = 0, 1, 2, \dots$ 
3:   Compute  $R[\rho_n^{\text{in}}]$  and store the pair  $\{\rho_n^{\text{in}}, R[\rho_n^{\text{in}}]\}$ 
4:   Solve equation (78) for  $\{c_i\}$ 
5:   if  $n \leq m$ 
6:     Set  $\rho_{n+1}^{\text{in}} = \sum_{i=1}^n c_i (\rho_i^{\text{in}} + \alpha R[\rho_i^{\text{in}}])$ 
7:   else
8:     Set  $\rho_{n+1}^{\text{in}} = \sum_{i=n-m+1}^n c_i (\rho_i^{\text{in}} + \alpha R[\rho_i^{\text{in}}])$ 
9:   if  $\|R[\rho_n^{\text{in}}]\|_2 < \text{tol}$ , 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 [56, 91, 109], a rearrangement of the optimisation problem in equation (78) reveals a close relationship between Pulay's method and type-II multiseant methods. A more detailed treatment of this correspondence is found in [56, 91, 109]; here, we simply state that the following unconstrained optimisation problem

$$\text{minimise } \|\mathbf{R}_n - \Delta\mathcal{R}_n\gamma\|_2, \quad (84)$$

for variations in $\gamma = (\gamma_{n-m+1}, \gamma_{n-m+2}, \dots, \gamma_n)$ is equivalent to Pulay's optimisation problem equation (78). The new coefficients $\{\gamma_i\}$ are related to the old coefficients $\{c_i\}$ such that the update in equation (83) now takes the form

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{in}} + \alpha\mathbf{R}_n - (\Delta\mathcal{P}_n + \alpha\Delta\mathcal{R}_n)\gamma, \quad (85)$$

where γ on iteration n is solved by

$$\gamma = (\Delta\mathcal{R}_n^\dagger \Delta\mathcal{R}_n)^{-1} \Delta\mathcal{R}_n^\dagger \mathbf{R}_n. \quad (86)$$

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

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{in}} + H_n \mathbf{R}_n, \quad (87)$$

$$H_n = \alpha\mathbf{I} - (\Delta\mathcal{P}_n - \alpha\Delta\mathcal{R}_n)(\Delta\mathcal{R}_n^\dagger \Delta\mathcal{R}_n)^{-1} \Delta\mathcal{R}_n^\dagger. \quad (88)$$

By comparison with the updated inverse Jacobian in equation (75), we can observe that Pulay's method is a type-II quasi-Newton step where the iterative inverse Jacobian is updated according to

$$\begin{aligned} &\text{minimise } \|\mathbf{H}_n - \mathbf{H}_0\|_f \\ &\text{subject to } \Delta\mathcal{P}_n = H_n \Delta\mathcal{R}_n. \end{aligned} \quad (89)$$

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. Note that $H_0 = \alpha\mathbf{I}$ is required in order to recover the update in equation (88). 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,

$$\begin{aligned} &\text{minimise } \|\mathbf{J}_n - \mathbf{J}_0\|_f \\ &\text{subject to } \mathbf{J}_n \Delta\mathcal{P}_n = \Delta\mathcal{R}_n, \end{aligned} \quad (90)$$

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 [91], has seen comparatively less application and testing in the context of Kohn–Sham codes [56]. 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 modifications to Pulay's method are now examined; although we note that these modifications are adaptable to all the secant-based methods discussed previously. First, the work in [110], based on [111], suggests that alternating between

Pulay and linear mixing steps can improve the robustness of the iterations over standard Pulay—the 'Periodic Pulay' method. Each new Pulay step utilises the history from the linear mixing and past Pulay steps to solve the optimisation subproblem equation (78). This is demonstrated to have a stabilising effect as the linear mixing history data is used well by the Pulay extrapolation. In [110], 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 6 with a damping parameter of $\alpha = 0.2$ for both the Pulay and linear mixing steps.

Second, [112] considers occasionally flushing the history every time a certain criterion is met, rather than iteratively overriding the history—'Restarted Pulay'. 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.

The final technique we discuss here is the 'Guaranteed Reduction Pulay' algorithm of [113]. The approach of Guaranteed Reduction Pulay involves ensuring that the Pulay predicted optimal residual $\|\mathbf{R}_{\text{opt}}\|_2$ decreases each iteration. This is achieved by rearranging the stored history of residuals $\{\mathbf{R}_i\}$ 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 equation (78). 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. All three of these techniques are benchmarked in section 6.

4.1.4. Modern multiseant-based algorithms. Here, we highlight one modern use of multiseant methods in particular, the methods outlined in [89, 92], which are now default self-consistency methods in WIEN2K [114]. These algorithms can be considered a sophisticated modern variant of the standard multiseant already methods discussed. Furthermore, they are

designed with the aim of converging Kohn–Sham calculations. The range of strategies utilised make [89, 92] an interesting case to isolate and examine here. The most recent published form of these algorithms is that given in [89] titled ‘multise-cant rank one’ (MSR1). However, this algorithm is designed to converge both the atomic (geometry optimisation) and elec-tronic degrees of freedom. Hence, we focus on the techniques that are relevant to the self-consistent field iterations, and the reader is referred to [89, 92] for a more in-depth treatment.

First, the updates considered in [89, 92] are defined in equations (90) and (89), which are of the form

$$H_n = \alpha I + (\Delta \mathcal{P}_n - \alpha \Delta \mathcal{R}_n)(\Delta \mathcal{R}_n^\dagger W_n)^{-1} W_n^\dagger. \quad (91)$$

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 [92], and further verified in section 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 combina-tion of updates,

$$W_n = Y_n + \beta S_n, \quad (92)$$

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 [89], this is similar to a technique used in [115] 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 feature of the methods in [89, 92] involve partitioning of the full update into a *predicted* and *unpredicted* component, which are now defined. Consider the update gen-erated from equation (91),

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{in}} + H_n R_n, \quad (93)$$

which is now split in two,

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{in}} + (H_n^p + H_n^u) R_n. \quad (94)$$

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

$$(H_n^u R_n)^\dagger \Delta \mathcal{R}_n = 0. \quad (95)$$

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 consid-ered to be the part of the update that the iterative subspace does describe. Equation (95) is shown in [92] to have the solution,

$$H_n^u R_n = -(I - \Delta \mathcal{R}_n (\Delta \mathcal{R}_n^\dagger \Delta \mathcal{R}_n)^{-1} \Delta \mathcal{R}_n^\dagger) R_n. \quad (96)$$

The update now takes the rearranged form

$$H_n^u R_n = (I - \Delta \mathcal{R}_n (\Delta \mathcal{R}_n^\dagger \Delta \mathcal{R}_n)^{-1} \Delta \mathcal{R}_n^\dagger) R_n, \quad (97)$$

$$H_n^p R_n = -S_n (\Delta \mathcal{R}_n^\dagger \Delta \mathcal{R}_n)^{-1} W_n^\dagger R_n. \quad (98)$$

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

$$\rho_{n+1}^{\text{in}} = \rho_n^{\text{in}} + (H_n^p + \alpha_n H_n^u) R_n. \quad (99)$$

In the implementation tested in section 6, the unpredicted direction is also Kerker preconditioned (see section 4.3). References [89, 92] 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. The precise greed controls for MSR1 are found in [89]; the controls used for section 6 are slightly modified for performance in the plane-wave pseudopotential setting. Furthermore, the matrix inverse in equation (91) is Tikhonov regularised [116] to prevent spurious behaviour due to rank-deficiencies, and the matrices involved are scaled appropriately. The method that has been described thus far is similar to ‘multisecant Broyden 2’ (MSB2) of [92]. The type-I variant, MSB1, can be derived in a similar fashion, and both are tested in section 6.

Both of these methods remain prone to charge slosh-ing due to the primary form of step length control being an implicit trust region, the greed controls. Hence, particularly ill-conditioned simulations can still lead to divergence of the self-consistency iterations. In this context, a further stabilis-ing measure is taken in the form of an explicit trust region. Given R_n as a descent direction on iteration n , the trust region subproblem can take a standard form

$$\text{minimise } \|R_n - H_n^{-1} X\|_2^2 \quad (100)$$

$$\text{subject to } \|X\|^2 - \delta^2 = 0 \quad (101)$$

for variations in X with some trust region radius δ . The vari-able 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 [83, 89]. As [89] also states, the trust region problem does not need to be defined in terms of the full update, $\|H_n R_n\|_2$. Instead, the trust region problem can be solved for the predicted compo-nent of the update, $\|H_n^p R_n\|_2$. This is motivated by the fact that the size of the step in the unpredicted direction is already restricted by the greed controls. Note that δ is also iteratively updated based on algorithmic progress. The inclusion of an explicit trust region for the predicted update, in conjunction

with the strategies described earlier, are able to significantly stabilise the self-consistency iterations [89, 92].

A second method we highlight is a global variant of the type-I Pulay update [56], which has not yet been tested in the context of self-consistency iterations. (However, note that this method utilises a selection of the techniques from section 4.1.3 that do perform well for self-consistency iterations). First, similarly to the previous method, the Jacobian is required to be non-singular, which is achieved here through a form of Powell regularisation [117]. The regularisation parameter that is introduced, in some sense, scales the update between an unregularised Pulay step, and a fixed-point step, and thus must be chosen appropriately as not to negatively impact the efficiency of the method. The remaining modifications are aimed at stabilising the iterations and preventing stagnation in the reduction of the residual norm. Type-I methods, more so than type-II methods, tend to suffer from stagnation due to rank-deficiency in the iterative subspace of density differences $\Delta\mathcal{P}_n$, see [56]. Rank-deficiency is avoided using a technique derived from the Restarted Pulay method [112] in the previous section, i.e. occasionally restart the history of stored iterates. Here, the restart condition can be triggered based on two separate criteria, the first of which is whenever the number of iterates stored reaches a maximum value m , similarly to [112]. The second of the restart criteria is based on ensuring the iterative subspace remains approximately of full rank. This is done by first using Gram–Schmit orthonormalisation on set of vectors used to construct $\Delta\mathcal{P}_n$ in order to generate a new set of vectors with the same span, $\Delta\hat{\mathcal{P}}_n$. Elements of this new set of vectors are identically zero if the original set had linear dependencies. The condition

$$\|(\Delta\hat{\mathcal{P}}_n)_i\|_2 < \tau \|(\Delta\mathcal{P}_n)_i\|_2 \quad (102)$$

is therefore used to quantify the degree of linear independence we require from the iterative subspace, parametrised by τ . If this condition is triggered, for some sensible value of τ , the iterative history is deemed too linearly dependent, and the history is reset.

Given that update Jacobian is now bounded, and the iterative history is linearly independent, global convergence is further guaranteed by using techniques similar to the Periodic Pulay method [110, 118, 119]. That is, linear mixing steps are included in between Pulay steps, and added to the history. However, instead of performing the linear mixing steps with a fixed period, the linear mixing steps are performed based on a criterion that ensures progress is made. Namely, the following inequality is defined

$$\|R_n\|_2 \leq D \|R_0\|_2 (n' + 1)^{-(1+\varepsilon)} \quad (103)$$

for parameters D and ε , where n' is the number of Pulay steps performed so far in the simulation. If the inequality is not satisfied, linear mixing steps are performed until the inequality is satisfied again, and then Pulay steps are resumed. In other words, if the Pulay steps are not making sufficient progress, then linear steps are performed until a certain amount of progress, defined by D and ε , has been made, and then Pulay steps

are continued. The proof of global convergence for this algorithm is given in [56].

4.2. Density matrix optimisation

The methods introduced previously define *density mixing schemes*, by which we mean the one-particle density is iteratively updated in order to find a fixed-point of the Kohn–Sham equations. Density mixing is common in implementations that utilise some form of delocalised basis set. In these implementations, both the Kohn–Sham Hamiltonian and the *density matrix* can be prohibitive to compute and store due to the size of the basis N_b . However, in the context of localised basis sets, it is common to formulate Kohn–Sham theory so that the density matrix equation (13) is the optimised variable, rather than the density¹². The Kohn–Sham energy functional has a closed-form expression in terms of the density matrix,

$$\begin{aligned} E_{\text{KS}}[D] = & -\frac{1}{2} \int_{\mathbb{R}^3} \nabla_{x'}^2 D(x, x')|_{x=x'} \\ & + \int_{\mathbb{R}^3 \times \mathbb{R}^3} \frac{D(x, x) D(x', x')}{|x - x'|} \\ & + \int_{\mathbb{R}^3} v_{\text{ext}} D(x, x) + E_{\text{xc}}[D(x, x)] \end{aligned} \quad (104)$$

which is now minimised over allowed variations in D ,

$$\int_{\mathbb{R}^3} D(x, x) = N, \quad (105)$$

$$D^\dagger = D, \quad (106)$$

$$\int_{\mathbb{R}^3} D(x, x'') D(x'', x') = D(x, x'), \quad (107)$$

for binary occupancy of the Kohn–Sham orbitals. These conditions define the set of density matrices \mathcal{D} which form the domain of $E_{\text{KS}}[D]$, i.e. the set of density matrices for which the Kohn–Sham functional is defined. These are necessary but not sufficient conditions for a density matrix to be so-called N -representable, meaning it can be constructed from an anti-symmetric wavefunction [120]. The first condition forces the density matrix to correspond to an N -particle system. The second condition forces the density matrix to be Hermitian, which is true by construction. The third condition is called *idempotency*, which leads to the requirement that the eigenvalues of discretised density matrix, the orbital occupancies, are $f_i = \{0, 1\}$ [28]. The Kohn–Sham Hamiltonian, sometimes referred to as the Fock matrix, can also be constructed and used to solve the Kohn–Sham equations for a self-consistent density matrix. As expected, a variety of self-consistent field techniques exist that are well suited to this formulation, such as leveling shifting and its modern variants [121–124], methods that minimise a local model energy functional [122, 125–127], and more [19, 60, 128–131]. As many of these

¹² Note that density mixing schemes are usually translated to density matrix mixing schemes relatively straightforwardly, e.g. [28].

methods are not readily adaptable to the plane-wave setting, we keep the discussion here brief, and instead refer readers to the following review articles [19, 130, 132, 133] and references therein. In fact, we highlight one class of methods in particular, the relaxed constraints algorithms given in [49, 60, 129, 131, 134]. A member of this class, ‘Energy DIIS’, is now the default self-consistency method in GAUSSIAN09 [18].

Relaxed constraints algorithms are introduced in the context of Hartree–Fock theory, where, unlike in extended Kohn–Sham theory discussed in section 3.6, the notion of fractional occupancy is not a part of the framework. Relaxed constraints algorithms operate by permitting fractional occupation of the Hartree–Fock orbitals as a tool to reach convergence. The binary occupation fixed-point solution is recovered at the end of the calculation. First, consider the set of allowed *discretised* density matrices,

$$\mathcal{D} = \{D \mid \text{Tr}D = N, D^\dagger = D, D^2 = D\}. \quad (108)$$

It is now apparent that the idempotency condition equation (107) requires that the eigenvalues of the density matrix, the orbital occupancies, are binary, $f_i \in \{0, 1\}$. The extension to fractional occupancy, as described in section 3.6, thus alters the set of allowed density matrices,

$$\tilde{\mathcal{D}} = \{D \mid \text{Tr}D = N, D^\dagger = D, D^2 \leq D\}; \quad (109)$$

i.e. the eigenvalues of the density matrix must now satisfy $f_i^2 \leq f_i$, which is the case for $0 \leq f_i \leq 1$. Relaxed constraint algorithms are now founded based on two theorems, the proofs of which are given in [135]. The first theorem states that the Hartree–Fock functional varied over \mathcal{D} has the *same* stationary points as the Hartree–Fock functional varied over $\tilde{\mathcal{D}}$. In other words, minimising the Hartree–Fock functional over $\tilde{\mathcal{D}}$ will always lead to physical (binary occupation) solutions. The second theorem states that the set $\tilde{\mathcal{D}}$ is convex. This means that, given a *convex combination* of density matrices $D_1, D_2 \in \tilde{\mathcal{D}}$,

$$D = (1 - \gamma)D_1 + \gamma D_2, \quad (110)$$

for $\gamma \in [0, 1]$, then $D \in \tilde{\mathcal{D}}$. This is not the case for elements of \mathcal{D} . Algorithms that utilise these theorems in conjunction with the set $\tilde{\mathcal{D}}$ belong to the class of relaxed constraints algorithms.

The question remains of how these theorems translate to the Kohn–Sham functional equation (104). In fact, the former theorem no longer holds [134], meaning that the functional varied over $\tilde{\mathcal{D}}$ does lead to different stationary points than in the case of binary occupation. This is expected, as these solutions correspond to solutions of the extended Kohn–Sham theory developed in section 3.6. Therefore, the absence of the former theorem does not pose a problem, as solutions that are obtained that are members of $\tilde{\mathcal{D}}$ retain meaning within extended Kohn–Sham theory. Furthermore, the set $\tilde{\mathcal{D}}$ remains convex. The *optimal damping algorithm* was the first of the relaxed constraints algorithms, as examined in [60]. This algorithm seeks, at each iteration, to find an ‘optimal’ damping parameter α_n in the linear mixing scheme equation (22). Consider the pair $\{D_n^{\text{in}}, D_n^{\text{out}}\}$ on iteration n . A trail density matrix is now constructed as a convex combination of this pair,

$$D^{\text{trial}} = (1 - \alpha)D_n^{\text{in}} + \alpha D_n^{\text{out}}, \quad (111)$$

for $\alpha \in [0, 1]$, where we note that $D_n^{\text{out}} - D_n^{\text{in}}$ is a descent direction. The Kohn–Sham energy functional $E_{\text{KS}}[D^{\text{trial}}]$ is then minimised along this line-segment by varying α . The value of the damping parameter that leads to the minimum energy is the so-called optimal damping value α^{opt} . The subsequent input density matrix is thus

$$D_{n+1}^{\text{in}} = D^{\text{trial}} = D_n^{\text{in}} + \alpha^{\text{opt}}(D_n^{\text{out}} - D_n^{\text{in}}). \quad (112)$$

Note that this method works precisely because the elements along the line-segment equation (111) remain in the domain of E_{KS} . Furthermore, it is much cheaper to evaluate the energy for a given density matrix, rather than construct and diagonalise the Fock matrix for a *new* output density matrix.

This method is improved by instead considering an (at most) m -dimensional iterative subspace of past density matrices, rather than just the most recent pair. In other words, similar to Pulay’s method, a trail density matrix is constructed

$$D^{\text{trial}} = \sum_{i=n-m+1}^m c_i D_i^{\text{in}} \quad (113)$$

for some unknown coefficients c_i . These coefficients are then fixed as the coefficients that minimise $E_{\text{KS}}[D^{\text{trial}}]$. An algorithm to accomplish this is provided in [131]. However, note that only *convex combinations* of past density matrices are permissible, otherwise the resulting density matrix may not be a valid according to the conditions imposed in $\tilde{\mathcal{D}}$. This observation necessitates the restriction $0 \leq c_i \leq 1$, which leads to the following constrained minimisation problem,

$$\text{minimise } E_{\text{KS}} \left[\sum_{i=n-m+1}^m c_i D_i^{\text{in}} \right] \quad (114)$$

$$\text{subject to } 0 \leq c_i \leq 1. \quad (115)$$

The coefficients that solve this problem form the optimal density matrix that is set equal to the subsequent input density matrix D_{n+1}^{in} . This method is titled ‘Energy DIIS’ (EDIIS) and has been demonstrated to perform well in a variety of cases [130, 131]. Furthermore, this method is global, as the energy functional is required to be minimised at each iteration in equation (114). The minimisation of the energy functional is an *interpolation* step in the space of past density matrices due to the constraint on the coefficients, and can be slow when the iterates are near convergence [131]. For this reason, EDIIS is commonly augmented with Pulay iterations (DIIS) when close to convergence, which demonstrably improves efficiency [130]. These methods, and similar methods, can be impractical in plane-wave codes as one is required to construct and store the density matrix.

4.3. 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

$$P^{-1}Ax = P^{-1}b, \quad (116)$$

where the $P^{-1}A$ has a lower condition number equation (43) 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

$$\|P(R[\rho_*])\|_2 = 0. \quad (117)$$

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 [58, 83, 97]. Therefore, a preconditioner should accomplish one or more 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 [83, 89, 136]. To simplify matters, we consider P to be a matrix *constant* with respect to the optimised variable, the density, such that $P(R) = PR$. 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. The strategy used in practice is to identify the source of divergent eigenvalues of the dielectric, as examined in section 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 [114], 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 [89, 92]. The following work assumes an entirely plane-wave basis set, although the preconditioners can be adapted for a variety of implementations.

Recall from section 3.8 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




Figure 8. The Kerker preconditioner for various values of $|G_0|^2$.

and isotropic, the dielectric eigenvalues of gapped and gapless phases are approximately determined by

$$\epsilon_0^{-1} = \left(1 + \frac{4\pi\gamma}{|G|^2}\right)^{-1}, \quad (118)$$

$$\epsilon_0^{-1} = (1 + 4\pi\gamma)^{-1}, \quad (119)$$

for some *a priori* unknown system-dependent constant γ . In the case of the homogeneous electron gas, this constant in equation (118) is identified with the square of the Thomas–Fermi screening wavevector k_{TF}^2 . Modest departures from homogeneity and isotropy remain accurately modelled by equations (118) and (119), particularly in the low $|G|$ limit [75, 88]. Therefore, these *model dielectrics* can be used to improve the condition of the residual map by allowing $P = \epsilon_0^{-1}$ for either equation (118) or (119) depending on whether one suspects the input to be metallic or insulating. Collecting and relabelling the unknown constants, the preconditioner becomes

$$P = \alpha \frac{|G|^2}{|G|^2 + |G_0|^2}, \quad (120)$$

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 $|G_0| = 1.5 \text{ Å}$ and $\alpha = 0.8$ [15, 31] demonstrably improves convergence, see section 6. The modified Kohn–Sham problem is now solved

$$\left\| \alpha \frac{|G|^2}{|G|^2 + |G_0|^2} R[\rho_*(G)] \right\|_2 = 0, \quad (121)$$

which is referred to as Kerker preconditioned equation (120) [84, 85]. The Kerker preconditioner suppresses charge sloshing, as defined in section 3.8, by damping eigenvalues of the dielectric corresponding to low $|G|$ components of the density, see figure 8. It is these components that have a generically amplified response due to the Coulomb kernel. Note that

the optimal damping algorithm detailed in section 4.2 can be regarded as a preconditioner that updates the value of α_n at each iteration. The adaptation of the Kerker preconditioner to real space implementations of Kohn–Sham theory more difficult. The dielectric response function of the homogeneous electron gas is non-local in real space, meaning the integral in equation (39) 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 [76]. An efficient real space implementation of the Kerker preconditioner is given in [137].

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 [73]. 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 [73, 138]. 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 [139], which involve more accurately modelling the dielectric response, e.g. [140]. However, in these examples, 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 [78–81, 86, 141]. 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, equation (117), 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 [78, 141]. 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 equation (49). In plane-wave codes, both the computation and storage of all eigenfunctions of the Kohn–Sham Hamiltonian is infeasible.

Past attempts [79, 80] 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 overhead¹³ [15, 31, 78]. Second, we highlight two strategies to remove the infinite summation over unoccupied eigenfunctions in the Adler–Wiser equation equation (49). The first strategy utilises density functional perturbation theory, and in particular the *Sternheimer equation* [142]. 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 residual—without requiring the full eigendecomposition of the Kohn–Sham Hamiltonian. This is utilised in, for example, the implementation of the GW approximation [143], and hence is available functionality within many Kohn–Sham codes. The density functional perturbation theory approach is used to construct a Newton step in [141] 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,

$$\delta(x - x') = \sum_{j=1}^{\infty} \phi_j^*(x) \phi_j(x'). \quad (122)$$

This identity cannot be applied to the Adler–Wiser equation equation (49) 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 equation (122) to transform the Adler–Wiser equation into a sum over $N_{\text{cut-off}} \ll N_b$ eigenfunctions, which is used [78, 81] 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 [79], 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.

Finally, methods are examined that attempt to include inhomogeneity and anisotropy through model response functions, rather than with the exact methods discussed prior. An extension to the Kerker preconditioner, which is based on Thomas–Fermi theory of the homogeneous electron gas, is considered [144]. Here, Thomas–Fermi–Von Weizsäcker¹⁴ theory is used, in combination with Pulay’s method, to generate the subsequent density. Crucially, this work does not restrict to the case of the homogeneous electron gas, and instead

¹³ In fact, a reduced set of plane-waves is also used for the mixing algorithms of section 4.1 in CASTEP [15, 31]. The components of the density corresponding to G vectors outside this reduced space are treated with the fixed-point algorithm.

¹⁴ The Von Weizsäcker kinetic energy extension to Thomas–Fermi theory serves to better model inhomogeneities in the density [87].

numerically minimises the relevant functional. That is, the modified *orbital-free* functional

$$E_{\text{TFW}}[\rho] = \int_{\mathbb{R}^3} |\nabla \rho(x)^{\frac{1}{2}}|^2 + \frac{3}{5} \int_{\mathbb{R}^3} \rho(x)^{\frac{5}{3}} - \int_{\mathbb{R}^3} v_{\text{eff}}[\rho](x) \rho(x) + \delta E[\rho_n^{\text{in}}, \rho] \quad (123)$$

is minimised with respect to variations in ρ , where $v_{\text{eff}} = v_h + v_{\text{xc}} + v_{\text{ext}}$, the effective potential used to solve the Kohn–Sham equations. The details of the modification term $\delta E[\rho_n^{\text{in}}, \rho]$ are given in [144], and is derived such that the minimiser of equation (123) can be used as the subsequent input density. Namely, the method calculates the response that is required to bring the current input density to self-consistency within the framework of Thomas–Fermi–Von Weizsäcker 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 [144], 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 ABINIT [10, 32].

Alternatively, given an input that can be transparently partitioned into metallic and insulating regions, such as an interface, inhomogeneity can be included explicitly by varying $|G_0|$ in each region. That is, set $|G_0| = 0$ for the insulating region, and have finite $|G_0|$ in the metallic region, which is considered in [73, 145]. This is non-trivial as the dielectric equation (39) becomes non-diagonal in both Fourier and real space, and hence becomes unfavourable to construct and store. However, note that one requires the dielectric applied to the residual vector R_n , rather than the dielectric itself. Therefore, inhomogeneity can be included using an algorithm that successively switches between Fourier and real space as to avoid constructing the non-diagonal dielectric; see [145]. Furthermore, if the potential is treated as the optimisation variable rather than the density, a modified Poisson equation can be solved for the updated potential residual, see [73]. Here, the inhomogeneity is specified *a priori* with two functions that are now inputs to the calculation. These methods thus have the drawback that they are not black-box, as one is required to specify the inhomogeneity using prior knowledge of one's input system. Nonetheless, for specific systems, these frameworks provide an expert user with an additional degree of freedom for optimising convergence.

4.4. Direct minimisation

Whilst self-consistent field methods are widespread, alternative techniques are available based on direct minimisation of the energy functional. These methods exploit the variational principle, and are thus global, varying $\{\phi_i\}$ to minimise $E_{\text{KS}}[\{\phi_i\}]$. The density in these schemes is a dependent quantity, always derived directly from $\{\phi_i\}$ with no history of previous densities. From an initial guess set of orbitals, $\{\phi_i^{(0)}\}$, a density




Figure 9. Comparison of the energy for a line-minimisation in the steepest descent direction with respect to the orbitals as a function of step-length α . The energy function is specified with fixed initial potentials (solid line) or with continually updated potentials along the line-minimisation (dashed line). The data is from the first step of a simulation of a conventional 8-atom *fcc* silicon cell, using the local density approximation.

$$\rho^{(0)}(x) = \sum_i |\phi_i^{(0)}(x)|^2 \quad (124)$$

is constructed, and the energy $E_{\text{KS}}[\{\phi_i^{(0)}\}]$ and Kohn–Sham eigenvalue estimates $\epsilon_i^{(0)}$ are evaluated, along with the energy gradient with respect to changes in $\{\phi_i^{(0)}\}$,

$$\left. \frac{\delta E_{\text{KS}}}{\delta \phi_i} \right|_{\phi_i = \phi_i^{(0)}} = H_{\text{KS}}[\rho] \phi_i^{(0)} - \epsilon_i^{(0)} \phi_i^{(0)}. \quad (125)$$

Since the energy gradients are the steepest *ascent* directions, the steepest descent direction is the negative of this. This steepest descent direction may be interpreted as a force acting on the degrees of freedom of the trial states $\{\phi_i^{(0)}\}$. If masses are assigned to these degrees of freedom, then the states may be evolved forward in time according to a suitable equation of motion, and this forms the foundation of the Car–Parrinello method [17, 146]. By damping the motion appropriately, the system evolves towards the ground state.

An alternative approach is to consider the search for the ground state as a minimisation problem. A candidate search direction $D_i^{(0)}$ may be constructed to minimise the energy as,

$$D_i^{(0)} = - \left. \frac{\delta E_{\text{KS}}}{\delta \phi_i} \right|_{\phi_i = \phi_i^{(0)}} \quad (126)$$

i.e. the steepest descent direction. In practice more sophisticated methods are used to construct a search direction, usually based on preconditioned quasi-Newton methods such as conjugate gradients.

Once a search direction has been obtained, an improved set of trial orbitals are constructed, e.g.

$$\phi_i^{(1)} = \phi_i^{(0)} + \alpha D_i^{(0)}, \quad (127)$$

where α is a scalar parameter, chosen to minimise $E_{\text{KS}}[\{\phi_i^{(1)}\}]$. Note that, in general, $\{\phi_i^{(1)}\}$ will not be orthonormal and must be orthonormalised explicitly first. The search for an optimal value of α is known as the line-minimisation step. It is also important to note that, in the evaluation of $E_{\text{KS}}[\{\phi_i^{(1)}\}]$, v_h and v_{xc} are always constructed from the density

$$\rho^{(1)}(x) = \sum_i |\phi_i^{(1)}(x)|^2. \quad (128)$$

This is the critical difference between the self-consistent field methods and the direct energy minimisation methods. In self-consistent field methods the corresponding optimisation of the orbitals is carried out using the original v_h and v_{xc} . Figure 9 shows a direct comparison between these two approaches for a simulation of silicon, using an 8-atom conventional unit cell. The effect of updating the Kohn–Sham potential at each step along the line minimisation is to increase the curvature of the energy with respect to the step-length α leading to higher energies and an energy minimum at a smaller value of α . In contrast, the over-estimation of α when the energy curve along the steepest descent direction is computed with the fixed initial potential (from $\alpha = 0$) can be considered one of the root causes of charge-sloshing instabilities in self-consistent field methods, see section 3.8.

The direct energy minimisation method discussed thus far is suitable for simulations with a band-gap. For metals and finite temperature insulators, however, it is not sufficient to consider only the lowest N eigenstates as occupied and there is an additional dependence on the partial occupancies f_i . The ground state must now be found by searching over both the Kohn–Sham states and their occupancies. One of the most robust methods of this form is the *ensemble density functional theory* (EDFT) method of Marzari *et al* [26]. In EDFT every update of the trial states $\{\phi_i\}$ is followed by a direct energy-minimisation over the density matrix in the basis of the trial states occupancies, $\{f_{ij}\}$. This density matrix is nothing more than a generalisation of the occupancies to the case when the trial states $\{\phi_i\}$ do not diagonalise the subspace Hamiltonian h directly, where

$$h_{ij} = \int dx \phi_i^*(x) H_{\text{KS}}[\rho] \phi_j(x). \quad (129)$$

5. Test suite

The test suite [147] presented here differs from available test suites (e.g. [35, 36]) as it is required to sample the range of sources of ill-conditioning discussed in section 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, as defined in section 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 56 systems. The geometries and relevant input parameters are given in [147] in the form of CASTEP input files, but are readily converted using, for example, [148].

The content of the test suite is compiled from a range of sources, for example, self-consistency articles (e.g. [73, 89, 92, 122, 123, 131, 144]), collaboration, and online databases. This content is now briefly motivated in terms of the theory in section 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 [47] 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 emphasis 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 [89]. 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).

6. Results and discussion

The aim of this section is to combine to the analysis of section 3, a sample of the methods presented in section 4, and the test suite of section 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 [131]. 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 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 [26], 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

$$\eta = \left(\frac{1}{N_{\text{conv}}} \sum_{i \in \text{conv}} n_i \right)^{-1}, \quad (130)$$

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.

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




Figure 10. Results of the benchmark tests of the algorithms in table 1. 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 11. 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.

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 $T = 300$ 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 2.6 GHz. 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 [147].

The results of the benchmark are given in table 1, and illustrated in figure 10. The first observation of note is that Pulay's algorithm, Kerker preconditioned using the default parameter set [15, 31], is Pareto optimal. In particular, Pulay's method significantly outperforms Broyden's methods in both the singlesecant and multisecant 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. From the algorithms tested here, these potential upgrades include certain parameterisations of Restarted Pulay [112], Periodic Pulay [110], and Marks & Lukes' MSB2 [92]. The parameters used for these methods, as detailed in table 1, 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. Figure 11 illustrates the results of calculations using eight different Kerker parameter sets for Pulay's method




Figure 12. A direct comparison of the efficiency measure of two algorithms. The plot is restricted to the range of 1–70 iterations for illustrative purposes.

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 [15, 31] are found to be approximately optimal. Reducing the history size to $m = 10$ rather than $m = 20$ had a slight stabilising effect.

As expected, EDFT [25] is able to converge the vast majority of the test suite—it is global by design. Note that the method is not 100% robust as two methods took over the maximum allowed time to converge. The cost of global 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 [56, 92, 131], see section 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, global self-consistent field methods. Note that a one-to-one comparison of algorithms can also be illustrated, figure 12. Here, we compare the efficiency of Pulay's method versus 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 figure 10, 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.

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 article is to elucidate these issues and provide a clearer path forward for algorithm development. We began in section 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 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.

The analysis of section 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 [110, 118, 119], and flushing the stored history of iterates after a given number of iterations [112]. Furthermore, considerable promise is shown by more sophisticated methods such as those in [56, 89, 92, 131]. 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 [15, 31] 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 article. Rather, the workflow used in section 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.

Acknowledgments

The authors would like to thank D Bowler, M J Smith, M Hutcheon, A Wade, C J Pickard, and L D Marks for many helpful discussions. NDW is supported by the EPSRC Centre for Doctoral Training in Computational Methods for Materials Science for funding under Grant No. EP/L015552/1. PJH is supported by an EPSRC RSE Fellowship, funded by EPSRC Grant ref EP/R025770/1. MCP acknowledges funding from EPSRC Grant EP/P034616/1.

ORCID iDs

N D Woods  <https://orcid.org/0000-0002-5913-2206>

P J Hasnip  <https://orcid.org/0000-0002-4314-4093>

References

- [1] Martin R M, Reining L and Ceperley D M 2016 *Interacting Electrons Theory and Computational Approaches* (Cambridge: Cambridge University Press)
- [2] Martin R M 2004 *Electronic Structure Basic Theory and Practical Methods* (Cambridge: Cambridge University Press)
- [3] McWeeny R 1992 *Methods of Molecular Quantum Mechanics* (New York: Academic)
- [4] Burke K 2012 *J. Chem. Phys.* **136** 150901
- [5] Kohn W and Sham L J 1965 *Phys. Rev.* **140** A1133
- [6] Hasnip P J, Refson K, Probert M I J, Yates J R, Clark S J and Pickard C J 2014 *Phil. Trans. R. Soc. A* **372** 20130270
- [7] Baerends E J and Gritsenko O V 1997 *J. Phys. Chem. A* **101** 5383–403
- [8] Mardirossian N and Head-Gordon M 2017 *Mol. Phys.* **115** 2315
- [9] Singh D and Nordstrom L 1994 *Planewaves, Pseudopotentials, and the LAPW Method* (Berlin: Springer)
- [10] Gonze X *et al* 2009 *Comput. Phys. Commun.* **180** 2582
- [11] Blaha P, Schwarz K, Sorantin P and Trickey S 1990 *Comput. Phys. Commun.* **59** 399
- [12] Madsen G K H, Blaha P, Schwarz K, Sjö E and Nordström L 2001 *Phys. Rev. B* **64** 195
- [13] Pople J A, Gill P M and Johnson B G 1992 *Chem. Phys. Lett.* **199** 557
- [14] Bowler D R and Miyazaki T 2012 *Rep. Prog. Phys.* **75** 036503
- [15] Kresse G and Furthmüller J 1996 *Comput. Mater. Sci.* **6** 15
- [16] Scuseria G E 1999 *J. Phys. Chem. A* **25** 4782
- [17] Payne M C, Teter M P, Allan D C, Arias T A and Joannopoulos J D 1992 *Rev. Mod. Phys.* **64** 1045
- [18] Frisch M J *et al* 2016 *Gaussian 09, Revision A.02* (Wallingford, CT: Gaussian, Inc.)
- [19] Kudin K N and Scuseria G E 2007 *ESAIM: Math. Modelling Numer. Anal.* **41** 281
- [20] Hartree D R 1928 *Math. Proc. Camb. Phil. Soc.* **24** 111132
- [21] Fock V 1930 *Z. Phys.* **61** 126
- [22] Slater J C 1930 *Phys. Rev.* **35** 210
- [23] Hohenberg P and Kohn W 1964 *Phys. Rev.* **136** B864
- [24] Gross E K U and Dreizler R M 1995 *Density Functional Theory* (Berlin: Springer)
- [25] Marzari N 1996 *Ab initio molecular dynamics for metallic systems PhD Thesis* University of Cambridge
- [26] Marzari N, Vanderbilt D and Payne M C 1997 *Phys. Rev. Lett.* **79** 1337
- [27] van Leeuwen R 2003 *Adv. Quantum Chem.* **43** 25
- [28] LeBris C 2003 *Computational Chemistry, Volume 10 (Handbook of Numerical Analysis)* (Amsterdam: Elsevier)
- [29] Lions P L 1987 *Commun. Math. Phys.* **109** 33
- [30] Heine V 1970 *Solid State Phys.* **24** 1
- [31] Kresse G and Furthmüller J 1996 *Phys. Rev. B* **54** 11169
- [32] Gonze X *et al* 2002 *Comput. Mater. Sci.* **25** 478
- [33] Giannozzi P 2009 *J. Phys.: Condens. Matter* **21** 395502
- [34] Clark S J *et al* 2005 *Z. Kristallogr.* **220** 567–70
- [35] van Setten M J *et al* 2015 *J. Chem. Theory Comput.* **11** 5665
- [36] Lejaeghere K, Van Speybroeck V, Van Oost G and Cottenier S 2014 *Crit. Rev. Solid State Mater. Sci.* **39** 1
- [37] Gill P M 1994 *Adv. Quantum Chem.* **25** 141
- [38] Harris J 1985 *Phys. Rev. B* **31** 1770
- [39] Farid B, Heine V, Engel G E and Robertson I J 1993 *Phys. Rev. B* **48** 11602
- [40] Zaremba E 1990 *J. Phys.: Condens. Matter* **2** 2479
- [41] Prodan E 2005 *J. Phys. A: Math. Gen.* **38** 5647
- [42] Kaiser H C, Neidhardt H and Rehberg J 2009 *Mon. hefte Math.* **158** 179
- [43] Cancès E, Chakir R and Maday Y 2010 *J. Sci. Comput.* **45** 90
- [44] Zeidler E 2011 *Nonlinear Functional Analysis and its Applications: I: Fixed-Point Theorems* (New York: Springer)
- [45] Banach S 1920 *Fundam. Math.* **3** 133
- [46] Latif A 2014 *Topics in Fixed Point Theory* (Cham: Springer)
- [47] Perdew J P, Burke K and Ernzerhof M 1996 *Phys. Rev. Lett.* **77** 3865
- [48] Roothaan C C J 1951 *Rev. Mod. Phys.* **23** 69
- [49] Cancès E 2000 SCF algorithms for HF electronic calculations *Mathematical Models and Methods for Ab Initio Quantum Chemistry (Lecture Notes in Chemistry)* vol 74 (Berlin: Springer)
- [50] Yang C, Gao W and Meza J C 2009 *SIAM J. Matrix Anal. Appl.* **30** 1773
- [51] Krasnosel'skii M A 1955 Two remarks on the method of successive approximations *Technical Report 1* (Uspekhi Mat. Nauk.)
- [52] Mann W R 1953 *Proc. Am. Math. Soc.* **4** 506
- [53] Ryu E K and Boyd S 2016 *Appl. Comput. Math.* **15** 3–43
- [54] Browder F E 1967 *Arch. Ration. Mech. Anal.* **24** 82
- [55] Mizoguchi N and Takahashi W 1989 *J. Math. Anal. Appl.* **141** 177
- [56] Zhang J, O'donoghue B and Boyd S 2018 Globally convergent Type-I Anderson acceleration for non-smooth fixed-point iterations *Technical Report* (arXiv:1808.03971v1)
- [57] Dederichs P H and Zeller R 1983 *Phys. Rev. B* **28** 5462–72
- [58] Boyd S and Vandenberghe L 2014 *Convex Optimization* (Cambridge: Cambridge University Press)

- [59] Davis J B A, Horswell S L and Johnston R L 2014 *J. Phys. Chem. A* **118** 208
- [60] Cancès E and Le Bris C 2000 *Int. J. Quantum Chem.* **79** 82
- [61] Schipper P R T, Gritsenko O V and Baerends E J 1998 *Theor. Chem. Acc.* **99** 329
- [62] Morrison R C 2002 *J. Chem. Phys.* **117** 10506
- [63] Katriel J, Roy S and Springborg M 2004 *J. Chem. Phys.* **121** 12179
- [64] Nesbet R K 1997 *Phys. Rev. A* **56** 2665–9
- [65] Ullrich C A and Kohn W 2001 *Phys. Rev. Lett.* **87** 037402
- [66] Mermin N D 1965 *Phys. Rev.* **137** A1441
- [67] Fu C L and Ho K M 1983 *Phys. Rev. B* **28** 5480
- [68] Verstraete M and Gonze X 2001 *Phys. Rev. B* **65** 035111
- [69] Doll K, Harrison N M and Saunders V R 1999 *J. Phys.: Condens. Matter* **11** 5007
- [70] Fowler A T, Pickard C J and Elliott J A 2019 *J. Phys. Mater.* **2** 34001
- [71] Pickard C J and Needs R J 2011 *J. Phys.: Condens. Matter* **23** 053201
- [72] Lee M, Leiter K, Eisner C, Crone J and Knap J 2015 *Comput. Theor. Chem.* **1062** 24
- [73] Lin L and Yang C 2013 *SIAM J. Sci. Comput.* **35** S277
- [74] Adler S L 1962 *Phys. Rev.* **126** 413
- [75] Wiser N 1963 *Phys. Rev.* **129** 62
- [76] Ashcroft N and Mermin N 1976 *Solid State Physics* (Philadelphia, PA: Saunders)
- [77] Lindhard J 1954 On the properties of a gas of charged particles *Mat. Fys. Medd. K. Dan* **28** 8
- [78] Anglade P M and Gonze X 2008 *Phys. Rev. B* **78** 045126
- [79] Ho K M, Ihm J and Joannopoulos J D 1982 *Phys. Rev. B* **25** 4260
- [80] Sawamura A and Kohyama M 2004 *Mater. Trans.* **45** 1422
- [81] Auer J and Krotscheck E 2003 *Comput. Phys. Commun.* **151** 265
- [82] Annett J F 1995 *Comput. Mater. Sci.* **4** 23
- [83] Nocedal J and Wright S J 1999 *Numerical Optimization* (Berlin: Springer)
- [84] Kerker G P 1981 *Phys. Rev. B* **23** 3082
- [85] Manninen M, Nieminen R, Hautojärvi P and Arponen J 1975 *Phys. Rev. B* **12** 4012
- [86] Auer J and Krotscheck E 1999 *Comput. Phys. Commun.* **118** 139
- [87] Lieb E H 1981 *Rev. Mod. Phys.* **53** 303
- [88] Ghosez Ph, Gonze X and Godby R W 1997 *Phys. Rev. B* **56** 12811
- [89] Marks L D 2013 *J. Chem. Theory Comput.* **9** 2786
- [90] Rohwedder T and Schneider R 2011 *J. Math. Chem.* **49** 1889
- [91] Saad Y 2008 *Numer. Linear Algebr. Appl.* **16** 197
- [92] Marks L D and Luke D R 2008 *Phys. Rev. B* **78** 075114
- [93] Borwein J M, Li G and Tam M K 2017 *SIAM J. Optim.* **27** 1
- [94] Lin L, Xu Z E and Ying L 2017 *Multiscale Modelling Simul.* **15** 29
- [95] Andrade X, Botti S, Marques M A L and Rubio A 2007 *J. Chem. Phys.* **126** 184106
- [96] Broyden C G 1965 *Math. Comput.* **19** 577
- [97] Dennis J E and More J J 1977 *Source: SIAM Rev.* **19** 46
- [98] Dennis J E Jr and Schnabel R B 1979 *SIAM Rev.* **21** 443
- [99] Sherman J and Morrison W 1949 *Ann. Math. Stat.* **20** 261
- [100] Srivastava G P 1984 *J. Phys. A: Math. Gen.* **17** L317
- [101] Eyert V 1996 *J. Comput. Phys.* **124** 271
- [102] Johnson D D 1988 *Phys. Rev. B* **38** 12807
- [103] Vanderbilt D and Louie S G 1984 *Phys. Rev. B* **30** 6118
- [104] Pulay P 1980 *Chem. Phys. Lett.* **73** 393
- [105] Pulay P 1982 *J. Comput. Chem.* **3** 556
- [106] Anderson D G 1965 *J. ACM* **12** 4
- [107] Del Ben M, Hutter J and VandeVondele J 2013 *J. Chem. Theory Comput.* **9** 2654
- [108] Artacho E *et al* 2008 *J. Phys.: Condens. Matter* **20** 064208
- [109] Walker H F and Ni P 2011 *SIAM J. Numer. Anal.* **49** 1715
- [110] Banerjee A S, Suryanarayana P and Pask J E 2016 *Chem. Phys. Lett.* **647** 31
- [111] Pratapa P P, Suryanarayana P and Pask J E 2016 *J. Comput. Phys.* **306** 43–54
- [112] Pratapa P P and Suryanarayana P 2015 *Chem. Phys. Lett.* **635** 69
- [113] Bowler D R and Gillan M J 2000 *Chem. Phys. Lett.* **325** 473
- [114] Blaha P *et al* 2001 *WIEN2k: An Augmented Plane Wave Plus Local Orbitals Program for Calculating Crystal Properties* (Wien: Technische Universität Wien)
- [115] Martinez J M 2000 *J. Comput. Appl. Math.* **124** 97
- [116] Tikhonov A 1963 *J. Phys. A: Math. Gen.* **151** 501
- [117] Powell M J D 1964 *Comput. J.* **7** 155
- [118] Henderson N C and Varadhan R 2018 (arXiv:1803.06673v2)
- [119] Lupo Pasini M 2019 *Numer. Linear Algebr. Appl.* **26** 2241
- [120] Pistol M E 2004 *Chem. Phys. Lett.* **400** 548
- [121] Thøgersen L, Olsen J, Yeager D, Jørgensen P, Sałek P and Helgaker T 2004 *J. Chem. Phys.* **121** 16
- [122] Høst S, Jansík B, Olsen J, Jørgensen P, Reine S and Helgaker T 2008 *Phys. Chem. Chem. Phys.* **10** 5344
- [123] Thøgersen L, Olsen J, Köhn A, Jørgensen P, Sałek P and Helgaker T 2005 *J. Chem. Phys.* **123** 074103
- [124] Saunders V R and Hillier I H 1973 *Int. J. Quantum Chem.* **7** 699
- [125] Zhou B and Wang Y A 2008 *J. Chem. Phys.* **128** 084101
- [126] Chen Y K and Wang Y A 2011 *J. Chem. Theory Comput.* **7** 3045
- [127] Wang Y A, Yam C Y, Chen Y K and Chen G 2011 *J. Chem. Phys.* **134** 241103
- [128] Francisco J B, Martinez J M and Martinez L 2004 *J. Chem. Phys.* **121** 10863
- [129] Cancès E and Le Bris C 2000 *ESAIM: Math. Modelling Numer. Anal.* **34** 749
- [130] Garza A J and Scuseria G E 2012 *J. Chem. Phys.* **137** 054110
- [131] Kudin K N, Scuseria G E and Cancès E 2002 *J. Chem. Phys.* **116** 8255
- [132] Garza A J and Scuseria G E 2015 *J. Chem. Phys.* **142** 164104
- [133] Chen X and Kelley C T 2019 *SIAM J. Sci. Comput.* **41** 365
- [134] Cancès E 2001 *J. Chem. Phys.* **114** 10616
- [135] Defranceschi M and Le Bris C 2000 *Mathematical Models and Methods for Ab Initio Quantum Chemistry (Lecture Notes in Chemistry vol 74)* (Berlin: Springer)
- [136] Kelley C T 1995 *Iterative Methods for Linear and Nonlinear Equations* (Philadelphia, PA: SIAM)
- [137] Wang J, Wang Y and Yu S 2008 *Modelling Simul. Mater. Sci. Eng.* **16** 035004
- [138] Hasnip P J and Probert M I J 2015 (arXiv:1503.01420)
- [139] Zhou Y, Wang H, Liu Y, Gao X and Song H 2018 *Phys. Rev. E* **97** 33305
- [140] Levine Z H and Louie S G 1982 *Phys. Rev. B* **25** 6310
- [141] Krotscheck E and Liebrecht M 2013 *J. Chem. Phys.* **138** 164114
- [142] Sternheimer R M 1954 *Phys. Rev.* **96** 951
- [143] Giustino F, Cohen M L and Louie S G 2009 *Phys. Rev. B* **81** 115105
- [144] Raczkowski D, Canning A and Wang L W 2001 *Phys. Rev. B* **64** 121101
- [145] Woods N 2018 (arXiv:1803.01763)
- [146] Car R and Parrinello M 1985 *Phys. Rev. Lett.* **55** 2471
- [147] Woods N The scf-xn test suite of Kohn-Sham input systems <https://bitbucket.org/NickWoods1/scf-xn-testsuite>
- [148] Rutter M 2018 *Comput. Phys. Commun.* **225** 174
- [149] Barzilai J and Borwein J 1988 *IMA J. Numer. Anal.* **8** 141