

Prediction and Inference of Dynamics in Quantum Plasmas



BRETT LARDER

CHRIST CHURCH

UNIVERSITY OF OXFORD

A Thesis Submitted for the Degree of

Doctor of Philosophy

Trinity 2020

Acknowledgements

I would like to thank my colleagues and collaborators at Oxford. Foremost of these, I'd like to thank Prof. Gianluca Gregori, my supervisor, for his invaluable guidance and patience. I would also like to thank Scott Richardson at AWE for his tireless support and countless helpful discussions.

Thanks, Mum.

Thanks, Emily, for everything.

Copyright Notice

This work has received support from AWE PLC, the Engineering and Physical Sciences Research Council (grant numbers EP/M022331/1 and EP/N014472/1), and the Science and Technology Facilities Council of the United Kingdom. The material of chapter 3 is partially based on work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Science under award number DE-SC0019268. Ministry of Defence © British Crown Copyright 2020/AWE.

Source Material and the Author's Contributions

This thesis describes work appearing in both journal publications and manuscripts that are currently being prepared for submission. The source material for each chapter – and the author's individual contribution to that source material – is clarified here.

Chapter 1 provides a brief, high-level description/abstract of the goals of the work developed in this thesis, alongside an outline of the thesis' structure. It contains no original work.

Chapter 2 examines the literature on computational methods for WDM. It contains no original work, but develops the theoretical foundation for the new work in chapters 3, 4 and 5.

Chapter 3 is an adapted and expanded version of the paper *Fast nonadiabatic dynamics of many-body quantum systems* (Science Advances)[1], of which the present author is the primary author. It has (very slight) overlap with the work presented in the paper *Modified Friedmann Equations via Conformal Bohm-de Broglie Gravity* (The Astrophysical Journal)[2], to which the present author provided some minor theoretical support (as a coauthor). The chapter concerns the development of a simulation method for quantum plasma dynamics, based on Bohmian Mechanics,

which retains electron-ion dynamic interactions while being sufficiently inexpensive to resolve the DSF.

Gianluca Gregori first suggested looking to Bohmian Mechanics to develop a new computational method for WDM, and he and Dirk Gericke provided many useful pointers on the corresponding computational work (and assisted the author in writing the paper itself). The reference VASP DFT-MD calculations exploited in section 3.5 were carried out by Scott Richardson, and further DFT results were provided by Paul Mabey and Thomas White for comparison.

The remainder of the theoretical and computational work of this chapter was carried out by the author (with the exception of cited external work).

Chapter 4 is an adapted and expanded version of the paper *Inverse Problem Approaches for Plasma Physics I: Static Properties* (in preparation), of which the present author is the primary author. It develops statistical techniques – for use in combination with quantum plasma simulations – that yield comprehensive insights into static plasma properties from experimental measurements.

The theoretical and computational work of this chapter was carried out by the author (with the exception of cited external work).

Chapter 5 is an adapted and expanded version of the paper *Inverse Problem Approaches for Plasma Physics II: Dynamic Properties* (in preparation), of which the present author is the primary author. It develops the theory of the DSF under stochastic conditions, and – by bringing together the core results from chapter 4 and the simulation methods of chapters 2 and 3 – it demonstrates a pathway to inference of dynamic plasma properties from experimental measurements.

The theoretical and computational work of this chapter was carried out by the author (with the exception of cited external work).

Chapter 6 summarises the overall contribution of the work in this thesis, and develops proposals for extensions of the work in chapters 3, 4 and 5. While some new algorithms are constructed in detail, validation and testing of the proposals in this chapter are left to future research.

The theoretical and computational work of this chapter was carried out by the author (with the exception of cited external work).

Contents

1 Introduction	15
1.1 Thesis Objective/Abstract	15
1.2 Thesis Layout	16
2 Background	19
2.1 Introduction	19
2.2 Modelling Quantum Plasma	22
2.2.1 Analysis Through the Dynamic Structure Factor	24
2.2.2 Linear Response Theory	26
2.3 Core Computational Approaches	38
2.3.1 Classical Molecular Dynamics	39
2.3.2 Classical Monte Carlo Methods	44
2.3.3 Integral Equation Methods	46
2.3.4 Quantum Monte Carlo Approaches for Static Structure	48
2.3.5 Density Functional Theory	54
2.3.6 Density Functional Theory Molecular Dynamics	57
2.3.7 Quantum Statistical Potentials	58
3 Predicting Dynamic Structure: The Bohmian Trajectory Method	65
3.1 Introduction	65
3.2 Theory	68
3.2.1 Bohmian Mechanics	69
3.2.2 Quantum Hydrodynamics and Linearisation	73
3.2.3 Linearisation for Bohmian Trajectories	77
3.3 Elements of the Bohmian Trajectory Method	80
3.3.1 Correlation Closure	81
3.3.2 Potential Corrections	82
3.3.3 Pseudopotentials	84
3.3.4 Setting λ	86
3.3.5 Fermi Thermostats	91
3.4 Implementation	96
3.4.1 Software Overview	99

Contents

3.4.2	Building Potentials	100
3.5	Validation and Predictions	101
4	Inferring Static Structure: Inverse Problem Approaches for Plasmas	107
4.1	Introduction	107
4.2	Markov Chain Monte Carlo	109
4.2.1	Basic Methods	111
4.2.2	Gradient-Driven Approaches	112
4.2.3	Automatic Differentiation	115
4.2.4	Stochastic Gradient Methods	116
4.3	Physical Parameter Inversion via SGMCMC	120
4.3.1	Unbiased Estimation Methods	121
4.3.2	Combining Multiple Simulation Fidelities	124
4.3.3	Variance Estimation and Corrections	132
4.3.4	Early Estimation and Conditioning	134
4.3.5	A (Very) Brief Note on Priors	136
4.4	Application: Proton Radiography	137
4.4.1	Problem Overview	138
4.4.2	Field Inversion via Linearised Maps	141
4.4.3	General Field Inversion	146
4.5	Application: Particulate Systems	148
4.5.1	Problem Overview	150
4.5.2	Configuration Inversion	151
4.5.3	General Error Function Inversion	156
5	Inferring Dynamic Structure: Analysis of DSF Moments via SGMCMC	161
5.1	Introduction	161
5.2	Moment Methods and Stochastic Calculus	162
5.2.1	Moments Under Smooth Reversible Dynamics	163
5.2.2	Stochastic Calculus	165
5.2.3	Langevin Moment Modifications	166
5.2.4	Practical Sampling of Moments	171
5.2.5	Memory Function Methods	173
5.2.6	Setting α Theoretically	182
5.2.7	Inferring α from Trajectory Data	187
5.3	Inference of Dynamic Properties via Moments	189
5.3.1	Static Sampling of Moment Gradients	190
5.3.2	Application: Yukawa+SRR with Langevin Noise Model	191
5.3.3	Extension to Arbitrary Properties: a Brief Outline	193
5.3.4	Extension to Mixed States: a Brief Outline	195
6	Future Directions and Summary	197
6.1	Introduction	197

6.2	Extending (or Overhauling) the Bohm Approach	198
6.2.1	Configuration Monte Carlo	199
6.2.2	Conditional Wavefunction Sampling	201
6.2.3	Other Possibilities	204
6.3	Extending the Inverse Problem Approach	206
6.3.1	Further Exploiting Connections to MD	207
6.3.2	Underlying Estimator Improvements	207
6.3.3	Other Possibilities	207
6.4	Summary	208

1

Introduction

1.1 Thesis Objective/Abstract

This thesis will focus on fast numerical calculations of properties of quantum plasmas. With this goal in mind, we will develop, implement and validate computational methods of two related types, which together provide a comprehensive insight into quantum plasma dynamics: methods that predict plasma properties from first principles physics, and methods that infer plasma properties from experimental measurements.

Quantum plasmas arise in a range of physical contexts, from planetary interiors to Inertial Confinement Fusion experiments. We require sophisticated numerical approaches to find the structural and dynamic properties of these complex, many-body quantum systems; diverse in accuracy and applicability, these computational approaches each have their own strengths and (often ignored) weaknesses.

By combining existing techniques with newly developed algorithms, we aim to address some weaknesses that are common to many state-of-the-art calculation methods, and open up paths to calculating plasma properties that were not previously accessible.

1.2 Thesis Layout

In chapter 2, we will review the literature on quantum plasmas and Warm Dense Matter, covering existing computational and theoretical approaches to calculating system dynamics (with emphasis on the Dynamic Structure Factor). Beginning with theoretical models via Linear Response Theory, the Random Phase Approximation and its extensions, we will go on to discuss modern numerical calculation techniques.

Chapter 3 then draws on the theory discussed in chapter 2, in combination with the theory of Bohm's interpretation of quantum mechanics, to develop a new trajectory-based approach to modelling quantum plasma dynamics. The chapter begins with an overview of Bohmian mechanics, develops a complete set of simulation algorithms, and then illustrates predictive results from a complete implementation.

Following this, we will consider the inverse problem; where chapter 3 develops a method for calculating experimental outputs from given plasma properties, chapters 4 and 5 will develop methods for calculating plasma properties from experimental outputs. Chapter 4 discusses Markov Chain Monte Carlo methods, domain specific extensions, and applications to plasma experiments; chapter 5 then extends these results to inference of dynamic quantum plasma properties, by applying the method of moments to stochastic models of ion dynamics and inferring dynamic parameters of a semi-classical plasma.

A final chapter, chapter 6, then discusses future directions for this research. It begins with a collection of detailed proposals for extending the Bohmian trajectory method, including derivations of new algorithms for this purpose. It then highlights how the Markov Chain Monte Carlo methods we have developed can be

1 Introduction

extended to further reduce simulation times in the future. Finally, the contributions to the field made by the range of theoretical and computational techniques we have developed are summarised.

2

Background

2.1 Introduction

Throughout this thesis, our focus will be on the physics of plasmas, and the calculation of plasma properties. We can define a plasma as an ionised gas, comprised of an interacting collection of charged particles[3]. Plasmas arise in high-temperature, high-pressure, and high-electromagnetic-field conditions, in which electrons have been partially or fully stripped from nuclei. Plasma structure is heavily influenced by the Coulomb interactions between its components, to the point that its properties are closer to those of fluids than those of standard gases. Long range correlations and a rich variety of collective wave dynamics and instabilities, in particular, distinguish plasma as its own state of matter.

We will be particularly interested in quantum plasma. At lower temperatures – before we obtain fully-ionised, effectively classical plasma – we find partially-ionised systems in which electrons exhibit quantum mechanical behaviour. The term quantum plasma can then be applied to these systems quite broadly – a quantum plasma is simply a plasma in which quantum mechanical effects significantly impact the structure and dynamics. As a loose quantification, this occurs when the De Broglie wavelength of the electrons is comparable the mean particle separation (or, equivalently, the Wigner-Seitz radius); the overlap of electronic wavefunctions means that Fermi-Dirac statistics come into play, shifting the system’s behaviour away from the classical regime[4].

To make our discussions more concrete, it will be useful to narrow our focus yet further, on a specific temperature/density domain – that of Warm Dense Matter. Warm Dense Matter (wDM) lies in the difficult middle ground between classical plasma and condensed matter. Spanning temperatures of order of a few electron volts (eV) and densities on the order of solid density, wDM plasmas form a key con-

stituent of large planet cores[5, 6], and occur as a transitory step during processes in which plasmas are formed from solids[7]. Further, Inertial Confinement Fusion (ICF) requires heating of a capsule target through the WDM regime on the path to ignition[8, 9]; as a result, understanding the dynamics of WDM will prove important to the optimisation of fusion reactors.

Under WDM conditions, ions can usually be treated as classical bodies (as we will discuss in detail in section 2.3); however, as a subset of quantum plasma, WDM requires a treatment of the quantum many-body problem for the electrons. The majority of modern WDM models that we will discuss begin from first principles – so-called *ab initio* methods. As WDM falls so far from both classical plasma on the one hand and from condensed matter on the other, perturbative techniques and linearised wave theories are ineffective; instead, we are usually forced to begin from the absolute basics – Newton’s laws and the Schrödinger equation – and build a microscopic description of the dynamics from there.

Ab initio methods have the advantage of including all of the rich physics of plasmas implicitly, without manual introduction of individual phenomena – that is to say, we needn’t introduce collective plasma behaviours and instabilities to our models artificially. The downside is that the basic equations of motion turn out to be incredibly difficult to tackle. Solving the Schrödinger equation exactly for thousands of interacting particles is beyond a lost cause – even the classical many-body problem is far from trivial. But – as we’ll show – if we frame the problem in the right way and choose the right approximations, we can describe the microscopic dynamics of these difficult systems with surprising accuracy.

This chapter lays out the groundwork for the rest of this thesis. We begin with a discussion of the Dynamic Structure Factor, a property of quantum plasmas/WDM that will be of particular importance in our modelling, and examine some of its key

2 *Background*

attributes via Linear Response Theory. From there, we survey existing ab initio simulation methods for WDM; the following chapters then extend these methods and apply them as part of improved prediction and inference schemes.

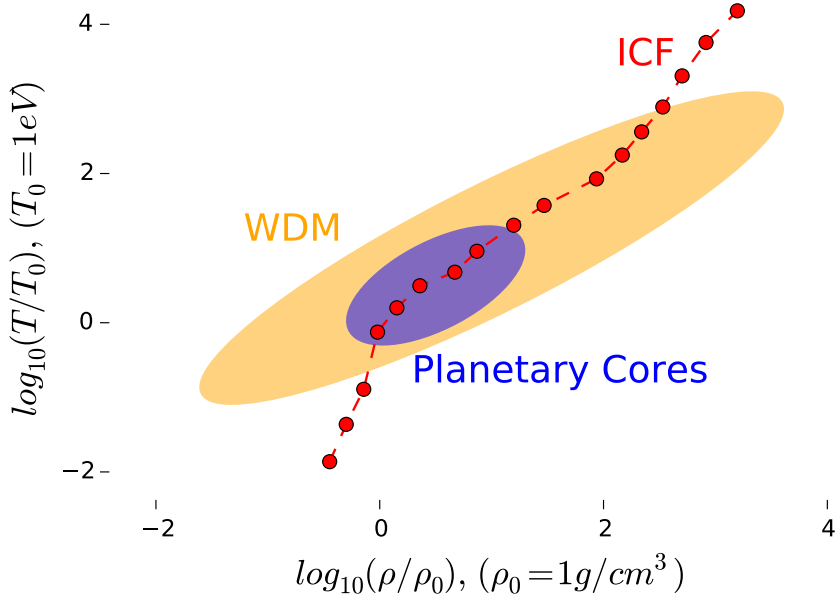


Figure 2.1: A schematic of the (very) rough region of temperature-density space occupied by Warm Dense Matter. During a typical inertial confinement fusion experiment - marked ICF - the deuterium-tritium fuel passes straight through wdm territory on its path towards ignition (data from ref. [10]). Planetary cores are also predominantly wdm.

2.2 Modelling Quantum Plasma

2.2.1 Analysis Through the Dynamic Structure Factor

Before delving into theory and computational methods, it will be helpful to isolate a specific target property of WDM systems that we wish to calculate. Through the lens of a specific dynamic quantity, our stated goal of modelling WDM dynamics becomes more precise, and our accuracy more measurable. Aside from standard properties such as the system pressure, temperature, density and so on, the Dynamic Structure Factor (DSF) of a quantum plasma is a strong candidate for such a target, fulfilling two key criteria: The DSF can be measured directly by experiment, allowing comparison between model and reality[11, 12, 13]; and key structural and dynamic quantities can be derived from the DSF, allowing a broad and detailed description of the physics of the plasma. These features have made the DSF a mainstay of the WDM literature, acting as a primary link between theoretical and experimental work[14, 15].

We can express the DSF via:

$$S(q, \omega) = \frac{1}{2\pi N} \int \exp(i\omega t) \langle \rho(q, \omega) \rho(-q, 0) \rangle dt, \quad (2.1)$$

where ρ is the Fourier space electron density. This is typically decomposed via the Chihara formula[16],

$$S(q, \omega) = |f_I(q) + \rho(q)|^2 S_{ii}(q, \omega) + Z_f S_{ee}(q, \omega) + S_{bf}(q, \omega), \quad (2.2)$$

where $f_I(q)$ and $\rho(q)$ are form factors describing the bound state and screening electron cloud surrounding each ion, $S_{ii}(q, \omega)$ is the dynamic ion-ion structure fac-

tor, $S_{ee}(q, \omega)$ is the free electron-electron structure factor (with a prefactor of the ion charge Z_f), and $S_{bf}(q, \omega)$ is a contribution due to bound-free transitions. Our primary interest throughout this thesis is the ion-ion term, which encodes information on the dynamic structure and modes of the ions. In particular, we concentrate on the low-frequency contribution to the ion-ion term – this contribution is neglected or poorly estimated by the most commonly employed computations, motivating our new methods.

To connect the DSF to experiment, we consider inelastic X-ray scattering from a WDM material, following the analysis of ref. [17]. The doubly differential cross section, which dictates the relative scattering intensity into solid angle range $d\Omega$ and frequency range $d\omega$, can be written as:

$$\frac{d^2\sigma}{d\Omega d\omega} = \frac{1}{N} \frac{j_1(r, k_1, \omega) \cdot \hat{l} r^2}{j_0}, \quad (2.3)$$

where N is the number of scatterers in the target, $j_1(r, k_1, \omega)$ is the current density of scattered particles at position r with wavevector k_1 and frequency ω , \hat{l} is a unit vector along r , and j_0 is the input current density (along wavevector k_0 , with energy ω_0). The numerator of this expression can be written in terms of the transition rate through:

$$N_0 \sum_{n_0, n_1} P_{n_0} w(n_0, k_0 \rightarrow n_1, k_1) \frac{V}{(2\pi)^3} k_1^2 d\Omega dk_1, \quad (2.4)$$

where n_0, n_1 are the values of a quantum number describing our scattering system before and after the scattering respectively, P_{n_0} is the probability of our system starting at n_0 , N_0 the number of particles starting with wavevector k_0 , and $w(n_0, k_0 \rightarrow n_1, k_1)$ is the transition rate between the given quantum numbers in a scattering event. In the Born approximation, this transition rate is proportional to

2 Background

the square of the matrix element of the interaction Hamiltonian, \hat{H}_{inter} , between the n_0, k_0 and n_1, k_1 states.

For inelastic X-ray scattering, the contributing term to \hat{H}_{inter} is quadratic in the photon field's vector potential (with linear contributions corresponding to absorption/emission). Through substituting an interaction Hamiltonian of this form and simplifying the result, our doubly differential cross section becomes proportional to:

$$\frac{1}{N} \sum_{n_0, n_1} P_{n_0} |\langle n_1 | \sum_j e^{iqr_j} | n_0 \rangle|^2 \delta(E_{n_0} - E_{n_1} + \hbar\omega - \hbar\omega_0), \quad (2.5)$$

where j enumerates the scatterers in the system, and E_n is the energy of the system when it has quantum number n . It is then just a matter of operator manipulation (see ref. [17]) to show that:

$$\frac{d^2\sigma}{d\Omega d\omega} \propto S(q, \omega), \quad (2.6)$$

where $S(q, \omega)$ is as defined in eq. 2.1.

While the principal contribution to the ionic part of $S(q, \omega)$ is due to direct Coulomb interactions between the ions, the details of the quantum interactions between electrons and ions - and electrons and electrons - have a substantial effect on the quantity. In turn, the DSF of a system embodies a great deal of dynamic and structural information. In the next section, our treatment of the direct relationship between the DSF and the linear response function will make this clear.

2.2.2 Linear Response Theory

Introduction and Definitions

We now have a target system - WDM - and a target property of this system that we wish to calculate - the DSF. Before turning to numerical methods, we will first attempt to make progress theoretically. Linear Response Theory allows us to do so,

providing both our first attempt at a calculation of the DSF, and also insight into its broader importance.

Linear Response Theory provides the general theoretical tools for establishing material response to perturbation. The results from this theory are important from two points of view. Firstly, we can produce an initial approximation to the structure of an electron fluid. Through this, we can establish how electrons behave under strong coupling and in the presence of ions, constituting our first attempt at a model of WDM. Secondly, the theory provides the links between how a material responds to stimuli and its unperturbed properties. This allows inference of properties of interest from simulation results and experiment, by providing the necessary relationships between measurable properties (for example, the DSF) and bulk properties such as viscosity, conductivity etc..

We follow Giuliani and Vignale[18] in introducing the elementary concepts of the theory. A linear response function $\chi_{yh}(t)$ can be defined in terms of a property of the examined system $y(t)$, and a perturbing input $h(t)$, via:

$$y(t) = \int_{-\infty}^t dt' \chi_{yh}(t-t')h(t') + \dots \quad (2.7)$$

where the ellipsis captures corrections to y non-linear in h (which are neglected for small h). In this sense χ_{yh} encodes the causal response of the system to the input, specific to the particular property y and perturbation h . In a quantum mechanical context, the perturbing input is usually a potential term added to the Hamiltonian.

We are interested in the response of a quantum system to an external potential that couples linearly to the density, as would be encountered in experiment, so that experimental measurements can be linked to the system properties. The general

2 Background

form for a Hamiltonian \hat{H}' that has been perturbed by such an external potential is:

$$\hat{H}' = \hat{H} + \int V_{ext}(x, t) \hat{n}(x) dx, \quad (2.8)$$

and the associated induced density change $n_1(x, t)$ can be written within linear response theory as:

$$n_1(x, t) = \int_0^\infty d\tau \int dx' \chi_{nn}(x, x', \tau) V_{ext}(x, t - \tau), \quad (2.9)$$

where we have introduced χ_{nn} as

$$\chi_{nn}(x, x', t) = \chi_{n(x)n(x')}(t) = -\frac{i}{\hbar} \Theta(t) \langle [\hat{n}(x, t), \hat{n}(x', 0)] \rangle_0, \quad (2.10)$$

with \hat{n} the density operator and Θ a step function. Due to our interest in translationally invariant systems (i.e. bulk plasmas) and driving by periodic potential (i.e. electromagnetic probes), we can simplify matters by examining linear response in momentum-frequency space. We define:

$$\begin{aligned} \chi_{nn}(q, q', \omega) &= \frac{1}{L^d} \int d^d x \exp(-iq \cdot x) \\ &\int d^d x' \exp(iq' \cdot x') \\ &\int dt \chi_{nn}(x, x', t) \exp(i\omega t). \end{aligned} \quad (2.11)$$

From this and equation 2.10 it is straightforward to see that

$$\chi_{nn}(q, q', \omega) = \frac{1}{L^d} \chi_{n_q n_{-q'}}(\omega), \quad (2.12)$$

where $n_q = \sum_{\text{positions}, i} \exp(-iq \cdot \hat{x}_i)$ is the Fourier-transformed density fluctuation operator. Translational invariance grants us $\chi_{nn}(q, q', \omega) = \chi_{nn}(q, \omega)\delta_{q, q'}$; our interest then lies in the calculation and measurement of this simplified quantity.

Connection to experiment - the Fluctuation-Dissipation Theorem

We now consider the connection between the linear response function and experiment, which arises from a relationship between the dynamic structure factor and the imaginary part of the linear response function: the Fluctuation-Dissipation theorem. This theorem - and related results - provide the key motivation for our interest in linear response: We find connections between experimental measurements, structural properties of simulations, and physical properties of the underlying system. By expanding the commutator in eq. 2.10 (and hence implicitly in eq. 2.11), we can determine the makeup of χ in terms of the density matrix elements and occupation probabilities, facilitating some insight into its structure. In terms of occupation probabilities P_n , the spectral representation of a general response function takes the form (for operators with matrix elements A_{ij}, B_{ij}):

$$\chi_{AB}(\omega) = \frac{1}{\hbar} \sum_{mn} \frac{P_m - P_n}{\omega - \omega_{nm} + i\eta} A_{mn} B_{nm}, \quad (2.13)$$

where $\eta \rightarrow 0^+$. For the density-density response this becomes

$$\chi_{nn}(q, \omega) = \frac{1}{\hbar L^d} \sum_{nm} \frac{P_m - P_n}{\omega - \omega_{nm} + i\eta} |(\hat{n}_q)_{nm}|^2. \quad (2.14)$$

Now we can compare this to the dynamic structure of the system. The DSF can similarly be written in a spectral form:

$$S(q, \omega) = \sum_{nm} P_m |(\hat{n}_q)_{nm}|^2 \delta(\omega - \omega_{nm}). \quad (2.15)$$

2 Background

Noting that inversion symmetry implies $S(-q, \omega) = S(q, \omega)$ (in combination with eq. 2.15), we can obtain further that $S(q, -\omega) = e^{-\beta\hbar\omega} S(q, \omega)$. Now, examining the imaginary part of χ and comparing its form to eq. 2.15, we can read off the relationship

$$S(q, \omega) = \frac{\hbar \operatorname{Im} \chi(q, \omega)}{\pi n (\exp(-\beta\hbar\omega) - 1)}, \quad (2.16)$$

which is the density form of the Fluctuation Dissipation theorem. The direct density correlations determining S (described by eq. 2.1) can then be extracted from simulation, or indirectly determined in experiment as a consequence of the scattering relation derived in eq. 2.6 (see for example ref. [19]), so that theory, simulation and experiment can be validated against one another.

Known Properties of the Linear Response Function - Sum Rules

It is useful here to very briefly discuss some known, exact properties of the linear response function for completeness, given that its detailed calculation (which we will first attempt in section 2.2.2) is involved and requires approximation. These exact properties then transfer into exact properties of the DSF through eq. 2.16, which will simplify our analysis going forward.

Firstly, the causality of the linear response function χ provides important relations between its real and imaginary parts (Kramers-Kronig relations):

$$\operatorname{Re} \chi(r, r', \omega) = \frac{1}{\pi} P \int d\omega' \frac{\operatorname{Im} \chi(r, r', \omega')}{\omega' - \omega}, \quad (2.17)$$

$$\operatorname{Im} \chi(r, r', \omega) = \frac{-1}{\pi} P \int d\omega' \frac{\operatorname{Re} \chi(r, r', \omega')}{\omega' - \omega}. \quad (2.18)$$

In combination with the density-density expectation expression

$$\langle [\hat{n}(r, t), \hat{n}(r', t')] \rangle = - \int \frac{d\omega}{\pi} \operatorname{Im} \omega(r, r', \omega) \exp(-i\omega(t - t')), \quad (2.19)$$

we can extract a series of properties of the response functions known as sum rules by taking successive derivatives with respect to t . In addition to providing insight into the response function structure, these rules give useful limits to which theoretical and computational models can be compared. From a theoretical perspective, the rules allow systematic construction of response functions by employing an ansatz representation constrained to satisfy the first n rules (ref. [20] illustrates this technique, and we will explore this process in detail in Chapter 5). Expanding the high frequency limit of the real part of the response function in reciprocal space, we have

$$\text{Re } \chi(q, \omega) = \frac{M_2^\chi(q)}{\omega^2} + \frac{M_4^\chi(q)}{\omega^4} + \dots, \quad (2.20)$$

and from this and the Kramers-Kronig relations we can see, as an example, the first sum rule (the f-sum rule):

$$M_1^\chi(q) = \frac{nq^2}{m}. \quad (2.21)$$

The simplicity of the first sum rule also allows its direct use in normalising scattering signals from experiments to obtain structure factors. Through this rule, and the infinite sequence of further rules produced by differentiation, we thus have an important set of validation and normalisation tools. In the following section, we consider the RPA and corrections to it in the form of local field factors – the sum rules provide important limits on these factors, and have been used to construct increasingly accurate analytic representations of the linear response function.

Calculating Electronic Linear Response – Random Phase Approximation

As a result of its connection to structural properties and experimentally measurable quantities – as outlined in the preceding sections – the utility of the linear response function is clear. We now complete the path between theory and experiment by calculating the quantity analytically for an electron gas, providing a baseline approximation to the electronic subsystem of a WDM plasma.

For the non-interacting electron gas, the density-density response function can be calculated readily from the spectral representation eq. 2.15. In this limit, we obtain the Lindhard response function:

$$\chi_0(q, \omega) = \frac{2}{V} \sum_k \frac{f(\epsilon_{k-q}) - f(\epsilon_k)}{\omega + i\delta + (\epsilon_{k-q} - \epsilon_k)}, \quad (2.22)$$

from which the structural properties of the non-interacting gas follow from eq. 2.16.

This function provides a strong starting point for extending to interacting electron systems (and is also valid for finite temperatures). Our work on DSF calculations then aims, indirectly, to calculate the linear response for the case of moderately to strongly interacting ion/electron fluids in the WDM regime.

The Lindhard response function cannot capture the more complex interacting electronic system, and the spectral representation does not allow us to immediately write down a simple form similar to eq. 2.22. The natural question to ask is: How, from a theoretical perspective, can we extend the simple non-interacting theory above into the interacting case? This is a fundamental problem in many-body quantum physics, and several routes of varying complexity have been developed. Historically, the first major step in modelling bulk electronic linear response came with the Random Phase Approximation (RPA).




Figure 2.2: Zeroth order diagram for the density-density response function, corresponding to the Lindhard response.




Figure 2.3: The diagrams that are summed to form the RPA response function. This corresponds to approximating the proper response function by the Lindhard function.

2 Background

The RPA, outlined originally in a series of papers of Bohm and Pines[21] in the early 1950s, provided a key step towards modelling quantum systems at high density. In terms of linear response theory, the simplest, most intuitive description of RPA is a formulation in terms of the stiffness theorem. Following Giuliani and Vignale[18], we consider for simplicity the zero temperature case (although the finite temperature generalisation follows simply). Consider a general system with expectation of an operator \hat{A} equal to 0 in its ground state. Under the constraint of a specific expectation value of \hat{A} , we define the minimum energy E via

$$E(A) = \min_{\psi \rightarrow A} \langle \psi | \hat{H} | \psi \rangle, \quad (2.23)$$

such that for small A , E can be expanded as

$$E(A) = E(0) + \alpha \frac{A^2}{2} + \dots, \quad (2.24)$$

where we assume E is an analytic function. The coefficient $\alpha > 0$ then defines the stiffness with respect to A . The stiffness theorem can be stated as:

$$\alpha = -\frac{1}{\chi_{AA}(0)}, \quad (2.25)$$

where the straight-forward proof can be found in ref. [18]. The RPA (firstly in the static case) can be formulated through approximating the energy term $\delta E = \alpha A^2/2$. For the case of density-density response, the stiffness theorem takes the form

$$\delta E = -\frac{n_q^2}{\chi_{nn}(q, 0)V}, \quad (2.26)$$

where V is the system volume. If we approximate this as the sum of terms due to electrostatic (Hartree) energy,

$$\delta E_C = \frac{v_q n_q^2}{V}, \quad (2.27)$$

and the stiffness of the non-interacting electron gas,

$$\delta E_0 = -\frac{n_q^2}{\chi_0(q, 0)V}, \quad (2.28)$$

then by comparing the quantities in the equation $\delta E = \delta E_C + \delta E_0$ we arrive at the RPA:

$$\chi_{nn}^{RPA}(q, 0) = \frac{\chi_0(q, 0)}{1 - v_q \chi_0(q, 0)}. \quad (2.29)$$

This expression can be shown to hold in the general finite frequency case, in addition (the diagrammatic formalism in the following section provides a clean route to this, for example). The RPA is thus just a time-dependent Hartree approximation to the linear response.

The principal success of the RPA beyond the Lindhard function is the inclusion of plasmon physics – that is to say, it exhibits poles corresponding to resonant driving of plasma oscillations in the electron liquid. This feature dominates the behaviour of the RPA response function at low frequency, and is easily demonstrated by examining the pole structure (i.e. examining the behaviour of χ_{nn}^{RPA} as $1 - v_q \chi_0(q, \omega) \rightarrow 0$, corresponding to a collective excitation). The RPA becomes increasingly accurate at higher densities/longer wavelengths, and despite its simplicity captures the correlation physics essentially exactly in these limits. However, there are multiple flaws to RPA's predictions, linked primarily to the failings of the Hartree approximation (principally due to the neglect of exchange effects). The breakdown at small wavelengths, arising due to complete neglect of

2 Background

short-ranged local field corrections/the exchange hole, manifests itself clearly in unphysical negative densities when the radial distribution function centred on an electron is calculated. Accurate prediction of correlations at small separation thus requires improvements over the basic RPA description.

Extensions to the RPA to correct its deficiencies away from the high density and long wavelength limits have been the subject of a great deal of study, and are typically formulated in terms of an approximate many-body local field factor G . Introducing additional terms to the stiffness energy (eq. 2.26) provides a natural way to extend the RPA - we define G via an additional exchange-correlation contribution to the energy:

$$\delta E_{exc} = -\frac{v_q G(q) n_q^2}{V}, \quad (2.30)$$

giving the more general linear response function of

$$\chi_{nn}(q, 0) = \frac{\chi_0(q, 0)}{1 - v_q(1 - G(q))\chi_0(q, 0)}, \quad (2.31)$$

with dynamic generalisation obtained by setting $\omega \neq 0$. Many of the properties of G can then be connected back to the sum rules, as highlighted at the end of section 2.2.2, allowing constraints to be imposed on its form. To calculate the quantity systematically, however, requires some additional mathematical apparatus. The diagrammatic Green function formalism can be employed to calculate χ through a series expansion, although this approach quickly runs into limitations. Individual perturbation terms are often divergent, and instead sets of diagrams that form finite sums must be considered collectively (this is known as 'partial summation'); the remaining errors in χ are non-trivial to characterise.

The Green function method relies on a broad range of related theorems (see ref. [18]), and so for brevity we only mention the mechanics of the calculations at a ba-

sis level. The density linear response can be determined by forming diagrams with two external points representing density operators (with momenta $q, -q$), connecting these points with combinations of solid (Green function) lines and dashed (interaction) lines, and conserving momentum/energy associated with these lines at vertices. The diagrams can then be translated into terms in the perturbation expansion of χ when equipped with a set of rules linking them back to mathematical expressions (the rules can be straightforwardly, if somewhat tediously, derived by simply examining terms in the perturbation expansion). The Lindhard function is just the simplest diagram of this type (see fig. 2.2).

Within the set of possible first-order diagrams, there exists the first example of an ‘improper’ diagram. When performing calculations, interior momenta values on lines are integrated over, in accordance with the rules mapping them back to mathematical expressions; however, in the first order, there is a diagram in which we find a single interaction line with momentum fixed to q (as opposed to a free variable). This diagram has an infinite contribution as $q \rightarrow 0$ as a result. The procedure to deal with this issue is to define the set of improper diagrams as those connected internally by such a fixed-momentum line. We can then focus on just calculating proper diagrams, and sum over them in all possible combinations on each side of the troublesome fixed-momentum line. This process leads to the Dyson equation, and several other decompositions of diagrams into classes lead to other integral equations, which together provide self-consistent equations for diagrammatic contributions (the Hedin equations). Solution to these equations provides an interesting alternative to calculating material properties through direct simulation of particles (see e.g. ref. [22] for the equations and possible methods of solution). In the case where we consider the Lindhard function as the only proper diagram, we obtain the sum in fig. 2.3, which is exactly the RPA expression.

2 *Background*

Even with partial summation methods going beyond order-by-order perturbation theory, however, calculating structural quantities of dynamic systems quickly becomes intractable analytically, and computational methods must then be considered to make further progress. Results from linear response theory can be then be exploited to refine more direct first-principles methods, as we will see.

2.3 Core Computational Approaches

2.3.1 Classical Molecular Dynamics

We have made considerable progress towards calculations of the linear response (and thus the DSF) of WDM systems; to move beyond the purely analytical methods above, we'll have to simulate the dynamics of WDM directly, extracting the DSF from the calculated time-dependent density profile. To do so, we'll have to consider how to take a collection of ions and atoms in space and update their properties through time. Our methods will often be based – at least in part – on Classical Molecular Dynamics.

Suppose we have a system of classical point particles, interacting through (for example) pair potentials. If we wish to find the dynamic properties of such a system (or, at least, the dynamic properties of a system that can be approximated in this manner), we can attempt to integrate Newton's equations of motion directly (i.e. $\ddot{x} = -\nabla U(x)$, where x is the position vector of all of our particles, and $U(x)$ the potential). Classical Molecular Dynamics, which we will often refer to as just MD, is exactly this approach. We place a number of particles in a virtual box, and sequentially update their positions and velocities in small, discrete time steps. The equations of motion are written as approximate finite difference equations, and numerically integrated over time.

This method can be applied to a huge variety of systems. In particular, by constructing appropriate potentials, we'll see in section 2.3.7 that WDM can be treated directly with MD as an effective ion-only system. Alongside pure MD approaches, we can also consider combining (classical) MD for ions with fully quantum mechanical methods for electrons – this kind of hybrid approach is very useful for

2 Background

treating quantum plasmas in general, with Density Functional Theory Molecular Dynamics (to be discussed in section 2.3.6) the most prominent example.

In order to construct an MD simulation in practice, we will have to bring together several numerical components. We use this section to very briefly discuss these, and their implications for realistic calculations. This will lay the foundation for our future work, particularly our implementation of a Bohmian trajectories model of WDM – the topic of chapter 3 – for which MD will be a core component. We base these outlines on standard MD results (which can be found in e.g. ref. [23]).

Discretised Equations of Motion (Integrators)

In order to effect MD, we must first discretise Newton's equations so that particle positions and velocities can be updated incrementally. The most direct discretisation, the Euler scheme, can be written as:

$$a_i = F(x_i) \tag{2.32}$$

$$v_{i+1} = v_i + a_i \Delta t \tag{2.33}$$

$$x_{i+1} = x_i + v_i \Delta t \tag{2.34}$$

where i enumerates each time step, x_i is the position vector of the system, v_i the velocity, a_i the acceleration, F is the force function, and Δt is the length of the time step. While possessing the advantage of simplicity and transparency, the Euler scheme has very poor numerical properties. As time progresses, errors in particle positions and velocities quickly accumulate, resulting in an inaccurate representation of the real continuous-time system.

Several options exist for updates with more robust properties – see ref. [23] for details of these, and the specific numerical properties that should be considered. For our purposes, we will rely on a Leapfrog approach:

$$a_i = F(x_i) \tag{2.35}$$

$$v_{i+1/2} = v_{i-1/2} + a_i \Delta t \tag{2.36}$$

$$x_{i+1} = x_i + v_{i+1/2} \Delta t \tag{2.37}$$

While this algorithm requires only the same number of function evaluations as the Euler algorithm, the staggering of the velocity and position time steps ensures it is a second-order rather than first-order method (i.e. it is equivalent to the exact equations of motion up to second order in Δt), and has an excellent balance between speed and accuracy.

Boundary Conditions

For the most part, we will be interested in simulating bulk properties of quantum plasma. However, we will only be able to model a finite number of interacting particles in a box with MD. In order to take into account the interaction of our tracked particles with an effective surrounding plasma, we will employ periodic boundary conditions: Rather than just our single simulation box, we consider our system as an array of identical boxes – each containing an identically arranged set of particles – extending in each cardinal direction. Together, these boxes form a simple cubic lattice. When calculating the force on a given particle in our original box, we consider its interaction with all particles throughout this lattice. Positions are updated for particles in our original box, with those in the other boxes ‘following along’. In the case of short range potentials, only the nearest boxes to our origi-

2 Background

nal need to be considered. For long-range potentials – in particular, the Coulomb potential – we need to be more careful.

In many cases when we employ classical MD, the effective potentials exerted by each particle are short range (i.e. effectively decay to zero faster than $1/r^3$). These cases emerge when we are able to approximately integrate out the electronic subsystem, resulting in effective ion-ion potentials that are Yukawa-like (i.e. exponentially decaying) at long range. However, when we wish to explicitly model the electrons as dynamic particles, we must account for long-range Coulomb interactions. In these cases, our lattice of simulation boxes gives rise to a problem – the potential associated with any individual particle cannot be neglected for far-off periodic images.

There are several approaches to dealing with this issue; the most widely-used of these is Ewalds summation. The essence of the method is to break up the Coulomb interaction into k -space and real space contributions, such that we can approximate the full Coulomb interaction with a finite, truncated sum over images in k -space, on top of a standard short-range potential in real space (again, see ref. [23]). We will tend to rely on this standard method whenever we effect MD with long-range potentials (particularly in chapter 3).

Thermostats

Equipped with the Leapfrog integration method (in addition to appropriate boundary conditions/long-range potential corrections), we are able to evolve our system of particles in time. However, our existing equations of motion implicitly assume adiabatic conditions – we assume that our particles have no interaction with the outside world. This is not necessarily an accurate representation of our real system. For example, we will often want to simulate the system at fixed temperature

or pressure instead. In order to do this, we must adjust the equations of motion given in section 2.3.1 to account for interaction with an external environment. We will focus on thermostats that maintain our model system at a fixed temperature for our purposes.

There are two key thermostats that will feature prominently in our work – those of the Langevin and Nosé-Hoover varieties. The Langevin thermostat emulates thermal contact through a stochastic contribution to particle velocity updates. At each time step, the driving forces are modified to:

$$F(x) = F_0(x) - \gamma v/m + f \quad (2.38)$$

where F_0 is the deterministic force, v is the velocity, m is the mass of each particle, γ is the Langevin friction parameter, and f is a mean-zero random variable with dimension equal to that of x and (diagonal) variance $\sigma^2 = 2m\gamma k_B T/\Delta t$, where T is the temperature and Δt the simulation time step. It is straightforward to show, via stochastic calculus (discussed at greater length in chapter 5), that the stationary distribution of the updated equation of motion is equal to the canonical distribution.

The Nosé-Hoover thermostat instead effectively couples the system of particles to a deterministic harmonic oscillator. The equations of motion become:

$$F(x) = F_0(x) - \gamma v/m \quad (2.39)$$

$$\dot{\gamma} = \frac{1}{Q}(mv^2 - (3N + 1)k_B T) \quad (2.40)$$

where the friction parameter γ now effectively models the position of an oscillator representing the external environment, N is the number of particles in the system,

2 Background

and Q is an adjustable parameter. Again, the stationary distribution associated with these updated equations of motion is the canonical distribution.

Choosing the correct thermostat for a given simulation is non-trivial, as the effects on the dynamics of the modified equations of motion can be significant. The Langevin thermostat, for example, can be used to model collisions with light particles (read electrons) that have been integrated out from the system, by scaling the Langevin friction parameter (see ref. [15]). The Nosé-Hoover thermostat, on the other hand, tends to have a less pronounced impact on the dynamics, but may need to be augmented to ensure adiabatic motion. The adjustable parameters must also be chosen carefully to minimise impact on the dynamics, while preventing temperature oscillations. We will consider the implications of thermostat choice in more detail in chapter 3, alongside some key extensions.

2.3.2 Classical Monte Carlo Methods

Classical Monte Carlo (CMC) is a statistical sampling approach to obtaining system properties. One of the many variants of Monte Carlo methods we'll consider in this thesis, CMC samples a sequence of static snapshots of a classical many-body system – or, again, a WDM system approximated by a classical one – allowing us to calculate its time-averaged properties. Calculations are restricted to static properties, but the results can often be obtained far more quickly than with MD and similar dynamic methods.

When we want to calculate a static system property, we can reduce the problem to calculating an integral over possible system states. Say we want to find an averaged property $\langle \lambda \rangle$, where $\lambda = \lambda(x)$ is a function of the system coordinates x . $\langle \lambda \rangle$ could be the pressure, for example, or a radial distribution function. The average

we're after can be written:

$$\langle \lambda \rangle = \sum_i \lambda(x_i) \exp(-\beta V(x_i)) / Z, \quad (2.41)$$

where i enumerates all possible states of the system x_i , Z is the partition function, $V(x_i)$ is the energy, and β is the inverse temperature. Classical Monte Carlo calculates this integral by sampling successive states of the system, ensuring the probability of sampling any individual state is properly weighted according to its energy. And, importantly, cmc only requires us to know the probability ratio of any two given states – as we don't need absolute probabilities, we can avoid calculating the partition function altogether.

We can illustrate this with an example, using the simplest cmc variant: the Metropolis Hastings algorithm[24]. We begin with a random arrangement of particles in a box (a la MD), with position vector x . We generate a proposal, y – a possible new position vector for the system – according to a proposal distribution, $Q(y|x)$. In the simplest case, we require this proposal distribution to be symmetric in x and y – a typical distribution would be a normal distribution in $|x - y|$.

After making this proposal, we either accept or reject it. In the case of Metropolis Hastings, we accept immediately if $V(y) < V(x)$; if this is not the case, we generate a uniform random number $u \in [0, 1]$. We then accept the proposal if $\exp(-\beta\{V(y) - V(x)\}) > u$, and reject otherwise. If accepted, we update the particle position to y , i.e. we update $x \rightarrow y$; otherwise, we leave it where it is, $x \rightarrow x$. We then calculate λ for the new particle positions. The sequence of λ values generated in this way, when averaged, yield the desired system property $\langle \lambda \rangle$.

Due to its relative numerical speed, cmc will prove important in developing self-consistent potential construction methods, as we will discuss in section 3.4.2.

2 Background

Together with MD, we then have the tools to determine a broad swathe of properties of classical particulate systems.

2.3.3 Integral Equation Methods

For some applications, first-principles application of MD and CMC may prove to be too computationally intensive; in these cases, we can turn to approximate, semi-analytical approaches. The most prominent of these (and the most important to our work in this thesis) are those based on Integral Equations.

Consider one particle in the midst of a many-body classical system. This particle's effect on others surrounding it can be decomposed into two parts – the direct effect due to the particle itself, and the indirect effects mediated by other particles in the system. For example, in a simplified three particle system, we can consider the effect of particle 1 on particle 2 as a direct interaction effect, alongside an effect due to particle 1 influencing particle 3, and then particle 3 influencing particle 2.

This insight allows us to write equations linking interaction potentials to structure – the Ornstein-Zernike equations (a detailed discussion of which can be found in ref. [25]). We define the radial distribution function $g(r_{12})$ as the ratio of the radial density of particles of species 2 when there is a particle of species 1 at the origin to the radial density of a uniformly distributed liquid of particles of species 2 (of the same mean density). Deviations of $g(r_{12})$ from unity therefore encode the influence of the particle at the origin on the surrounding particles. It is useful to additionally define the total correlation function $h(r_{12})$ via:

$$h(r_{12}) = g(r_{12}) - 1, \tag{2.42}$$

such that for classical particles with no interactions, $h(r_{12}) = 0$. Now, the splitting of g into direct and indirect contributions can be performed by further defining the direct correlation function $c(r_{12})$, which represents the component of h that would remain were the indirect effect mediated by additional particles to be switched off. The principal Ornstein-Zernike equation then links these quantities by decomposing the indirect effect as a sum (convolution) of direct effects:

$$h(r_{12}) = c(r_{12}) + \rho \int dr_3 c(r_{13})h(r_{32}). \quad (2.43)$$

In reciprocal space, this takes a simple multiplicative form:

$$\hat{H}(k) = \hat{C}(k) + \rho \hat{H}(k)\hat{C}(k). \quad (2.44)$$

We may then write the following implicit Ornstein-Zernike equations for the correlation functions in reciprocal space as:

$$\hat{C}(k) = \frac{\hat{H}(k)}{1 + \rho \hat{H}(k)}, \quad (2.45)$$

$$\hat{H}(k) = \frac{\hat{C}(k)}{1 - \rho \hat{C}(k)}. \quad (2.46)$$

It remains now to close the system of equations in terms of the inter-particle potentials. For simplicity, we consider first a system with only one species of particle. For a pair potential $u(r)$, the Ornstein-Zernike equations are closed with the exact relation:

$$g(r) = \exp(-\beta u(r) + h(r) - c(r) + B(r)), \quad (2.47)$$

where, in the limit of weak to moderate coupling, the so-called bridge term $B(r)$ can be neglected. In this limit ($B(r) = 0$), the equation 2.47 combined with the

2 Background

Ornstein-Zernike equations 2.45, 2.46 are closed in the Hypernetted Chain (HNC) Approximation (alternative approximations may also be made for this term – for example, the Percus-Yevick closure is useful when analysing hard spheres, as it provides an analytical solution in this case[25]). The corresponding system of integral equations can then be solved for known pair potentials and system temperature. The inverse relationship can also be exploited to obtain an effective potential $u_{eff}(r)$ as a functional of the pair correlation functions[26].

As an alternative to direct first-principles approaches, integral equation methods have a key desirable property – they are extremely fast. Although not analytically tractable, the Ornstein-Zernike equations can be solved self-consistently very quickly with standard numerical methods (see e.g. [27]). The downside is the introduction of a further approximation, beyond any involved in constructing inter-particle potentials – the closure. In general we will prefer cMC to Ornstein-Zernike methods to avoid this approximation, but we will see that HNC methods, in particular, are very useful tools for analysing electron gases and WDM.

2.3.4 Quantum Monte Carlo Approaches for Static Structure

We now turn our attention to numerical methods that have been developed specifically for quantum systems; in combination with the classical methods we have discussed, they will allow us to establish a comprehensive picture of WDM numerically. First, we consider the direct Quantum analogue of cMC – Quantum Monte Carlo.

Quantum Monte Carlo[28] (QMC) methods enable state-of-the-art calculations of the static structure of quantum plasmas. Encompassing a range of related methods for both zero and finite-temperature systems, QMC has been exploited to provide benchmark properties of electronic gases across a range of conditions (includ-

ing conditions relevant to wDM); through this, it has provided highly accurate inputs for constructing density functionals (discussed in section 2.3.5). We outline here a subset of the more prominent QMC methods and their key results, alongside the limitations that must be addressed by companion simulation methods.

Zero Temperature – Variational and Diffusion Monte Carlo

The simplest form of QMC, emerging in the very early days of quantum mechanics, is the variational method – a standard approach in any introductory quantum mechanics text. By beginning with a trial wavefunction parametrised with an array of variables a , the ground state wavefunction may be approximated by $\psi_0 = \psi(a_0)$, where

$$a_0 = \operatorname{argmin}_a \frac{\langle \psi(a) | \hat{H} | \psi(a) \rangle}{\langle \psi(a) | \psi(a) \rangle}$$

and *argmin* is the minimum of the argument function (taking the value of a at which the energy argument is a minimum). The Monte Carlo aspect of this approach comes down purely to how the energy integral is evaluated. Writing the energy expectation as

$$\frac{\langle \psi(a) | \hat{H} | \psi(a) \rangle}{\langle \psi(a) | \psi(a) \rangle} = \frac{\int |\psi(x, a)|^2 \frac{H\psi(x, a)}{\psi(x, a)} dx}{\int |\psi(x, a)|^2 dx}, \quad (2.48)$$

the quantity

$$\frac{|\psi(x, a)|^2}{\int |\psi(x, a)|^2 dx} \quad (2.49)$$

can be interpreted as the probability distribution, which can be sampled in effecting the random Monte Carlo integral. This approach becomes exact when the ansatz wavefunction form incorporates the ground state for some value of a (and provides an upper bound to the ground state energy otherwise); however, the sen-

2 Background

sitivity to the quality of the employed ansatz limits the accuracy of this simple approach, even before numerical concerns are taken into account.

Diffusion Monte Carlo and its various extensions provide an alternative, more accurate and more general approach at zero temperature. In principle the method is numerically exact, but the time complexity for fermion systems is exponential (this is related to the ‘fermion sign problem’, which must be dealt with when constructing a QMC method in order to obtain an efficient simulation technique – the problem manifests itself as a fluctuating sign in the weights of different contributions to the required integrals). To reduce this complexity, the method is normally employed alongside additional simplifying approximations when treating systems with more than a couple of electrons (see e.g. ref. [29], which employs the common ‘fixed-node’ approximation, discussed further in section 2.3.4 in the context of path integrals). We outline just the basics of the algorithm, following ref. [30].

To make contact with classical diffusion, the Schrödinger equation is first transformed into imaginary time via $\tau = it$, yielding:

$$\frac{\partial\psi(x, \tau)}{\partial\tau} = \frac{1}{2} \frac{\partial^2\psi(x, \tau)}{\partial x^2} - (V - V_{ref})\psi(x, \tau), \quad (2.50)$$

where V and V_{ref} are the local potential and (arbitrary) reference potential; the evolution of the eigenstates under this mapping is dictated by:

$$\psi(x, \tau) = \psi(x)e^{-(E - V_{ref})\tau}. \quad (2.51)$$

Equation 2.50 can be iteratively solved by evolving through time a set of imaginary particles, referred to as ‘psips’. The time evolution of the concentration of psips, C ,

is controlled by a diffusion equation of the form:

$$\frac{\partial C}{\partial t} = D \frac{\partial^2 C}{\partial x^2} - kC, \quad (2.52)$$

with diffusion coefficient D and rate constant k ; choosing these constants properly leads to the equivalence of eq. 2.50 and eq. 2.52 equivalent. At each time step, the psips are moved randomly along x with distance determined by the magnitude of D ; following this, each psip may multiply or be destroyed, according to a probability that is a function of k . When a fluctuating equilibrium is reached in terms of the spatial psip distribution, the energy eigenvalue of the ground state can be extracted. The explicit algorithm steps and probabilities as functions of D , k can be found in ref. [30] and the references therein.

This algorithm, and a family of highly related zero-temperature methods, have provided much of the data that has been used in the construction of density functionals for DFT methods[31].

Finite Temperature - Path-Integral Monte Carlo

In the case of WDM, we are principally interested in finite temperature systems and methods - as a consequence of this, we will need to consider how the general diffusive algorithms above can be extended to this regime. Path-Integral Monte Carlo offers one such extension.

To highlight the power of Path-Integral Monte Carlo for our purposes, we focus on a more realistic physical system than those considered previously: the electron gas in the WDM regime. In the case of finite-temperature systems, we take interest in the density matrix directly as opposed to considering details of individual eigenstates. The approach in this case is to exploit Feynman's path integral formulation

2 Background

of quantum mechanics to write the density matrix as (following ref. [32]):

$$\rho_F(R, R', \beta) = \hat{A}\rho_0(R, R', \beta) \left\langle \exp \left[- \int_0^\beta d\tau V(R(\tau)) \right] \right\rangle, \quad (2.53)$$

where ρ_F is the full density matrix, ρ_0 the non-interacting density matrix, \hat{A} the antisymmetrisation operator, β the inverse temperature, R the set of electron coordinates, $\tau = \beta/M$ (where M is the number of imaginary time discretisations), and the angle brackets indicate an average over Brownian paths of the system from R' to R . The non-interacting density matrix appearing in this equation is known exactly:

$$\rho_0(R, R', \tau) = \left(\frac{r_s^2}{4\pi\tau} \right)^{\frac{3N}{2}} \exp \left(- \frac{(r_i - r'_j)^2}{4\tau/r_s^2} \right). \quad (2.54)$$

Using these expressions, an algorithm can be constructed to sample the average in eq. 2.53. This can be carried out by constructing a path in configuration space parametrised by the imaginary time τ , and developing different Monte Carlo steps that sequentially modify the path (with acceptance probability determined by energy considerations - see ref. [32] and references therein for explicit details). This procedure can be conducted at a relatively high temperature and, subsequently, the density matrix at a desired temperature obtained using the pair product approximation, which decomposes the full density matrix into a product of higher-temperature electron pair density matrices[33].

As alluded to in the previous section on Diffusion Monte Carlo, there is an exponential time scaling for fermion systems, as weights from different paths have fluctuating signs and equal orders of magnitude. In the Path Integral case, the ‘fixed node’ approximation circumvents this problem by assuming the nodal structure of the interacting density matrix can be approximated by that of the non-interacting density matrix in eq. 2.54. Armed with this assumption, paths consid-

ered in the average in eq. 2.53 can be restricted to those that are wholly contained in a positive definite region of ρ_0 , removing issues of cancelling signs.

The resulting method is just one of many finite temperature QMC approaches (known as Restricted Path integral Monte Carlo – RPIMC). The principal limitation of the discussed QMC methods is that – despite their accuracy for static properties – they are incapable of directly determining the dynamics of the system. The time-dependence appearing in the associated equations is explicitly a dependence on imaginary time – as constructed for calculation purposes – and any genuine time-dependence of the system is not accounted for. As such, QMC must be paired with additional methods when dynamic properties are required (or, alternatively, one must invoke additional approximations in order to analytically continue from imaginary to real time)[34].

It should also be noted that the computational complexity of QMC imposes a restrictive limit on the number of particles it can treat. This is especially true when ions and electrons are considered simultaneously; in these cases, the separation of time-scales issue of classical simulation appears as the separation of imaginary-time scales for the path integrals (see for example ref. [35] for a recent state-of-the-art PIMC simulation of Sodium, restricted to an 8-atom cell). As a result, there is a finite-size error that must be considered when attempting to generalise to the thermodynamic limit. Considerable effort has gone into methods of reliably overcoming this error, which can often be the largest present in QMC results. The simplest approach to this correction is to perform simulations with a range of box sizes, and then to attempt to extrapolate from the results up to the thermodynamic limit.

2.3.5 Density Functional Theory

Once we require dynamical information for large systems, the limitations of QMC usually render it impractical (or, at least, incomplete). By incorporating results from QMC into a reliable and scalable formulation of quantum mechanics, though, we can turn to more general-purpose Density Functional Theory methods.

Density Functional Theory (DFT) has established itself as the most prominent method for modern numerical calculations of electronic structure. It is a broad subject with enormous representation in the literature, with applications as diverse as chemistry, condensed matter and dense plasma[14, 36, 37, 38, 39, 40, 41]; here we discuss only the very basic elements of the theory for context. DFT simulations exploit the theorem due to Hohenberg and Kohn[42] that:

- The energy of an electronic system can be expressed as a (unique) functional of the electronic density, without direct recourse to the exact many-body wavefunction (i.e. $E = E[\rho]$).
- The ground state electronic density can be obtained by finding the function ρ that minimises the energy functional, via $\rho_0 = \operatorname{argmin}_{\rho} E[\rho]$.

In general, then, the central object of DFT is the electronic density, which combined with an energy functional yields ground state structure and energies. A procedure for performing practical calculations using these theorems was presented by Kohn and Sham[43] shortly after the initial Hohenberg-Kohn paper. A fictitious non-interacting electron system is considered; for such a system, the kinetic energy can be written exactly in terms of the occupied orbitals. The total energy functional has the form (we examine the zero-temperature formalism for simplicity):

$$E[\rho] = T_s[\rho] + \int dr \{v_{ext}(r)\rho(r) + E_H[\rho] + E_{xc}[\rho]\}, \quad (2.55)$$

where the kinetic energy contribution can be written in terms of the fictitious orbitals as:

$$T_s[\rho] = \sum_i \int dr \left\{ \phi_i^*(r) \left(\frac{-\hbar^2}{2m} \nabla^2 \right) \phi_i(r) \right\}, \quad (2.56)$$

and E_H is the standard Hartree contribution:

$$E_H = \frac{e^2}{2} \int dr \int dr' \frac{\rho(r)\rho(r')}{|r - r'|}. \quad (2.57)$$

The remaining part of the functional, E_{XC} , then contains contributions not accounted for by this model system - namely the exchange interaction, and the difference between kinetic energies of the interacting and model systems. Formally this can be written as:

$$E_{XC}[\rho] = T[\rho] - T_s[\rho] + E_{ee}[\rho] - E_H[\rho], \quad (2.58)$$

where $T[\rho]$ is the exact kinetic energy functional, and $E_{ee}[\rho]$ the exact electron-electron interaction energy. In practice, this is the component that must be approximated, but years of intense research have produced XC functionals accurate for a wide range of systems.

The Kohn-Sham equations for the fictitious electron orbitals are then:

$$\left(-\frac{\hbar^2}{2m} \nabla^2 + v_{eff}(r) \right) \phi_i(r) = \epsilon_i \phi_i(r), \quad (2.59)$$

where the effective, non-interacting potential experienced is:

$$v_{eff}(r) = v_{ext}(r) + e^2 \int \frac{\rho(r')}{|r - r'|} dr' + \frac{\delta E_{xc}[\rho]}{\delta \rho(r)}. \quad (2.60)$$

2 Background

These equations considerably simplify the tremendous complexity involved in solving for an N -body wavefunction for large N . The wavefunctions can be solved for iteratively to achieve self-consistency between the orbitals and the density-dependent potential terms. While the calculated wavefunctions correspond to the fictitious non-interacting system, the Hohenberg-Kohn theorem guarantees the equivalence of the electronic density of the fictitious system and the real physical system. As a consequence, a great deal of research has gone into deriving and evaluating different approximations for the exchange-correlation term E_{XC} , as it is the only unknown element of what is a fully exact theory.

Computational difficulties remain for large systems, however, and for modelling such systems Orbital-Free Density Functional Theory[44] (OFDFT) provides a less expensive approach. The key step that is circumvented through OFDFT is the self-consistent determination of the fictitious orbitals. This can become prohibitively expensive for a large number of orbitals (which must be orthogonalised, and iteratively optimised); the problem becomes even greater when the electronic system is treated at finite temperature, where potentially large numbers of orbitals beyond those in the ground state are occupied with non-negligible probability.

The orbital-free approach does away with individual wavefunctions, instead considering only the electronic density – as such, the kinetic energy of the system in eq. 2.56 becomes a further unknown, which must be approximated alongside the exchange-correlation contribution to the energy. The accuracy of OFDFT is therefore reduced relative to full Kohn-Sham DFT (energy errors of up to order 10% for certain material properties are fairly typical), but the computational speed-up for finite systems often makes it a more appropriate tool. Approximations to the orbital-free energy functional (inclusive of the kinetic term) have a long history – in fact longer than DFT itself – dating back to the Thomas-Fermi model of the ki-

netic energy:

$$E_{TF} = \frac{3}{10}(3\pi^2)^{\frac{2}{3}} \int n(r)^{\frac{5}{3}} d^3r, \quad (2.61)$$

and its subsequent extension by Dirac to include exchange effects. Alongside Kohn-Sham methods, OFDFT remains a promising route with active research into improved functionals/associated algorithms.

2.3.6 Density Functional Theory Molecular Dynamics

With the tools of DFT to hand, there are multiple approaches to translate the associated electron structure calculation to dynamic simulation. Perhaps the most immediately obvious route is to introduce time-dependence into the governing equations directly. To do so, one can exploit the more theoretically involved time-dependent Density Functional Theory (TDDFT), allowing effective electron wavefunctions to evolve in time. This is certainly a valid route, and much research has been done in this area (a broad summary can be found in ref. [45] and the references therein). While appealing from the point of view that it is the natural time-dependent extension to DFT, TDDFT is a complex and computationally very expensive procedure; it provides an important benchmark for systems with relatively small numbers of electrons due to its high accuracy in predicting dynamic behaviour, but is currently not a practical, general purpose method (although this state of affairs is improving with robust publicly available TDDFT codes such as Octopus[46]). Recent state of the art TDDFT results have been able to accurately construct DSF estimates, but have proved too expensive to resolve the low-frequency ion contribution (see ref. [47], where even vast computational resources could not resolve this feature in the all-electron case).

A computationally simpler approach, within the theoretical bounds of standard time-independent DFT, is to exploit the separation of time-scales between

2 Background

ions and electrons in a system via the Born-Oppenheimer approximation. Ion positions are updated via classical MD, and at each ion time step new electron densities – and associated forces – are calculated. This approach, usually referred to as Density Functional Theory Molecular Dynamics (DFT-MD) represents a reasonable compromise between speed and accuracy for the WDM regime, and there are numerous examples in the literature of accurate predictive simulations of this kind (see, for example, ref. [14]).

While offering strong accuracy for many properties of WDM systems, DFT-MD has its limitations; in particular, the Born-Oppenheimer approximation misses some elements of the physics – that is, the impact on the ions of the electronic dynamics. As electrons are treated as a static background at each ion time step, any effects of the electron motion on the ions are lost. An approximate approach to determine these effects (alongside DFT methods for the many system properties it is able to predict accurately) is clearly needed to establish the full picture (this will be discussed in greater detail in chapter 3). Further, the incorporation of finite temperature electrons has proved non-trivial for DFT-MD methods – the production of high-quality finite-temperature functionals is still underway, and many off-the-shelf DFT-MD packages do not yet incorporate finite temperature physics properly (many use fractionally occupied orbitals, but ignore the temperature dependence of the energy functional itself).

2.3.7 Quantum Statistical Potentials

While DFT-MD is in many ways the state-of-the-art method for WDM calculations, we should also consider alternative approaches that address some of its flaws – in doing so, we open up avenues for developing new, more effective simulation techniques. Quantum Statistical Potentials (QSPs) should be considered from this point

of view; by treating quantum systems with effective classical potentials, they will provide a useful reference point for our work going forward.

Treating quantum systems with classical simulation methods – with the intent of reducing simulation complexity – has a long history in many-body physics. In order to utilise classical simulation methods, a reliable map between the properties of the real quantum system and the model classical system must be constructed. Of particular concern is the construction of classical potentials that somehow take into account the quantum effects present. In this section we will summarise some of the approaches that have been taken to this problem.

We'll begin with a collection of methods that – in a sense – do away with electrons altogether. Effective potential methods take the route of obtaining ion-ion potentials that account for the electron structure, allowing calculations (typically using MD) with ions only – with appropriately adjusted potentials. By absorbing electrons into ion-ion potentials, we can obtain classical systems that are easy to simulate. Better still, the time scale problem resulting from differing electron and ion masses is neatly swept under the rug, in a manner akin to direct use of the Born-Oppenheimer approximation. Of course, this means electron dynamics are ignored – but the resulting calculations are fast and straightforward, and remain accurate for many properties of interest.

The approach is neatly exemplified by the neutral pseudoatom method. This method first calculates the electron density surrounding an average ion in the system – via some form of quantum mechanical calculation – and thereby constructs effective ion-ion potentials that take electron structure and response into account. There are a few different variants of this approach, but the majority are in the spirit of that in ref. [48] (the references therein highlighting some alternatives).

2 Background

Initially, a single ion in a spherical box is considered, with some initial trial value for the ion-ion pair correlation function, $g_{ii}(r)$, and ion-electron pair correlation function, $g_{ie}(r)$. Using the charge distribution around the ion implied by these correlation functions, a new electron distribution is determined using DFT. In turn, this electron distribution establishes an effective ion-ion potential via linear response theory, which can be used to recalculate $g_{ii}(r)$ (usually via an integral equation approach). This loop can be iterated to self-consistency.

With this, the effective potential between ions takes into account the electron distributions ‘carried around’ by each. Another similar (and even simpler) approach can be found in the Yukawa+SRR model[49] (Yukawa plus short-range repulsion). In this case the ion-ion potential V_{ii} is written

$$V_{ii} = \sigma^4/(\beta r^4) + \exp(k_s r)/r, \quad (2.62)$$

Here, the usual Yukawa potential appears in the second term, with the first term accounting for short-range repulsion between ions. The additional parameter σ can then be fitted based on a reference calculation, resulting in a pair potential that implicitly takes into account an adiabatic electronic background.

These methods, and many others along similar lines, offer a route to subsequent classical MD. We are then able to perform fast calculations of dynamic structure – and, as a result, effective potential methods are very valuable (and often surprisingly accurate) tools. Clearly, however, some severe dynamical approximations have been made. As mentioned above, effects due to electron dynamics have been completely ignored, with electrons being effectively integrated out in the process of constructing the potentials – at face value, this is at least as bad as the Born-Oppenheimer approximation. However, the approximation we’ve made with these effective potentials is actually far more severe: the electron static den-

sity is treated in an averaged way. The real system cannot be exactly reduced to one with pair potentials between neutral bodies, as electrons do not uniformly superimpose themselves on each ion irrespective of the behaviour of surrounding ions.

To go beyond the effective potential methods above, we can try to construct models that retain dynamic electrons explicitly – without sacrificing our ability to use classical simulation methods. To do so, we'll need electronic potentials that implicitly account for their less-than-classical behaviour.

The classical potentials we're seeking – capable of (partially, and approximately) reproducing quantum effects – can be obtained by examining contributions to the full quantum density matrix. Using known probabilistic distributions of the quantum system (via the density matrix), it is relatively straightforward to determine classical potentials capable of reproducing these distributions for a model classical system. That is to say, provided the solution to the full quantum system is known, one may produce a classical system with effective potentials that can reproduce statistical properties of the quantum system.

Of course, this observation alone does not provide a route to a useful classical model, as the solution to the quantum system is required as input. Instead, we can use the properties of a simplified quantum system that can be solved for exactly (or accurately) to produce effective potentials, and use these effective potentials as partial components of potentials in a fully classical model.

Consider first the interaction of two isolated electrons, which can be used to yield an effective potential for low density systems. In the low density limit, we can use first order perturbation theory to approximately solve the two-particle Bloch

2 Background

equation for the density matrix, giving:

$$\rho_{ij} = \frac{(m_i m_j)^{\frac{3}{2}}}{(2\pi\hbar\beta)^3} f_{ii'} f_{jj'} \exp(-\beta\Phi^{ij}), \quad (2.63)$$

where

$$f_{ab} = \exp\left(-\frac{m_a}{2\hbar^2\beta}(r_a - r_b)^2\right). \quad (2.64)$$

For Coulomb systems, the resulting potential matrix entering into equation 2.63 is:

$$\Phi^{ij}(r_{ij}, r'_{ij}) = e_i e_j \int_0^1 \frac{d\alpha}{d_{ij}(\alpha)} \operatorname{erf}\left(\frac{d_{ij}(\alpha)}{2\lambda_{ij}\sqrt{\alpha(1-\alpha)}}\right), \quad (2.65)$$

where $d_{ij}(\alpha) = |\alpha r_{ij} + (1-\alpha)r'_{ij}|$, $\lambda_{ij}^2 = \frac{\hbar^2\beta}{2\mu_{ij}}$, the reduced mass $\mu_{ij}^{-1} = m_i^{-1} + m_j^{-1}$, and β is the usual inverse temperature. This allows capture of quantum diffractive effects, seen explicitly in the finite potential limit when $r_a - r_b \rightarrow 0$; in low density scenarios, this approach then yields an effective electron-electron potential that can be used within classical simulation, yielding statistical correlations equivalent to the quantum case (becoming exact in the low density limit).

Suppose now we wish to extend this result to obtain an electron-electron potential valid also for higher densities. We now require a more robust starting point than the perturbative result of equation 2.63. This can be accomplished at high temperature with techniques such as matrix squaring (which we earlier saw exploited in Path-Integral Monte Carlo approaches – see e.g. ref. [50]); in general terms, we can use quantum mechanical techniques to approximate the full probability distribution, then obtain an exact n-body potential through inverting:

$$\rho(\{r_i\}; \beta) \propto \exp(-\beta U_p(\{r_i\}; \beta)), \quad (2.66)$$

where d is the dimension of the system, and U_p is the full system potential.

The quality of the potential depends strongly on how accurately the calculated density matrix ρ reflects the full system, however. Obtaining a highly accurate solution to the full Hamiltonian in order to use this inversion is extremely difficult. Rather than approximating the full density-matrix quantum mechanically, then, we can alternatively consider the potential extracted from equation 2.66 as a component of the overall potential, and approximate the remainder classically.

We can turn then to the Ornstein-Zernike equation (as we discussed in section 2.3.3) for this purpose. A simple procedure for simulation using effective quantum statistical potentials, in the same spirit as the density matrix inversion used above, can be summarised as follows:

1. Solve for the pair correlations of the quantum system, with interaction potentials $V_{approx.}$ that permit accurate or exact solution (through the Bloch equation/Schrödinger equation as appropriate)
2. Invert the Ornstein-Zernike equations to obtain classical potentials that reproduce these pair correlations, $V_{eff.}$
3. Perform a classical simulation using potentials equal to $V_{QSP} = V_{eff.} + V_{exact} - V_{approx.}$

This procedure is formalised by Classical Map methods (see e.g. ref. [51], [26] for various implementations and overviews). The electron fluid, in particular, is a favourite target – a potential is constructed that replicates the spin statistics of a non-interacting electron gas at the pair correlation level by inverting the Hypernetted Chain Equations, and is added to a quantum statistical Deutsch potential (which accounts for Coulomb interactions and quantum diffraction). The Classical Map procedure then adjusts the effective system temperature in order to match

2 Background

known pressure information. The result is a model system that attempts to match both dynamic and static properties of the real quantum system with a classical one.

While certainly an effective computational tool, the Classical Map technique has drawbacks – by reducing a quantum system to a purely classical one, inevitably some properties of the real system cannot be replicated. However, the core strategy behind Classical Maps is a very useful one; a model system is chosen, and parameters are chosen to match known, desired properties under restricted conditions. By combining this idea with a quantum formulation of the underlying system (and non-classical electronic dynamics), we now aim to construct a new method for wDM modelling that completely avoids Born-Oppenheimer.

3

Predicting Dynamic Structure: The Bohmian Trajectory Method

“You were torturing a cat,” she says. “With a freaking prod.”
“A prod I built myself in metal shop,” he says. “But of course you never mention that.”

Pastoralia
George Saunders

3.1 Introduction

In the previous chapter, we discussed several prominent numerical methods for calculating the properties of quantum plasmas. Amongst the methods discussed, we observed a near-universal flaw: Dynamical coupling effects between electrons and ions were neglected, either in part or – more often – in their entirety. There is good reason for this: Electrons and ions evolve on very different timescales due to their mass difference, and modelling the dynamics of both simultaneously is a time-consuming task. Taking a standard quantum-mechanical method for the electron dynamics and tacking on a classical method for the ions – the most commonplace approach to quantum plasma – will only get you so far. On an electronic timescale, long-time ion-ion correlations rapidly become too expensive to calculate.

The issue is seen clearly on examination of DFT methods and results. The time-dependent formulation of DFT, TDDFT, *does* allow us to model both the electrons and ions simultaneously, giving a complete and very accurate insight into the dynamics; however, this method remains too computationally intensive when we wish to calculate long-term correlations (particularly the ion modes). A recent, state-of-the-art TDDFT calculation exemplifies this. Ref. [47] models a system of 128 electrons for approximately 0.001 attoseconds per CPU-core and second; despite leveraging enormous computational resources, the dynamic ion-ion parameters are left unresolved. The simplest ‘solution’ – as exploited by DFT-MD – has been to decouple the electron and ion dynamics with the Born-Oppenheimer approximation. Rather than considering fully dynamic electrons, the ions are updated with classical equations of motion, and the electron density is calculated statically at every ion time step.

The work in this chapter is motivated by experimental and numerical evidence that this approach is often not good enough. Dynamical coupling between electrons and ions cannot be neglected altogether when calculating dynamic ion properties. Ref. [15], in particular, explores the effect of electron dynamics on the ions via a simple noise model – the data indicates that Born-Oppenheimer simulations are missing much of the central diffusive peak of the DSF, and are overestimating the acoustic mode frequency – the location of the side peaks – as a result. Clearly, electron-ion dynamic coupling – as a source of randomisation for the ion dynamics – is crucial to accurate calculations of the Dynamic Structure of the ions.

This is a particular case of a more general problem. When we coarse-grain a system – that is, remove problematic (usually fast) degrees of freedom to focus on those of interest – we inevitably face a trade-off between speed and accuracy. This is clear in much earlier foundational work on Dissipative Particle Dynamics – the original treatment in ref. [52] illustrates the difficulty of capturing full dynamic correlations in coarse-grained systems. The Born-Oppenheimer approximation introduces analogous flaws that must be accounted for.

Correcting for missing electron-ion interactions is a non-trivial task, however: It is usually not computationally feasible to consider full quantum electronic dynamics alongside standard ion models. And, while ref. [15] highlights the issue via variation of an artificial noise model, it does not offer a complete answer to the core problem. The magnitude of the noise term cannot be determined with good accuracy ahead of time, and must instead be matched to experiment retrospectively – so predictive simulations of material properties are not possible. A clean solution for calculating ion dynamics without neglecting electron-ion dynamic coupling will require a new approach entirely.

3 *Predicting Dynamic Structure: The Bohmian Trajectory Method*

Classical Maps, examined at the tail end of the previous chapter, offer a step towards a solution to our problem – by modelling quantum systems with property-matched classical systems, we saw that it is possible to avoid Born-Oppenheimer via classical electron trajectories. However, Classical Maps strictly map quantum systems onto toy, representative classical systems; our goal will be to instead develop an approximate quantum model – with the desirable computational properties inherent to Classical Maps – beginning from a fully quantum formulation of the system and without restricting ourselves to purely classical dynamics.

In summary, then, the goal of this chapter is the development of a new simulation technique, incorporating coupled electron and ion dynamics while – crucially – being fast enough to examine long-time ion correlations. The work herein is a combination of theoretical and computational advances towards this aim, culminating in the theory, implementation and testing of a new trajectory-based simulation method. We begin the chapter with an overview of the theoretical basis of the new method (in particular, Bohmian Mechanics). We then develop the various components of our simulation scheme, before validating our complete implementation with a collection of `wDM` test cases.

3.2 Theory

3.2.1 Bohmian Mechanics

Basics and Derivation

Bohm's interpretation of quantum mechanics treats a quantum system as a classical ensemble of trajectories in the full configuration space (containing all the positional degrees of freedom); in the non-relativistic case, this classical ensemble then reproduces the probabilistic result of the Schrödinger equation exactly[53]. That is to say, by following an ensemble of trajectories through configuration space, each with clearly defined physical properties, we can calculate anything that could in principle be calculated from knowledge of the system's wavefunction. In practice, there are two equivalent calculation strategies for obtaining these 'quantum trajectories' – either we use a modified classical equation of motion to track the time evolution of a large number of trajectories that are dense in the configuration space (and thus have well-defined probability distribution), or we track a smaller subset of trajectories alongside solving for the traditional wavefunction. Quantum trajectories can then either be considered as interacting with each other through a density-dependent potential (the Bohm potential), or as guided by a wavefunction ('pilot wave' theory). We consider both of these points of view in constructing algorithms, although of course they are equivalent, merely differing in whether the configuration space density is tracked directly or inferred from a particle distribution.

This trajectory-based approach to quantum mechanics has proved controversial from a philosophical point of view – particularly with regards to the reality of the quantum trajectories and their physical meaning. However, here we use

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

Bohm's approach purely as a computational tool. The equivalence in prediction of probability distributions from Bohm's theory and 'traditional' quantum mechanics is known, and thus the philosophical correctness is not relevant for our purposes – we hence do not discuss it further in this work.

Although Bohmian mechanics can cast a quantum system in terms of classical trajectories, it does not circumvent the complexity of many-body quantum theory. Even when we make no recourse to a wavefunction, in principle we must construct an infinite ensemble of classical trajectories to exactly reproduce a quantum system. Further, the trajectories themselves exist in the full, high-dimensional configuration space. Advantages of the theory are therefore not free, and emerge only in the form of different approximation strategies, which we will discuss.

We begin with a very brief outline of the underlying mathematics. In the case of a single particle, the formalism follows from writing the wavefunction ψ as:

$$\psi(x, t) = R(x, t) \exp(iS(x, t)/\hbar). \quad (3.1)$$

From this form for ψ , the Schrödinger equation yields:

$$\frac{\partial R^2}{\partial t} + \nabla \cdot \left[R^2 \frac{\nabla S}{m} \right], \quad (3.2)$$

$$\frac{\partial S}{\partial t} + V_{ext} + V_B + \frac{(\nabla S)^2}{2m} = 0, \quad (3.3)$$

where V_{ext} is the external potential, and the Bohm potential V_B is given by:

$$V_B(x, t) = \frac{-\hbar^2 \nabla^2 R}{2m R}. \quad (3.4)$$

Here, eq. 3.2 represents a probability continuity equation, and eq. 3.3 is of the form of a Hamilton-Jacobi equation, with the additional quantum potential of eq.

3.4. The term R^2 then represents the probability density of classical trajectories in configuration space. Trajectories within the ensemble, or equivalently trajectories guided by the wavefunction, then evolve via:

$$v_\alpha(x, t) = \frac{1}{m_\alpha} \frac{\partial S}{\partial x_\alpha}. \quad (3.5)$$

This velocity relationship allows sample trajectories to be evolved when ψ is known, or a large ensemble of trajectories to be evolved concurrently. The statistical properties of the classical trajectories then reproduce all quantum mechanical results, provided their initial positions are sampled from the initial probability distribution.

Quantum Trajectories in Bohmian Mechanics

In order to produce practical simulation methods, approximations must typically be made to avoid the high dimensionality associated with the full configuration space. In DFT, for example, this approximation enters through the xc functional, allowing reduction of dimensionality to that of real space in exchange for a non-exact contribution to the energy. In applying Bohmian mechanics, similar forms of approximation are also possible; a prominent example is the use of so-called conditional wavefunctions, in which single particles can be treated as having their own separable wavefunctions with modified (complex) potential terms in the Hamiltonian. We follow the original formulation due to Oriols[54] here, in which we focus on an individual particle with label a , and treat the positions of all other particles in the system as components of the vector x_b .

We consider a system with overall wavefunction $\phi(x_a, x_b, t)$, and an associated set of Bohm trajectories with positions $x_a[t], x_b[t]$, where throughout this section we distinguish the trajectory coordinates with their explicit t argument. The con-

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

ditional wavefunction for particle a , $\psi_a(x_a, t) = r_a(x_a, t) \exp[is_a(x_a, t)/\hbar]$, is then defined as the full wavefunction evaluated at the trajectory positions $x_b[t]$ of the other particles, $\psi(x_a, t) \equiv \phi(x_a, x_b[t], t)$. Starting with a single-particle Schrödinger equation for the conditional wavefunction, we can construct an effective potential energy with which $\psi(x_a, t)$ can be evolved in time. This effective single-particle potential profile is written as a sum of terms

$$u_a(x_a, t) = U_a(x_a, x_b[t], t) + G_a(x_a, x_b[t], t) + iJ_a(x_a, x_b[t], t), \quad (3.6)$$

where the overall potential term in the Hamiltonian, U , has been split into components associated with a , and those associated with particle interactions not including a , via $U = U_a + U_b \equiv U_a(x_a, x_b, t) + U_b(x_b, t)$. Here, G_a , J_a are defined as:

$$G_a = U_b + \sum_{k \neq a} \left\{ K_k + V_{B,k} - \frac{\partial S}{\partial x_k} v_k \right\}, \quad (3.7)$$

$$J_a = \sum_{k \neq a} \frac{\hbar}{2R^2} \left\{ \frac{\partial R^2}{\partial x_k} v_k - \frac{\partial}{\partial x_k} \left(\frac{R^2}{m_k} \frac{\partial S}{\partial x_k} \right) \right\}, \quad (3.8)$$

where v_α is the trajectory velocity as defined by eq. 3.5, and $V_{B,k}$ is equal to the previous Bohm potential definition of eq. 3.4, where the Laplacian operator is over the coordinates of particle k , and $K_a = (\partial S / \partial x_a)^2 / 2m_a$. The effective single-particle Schrödinger equation for the reduced wavefunction associated with particle a is then just:

$$i\hbar \frac{\partial \psi_a}{\partial t} = \left\{ -\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x_a^2} + u_a \right\} \psi_a. \quad (3.9)$$

For a single particle the summations over $k \neq a$ in eqs. 3.7, 3.8 necessarily vanish, yielding the normal single-particle Schrödinger equation. G_a , J_a are analogous to the exchange-correlation contribution of DFT, in that an estimate must

be constructed for these terms, but in principal the time evolution equation 3.9 becomes exact when they are known. In combination with the general theory of Bohmian trajectories, the concept of the conditional wavefunction allows for new strategies to modelling quantum dynamics, provided good estimates of G_a , J_a can be obtained.

For our purposes, we will construct a slightly different, but related, route to a practical simulation method, based on the same foundation of Bohmian mechanics. We will combine the theory above with the pragmatic elements of Classical Maps (i.e. systematic property matching, as per section 2.3.7), allowing us to construct effective potentials that both closely replicate known system properties, and are also derived from a fully quantum formulation of the problem. Before we can construct such a method – and indeed, before we can make progress on the related inverse problem, of obtaining system properties from experimental data – we will need to build up some further theoretical elements.

Most pressingly, we will require a scheme to simplify the evolution equations of our many-body quantum system – that is to say, we must develop an approximation that removes the necessity of an infinite ensemble of interacting trajectories. To do so, we can look to effective approximations employed by the Quantum Hydrodynamics method.

3.2.2 Quantum Hydrodynamics and Linearisation

From a computational perspective, a hydrodynamic approximation to the electron fluid offers considerable speed-ups over the full kinetic formalism, particularly when good approximations for the electronic dynamics can be established that incorporate quantum effects. The bulk treatment of fluid elements, each containing large numbers of electrons, clearly has computational advantages over propagat-

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

ing individual trajectories, and is considerably faster than fully quantum mechanical DFT calculations. Quantum Hydrodynamics has thus emerged as a fast alternative to DFT simulations when studying strongly coupled plasmas – however, there are additional approximations invoked by this method, and associated accuracy penalties. In order to construct this theory, we follow Manfredi et al. (ref. [55]) by beginning with decoupled individual particle wavefunctions ψ_α for the electrons. The time evolution of these wavefunctions can be expressed (in the Hartree approximation) by coupling independent Schrödinger equations for each electron with Poisson's equation:

$$i\hbar \frac{\partial \psi_\alpha}{\partial t} = -\frac{\hbar^2}{2m} \frac{\partial^2 \psi_\alpha}{\partial x^2} - e\phi \psi_\alpha, \quad \alpha \in \{1, \dots, N\}, \quad (3.10)$$

$$\frac{\partial^2 \phi}{\partial x^2} = \frac{e}{\epsilon_0} \left(\sum_\alpha p_\alpha |\psi_\alpha|^2 - n_0 \right), \quad (3.11)$$

where p_α is the occupation probability of state ψ_α . Manfredi et al. then suggest interpretation of these equations as analogous to the classical multistream model of Dawson, which assumes the distribution function can be expressed as a probabilistic sum of streams α :

$$f(x, v, t) = \sum_\alpha p_\alpha n_\alpha(x, t) \delta(v - u_\alpha(x, t)) \quad (3.12)$$

The Vlasov equation then implies the following properties (representing continuity and momentum equations) for each α :

$$\frac{\partial n_\alpha}{\partial t} + \frac{\partial}{\partial x} (n_\alpha u_\alpha) = 0 \quad (3.13)$$

$$\frac{\partial u_\alpha}{\partial t} + u_\alpha \frac{\partial u_\alpha}{\partial x} = \frac{e}{m} \frac{\partial \phi}{\partial x} \quad (3.14)$$

The proposed quantum extension to these equations is made by linking the classical stream densities to quantum amplitudes - an approach in clear correspondence with the Bohmian description we have already discussed. Associating with each stream an amplitude

$$\psi_\alpha = R_\alpha \exp(iS_\alpha/\hbar), \quad (3.15)$$

we have densities and velocities defined via

$$n_\alpha = |\psi_\alpha|^2 = R_\alpha^2, \quad (3.16)$$

$$u_\alpha = \frac{1}{m} \frac{\partial S_\alpha}{\partial x}. \quad (3.17)$$

Inserting these quantities into the Hartree equations (eqs. 3.10, 3.11) yields a multistream quantum hydrodynamics model:

$$\frac{\partial n_\alpha}{\partial t} + \frac{\partial}{\partial x}(n_\alpha u_\alpha) = 0, \quad (3.18)$$

$$\frac{\partial u_\alpha}{\partial t} + u_\alpha \frac{\partial u_\alpha}{\partial x} = \frac{e}{m} \frac{\partial \phi}{\partial x} + \frac{\hbar^2}{2m^2} \frac{\partial}{\partial x} \left(\frac{\partial^2(\sqrt{n_\alpha})/\partial x^2}{\sqrt{n_\alpha}} \right). \quad (3.19)$$

The new term emerging on the right hand side of eq. 3.19 can be associated with the Bohm potential of eq. 3.4. These equations, as they stand, dictate the evolution of each individual stream, and as a result still remain a complex system of $2N$ (where N is the number of streams) equations to be solved. We can, however, use these equations to assess the evolution of the total density and velocity, defined through:

$$n = \sum_{\alpha} p_{\alpha} n_{\alpha}$$

$$u = \sum_{\alpha} p_{\alpha} \frac{n_{\alpha}}{n} u_{\alpha}$$

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

Multiplying the multistream equations by p_α and summing over α we obtain

$$\frac{\partial n}{\partial t} + \frac{\partial}{\partial x}(nu) = 0 \quad (3.20)$$

$$\begin{aligned} \frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} &= \frac{e}{m} \frac{\partial \phi}{\partial x} + \\ &\frac{\hbar^2}{2m^2} \frac{\partial}{\partial x} \sum_{\alpha} p_{\alpha} \left(\frac{\partial^2(\sqrt{n_{\alpha}})/\partial x^2}{\sqrt{n_{\alpha}}} \right) \\ &- \frac{1}{mn} \frac{\partial P}{\partial x} \end{aligned} \quad (3.21)$$

where P , a pressure-like term, has been introduced as:

$$P = mn \left[\frac{\sum_{\alpha} p_{\alpha} n_{\alpha} u_{\alpha}^2}{n} - \left(\frac{\sum_{\alpha} p_{\alpha} n_{\alpha} u_{\alpha}}{n} \right)^2 \right] \quad (3.22)$$

Simplification of the equations is now possible, provided approximations can be made for the two terms involving summation over streams. Manfredi et al. suggest a classical equation of state, $P = P(n)$, for P , and a linearisation of the Bohm potential[55] – it is this linearisation (which can be shown to hold for long wavelengths) that will form the basis for our core approximation in the Bohmian approach we will develop. The approximation takes the form:

$$\sum_{\alpha} p_{\alpha} \left(\frac{\partial^2(\sqrt{n_{\alpha}})/\partial x^2}{\sqrt{n_{\alpha}}} \right) \rightarrow \frac{\partial^2(\sqrt{n})/\partial x^2}{\sqrt{n}} \quad (3.23)$$

with which the simple pair of equations 3.20, 3.21 is closed, and can be readily solved to approximate the more complex multistream system.

3.2.3 Linearisation for Bohmian Trajectories

Having laid out the theory behind Bohmian mechanics, we now want to develop this theory into a new computational method that:

- treats the full dynamics of *both* electrons and ions;
- is fast enough to fully resolve the ion-ion dynamic structure in a reasonable simulation time;
- is applicable to systems of particles at finite temperature.

To do this, we will need to make some simplifying approximations. For computational purposes, the most important of these is the use of a thermally linearised Bohm potential. First, let's recall the full (and exact) non-linearised method for a pure quantum system. In the Bohmian picture, we saw that a quantum system can be replicated by an ensemble of classical systems. In this context, we will refer to the classical systems as N -trajectories, each of which moves through $3N$ -dimensional phase space. The equations of motion of these N -trajectories contain an additional Bohm potential, which is a functional of the density of the N -trajectories in phase space. The Bohm potential introduces non-local quantum effects to the motion of the N -trajectories, which map out the probability density of the quantum system.

Now, let's consider a similar quantum system, this time in thermal equilibrium (i.e. the actual kind of system we need to simulate). The crux of our approximation is this: we model the thermal system as an ensemble of N -trajectories – in the same manner as the pure case – *but with each N -trajectory's temperature set independently by a thermostat*. A schematic of this process is shown in fig. 3.1.

This is an approximation to the exact dynamics as we're coupling the individual N -trajectories to a heat bath, as opposed to the full quantum system. We model




Figure 3.1: A schematic of the linearisation approximation. In (a), we show the exact system. Each box represents a pure quantum state, with its own N -trajectory ensemble (represented by filled circles). Equilibrium statistical properties of the finite temperature system can be extracted through a thermal average over the properties of many pure states – that is, the system as a whole is coupled to a heat bath. In (b), we show the linearised system. Each N -trajectory is coupled to a thermostat independently – rather than summing over many pure states, we evolve an averaged thermal state. As discussed in the text, this allows us to simplify calculations to a single N -trajectory, which generates its own Bohm potential self-consistently.

a thermal average of pure states – each comprised of an independent N -trajectory ensemble – with an ensemble of effective N -trajectories that are each thermally coupled. As a result, in this approximation, each N -trajectory feels a Bohm potential that is a form of linearised average over states. In section 3.2.2, we discussed how quantum hydrodynamics deals with the Bohm potential through an effective average over states (the mapping described in eq. 3.24). Conceptually, the Bohm potential in our model is averaged over in the same way. Just as in the hydrodynamic approach, rather than tracking a set of independent phase space densities and Bohm potentials for an array of pure quantum states, we consider an averaged Bohm potential that is a functional of the total phase space density.

At first glance, this approximation still leaves us with a complex (and extremely expensive to simulate) ensemble of N -trajectories. The problems of the exact calculation appear to remain – most crucially, the need for an intractable number of N -trajectories (which must be dense in $3N$ -dimensional space). However, the fact that the individual N -trajectories are connected to thermostats allows us to simplify our calculations. This follows from two observations:

1. *Properties of a classical system with correlation-dependent potentials can be determined self-consistently.* For an arbitrary system of N well-localised particles, the N -particle correlation function can be written as $g(\mathbf{x}) = P(\mathbf{x})/P_0(\mathbf{x})$, where P denotes the joint positional probability distribution of the particles, and P_0 is the distribution for a non-interacting classical system with equal particle densities. Here, the variable \mathbf{x} is the set of particle positions, $\mathbf{x} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$. We may construct inter-particle potentials that are functionals of g : $V(\mathbf{x}) = V_0(\mathbf{x}) + V_g(\mathbf{x}|g)$, where V_0 denotes pair interactions and external forces, and V_g is a contribution that varies with g . The equilibrium properties of a system in this potential can be found self-consistently. Starting from an initial guess for $V(\mathbf{x})$, we can calculate

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

$g(\mathbf{x})$, through a Monte Carlo simulation or a similar method [23]. This value of the N -particle correlation function g then gives rise to a new approximation for the potential. Iterating this procedure allows for both g and V to be found.

2. *Our linearised Bohmian system is equivalent to a classical system with correlation-dependent potentials.* Consider a number of coupled thermal N -trajectories representing our linearised Bohmian system. Assuming that the system is in a temperature regime in which the particle motion is ergodic, we find that each N -trajectory has the same time-integrated correlations; that is, each has the same g . In the limit of infinite N -trajectories in our ensemble, it follows that the configuration space density of N -trajectories Φ is exactly proportional to this common g . As a result, each N -trajectory moves in a common static potential, and, as this static potential is a functional of configuration space density, Φ , it is equivalently just a functional of g .

Combining the results above, our linearisation approximation becomes a simple mapping:

$$V_B(\mathbf{x}|\Phi) \mapsto V_g(\mathbf{x}|g) = -\frac{\lambda\hbar^2}{2\sqrt{g}} \sum_{i=1}^N \frac{\nabla_i^2}{m_i} \sqrt{g}, \quad (3.24)$$

where m_i is the mass of particle i and λ is a linearisation factor to be determined. As g is common to all N -trajectories, our approximation scheme allows us to consider just a single N -body classical system (fig. 3.1b). The required simulation is thus amenable to (computationally cheap) classical Molecular Dynamics [23]. The classical particle trajectories simulated then approximate the statistics of the full quantum system.

3.3 Elements of the Bohmian Trajectory Method

The theory outlined above defines an approximate approach to modelling quantum systems via effectively classical trajectories; however – as with classical MD, or DFT, or any other modern modelling approach – there are multiple simulation elements and approximation schemes that must now be constructed to obtain a practical simulation method. This section describes each of these in turn.

3.3.1 Correlation Closure

The full correlation function, g , as it appears in eq. 3.24, cannot feasibly be calculated for many-body systems. The situation is the same as that associated with the many-body wavefunction (or density matrix) – as the number of system coordinates grows, the space over which the wavefunction is defined grows geometrically. Numerical calculations of the correlation function become impractical.

To remedy this issue, we seek a closure approximation to the full correlation function g , appearing in eq. 3.24, in such a way that the Bohm potential can be constructed for the system. The simplest, non-trivial approximation of this kind is a direct product representation in terms of pair correlation functions

$$g(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) \simeq \prod_{j>i} g(\mathbf{x}_i, \mathbf{x}_j), \quad (3.25)$$

where the pair correlation $g(\mathbf{x}_i, \mathbf{x}_j)$ is the full correlation function with all other coordinates integrated out:

$$g(\mathbf{x}_i, \mathbf{x}_j) = \frac{\left(\prod_{k \neq i, j} \int d\mathbf{x}_k \right) g(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N)}{\Omega^{N-2}}. \quad (3.26)$$

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

Here, Ω is a normalisation volume. Combining eq. 3.24 and eq. 3.25 we arrive at an easily calculable expression for the Bohm potential,

$$V_B(\mathbf{x}|g) \simeq -\frac{\hbar^2}{2} \sum_{j>i} \frac{\lambda_{ij} \left(\frac{\nabla_i^2}{m_i} + \frac{\nabla_j^2}{m_j} \right) \sqrt{g(\mathbf{x}_i, \mathbf{x}_j)}}{\sqrt{g(\mathbf{x}_i, \mathbf{x}_j)}}, \quad (3.27)$$

in which we have retained interactions up to pair terms only, and have generalised the dependence on λ to a set of λ_{ij} to accommodate different particle species.

At first glance, it may appear that a simple pair-product relationship of this kind is overly restrictive – clearly, effects beyond two-body interactions are crucial to the underlying dynamics of the systems we wish to treat (hence our need for a new method). However, it is important to note that our invocation of this approximation does not preclude a treatment of many-body correlations, of third order and beyond. Throughout we will calculate pair correlations using full N -body simulations (either CMC or MD), such that many-body correlations are implicitly included – our use of this closure is then akin to the usual approximate treatment of many-body systems via pair potentials only, rather than a wholesale abandonment of correlations beyond second order. Ultimately, the validity of the closure as an approximation tool can be assessed – at least to some degree – by the quality of the results it produces: This assessment will be carried out in section 3.5.

3.3.2 Potential Corrections

At the base level of the linearised Bohmian approach, we model the system within the Hartree approximation – effects due to the electron exchange interaction are not directly accounted for. To rectify this, we introduce an additional potential term. By inverting the known pair correlations of the non-interacting electron gas using inverse Monte Carlo (IMC, to be discussed in greater detail below) [56] or the

Hypernetted Chain Equation (as discussed in section 2.3.3), we can find pair potentials that reproduce the quantum correlations of the free electron gas exactly. We can then introduce these pair potentials as additional contributions to the system Hamiltonian.

This procedure follows similar lines to the potential inversion method of Lado [57], which subsequently became the basis for the Classical Map method of Perrot and Dharma-wardana for numerous applications [58, 59, 60]. This technique has also proven to be successful in related semi-classical correlation studies [26, 61, 62] (as alluded to in section 2.3.7) and can be seen as analogous – in terms of correcting for exchange effects – to the Pauli pressure correction emerging in Quantum Hydrodynamics (eq. 3.22).

By insisting that the total potential is equal to the IMC result V_{IMC} , when applied to a non-interacting electron gas, we arrive at an expression for the additional Pauli potential

$$V_P(\mathbf{x}) = -V_B(\mathbf{x}|g_0) + \sum_{i,j} V_{\text{IMC}}(\mathbf{x}_i, \mathbf{x}_j|g_0). \quad (3.28)$$

Here, g_0 is the ideal electron gas pair correlation, and $V_B(x|g_0)$ is the Bohm potential evaluated at g_0 . The full system potential V is then given by

$$V = V_{ext} + V_{int} + V_B + V_P, \quad (3.29)$$

where V_{ext} is an external potential (which is set to zero throughout our calculations in this chapter), V_{int} is the contribution from direct pair interactions of the simulated particles (e.g., the Coulomb potential), V_B is the Bohm potential, and V_P is the Pauli potential of eq. 3.28.

3.3.3 Pseudopotentials

Thus far, all electrons have been treated on equal footing. However, at the level of our pair-product closure approximation, the core electrons of the ion species require an explicitly separate treatment from the (effectively free) valence electrons. The pseudopotential approximation, commonly used in implementations of DFT, provides a simple alternative to modelling the core electrons directly. Core electrons are removed from the system to be simulated and ion-electron potentials are constructed in such a way that valence electrons form the correct density profile in a reference calculation. For our purposes, we use a reference single ion DFT calculation to determine an electron density profile and construct the pseudopotential.

The Troullier-Martins approach (also used for DFT pseudopotentials [63]) is a particularly simple choice for pseudopotential construction. We first perform scalar-relativistic electron density calculations within DFT for an isolated ion with both n_{core} and $n_{core} + 1$ electrons. The difference of these densities is used as an input valence electron density. We then construct the pseudopotential by requiring the following properties:

1. The radially integrated electron density up to a cutoff, r_c , is the same for both the full ion system and pseudosystem. When solving for the two-body electron-ion pseudosystem via the Schrödinger equation, we must have the same total electron density within r_c as we found with the full DFT calculation.
2. The pseudopotential for $r > r_c$ is equal to the (screened) Coulomb potential
3. The pseudopotential is smooth and continuous at r_c (ref. [63] discusses the specifics of the smoothing requirement)

3.3 Elements of the Bohmian Trajectory Method

Due to the direct connection to the approach used in standard DFT implementations, it should be possible to re-purpose existing pseudopotential generation codes for the method explained here. This should prove useful for materials with complex bound state structures where more sophisticated pseudopotential construction schemes have been shown to provide greater accuracy.

While the Troullier-Martins scheme above is easily implemented, we can also consider a more general-purpose pseudopotential generation technique, beginning from first-principles data for a full electron-ion system. Given ion-ion and electron-ion radial distribution functions for a system, we can combine traditional Inverse Monte Carlo and the generalised IMC technique that we use to set λ values (to be discussed in section 3.3.4) to obtain best-fit values of both the electron-ion potential and λ simultaneously. The pseudopotential is then determined as the difference between the IMC result and the Bohm potential evaluated for the given RDFs, as per eq. 3.28. This fully determines the linearised Bohm potentials for the given system conditions, and generates the pseudopotential as a side effect for further use. This ‘on-the-fly’ method guarantees optimal static property matching for the Bohm model for the input configuration.

As a future extension beyond the present work, we propose application of this on-the-fly method for every density and temperature we seek to model, rather than relying on the transfer of pseudopotentials from other conditions. This renders potential transferability – that is, the effectiveness of potentials in conditions outside of those in which they were constructed – unimportant. The associated cost is the need for more detailed DFT inputs – both ion-ion and electron-ion static correlations are required at each temperature/density under consideration. The potential construction procedure then becomes an extension of typical Classical Map methods, in which potentials are moulded to replicate known properties of

the exact system, but with electronic potentials constrained to our Bohmian potential form. Together with our quantum thermostat for the dynamics, the constraint on potentials to be consistent with Bohmian mechanics ensures our system is consistent with our quantum approximation model, as opposed to emerging from a purely classical statistical approach.

Accompanying the application of pseudopotentials – of any kind that we have discussed – is the sacrifice of a full treatment of bound-free electron transitions. As the number of valence electrons must be known a priori, as in implementations of orbital free DFT (OFDFT), we restrict our method here to cases where these numbers are both calculable/known and reasonably well-defined.

3.3.4 Setting λ

With this, the static properties of the model are completely described – it remains only to set the linearisation parameter(s) λ_{ij} . In this work we set $\lambda_{ij} = 1$ for the ion-ion and ion-electron contributions, and seek to fix the electron-electron contribution $\lambda := \lambda_{ee}$. Providing we can set this value with only *static* knowledge of the system, we can perform subsequent *dynamic* simulations of wDM with no free parameters. To approach this, we re-examine our core approximation. In addition to its similarities to the linearisation used in Quantum Hydrodynamics, our approximation is connected to the von-Weizsacker correction to Thomas-Fermi DFT [64] for the electronic case. This is clear if we approximate R^2 in the Bohm potential by the real-space electron density: our linearised potential becomes equal to von-Weizsacker’s density functional when $\lambda = 1$. This functional corresponds to the high-density limit – kinetic theory corrections yield equality when $\lambda = 1/9$ for the low-density case[65]. These correspondences suggest that the optimal choice of λ has a non-trivial dependence on the system density.

3.3 Elements of the Bohmian Trajectory Method

In general, λ also exhibits a temperature dependence. We can demonstrate this explicitly, and support our use of this potential form, by treating the simplified case of the harmonic oscillator. Here it becomes possible to choose $\lambda = \lambda(\beta)$, with $\beta = 1/k_B T$ the inverse temperature, such that the exact real space probability density of the quantum system is reproduced. Noting that the diagonal element of the oscillator's density matrix is:

$$\rho(x) = \left[\frac{M\omega}{\pi\hbar} \tanh\left(\frac{1}{2}\beta\hbar\omega\right) \right]^{\frac{1}{2}} \exp\left[-\frac{M\omega}{\hbar} \tanh\left(\frac{1}{2}\beta\hbar\omega\right) x^2\right], \quad (3.30)$$

where M is the particle mass and ω its characteristic frequency, one can show that the choice:

$$\lambda = \frac{y - \tanh(y)}{y \tanh^2(y)}, \quad (3.31)$$

where $y = \beta\hbar\omega/2$, produces equality in real-space density for finite β (see fig. 3.2 - we recover $\rho(x) = \rho_0 \exp[(-\beta V(x|\rho))]$ by setting λ as this value in eq. 3.27).

For a general system, it is clear that determining λ theoretically is not a trivial task. To set the value accurately, we take the approach of coupling the Bohm calculation to known static ion correlations, by fixing a system-dependent λ such that the static ion-ion pair correlation is accurately reproduced. This means that the optimal value of λ is determined via:

$$\lambda = \lambda(\rho, \beta) = \operatorname{argmin}_{\lambda'} \left[\int r^2 (g_{\lambda'}(r) - g_{DFT}(r))^2 dr \right], \quad (3.32)$$

where $g_{\lambda'}$ is the ion-ion pair correlation using the Bohm approach with $\lambda = \lambda'$, and g_{DFT} is the ion-ion pair correlation as determined by DFT simulation. The result of this combined DFT-Bohm approach is then the removal of our only unknown parameter.

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

We now set out how the search for the optimal value of λ , as defined as by eq. 3.32, can be carried out numerically. Due to the low computational demands associated with calculating $g_{\lambda'}$ for each λ' , we note that a brute force line-search should be adequate in most cases. However, a more robust approach (applicable to potential forms with arbitrary numbers of free parameters) is possible through a generalisation of Inverse Monte Carlo.

Basic Inverse Monte Carlo Method

Much as it is possible to obtain particle correlations from inter-particle potentials via cmc, as shown in section 2.3.2, it's also possible to do the inverse – that is, obtain effective potentials that reproduce known particle correlations. This approach – Inverse Monte Carlo (imc) – was developed to establish coarse-graining of complex systems, by absorbing uninteresting degrees of freedom into effective potentials.

By adapting this imc algorithm, we'll be able to develop a general approach to creating semi-classical models of a system with partial knowledge of its potentials and correlations. In passing, we have mentioned our usage of imc for obtaining potential contributions that reproduce Pauli repulsion in an electron gas; this section will describe the mechanics of this process, and – with suitable modification – describe how imc can be employed to search over parametrised potential functionals.

We'll start by outlining the original imc formulation. For a system with pair interactions only, we can construct the algorithm (established in ref. [56]) by writing the Hamiltonian of the system as:

$$H = \sum_{\alpha} S_{\alpha} K_{\alpha}, \quad (3.33)$$

where K_α represents the potential between two bodies at a particular distance, and S_α is the number of pairs of bodies separated by that distance. We assume knowledge of the desired mean values of $\{S_\alpha\}$, these being the correlations that we wish to replicate with effective potentials. The set of potentials at each point $\{K_\alpha\}$ are sequentially adjusted until the the desired $\{\langle S_\alpha \rangle\}$ are obtained.

The adjustments of $\{K_\alpha\}$ are found by using a Monte Carlo simulation to obtain the matrix $\langle S_\alpha S_\gamma \rangle$ and the vector $\langle S_\alpha \rangle$. Using an expression for the Boltzmann expectation, we have the relationship (with summation over repeated indices implied):

$$\delta \langle S_\alpha \rangle = \frac{\partial \langle S_\alpha \rangle}{\partial K_\gamma} \delta K_\gamma + O(\delta K^2), \quad (3.34)$$

in which we have the matrix:

$$\frac{\partial \langle S_\alpha \rangle}{\partial K_\gamma} = -\beta (\langle S_\alpha S_\gamma \rangle - \langle S_\alpha \rangle \langle S_\gamma \rangle). \quad (3.35)$$

This set of linear equations can then be inverted to obtain the required (Newton's method) adjustments $\{\delta K_\alpha\}$ in terms of the correlation errors $\{\delta S_\alpha\}$, which is formally:

$$\delta K_\alpha = \left(\frac{\partial \langle S_\alpha \rangle}{\partial K_\gamma} \right)^{-1} \delta \langle S_\gamma \rangle. \quad (3.36)$$

These additions to $\{K_\alpha\}$ are made, and the procedure repeated until convergence (when $\{\delta K_\alpha\}$ are sufficiently small).

Potential Functional Inverse Monte Carlo

To adapt this method to our search for λ values, we begin along the original lines of IMC by writing the positional Hamiltonian of the system as a sum over pair energies, via eq. 3.33. Our goal is to obtain a set of updates to a set of potential pa-

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

rameters, $\Delta\lambda_i$, so that we can effectively adjust the mean bin counts to match our desired pair correlations.

In our case, unlike the original formulation of IMC, the potential points K_α are functionals of the thermally averaged correlations of the system, and are also functions of the chosen λ values, via $K_\alpha = K_\alpha(\{\langle S_\alpha \rangle\}, \{\lambda_i\})$. Here the thermal expectation $\langle X \rangle$ (of an arbitrary function of the system coordinates $X(q)$) is defined by:

$$\langle X \rangle = \frac{1}{Z} \int X(q) \exp(-\beta H) dq, \quad (3.37)$$

where Z is the partition function and q is the set of all positional coordinates of the system. That is to say, the potentials of the system depend on both the set of λ linearisation parameters, and also the thermally averaged set of pair correlations between bodies (as per eq. 3.27).

Applying eq. 3.37 to S_α and differentiating with respect to λ_i , we obtain:

$$\frac{\partial \langle S_\alpha \rangle}{\partial \lambda_i} = -\beta \sum_\gamma M_{i\gamma} (\langle S_\alpha S_\gamma \rangle - \langle S_\alpha \rangle \langle S_\gamma \rangle), \quad (3.38)$$

where the total λ derivative matrix M is given by:

$$M_{i\gamma} = \frac{\partial K_\gamma}{\partial \lambda_i} + \sum_\delta \frac{\partial K_\gamma}{\partial \langle S_\delta \rangle} \frac{\partial \langle S_\delta \rangle}{\partial \lambda_i}. \quad (3.39)$$

Combining eqs. 3.38, 3.39, we find a separate set of linear equations for each λ value (each value of i):

$$M_{i\gamma} = \frac{\partial K_\gamma}{\partial \lambda_i} - \beta \sum_\epsilon M_{i\epsilon} \sum_\delta \frac{\partial K_\gamma}{\partial \langle S_\delta \rangle} (\langle S_\delta S_\epsilon \rangle - \langle S_\delta \rangle \langle S_\epsilon \rangle). \quad (3.40)$$

Using these relationships, we can now perform a directed search for the optimal λ values. To leading order, we can write an equation for changes in $\langle S_\alpha \rangle$ in terms of

changes in λ_i :

$$\Delta \langle S_\alpha \rangle = \sum_i \frac{\partial \langle S_\alpha \rangle}{\partial \lambda_i} \Delta \lambda_i. \quad (3.41)$$

If we perform a calculation of g with a given set of λ values, we can simultaneously calculate values of $\langle S_\alpha \rangle$, $\langle S_\alpha S_\gamma \rangle$ numerically. We can then solve the linear equation set 3.40, and in turn have an overdetermined set of equations (eq. 3.41) for the required changes in the λ parameters to match the desired correlations. The optimal changes in λ can then be determined by least-squares inversion of this equation. Defining the matrix of values $A_{\alpha i} = \partial \langle S_\alpha \rangle / \partial \lambda_i$, the formal solution is:

$$\Delta \lambda = (A^T A)^{-1} A^T \Delta \langle S \rangle. \quad (3.42)$$

This can be determined directly, or implicitly through QR factorisation for computational efficiency. This update procedure can then be applied iteratively to arrive at the optimal λ .

3.3.5 Fermi Thermostats

As highlighted in section 2.3, most modern DFT-MD simulations rely on either the Nosé-Hoover or Langevin thermostat to establish ion dynamics at a given temperature [14, 15, 36]. To recap: In the case of the Nosé-Hoover thermostat, an additional dynamic variable, coupled linearly to particle momenta through a friction term, is introduced to the equations of motion [66]. The parameters for the thermostat can be chosen to ensure a balance between temperature stability and the equilibration time, but in general the dynamics are parameter-insensitive.

The Langevin thermostat, on the other hand, adds both a frictional term and a stochastic noise term to the equations of motion [67]. Again, this ensures the canonical distribution is sampled. The magnitude of the friction and noise are con-

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

trolled by a free parameter, the Langevin friction σ , which must be chosen to be sufficiently small to minimise spurious effects on the particle dynamics, while remaining large enough that the correct distribution is sampled in a reasonable time-frame. We have also seen that σ can be scaled up to a larger value as a first approximation to electron-ion collision effects in Born-Oppenheimer simulations[15].

As our method does not rely on the Born-Oppenheimer approximation we can employ the parameter-insensitive Nosé-Hoover thermostat for the ions. This represents a key advantage of our approach: The electron dynamics are modeled directly, and recourse to an additional noise approximation is not necessary. The ion thermostat, and accordingly the ion dynamics we wish to calculate, therefore do not rely on the unknown free parameter σ to mimic the dynamic electrons.

In dealing directly with dynamic electron trajectories, however, in contrast to Born-Oppenheimer DFT-MD, we must develop an appropriate thermostat for the electrons as well. In the degenerate and semi-degenerate cases, the electronic velocity distribution is known to be far from the Boltzmann distribution. While a standard thermostat allows us to move past the Born-Oppenheimer approximation and include dynamic electrons, it cannot capture the effect of Fermi-statistics on the electron motion.

An exact treatment of our quantum system should, in principle, fully capture a Fermi distribution of velocities, assuming one can incorporate properly symmetrised wavefunctions into dynamical calculations. However, our series of approximations – in particular, our use of a linearised potential, required in order to construct a fast simulation scheme – preclude an exact treatment of exchange effects on the dynamics. Instead, we must consider a leading order correction to the electron equilibrium velocity distribution via the equations of motion.

In summary, we desire a thermostat that ensures static correlations are maintained, while enforcing a non-Boltzmann distribution of velocities. This will act as a dynamical correction for exchange effects – much as the Pauli potential described in section 3.3.2 acts as a static correction. With this in mind, we will develop modifications of both standard Langevin and Nosé-Hoover thermostats.

By building these thermostats – in addition to our quantum-derived semiclassical potentials – we move past the traditional Classical Map paradigm. While Classical Map methods adjust the effective temperature of electrons to match a desired mean electron velocity – partially accounting for the non-classical dynamics of the electrons[58] – we can match a target electron velocity distribution in much greater detail through a modified thermostat. Rather than using a strictly classical system model (with modified *static* properties), we modify the full *dynamic* evolution of the system.

Langevin-Style Fermi Thermostat

The equations of motion of a particle in a stochastic force field can be written quite generally as:

$$d\mathbf{X} = \boldsymbol{\mu}dt + \boldsymbol{\sigma} \cdot d\mathbf{W}, \quad (3.43)$$

where

$$\mathbf{X} = (X_1, X_2, X_3, V_1, V_2, V_3)^T, \quad (3.44)$$

and

$$\boldsymbol{\mu} = (V_1, V_2, V_3, \mu_1(\mathbf{X}), \mu_2(\mathbf{X}), \mu_3(\mathbf{X}))^T. \quad (3.45)$$

Here, μ_i is the i^{th} component of the deterministic part of the particle's acceleration. \mathbf{X} is a vector containing the particle's position $(X_1, X_2, X_3)^T$ and velocity $(V_1, V_2, V_3)^T$, $\boldsymbol{\sigma}$ represents the stochastic collision frequency, and \mathbf{W} is a multidimensional Wiener process.

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

dimensional Wiener process [68] (for the time being, it is sufficient to note that the Wiener process acts as a white noise contribution to the force – a slightly more complete description of the Wiener process and Stochastic Differential Equations is deferred to section 5.2.2). We take the driving noise to be uniform and isotropic in velocity space. Hence,

$$\boldsymbol{\sigma} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ \sigma & 0 & 0 \\ 0 & \sigma & 0 \\ 0 & 0 & \sigma \end{pmatrix}. \quad (3.46)$$

Defining

$$\mathbf{V} = (V_1, V_2, V_3)^T, \quad (3.47)$$

$$\mathbf{R} = (X_1, X_2, X_3)^T, \quad (3.48)$$

$$\boldsymbol{\eta} = (\mu_1(\mathbf{X}), \mu_2(\mathbf{X}), \mu_3(\mathbf{X}))^T, \quad (3.49)$$

we can then write a Fokker-Planck equation for the equilibrium particle probability density, p , in configuration space at long times [68]. We have

$$\frac{\partial}{\partial \mathbf{X}} \cdot (\boldsymbol{\mu}p) + \frac{\sigma^2}{2} \left(\frac{\partial}{\partial \mathbf{V}} \right)^2 p = 0, \quad (3.50)$$

which reduces to

$$-\mathbf{V} \cdot \frac{\partial p}{\partial \mathbf{R}} - \frac{\partial}{\partial \mathbf{V}} \cdot (\boldsymbol{\eta}p) + \frac{\sigma^2}{2} \left(\frac{\partial}{\partial \mathbf{V}} \right)^2 p = 0. \quad (3.51)$$

We approximate the probability density distribution of the particles as decoupled in momentum and position space, via

$$p \equiv p_L = A \frac{1}{1 + e^{\beta(m\mathbf{V}^2/2 - \mu_c)}} \cdot e^{-\beta U(\mathbf{R})}, \quad (3.52)$$

where U is the potential energy, A is a normalisation constant, and μ_c is the chemical potential. By inserting p from eq. (3.52) into eq. (3.51) we can solve for η ; the equations of motion in eq. (3.43) can then be written as

$$d\mathbf{R} = \mathbf{V} dt, \quad (3.53)$$

$$d\mathbf{V} = \frac{\mathbf{F}}{m} (1 + e^{\beta E}) [\log(1 + e^{\beta E}) - \beta E] dt - \frac{\sigma^2}{2} m \beta \left(\frac{e^{\beta E}}{1 + e^{\beta E}} \right) \mathbf{V} dt + \boldsymbol{\sigma} \cdot d\mathbf{W}, \quad (3.54)$$

where m is the particle's mass, $E = m\mathbf{V}^2/2 - \mu_c$ and $\mathbf{F} = -\partial U/\partial \mathbf{R}$ (and the boundary conditions were set to ensure that the standard Langevin equation is recovered at high energy). These represent the modified equations of motion, solvable with standard methods, which have the desired equilibrium distribution function under the assumption of a uniform noise (frequency) distribution.

Nosé-Hoover-Style Fermi Thermostat

We can now consider a Nosé-Hoover-Style thermostat with the same equilibrium distribution. By analogy with the Langevin case, we begin with the equations for the dynamics

$$\dot{\mathbf{R}} = \mathbf{V}, \quad (3.55)$$

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

$$\begin{aligned} \dot{\mathbf{V}} = & \frac{\mathbf{F}}{m}(1 + e^{\beta E}) [\log(1 + e^{\beta E}) - \beta E] \\ & - \frac{\alpha}{2} m \beta \left(\frac{e^{\beta E}}{1 + e^{\beta E}} \right) \mathbf{V}, \end{aligned} \quad (3.56)$$

where α is a new dynamic variable coupling the system to the heat bath. In equilibrium, our deterministic system must now satisfy the generalised Liouville equation [69]

$$\begin{aligned} \frac{\partial p_{NH}}{\partial \mathbf{R}} \cdot \dot{\mathbf{R}} + \frac{\partial p_{NH}}{\partial \mathbf{V}} \cdot \dot{\mathbf{V}} + \frac{\partial p_{NH}}{\partial \alpha} \dot{\alpha} \\ + p_{NH} \left(\frac{\partial}{\partial \mathbf{R}} \cdot \dot{\mathbf{R}} + \frac{\partial}{\partial \mathbf{V}} \cdot \dot{\mathbf{V}} + \frac{\partial \dot{\alpha}}{\partial \alpha} \right) = 0, \end{aligned} \quad (3.57)$$

where the probability distribution p_{NH} is given by $p_{NH} = p_L \exp(-\beta\alpha^2/2m_\alpha)$ - in which we force α to have a Gaussian distribution, as per the standard Nosé-Hoover scheme - and m_α is a thermostat mass. Again, asserting that we must recover the standard Nosé-Hoover thermostat in the classical limit, we can solve this equation for $\dot{\alpha}$. We obtain

$$\frac{\dot{\alpha}}{m_\alpha} = -\frac{m}{2} \left(\frac{e^{\beta E}}{1 + e^{\beta E}} \right) \left[\beta m \mathbf{V}^2 \left(\frac{1 - e^{\beta E}}{1 + e^{\beta E}} \right) + 3 \right], \quad (3.58)$$

which, in combination with eqs. (3.55) and (3.56), forms a closed dynamical system with the desired equilibrium distribution. For multiple particles, $\dot{\alpha}$ is just a sum over terms of this form for each particle. fig. 3.3 demonstrates the recovery of the desired momentum distribution for a sample system.




Figure 3.2: Values of λ producing exact correspondence (in position-space density) for a quantum harmonic oscillator.




Figure 3.3: Reproduction of a Fermi kinetic energy distribution using our modified thermostat. The system considered comprised particles at the same density/temperature as valence electrons in twice-compressed-Al at 3.5 eV, subject to a simple exponential pair potential $V = V_0 \exp(-\kappa r)$ with $V_0 = 1 \text{ Ha}$ and $\kappa = 1 \text{ a}_B^{-1}$. Distributions are normalised to have maxima of one.

3.4 Implementation

3.4.1 Software Overview

For the purposes of implementing the method(s) established in sections 3.2 and 3.3, a new code has been developed. It has been constructed with an extensible and loosely-coupled object-oriented design, allowing the various elements of simulation such as Verlet integrators, force calculators, and numerical thermostats (including the novel thermostats of section 3.3) to be individually applied to produce a simulation. This will allow for future extension methods (such as those described in chapter 6) to be implemented straightforwardly. The code is also inclusive of related static methods (Monte Carlo, HNC solver, HNC inverter and IMC solver for obtaining effective potentials).

The ‘glue code’ and user-facing interface are written in Python (3), to allow for rapid prototyping and extension, and the computationally demanding bottlenecks are implemented in C/C++ and wrapped with Cython[70] (a statically typed extension to Python, with speeds at the same order of magnitude as C achieved through compilation). The code is natively parallel in performance-critical sections using language-level concurrency features in C++/C (i.e. there is no reliance on OpenMP). Several usability features have also been incorporated, including video output of dynamic simulations, storage of data of arbitrary size in the transferable HDF5 format, and notebook compatibility to allow interaction with simulations via Jupyter notebooks.

Alongside the core code, several analysis tools have been developed, including a highly-optimised DSF calculator for use with large data files (that may not fit in memory), and a detailed logging system. The DSF calculator uses particle position outputs over time to explicitly calculate the integral in eq. 2.1; the in-

intermediate scattering factor is constructed and Fourier transformed following a smoothing procedure (following the simple procedure outlined in Chapter 4 of ref. [71]). Further, a general purpose ‘runner’ module has been implemented for starting/monitoring/validating multiple simulations, to allow for searches over parameter space to be set up quickly and cleanly. Interoperability with data formats common to other related codes (LAMMPS[72], ABINIT[73]) has also been implemented.

3.4.2 Building Potentials

From an algorithmic viewpoint, the linearised structural method described in section 3.2.3 is implemented through a procedure of static structure determination, followed by dynamical simulation. As a first step, the equilibrium Bohm potential and static correlation functions for the system are found self-consistently via cmc. The Monte Carlo update employed is taken as a single-particle move, with acceptance probability given by $p = \min(1, \exp(-\beta\delta E))$, where δE is the potential energy change due to the move. The magnitude of particle movements is dynamically adjusted to obtain $\langle p \rangle \approx 0.5$, and a simple admixture procedure is used to update the Bohm potential as a functional of the pair correlations.

The explicit stages of the algorithm are as follows:

1. Initialise pair correlations g to unity, and Bohm potentials V_B to 0
2. Perform a standard classical Metropolis update, recording pair correlations for each updated particle to form g'
3. Update g through $g \rightarrow \gamma_g g' + (\gamma_g - 1)g$
4. Update V_B through $V_B \rightarrow \gamma_B V'_B + (\gamma_B - 1)V_B$, where $V'_B = V_B(x|g)$

5. Complete if $|g - g'| < \epsilon_g$ and $|V_B - V'_B| < \epsilon_B$; else return to 2.

Here, ϵ_g and ϵ_B are (small) convergence parameters, γ_g and γ_B are mixing parameters, and $|f - f'|$ represents a difference metric, which we define as simply

$$|f - f'| = \int (f(x) - f'(x))^2 dx. \quad (3.59)$$

This produces mean pair potentials for each of the species involved, which are then used in a standard MD simulation (inclusive of our modified thermostats) to produce dynamic properties. This is a special case of a more general approach, in which the potentials of mean force implied by static structure (via inversion of the HNC approximation or IMC) can be extracted for use in dynamic simulation. We therefore expect that a traditional Path-Integral Monte Carlo simulation – which produces static results with high accuracy – could be used to obtain a more accurate baseline (albeit with a higher computational cost, and under an appropriate approximation for the dynamical electron temperature).

3.5 Validation and Predictions

To demonstrate the strength of the method developed in this chapter, we apply it to several warm dense matter test cases. In addition to being the principal target of this thesis, WDM is an ideal testbed for quantum simulations more generally [74], as it combines the need for quantum simulations of degenerate electrons with the description of a strongly interacting ion component.

The primary case we consider is compressed liquid aluminum with a density of 5.2 g cm^{-3} and a temperature of 3.5 eV, which allows for direct comparison with previous results from more traditional methods. The linearisation factors and corresponding Bohm potentials were constructed according to the core algorithm in section 3.4.2. The information about static correlations needed as input – specifically the radial distribution function of the ionic system – was produced with OFDFT calculation running the using VASP package [75, 76, 77, 78] for short times. For this purpose, we used plane wave and augmentation energies of 250 eV and 500 eV, respectively, the 3e Al pseudopotential provided with VASP, and a simulation box containing 256 atoms. A total of 2680 electron bands were included, such that the occupations of the highest energy bands were less than 10^{-6} . This simulation was run for 5000 time steps of 1 fs.

Electron density profiles required to produce pseudopotentials were calculated using the GPAW DFT package [79, 80, 81] with the PBE functional [82]. The MD simulation step was carried out with the MD code described in 3.4 (inclusive of our Fermi thermostats). Periodic boundary conditions were employed, and the simulation box contained 256 aluminum ions. The simulations were run with a time step of 0.1 a.u. All properties were extracted by averaging over multiple simulation runs of 10^7 time steps.

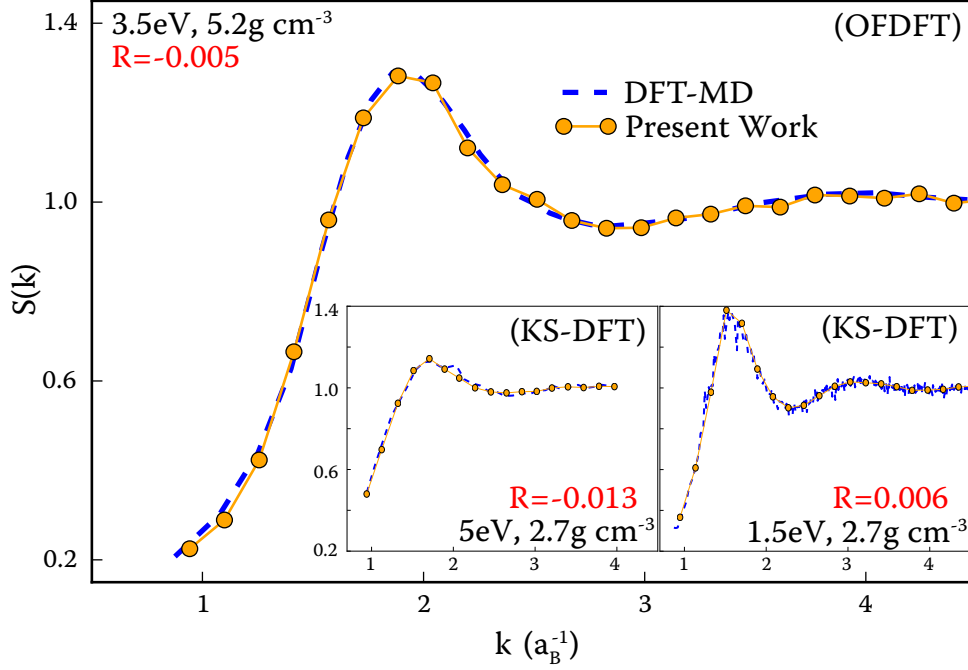


Figure 3.4: Static ion-ion structure factors for aluminum. The static structure factor is defined as $S(k) = \int S(k, \omega) d\omega$. The main graph compares our results from Bohmian dynamics with data obtained by orbital free DFT [15] for a density of 5.2 g cm^{-3} and a temperature of 3.5 eV . The lower insets compare our results to data from full Kohn-Sham DFT simulations at solid density and two different temperatures. The excellent agreement of the methods is also demonstrated by the very small differences in pressure as quantified by the parameter R : These values give the difference in ionic pressure between the methods normalised to the difference of the DFT pressures and the pressure of an ideal gas; that is $R = (P_{Bohm} - P_{DFT}) / (P_{DFT} - P_0)$.

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

The validity and accuracy of our implementation of Bohmian dynamics are supported by the excellent reproduction of static ion-ion correlations from DFT simulations. Fig. 3.4 illustrates the static ion-ion structure factor obtained with the Bohmian trajectories technique. The comparisons with orbital free DFT and the computationally more intensive Kohn-Sham DFT both yield agreement within the statistical error of the simulations. This match was achieved by a single parameter fit defining λ_{ee} (as described in section 3.3.4). The different simulation techniques predict almost the same thermodynamics as shown by the small pressure difference.

Fig. 3.5a shows calculations of the fully frequency-dependent DSF. We see the appearance of side peaks in the DSF that correspond to ion acoustic waves. Their dispersion for smaller wavenumbers, and the corresponding sound speed, are very sensitive to the interactions present in the system. As a result, they reflect the screening of ions by electrons as well as dynamic electron-ion collisions. For larger wavenumbers k , these modes cease to exist due to increased damping.

The dispersion relation of the ion acoustic modes is displayed in fig. 3.5b, which also shows results from the Langevin approach discussed in ref. [15]. The latter approach requires *ad hoc* friction parameters, which were chosen to cover the range between the classical and quantum Born limits. The Bohmian approach, however, does not require a free parameter, thereby allowing access to a self-consistent prediction of the sound speed. This comparison may also be used to assess the quality of the friction parameter applied in the Langevin approach. For the case considered, one finds that neither the classical nor the weak coupling Born limit are applicable – a finding that is typical of the WDM regime.

Collectively, our results demonstrate the strength of the Bohmian approach in modelling quantum systems with strong interactions and nonlinear ion dynamics.

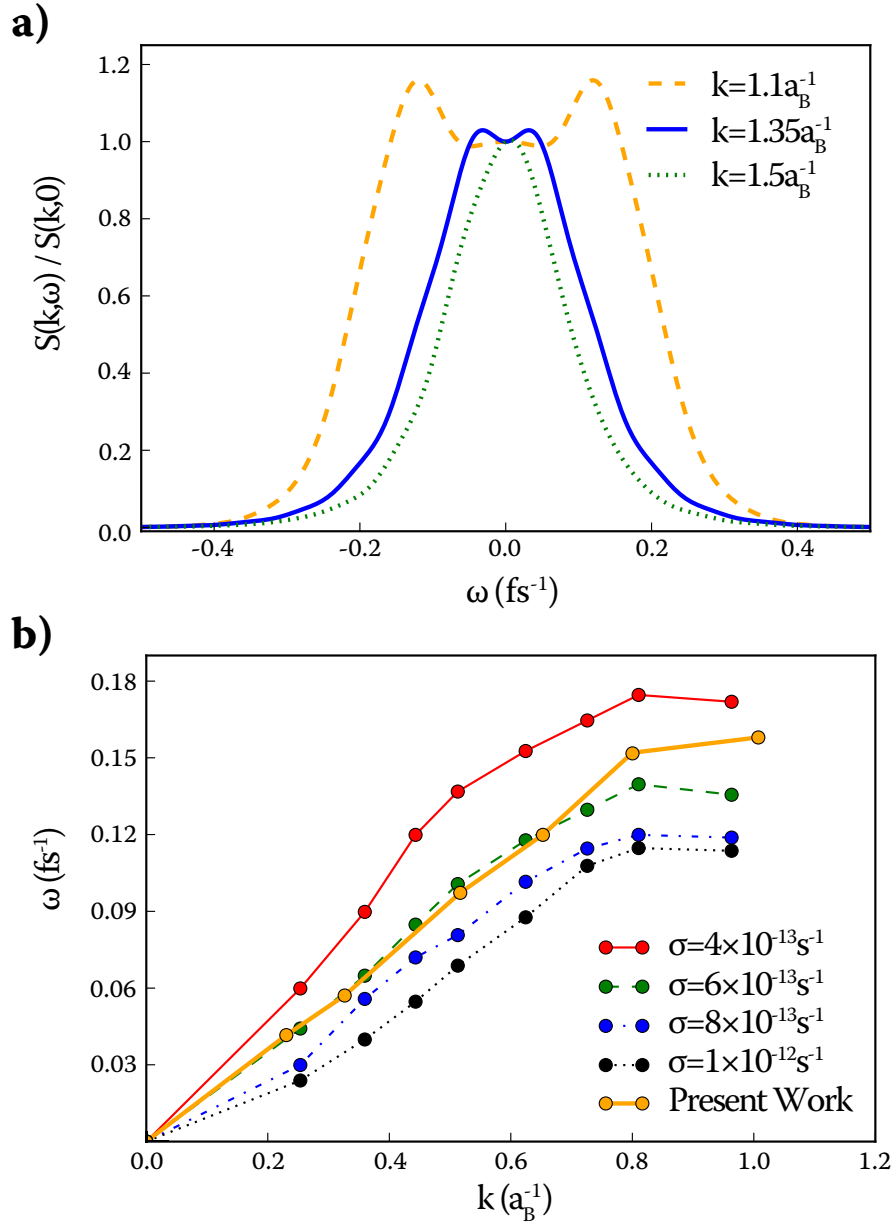


Figure 3.5: Results for the dynamic ion structure for aluminium at 3.5eV and 5.2gcm^{-3} . The upper panel shows the frequency-resolved DSF from the Bohmian dynamics. The lower panel provides a comparison of the dispersion relation of the ion acoustic modes from our Bohmian approach with the data from the Langevin model of ref. [15].

3 Predicting Dynamic Structure: The Bohmian Trajectory Method

For static and thermodynamic properties, we obtain results in very close agreement with DFT simulations. Further – as we have discussed – the standard implementation of DFT-MD invokes the Born-Oppenheimer approximation, while the Bohm approach treats electrons and ions non-adiabatically, retaining the full coupling of the electron and ion dynamics. We are able to examine changes to the ion modes due to dynamic electron-ion correlations that are inaccessible to standard DFT.

The core advantage of our approach, then, is the combination of its relative numerical speed, which allows for the modelling of quantum systems with large numbers of particles, and its retention of electronic dynamics. Calculations dependent on these properties are relevant not just for WDM, but also address core problems in chemical and biological systems (e.g., protein folding), as well as radiation damage of materials [83, 84, 85].

Finally, we note that while our initial implementation of the Bohmian trajectory method focuses on establishing dynamic correlations of systems in thermal equilibrium, generalisation to non-equilibrium systems is also possible. This will involve dynamically updating the system potentials to account for local thermodynamic conditions that evolve in time. In particular, the electron-ion or electron-phonon energy exchange in two-temperature systems is amenable to this approach – we discuss this further as a future path for research in chapter 6.

4

Inferring Static Structure: Inverse Problem Approaches for Plasmas

4.1 Introduction

The methods we have examined and developed thus far have focused on predicting experimental observables from given physical parameters. Such predictions are invaluable for experimental design; however, the inverse problem is often what we are really interested in. That is: Given experimental measurements, what are the physical properties of the measured system?

It's rare that we're able to directly calculate a complete set of system properties from experiment measurements. Typically we have access only to simulations of the type we've considered in earlier chapters – forward simulations, such as our Bohmian trajectories approach, which take system properties as inputs. Equipped with only these tools, the naive solution to the inverse problem is to run a broad series of forward simulations, in the hope of finding one whose outputs match up with experiment. In fact, this solution is by far the most common in the practice, and for many situations it provides an adequate insight into the system under consideration. However, blind application of this kind of simple parameter search can miss nuances in parameter distributions and, worse, can lead to completely incorrect conclusions.

The lack of a formal treatment of the inverse problem is somewhat understandable in many cases, as a complete solution can be extremely computationally demanding; for many complex models, with time-consuming forward simulations, an exact treatment is not feasible. Considerable care must be taken, though, to avoid assigning too much value to a simple search solution. Kasim et. al. highlight the problem with the naive approach, particularly in the domain of plasma physics[86]. A given set of experimental measurements and error bars can 'match'

with a huge range of physical system parameters. The naive search approach only selects one of these systems. It cannot, therefore, offer robust conclusions.

Even when hedged with error propagation methods, the conclusions drawn from simply matching up results are often incomplete – standard error propagation methods assume that all feasible physical parameter sets are neatly clustered around their best fit values; this is plainly not the case in general, as the sample distributions found in ref. [86] demonstrate. Determining the bounds of the physical parameter space that fit an experimental data set is an important and challenging problem: It should be treated with the same attention as the forward problem that underlies it.

This chapter develops and applies a range of Markov Chain Monte Carlo approaches to the inverse problem for plasma physics. We begin with a brief summary of Markov Chain Monte Carlo, and go on to develop some general extensions to standard methods that will prove useful for our application domain. The chapter rounds out with multiple example cases.

4.2 Markov Chain Monte Carlo

Many of the Monte Carlo methods we've discussed in previous chapters (i.e. Classical Monte Carlo and its variants) form a subset of a broader family of techniques: Markov Chain Monte Carlo (MCMC) methods[87]. MCMC is essentially an approach to sampling probability distributions – particularly probability distributions for which it is possible to calculate a likelihood value for a given sample (i.e. a value proportional to, but not necessarily equal to, the probability of that sample).

To clarify this, we return to the Metropolis Hastings algorithm discussed in section 2.3.2. This algorithm is used to sample the canonical distribution of a system, in order to calculate its integrated properties. Let W be a vector containing the coordinates of the system – in the Classical Monte Carlo case, we can initially take W as the collection of the real-space coordinates of all of the system's particles, $W = x$. We refer to a single set of values of W_i (indexed by i) as a configuration.

Roughly speaking, MCMC methods construct a sequence of configurations such that, under appropriate asymptotic conditions, the distribution of the sampled configurations converges to a desired probability distribution. To decide whether to move from one configuration to another, we calculate the likelihood function $\exp(-\beta E)$ for each, where $E = E(W)$ is the energy of the configuration corresponding to each sample point. Clearly, this likelihood is proportional to the probability of the configuration, $\exp(-\beta E)/Z$, where Z is the system partition function. The advantage of the Metropolis Hastings algorithm, and other classes of MCMC algorithms, is that the probability value itself does not need to be known – in this case, for example, we don't need to calculate the full partition function in order to sample the canonical distribution.

However, MCMC approaches are far more general than the direct physical application we considered above. Experimental data points have error bars, and via these error bars we can assign likelihoods to simulations themselves (and their corresponding physical parameters) based on their predictions. Whereas we have already seen that MCMC can calculate physical observables from system parameters, we will now discuss methods for the inverse – calculating distributions over system parameters from physical observables. As a concrete example, we might want to take a DSF measurement (and its error bars) from an experiment and use it to deduce an implied temperature probability distribution. The methods developed here will allow experimental data to be rigorously analysed, in combination with the forward simulation methods discussed and developed in chapters 2 and 3.

4.2.1 Basic Methods

As mentioned in the preceding section, the Metropolis Hastings algorithm is a particularly simple MCMC method. We need only to cast the inverse problem as a physical problem in order to apply Metropolis Hastings precisely as we have done in the context of Classical Monte Carlo.

This is a simple procedure – we replace occurrences of $-\beta E$ in the algorithm described in section 2.3.2 with the more general expression $\log \pi(W)$, where π is the likelihood of the physical parameters W . For example, we can consider a case where we have a single experimental observation, X_{exp} , with Gaussian error bars of width Δ . Our forward simulation can be treated as a function mapping from the domain of possible physical parameters, $\{W\}$, onto the range of a predicted

observation X_{pred} . Then we can write:

$$\log \pi(W) = -\frac{\log 2\pi}{2} - \frac{\log \Delta^2}{2} - \frac{(X_{exp} - X_{pred}(W))^2}{2\Delta^2}. \quad (4.1)$$

As with our application of Metropolis Hastings to particulate systems, we construct a chain of parameter values by proposing moves (i.e. changes in physical parameters), and accepting or rejecting these moves according to the ratio of log-likelihoods of the proposed and current parameters.

Beyond the simple Metropolis Hastings method, there are several variants of MCMC that provide faster convergence and require less hand-tuning - DRAM[88] (an adaptive sampler with delayed rejection) and EMCEE[89] (an affine-invariant ensemble sampler), for example. However, in order to tackle problems over large parameter spaces with expensive forward simulations, we will need to consider MCMC methods that exploit not only the log-likelihood, but also its gradient with respect to W .

4.2.2 Gradient-Driven Approaches

In principle, the Metropolis Hastings method (or a variant along the line of EMCEE or DRAM) is all we need to robustly map experimental error bars onto uncertainties in physical parameters. In practice, however, we are often limited by computational costs. For each evaluation of the log-likelihood, we have to run a (potentially very expensive) forward simulation; it is this issue that, in part, motivated our development of a fast forward simulation method in chapter 3. Clearly, care must be taken to minimise the number of simulations required to effectively sample our physical parameters.

We can achieve far more effective sampling if we are able to exploit the geometry of our parameter space. Suppose, in particular, that we are able to determine derivatives of the log-likelihood with respect to the parameters we are sampling. Heuristically, we are then able to make more informed parameter proposals, and fully explore the space of valid parameters with fewer simulations/log-likelihood evaluations.

Initially, we consider a direct extension to Metropolis Hastings that incorporates the gradient of the log-likelihood. Under the assumption that we are able to calculate this gradient, we can calculate the analogue of physical forces in parameter space (just as the log-likelihood is analogous to the energy). We can then draw on the physical MD methods highlighted in chapters 2, 3 to write equations of motion for W with the desired equilibrium distribution. As a simple first step, we can apply an overdamped form of the dynamical equations implied by eq. 2.38 to parameter space dynamics. This corresponds to the equation of motion for W :

$$dW = \frac{1}{2} \nabla \log \pi(W) dt + dB, \quad (4.2)$$

where π is, as before, the likelihood distribution over the parameter space, and B is a standard Brownian motion/Wiener process. Following this equation of motion, we know – from the properties of Langevin dynamics – that the path of W through time samples the effective canonical distribution over parameters implied by the experimental data. A first order (Euler-Maruyama) discretised update for proposals, for the sequence of values W_k , can be constructed as:

$$\tilde{W}_{k+1} = W_k + \tau \nabla \log \pi(W_k) + \sqrt{2\tau} \zeta_k, \quad (4.3)$$

where τ is a time step, and ζ_k an independently draw from a normal distribution.

4 Inferring Static Structure: Inverse Problem Approaches for Plasmas

As with classical MD, we must consider the implications of using a finite time step; direct application of this update scheme will only *exactly* sample our desired posterior distribution as τ tends to zero. However, we must simultaneously contend with the fact that allowing the time step to become very small will result in a slower sampling scheme. This is where a connection to Metropolis Hastings can be exploited: by employing an accept/reject step after each step of our discretised update scheme, we can ensure that we preserve our desired equilibrium distribution with an arbitrary time step length. After each time step – in which we begin at parameter positions W_k – we accept/reject the newly proposed parameter positions \tilde{W}_{k+1} by first calculating α :

$$\alpha = \min \left\{ 1, \frac{\pi(\tilde{W}_{k+1})q(W_k|\tilde{W}_{k+1})}{\pi(W_k)q(\tilde{W}_{k+1}|W_k)} \right\}, \quad (4.4)$$

where

$$q(w_2|w_1) \propto \exp \left(-\frac{1}{4\tau} |w_2 - w_1 - \tau \log \pi(w_1)|^2 \right). \quad (4.5)$$

We compare this value to a sample from the uniform distribution, $\gamma \sim U(0, 1)$. If $\alpha > \gamma$, the stochastic update is performed; otherwise, the parameter set is not updated over the time step. Our path through parameter space then properly samples the posterior distribution over physical parameters. The method brings together Classical Monte Carlo and Molecular Dynamics for particulate systems in a more general statistical inference context.

This approach, a form of Metropolis Adjusted Langevin Algorithm (MALA), is substantially more efficient than black-box Markov Chain Monte Carlo approaches (see refs. [90, 87] for detailed analysis and numerical properties), particularly when W exists in a high-dimensional space (as we will see in section 4.4). Gen-

eralisation to other driving equations of motion with the correct equilibrium distribution is straightforward, although we do not pursue it further here.

4.2.3 Automatic Differentiation

The advantages of gradient-driven MCMC are significant; however, we must consider whether it is feasible – and computationally efficient – to calculate the required gradients. In most of the specific situations we will consider, the calculations of gradients will be direct extensions of the associated forward simulation methods, derivable through analysis of the underlying physics – we’ll treat these cases as we discuss them. But it is worth mentioning that, in principle, the gradients we require are *always* accessible without an increase in time complexity of the underlying simulation.

A simulation – regardless of its underlying steps and complexity – can be considered as a parametrised mapping from a set of inputs X to a set of outputs y . The parametrisation is accomplished through a set of weights, W , which dictate the fixed inputs that characterise the simulation – for an MD simulation, for example, these weights could specify interparticle potentials, temperature, density, and so on. The effect of the simulation can be written $y = f(X; W)$. The dependence of the function f on its parameters is set by the simulation architecture, i.e., by the series of mappings from inputs values to outputs.

By taking this viewpoint on simulations as general functions, Automatic Differentiation allows explicit access to gradients. The key observation exploited is that any computer program, regardless of its complexity, can be reduced to a set of elementary, differentiable operations. As a result, it is possible to determine the variation of simulation output y with respect to our parameters W .

Fig. 4.1 illustrates this concept, viewing a simulation as a sequence of functions applied sequentially to an input state, with each function in the sequence a function of W (see caption). Automatic Differentiation works by a simple application of the chain rule to this sequence of functions – this is shown in fig. 4.2. By propagating gradients through each step of the simulation, it becomes possible to obtain $\partial y/\partial W$ alongside y .

AD is a very general computational tool, and has seen a range of applications throughout the sciences in numerical simulations – examples of physical applications include (but are far from limited to) force calculation in Quantum Monte Carlo[91], functional calculation in Density Functional Theory[92], Hartree-Fock calculations[93], and uncertainty propagation in Atmospheric Physics and fluid dynamics[94]. While we will focus predominantly on domain-specific gradient calculations for quantum plasma systems, AD will prove a useful tool for analysis of plasma diagnostics in section 4.4.

4.2.4 Stochastic Gradient Methods

In the case of physical simulations, we are often in a position where calculations have an associated statistical noise. In particular, physical correlation properties are usually expressed as expectation values for the system, so that finite numbers of simulation outputs cannot allow exact correlation functions to be calculated. We could consider, for example, the positional correlation between two particles in a system, as a function of time (or perhaps more usefully, the SSF and DSF). A calculation of a form of particle correlation could be performed by averaging over many realisations of the system, but it may only be possible to achieve an exact result when the number of realisations used tends to infinity. Application of MCMC techniques is more subtle in this case, as we must ensure that the applied algo-




Figure 4.1: Schematic of a generic simulation interpreted as a differentiable function. A set of input variables $x = x_0$ are generated according to simulation requirements/known input distributions. The core simulation network then maps these inputs onto measurements via a series of steps parametrised by network weights W . These measurements are then combined into a metric that quantifies the network error, $y = f(x; W)$, where $f = f_1 \circ f_2 \circ \dots \circ f_n$.




Figure 4.2: Schematic of Automatic Differentiation, as applied to the simulation of fig. 4.1. By propagating gradients backwards sequentially via the chain rule, gradients of y with respect to W can be extracted. At each layer of the graph we calculate $y_i = \partial y / \partial x_i | W$ recursively, using the fact that $y_{i-1} = y_i f_{i,i-1}(x_{i-1}; W)$. The functions $f_{i,i-1}(x_{i-1}; W) = \partial x_i / \partial x_{i-1} | W$ are straightforwardly determined through differentiation of the elementary forward functions. The derivative with respect to W is then calculating by accumulating the values of y_i , via $\partial y / \partial W | x_0 = \sum_n y_n \partial x_n / \partial W | x_{n-1}$.

rithm converges to the correct posterior distribution when noisy measurements are provided as input.

Stochastic Gradient Langevin Dynamics acts as an extension to MALA in this sense, updating parameters using an estimator of the log-likelihood gradient (as opposed to its exact value). This scheme eliminates the accept/reject stage of MALA, and provides the correct posterior distribution using physical inputs that are random in the statistical sense. Updates to parameters are carried out with a modified equation of motion:

$$dW = \frac{1}{2} [\nabla \log \pi(W)]_{est.} dt + dB, \quad (4.6)$$

where $[\nabla \log \pi(W)]_{est.}$ is an unbiased estimator of the gradient $\nabla \log \pi(W)$ appearing in the MALA equation of motion. This allows us to take samples when it is not possible to determine $\nabla \log \pi(W)$ exactly – this is true in the case of many physical simulations, where only a statistical estimate of a physical property can be constructed.

While SGLD has been applied in a variety of statistical inference applications, it is overwhelmingly used alongside a batch mean estimator for the log-likelihood gradient. That is to say, it is almost exclusively used in the context in which the exact log-likelihood gradient is a simple sum over a finite number N terms, through $\nabla \log \pi(W) = \sum_i^N F_i$, and an unbiased estimator can be trivially constructed as $\sum_{\{i_{sub}\}} F_i$ for some subset of indices $\{i_{sub}\}$. Our (slightly) more general use of the algorithm in this work is motivated by ref. [95] – the analysis therein confirms the favourable properties of SGLD with only the (more general/less restrictive) assumption that an unbiased estimator of the gradient is obtainable.

By sequentially reducing the step size over the course of sampling, it can be shown that the SGLD chain's distribution converges to the posterior distribution over parameters as required. The accept/reject step is avoided by ensuring the time

4 Inferring Static Structure: Inverse Problem Approaches for Plasmas

steps used at each step of the algorithm, τ_k , satisfy:

$$\sum_{k=1}^{\infty} \tau_k = \infty, \quad \sum_{k=1}^{\infty} \tau_k^2 < \infty. \quad (4.7)$$

SGLD is just one of many possible Stochastic Gradient Markov Chain Monte Carlo (SGMCMC) methods (see ref. [96] in particular, which highlights the general relationships between these methods); in the following section, we consider extensions to SGLD, and SGMCMC more generally, for application to quantum plasma data.

4.3 Physical Parameter Inversion via SGMCMC

While we've discussed the overarching principles of SGMCMC, there are many ancillary methods we will rely on in order to effect physical parameter inversion over practical time frames. This section brings together concepts from the literature and general-purpose, new developments that we will exploit in SGMCMC applications.

4.3.1 Unbiased Estimation Methods

In order to apply SGMCMC methods, we'll have to consider the domain-specific problem of extracting gradients of log-likelihoods with respect to parameters of interest. In doing so, we must take considerable care when exact gradients cannot be extracted – we must ensure that any gradient estimates we construct from noisy simulations are unbiased (to ensure our SGMCMC method converges to the true posterior distribution).

Cumulant-Based Methods for Polynomials of Expectations

In many of the scenarios we will consider, the make-up of the log-likelihoods and gradients we want to calculate will make unbiased estimation fairly straightforward. Specifically, we will see several cases in which we need to form unbiased estimators of polynomials of expectations emerging from underlying Classical Monte Carlo simulations. These expectations will have the form:

$$E_W[f^i(x)] = \frac{\int \exp(-\beta V(x; W)) f^i(x) dx}{\int \exp(-\beta V(x; W)) dx} \quad (4.8)$$

4 Inferring Static Structure: Inverse Problem Approaches for Plasmas

for various functions f^i of the configuration coordinates x , at a given set of system properties W (with inverse temperature β and potential V). Assuming there is no bias associated with the start point of our Classical Monte Simulation (more on this later), $E_W[f^i(x)]$ is straightforward to estimate without bias: we draw a set of N configuration samples $\{x_j\}$, and have estimate the expectation as $\sum_j f^i(x_j)/N$.

The problem we face is how to estimate, without bias, polynomials of expectations of this type, along the lines of:

$$\prod_i E_W[f^i(x)]^{n_i}. \quad (4.9)$$

When we deal with Gaussian error bars, in particular, we will typically see that all of the terms we need to estimate will be of this form. In such cases, we can apply a straightforward recursive procedure to arrive at an unbiased sampling strategy: We can rewrite our desired expectations in terms of multivariate cumulants, and use standard results for unbiased estimators of cumulants (for example, those derived in ref. [97]). While this does not necessarily simplify individual calculations, the general procedure can be carried out mechanically in combination with standard statistical software, removing the need for manual derivation of contributions from individual terms.

As an example - in fact, the simplest non-trivial example - we may wish to determine an unbiased estimator of $Z = E_W[f^1(x)]E_W[f^2(x)]$ for functions of the coordinates f^1, f^2 . In this case we see immediately from the form of the second multivariate cumulant,

$$C_2(f^1(x), f^2(x)) = E_W[f^1(x)f^2(x)] - E_W[f^1(x)]E_W[f^2(x)], \quad (4.10)$$

that $Z = E_W[f^2(x)f^2(x)] - C_2(f^1(x), f^2(x))$. For N independent samples of x , enumerated by j , we can then apply a standard estimator for C_2 and the single expectation term to arrive at an unbiased estimator of Z , \bar{Z} :

$$\bar{Z} = \frac{1}{N^2} (m \langle f^1(x)f^2(x) \rangle + N(N-1) \langle f^1(x) \rangle \langle f^2(x) \rangle), \quad (4.11)$$

where $\langle g(x) \rangle = \sum_{j=1}^N g(x_j)/N$. More complicated expressions of the form of eq. 4.9 can be approached in an analogous way, in general relying on unbiased estimators of various cumulants C_n of (polynomials of) the functions f^i .

Russian Roulette Estimators and General Approximate Methods

Beyond simple polynomials, we can take a more general approach to forming unbiased estimators – that of Russian Roulette. Suppose we want an unbiased estimator of a quantity Y , which cannot be calculated in finite time; further, suppose that we can establish an infinite sequence $Y_n, n = 0, 1, 2, 3, \dots$ such that $\lim_{n \rightarrow \infty} Y_n = Y$. We define Δ_n via $\Delta_n = Y_n - Y_{n-1}$, where we take $Y_{-1} = 0$. We then consider an additional finite-valued, non-negative, integer-valued random variable N – independent of Y_n – with $P(N \geq n) > 0$ for all $n \geq 0$. Defining \bar{Y} as:

$$\bar{Y} = \sum_{l=0}^N \Delta_l / P(N \geq l), \quad (4.12)$$

\bar{Y} is an unbiased estimator of Y as required (see ref. [98] for derivation/proof). This technique allows us to form unbiased estimators from (truncated) sequences of biased estimators (see e.g. the application in ref. [99]). While this Russian Roulette approach has the advantage of generality, its utility is predicated on the assumption that the random variable defined by equation 4.12 has finite variance. A priori, it may not be possible to assess whether this assumption is valid – we re-

quire knowledge of how quickly the variances and biases of each term in our series decay with n . We can adjust these decay rates by varying the number of samples averaged over at each value of n (for example, we may require that the number of samples taken grows exponentially with n) – but there is trade-off here with sampling speed.

We can assure finite variance by truncating the series at some finite value of n ; however, we must then make sure that the bias of the final term in our series is small enough to be neglected. In short, the Russian Roulette approach is not a silver bullet – we must ensure fast decay of biases, or alternatively allow the series itself to have a finite bias in exchange for a finite variance.

In the case of Classical Monte Carlo sampling, we can rely on a simpler (although perhaps less robust) approach. Bias associated with our choice of initial configuration can first be made negligible (although not completely eliminated) by using a burn-in period – we allow the Classical Monte Carlo algorithm to run for a number of steps, but discard the associated samples. For a sufficiently large burn-in period, we can then begin sampling from a configuration that is approximately independent of our initialisation scheme. Following a burn-in period, we can separately debias expectation polynomials via multivariate cumulants as above.

4.3.2 Combining Multiple Simulation Fidelities

During the process of sampling, we may find that our dynamics in parameter space (dictated by *SGLD* or similar) retreads old ground. At the very least, we would expect our Markov chain to eventually reapproach high-likelihood regions that have been previously sampled. It is natural, then, to consider the possibility of exploiting the gradient information we have already gathered, rather than performing new expensive gradient calculations.

Further, we may have access to multiple different ways (of varying accuracy) of calculating gradients. We could ask, in addition – do we always need to use our most accurate and expensive simulation method? Or can we partially lean on faster less accurate methods along our SGMCMC chain?

These two questions relate to multifidelity sampling: combining approximations of different speeds and accuracies to more effectively sample parameter space (modern multifidelity techniques are summarised in ref. [100] and references therein). We can approximate local gradients using information from previous samples – interpolating between points or using more sophisticated regression schemes – or we can use different techniques altogether – for example, we might calculate system properties with the Hypernetted Chain Equation, as opposed to using Molecular Dynamics or Monte Carlo methods. If our faster and rougher estimates of gradients are reasonably accurate, we can reduce our reliance on heavy duty ‘exact’ calculations. Through this observation, the methods developed in this section provide a pathway to reducing effective sampling times, which we can exploit in our future error analysis.

Finite Sequence Russian Roulette with SGLD

The simplest scenario we can consider is one in which we have, a priori, multiple methods for calculating a physical quantity (and, therefore, multiple methods for calculating the log-likelihood). If we are targeting experimental measurements of the static properties of a particulate system – the radial distribution function, for example – we have HNC, MD, CMC, etc.. We can construct a sampling algorithm that utilises all of these approaches in combination.

We define the total number of simulators we wish to use as $N_{sim} + 1$. We now order our simulators in approximate increasing order of accuracy and simulation

4 Inferring Static Structure: Inverse Problem Approaches for Plasmas

time, enumerating them with i , with $0 \leq i \leq N_{sim}$. The simulator with $i = N_{sim}$ is then our ‘best’ simulator: to accomplish exact parameter sampling, we require that this simulator gives an unbiased estimate of the gradient of the log-likelihood. For a given W , we write the outputs of our simulations as:

$$X_i(W) = \bar{X}(W) + b_i(W), \quad (4.13)$$

where $\bar{X}(W)$ is the exact value of the gradient of the log-likelihood, and $b_i(W)$ is a random variable containing all the error and noise associated with simulation i . By assumption, $E[b_{N_{sim}}(W)] = 0 \forall W$.

We want to construct an unbiased estimator of $\nabla \log \pi$ by running some combination of our simulations. Clearly, just running our best simulation yields such an unbiased estimator; however, we work under the assumption that the best simulator takes a long time to run relative to our other less accurate options. We want to avoid running more expensive simulations where possible, striking a balance between estimator variance and simulation time.

To do so, we consider a finite-series Russian Roulette approach, as per the previous section. First, we define $\Delta X_i = X_i - X_{i-1}$, and choose a sequence of decreasing, positive probabilities, Q_i , for $1 \leq i \leq N_{sim}$. We then construct Y via:

$$Y = X_0 + \sum_{i=1}^N \frac{\Delta X_i}{Q_i}, \quad (4.14)$$

where N is a random variable, independent of our X_i , with probability distribution established by requiring $P(N \geq n) = Q_n$. Then we have:

$$E[Y] = E[X_0] + \sum_{n=1}^{N_{sum}} E[X_n - X_{n-1}] = \bar{X}, \quad (4.15)$$

such that Y is an unbiased estimator of the exact gradient of the log-likelihood. Let us assume our i^{th} simulator requires a time t_i to run. Then the expected total calculation time for Y , τ , can be expressed as:

$$\tau = E \left[t_0 + \sum_{i=1}^N t_i \right] = E[t_0] + \sum_{i=1}^{N_{sum}} Q_i E[t_i]. \quad (4.16)$$

Provided the values of Q_i are chosen such that $\tau < E[t_{N_{sum}}]$, our estimator is, on average, faster to calculate than an output from our best simulator. The drawback of calculating log-likelihood gradients in this way is the corresponding increase in variance of the estimator (relative to direct use of our best simulator). The tradeoff between variance and speed can be adjusted through judicious choice of the probabilities Q_i ; for applications in which cheaper simulation methods do not suffer from large accuracy penalties, a substantial speed-up can be obtained with little sacrifice.

In our case, our primary interest in the Russian Roulette estimator will be as an input for SGMCMC. More specifically, we focus on SGLD, for which we can assess the variance-speed tradeoff semi-quantitatively. In very broad terms, SGLD transitions from optimisation to posterior sampling when the magnitude of the injected stochastic noise begins to exceed the magnitude of the noise in the gradient term[101]. Writing the variance of our estimator Y at W as $V(W)$, this condition can be expressed as:

$$\epsilon_{max} \gg \left(\frac{\epsilon_{max}}{2} \right)^2 V(W), \quad (4.17)$$

or:

$$\frac{1}{V(W)} \gg \epsilon_{max}, \quad (4.18)$$

where ϵ_{max} is the maximum time step value at which we can reliably sample.

4 Inferring Static Structure: Inverse Problem Approaches for Plasmas

In addition, the maximum time step employed during sampling is limited by the geometry of the parameter space (i.e. how smoothly the log-likelihood gradient varies with changes in parameters). This is a restriction analogous to that on the time step in Molecular Dynamics – we must select a time step small enough for smooth dynamics based on the temperature, potentials, etc.. As such we can write:

$$\epsilon_{top} \gg \epsilon_{max}, \quad (4.19)$$

where ϵ_{top} is the maximum time step permissible for smooth dynamics (in general, we would expect ϵ_{top} to be a function of position in parameter space, $\epsilon_{top} = \epsilon_{top}(W)$). Now, the number of effectively independent samples drawn by SGLD over a fixed number of time steps is (approximately) proportional to the current magnitude of the time step. As such, to maximise the number of independent samples per simulation time, we want to maximise:

$$\frac{\epsilon_{max}}{\tau} \propto \frac{1}{\max[V(W), \epsilon_{top}^{-1}]\tau} \quad (4.20)$$

We can then use this condition to set our values of Q_i . One simple approach to this is to construct an estimate of the variance in terms of previously calculated quantities. For example, it can be shown that \hat{V} , where

$$\hat{V} = \sum_{n=1}^N \frac{(\Delta X_n)^2}{Q_n^2} (1 - Q_n) + 2 \sum_{j=1}^N \sum_{n=j+1}^N \frac{\Delta X_n \cdot \Delta X_j}{Q_j^2 Q_n} (1 - Q_j), \quad (4.21)$$

is an unbiased estimator of V [98] when our final simulator is nearly exact (and this expression from ref. [98] can be extended to more general cases). We can then estimate the variance at any value of W for which we have already calculated Y , using the terms calculated in order to construct Y . In turn, these variance es-

imates can be used to tune our values of Q_i going forward, in order to ensure our SGLD sampler is as efficient as possible.

We require an estimate of $V(W)$ before running our simulations at a given W , however, so that we can calculate our values of Q_i , which in turn allow us to sample N (and thus determine how many simulations we need to run). We have several options for constructing such an estimate, which we will label $\tilde{V}(W)$. In particular, we can use a range of regression methods (linear, K-nearest neighbours, etc.) using variances calculated at previous time steps. The simplest approximation is just to use the estimated variance at the previous time step, via $\tilde{V}(W_n) = \hat{V}(W_{n-1})$. Note that a sampling step is in-principle exact regardless of how good this estimate is, but a good choice will allow SGLD to converge more quickly.

Our full algorithm is then just SGLD with an optimally chosen gradient estimator:

1. Set $i = 0, Q_n = 1, 1 \leq i \leq N_{sim}$, initial parameters W_0
2. Sample a value of N , given the set of probabilities Q_n
3. Run the simulators with labels from 0 to N , obtaining their log-likelihood gradient estimates X_0, \dots, X_N
4. Construct the unbiased estimator Y using eq. 4.14
5. Set $W_{i+1} = W_i + \frac{\epsilon_i}{2} Y(W_i) + \eta_i$, where $\eta_i \sim N(0, \epsilon_i I)$
6. Via the (nearly) unbiased estimator $\operatorname{argmin}_{\{Q_n\}, \max[\tilde{V}(W_{i+1}), \epsilon_{top}^{-1}]} \tau$, Set the values of $\{Q_n\}$ subject to the constraint that Q_n are positive and decreasing with n
7. Set $i = i + 1$
8. if $i > i_{max}$, end; else, go to 2

The minimisation step setting Q_i should be very fast, as it is typically optimisation of a simple polynomial in a small number of variables. For the choice of $\tilde{V}(W_n) = \hat{V}(W_{n-1})$, the general solution can likely be found symbolically ahead of time.

Clearly this approach to multifidelity sampling is quite general, reaching far further than our prototypical radial distribution function example. Simulations that rely on finitely-spaced grids, in particular, are natural candidates for this kind of sampling. For example, we could consider a fluid dynamics simulation with a range of difference resolutions/grid spacings. Using this approach, we could rely on cheap coarse-grained simulations as far as is possible, turning to expensive high-resolution simulations comparatively rarely. As the level of resolution used is effectively a continuous variable for this type of simulation, the scope for multifidelity sampling is expanded: Beyond just selecting probabilities Q_i that minimise variance, we can attempt to optimally select which grid-sizes/resolutions should be incorporated into our Russian Roulette sum. Assuming one can construct a simple analytic model of the relationship between simulation fidelity and accuracy, choosing which resolutions to build into the telescoping sum becomes a fairly straightforward optimisation problem.

Machine Learning Estimators and Emulation

As a final note on multifidelity sampling, we consider in slightly more detail the case in which we can build an additional data-driven simulator, for use in conjunction with existing simulation methods (this general approach is highlighted by ref. [100]'s review and references therein). Such a simulator takes a collection of inputs and outputs from existing simulation methods and constructs an approximate mapping between them; there are a huge quantity of different options for this task that have emerged from Machine Learning research[102].

As we take samples from a system of interest, we can simultaneously provide results to data-driven estimators (online). We then obtain an additional simulation method that can slot into our telescoping Russian Roulette estimator, whose accuracy should, in principle, increase as we take more and more samples. Incorporating such an estimator into the sum is especially useful when we sample values very close in parameter space to those sampled previously – the estimator can provide a very accurate value for a very low simulation cost in this case.

The other great advantage of building a data-driven estimator is its flexibility. In general terms, estimators are constructed according to a set of hyperparameters, which allow a fine-tuned control over model complexity, and thus over estimator bias and variance. An increase in model complexity tends to lead to greater bias and lower variance (and the converse also holds). While a non-trivial task (and beyond the scope of this brief discussion), one can in principle control the bias-variance tradeoff of the data-driven estimator to ensure that, once combined into the Russian Roulette telescoping sum, the resulting combined unbiased estimator has minimal variance.

In many low-dimensional cases, it is even feasible to pretrain a data-driven estimator. For simulations with small numbers of input parameters, one can construct a general-purpose ‘emulator’ of the simulation with a data-driven estimator that is sufficiently accurate across the parameter space. The MCMC process can then be accelerated by several orders of magnitude, relying on emulator outputs rather than full simulation runs for each sample. Ref. [103] demonstrates effective application of this technique to a range of physical simulations using Neural Architecture Search – integration of these emulators of this kind with multifidelity sampling is an obvious next step for the future.

4.3.3 Variance Estimation and Corrections

In addition to constructing more sophisticated estimators of the log-likelihood gradient, we can refine the SGLD algorithm itself with various techniques that reduce sampling variance (and thereby reduce the number of sampling points required to approximate our underlying probability distribution).

First, we can consider the method of control variates – a simple, widely used variance-reduction method. Suppose we wish to calculate an expectation of a function $g(W)$ of our parameters W , $E[g(W)]$. We could instead consider the alternative function $\tilde{g}(W) = g(W) + h(W)$; if $E[h(W)] = 0$ and g and h are negatively correlated, the variance of $\tilde{g}(W)$ can be substantially lower than that of $g(W)$. Baker et. al. discuss multiple implementations of this idea for SGLD in ref. [104]. The SGLD-ZV algorithm detailed in that paper, in particular, can be applied as a simple post-processing step after running a standard SGLD simulation in order to obtain reduced-variance averages. Similar variance reduction strategies are also described in ref. [105].

More generally, though, we can also consider modifications to the equation of motion of SGLD itself. As we briefly discussed in section 4.3.2, SGLD essentially functions as an optimisation algorithm while the variance of our unbiased estimator greatly exceeds the variance of the injected noise at each time step. It is only when the time step is sufficiently small – when the injected noise dwarfs that associated with our estimator – that we effectively sample our desired probability distribution; it is only in this case that our discrete equation of motion well-approximates continuous Langevin dynamics. However, we can slightly modify the SGLD equation of motion to improve this state of affairs, allowing us to proceed to the sampling stage earlier.

To do so, we can adjust the covariance of the injected noise such that the injected noise/estimator noise combination is of the correct form for sampling. Assuming we can form some basic estimate of the variance of our stochastic gradient – this could be estimated, for example, from previous samples at nearby parameter sets – we need only to deduct this from the variance of the applied noise (as shown in e.g. ref. [96]). We then effectively enter the sampling stage of the SGLD algorithm at a larger time step. In essence, we replace an injected noise with covariance C_0 with an injected noise with covariance $C_0 - C_{est}$, where C_{est} is the approximate covariance of our gradient estimate.

In the case of Monte Carlo estimates of log-likelihood gradients, we are often in a position where we can manually control our estimator variance by altering the number of samples averaged over – this will be the case in the applications we consider, and provides us with the opportunity for another significant algorithmic improvement. In combination with SGLD, we can dynamically vary the number of samples used by our Monte Carlo estimator at each time step.

Consider the case where our gradient estimate is the mean over n simple Monte Carlo samples. As we collect these samples, we can track the sample covariance matrix, C , of the gradient online (i.e. we can update an estimate of the covariance matrix on the fly after each sample is taken). As a consequence of the multivariate central limit theorem, the noise associated with our mean over samples is approximately a multivariate Gaussian with covariance $\tilde{C} = C/n$ for large n . Broadening this slightly, we can turn to the Markov Chain Central Limit Theorem when our inner Monte Carlo estimator produces samples with finite covariance – in this case the estimated covariance can be calculated online as well, with an additional lag term to account for correlations between the points traversed by the chain (see

e.g. ref. [87] for discussions of these limit theorems). Again, we can calculate and update \tilde{C} as we draw samples.

We continue to take samples until the additional covariance $C_0 - C_{est}(\tilde{C})$ becomes positive semidefinite; the injected noise entering into the SGLD update can then be constructed with covariance $C_0 - C_{est}$. By approaching gradient estimation in this way, we take the minimum number of Monte Carlo samples required to advance SGLD – in its sampling phase, and at the time step required – thereby minimising the time taken for effective sampling. At each iteration i of the SGLD procedure, we have a time step given by ϵ_i ; to calculate the parameters at the next time step, W_{i+1} , a rough outline of the algorithm can be constructed via:

1. Work with a gradient estimate Y , gradient covariance \tilde{C} , and set our target covariance $C_0 = \epsilon_i I$
2. Draw a few samples from our inner Monte Carlo simulation, set Y as the mean of all values drawn so far, and set \tilde{C} according to the Markov Chain Central Limit Theorem
3. Set $C_{est} = \epsilon_i^2 \tilde{C} / 4$
4. If $C_0 - C_{est}$ is positive semidefinite, then we have $W_{i+1} = W_i + \frac{\epsilon_i}{2} Y + \eta_i$, where $\eta_i \sim N(0, C_0 - C_{est})$; else, go to 2

4.3.4 Early Estimation and Conditioning

From a physical point of view, SGLD tracks the motion of a stochastic particle in parameter space – the path of this particle is then used to infer properties of our desired probability distribution over parameters. Even in the early stage of its motion, however, it is often possible to make inferences about the potential landscape in which the particle is moving. These inferences allow us to obtain early estimates

of the full probability distribution; in turn, these early estimates allow us to adjust our equation of motion for the particle to ensure the parameter space is explored efficiently.

As a concrete example, we can assume that our probability distribution is approximately a multivariate Gaussian (as it almost certainly will be, approximately, in the immediate region about the mode) – in physical terms, this corresponds to a generally non-isotropic harmonic oscillator potential about the mode. We can then use early measurements to estimate this Gaussian’s parameters. Note that we’ll just be using this approximation to set some initial points of our Markov Chains – even if the approximation turns out to be poor, our chains will still converge to the correct distribution (albeit with suboptimal efficiency).

The (maximum likelihood) parameter estimation procedure just requires a simple multivariate linear regression, fitting the constant, positive-definite matrix A and vector B in $Y = -A(W - B)$, where Y is the log-likelihood gradient and W is the position in parameter space. This can be done based on the (noisy) samples of the gradient of the log-likelihood we acquire during the early stages of the sampling process. We then arrive at an early estimate of our desired probability density, given by

$$P(W) = \frac{\exp\left(-\frac{1}{2}(W - B)^T A(W - B)\right)}{\sqrt{(2\pi)^k |A^{-1}|}}, \quad (4.22)$$

in which k is the dimension of W . Using this information, we can then attempt to accelerate sampling. To do so, we consider a linear-mapped parameter space, $\phi = \phi(W)$, in which our approximate probability density is a Gaussian with covariance matrix equal to 1; this space can then be explored with optimal efficiency by the basic SGLD algorithm (see ref. [101]). The mapping in question is straightforward to find from our values of A, B , via standard linear algebraic techniques –

we require the square root of A and its inverse to apply (and invert) the mapping $\phi = A^{\frac{1}{2}}(W - B)$. In cases in which we are able to calculate the Hessian matrix at the modal point of our probability distribution, we immediately have values for A and B without requiring our multivariate fit. Either way, we have a fast approximation to our desired probability distribution, and an accelerated procedure for exactly sampling from it.

Working in ϕ space is straightforward – we just use our usual SGLD dynamic update equation for ϕ at each time step, convert to W space for our actual log-likelihood gradient calculation, and then scale the log-likelihood gradient to obtain its value in ϕ space for the next update. This can be viewed as a fairly straightforward form of (global) preconditioning, a general tool applied to a variety of SGMCMC to accelerate convergence[96]. Alongside improved dynamics, we can set initial positions at the start of our SGLD algorithm in line with our approximate probability density (i.e. initial positions ϕ_0 can be drawn from a multivariate normal distribution with mean 0 and covariance 1). The net result is an SGMCMC procedure that handles awkward covariances with much greater efficiency.

4.3.5 A (Very) Brief Note on Priors

Throughout the applications consider from this point onward, we will rely on uniform (objective) prior distributions over physical parameters of interest. That is to say, we will assume – before we take into account our experimental measurements – that all values of the parameters we wish to analyse are equally likely (within some arbitrary, very large range).

This is not necessarily as benign a statement as it appears on the surface. In the potential parameter case in section 4.5, for example, we will assume a form for a pair interaction potential and then use experimental data from a system to

sample its parameters. In doing so, we will be asserting that all values of the potential parameters are equally likely a priori – this is actually a strong assumption about the form of the potential. If we reparametrised the potential, for example, equivalent uniform priors would lead to very different error bars on our final posterior distribution over the complete potential. A more objective observer – with fewer assumptions about the underlying potential – might instead use a prior that is invariant under reparametrisation, such as that in ref. [106].

In short, it is important to keep in mind that the parameter distributions we will sample going forward are based on assumptions which will not hold universally, but only for our particular cases – information from other diagnostics, and physical limitations on measurements, can (and should) be taken into account via non-uniform priors in general.

4.4 Application: Proton Radiography

4.4.1 Problem Overview

We now apply gradient-based MCMC methods to plasma experimental design and analysis. To start, we'll consider the use of proton radiography for inference of plasma properties, via analysis of data emerging from a prototypical experimental diagnostic.

Proton radiography is an experimental technique for extracting properties of electric and magnetic fields. A proton beam is projected through the fields to be measured and onto output films. The spatial density profile of proton impacts on the films can then be used to infer the fields' properties [107] (fig. 4.3). Specific experimental implementations vary, but in the majority of cases the analysis process is much the same. We give a brief overview here of this analysis framework, based on the more comprehensive treatment in refs. [107, 108].

The images produced on the output films can be determined based on the deflection of the proton beam at each spatial position in the image plane. For a proton with position in the object plane (x_0, y_0) , we can write linearised maps to the image plane position (x, y) via:

$$x = x_0 + \frac{x_0}{l}L + \alpha_x L, \quad (4.23)$$

$$y = y_0 + \frac{y_0}{l}L + \alpha_y L, \quad (4.24)$$

where we have neglected higher order terms in the small quantities a/l , and α_x, α_y are the angular deflections of the beam in the x and y directions respectively (we note - as in ref. [107] - that the lateral beam displacement can be neglected, it being of order $\alpha a \ll \alpha L$). The corresponding α functions dictate how the beam is distorted, and can be determined from the first-principles physics. In the cases

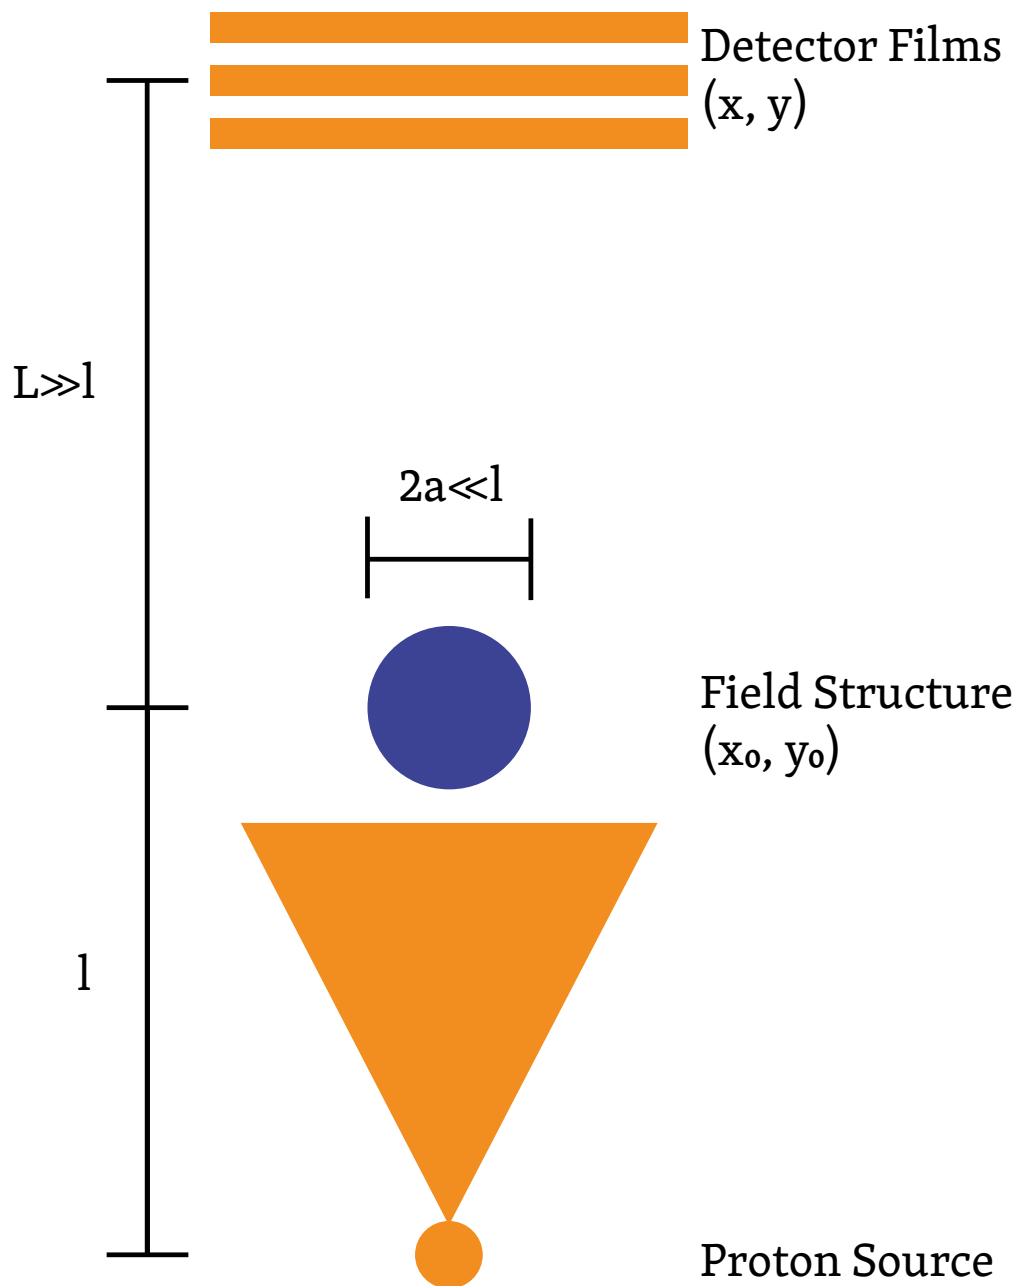


Figure 4.3: Typical proton radiography diagnostic setup. We assume a point source of protons propagating through field structures in the object plane, (x_0, y_0) , to films in the image plane (x, y) . Assumptions regarding the relative length scales are highlighted in the figure, although in principle these restrictions (as well as the point source approximation) can be relaxed with slightly more complicated numerical models.

4 Inferring Static Structure: Inverse Problem Approaches for Plasmas

we'll treat here, we will assume magnetic field effects are dominant (this can be shown to be the case for high proton energies[108]).

With appropriate gauge choice, it is straightforward to show from the Lorentz force that the deflection angles can be written as:

$$\alpha_x = \frac{e}{c\sqrt{2m_p E_p}} \frac{\partial}{\partial x_0} \int_{-\infty}^{\infty} A_z(x_0, y_0, z_0) dz_0, \quad (4.25)$$

$$\alpha_y = \frac{e}{c\sqrt{2m_p E_p}} \frac{\partial}{\partial y_0} \int_{-\infty}^{\infty} A_z(x_0, y_0, z_0) dz_0. \quad (4.26)$$

Here E_p is the proton beam energy, e is the electronic charge, c the speed of light, m_p the proton mass and A_z the longitudinal component of the magnetic vector potential. We can determine the resulting image on the screen in terms of these functions. A surface element dS_0 in the object plane can be mapped to the image plane via:

$$dS = |D(x_0, y_0)| dS_0, \quad D(x_0, y_0) = \frac{\partial(x, y)}{\partial(x_0, y_0)}, \quad (4.27)$$

where dS is the corresponding image plane surface element and D is the determinant of the mapping in equations 4.25, 4.26. The image plane intensity can then be written as:

$$I = \frac{I_0(1 + \epsilon)}{\left| \frac{\partial(x, y)}{\partial(x_0, y_0)} \right| + \epsilon} = \frac{I_0(1 + \epsilon)}{|D(x_0, y_0)| + \epsilon}, \quad (4.28)$$

where I_0 is the beam intensity in the absence of any fields and ϵ is a small, positive parameter arising from the finite resolution of the detecting film. ϵ is only important in the so-called caustic regime, where D becomes close to zero[108] – in these cases, it can be shown that there are multiple possible field values that give rise to the same image. We will focus initially on the more useful cases in which there is a unique, reversible mapping from fields to image – in which this parameter is unimportant – before discussing the general case.

The equations 4.23-4.28 together provide a simple, direct mapping from fields to image. The inverse problem of obtaining fields *from* an image is more involved, however. Several methods have been developed for field calculations, with the most prominent – which solves the integral equations implied by equations 4.23-4.28 – having been discussed in detail in ref. [108]. The issue with ref. [108]’s method is that it is limited by some major simplifying assumptions – in particular, it applies only to monoenergetic proton beams, failing to take into account both the input energy spectrum and the energy-dependent response of the output films. Further, it cannot take into account multiple film readings with different energy responses[107, 108] – this information could in principle be exploited to obtain better field estimates, or to constrain error bars. While successful for inferring broad system properties, this existing approach cannot, then, be applied for *exact* experimental inversion.

4.4.2 Field Inversion via Linearised Maps

In the regime free from caustics, we can make substantial progress towards solving the inverse problem before even turning to MCMC: We can discard several simplifying approximations inherent to existing methods merely by casting the simple, forward mapping as a differentiable simulation – via Automatic Differentiation – as discussed in section 4.2.3. Throughout, we will work in dimensionless units, with an artificial target magnetic potential and corresponding image generated as the sum of 100 randomly placed Gaussian distributions (the generated target image is shown in fig. 4.4). We will also assume that zero-field boundary conditions are applicable far from the imaged region. To do this, we embed film images – each 50 pixels by 50 pixels – at the centre of 100 by 100 pixel grids, and constrain the field on the edges of the larger grids to be zero. This is a commonly

used technique for establishing boundary conditions for experimental images, although we note that care should be taken to ensure applicability when considering a specific experiment.

Initially, we consider the optimisation problem of determining field distributions from images – these cases illustrate the strength of AD as applied in this context, and also provide starting points for our ultimate applications of MCMC. To do so, we employ Stochastic Gradient Descent – we apply the equations of motion of SGLD with the injected noise term set to 0. For the case of a single film and monoenergetic beam – the case in which existing methods work well – fig. 4.4 demonstrates that our approach is equally accurate (see caption for specifics). But the true advantage of our approach is shown in fig. 4.5. Here, we’ve constructed an example in which the beam input is across a range of energies – the simulation we’ve built accounts for this by sampling randomly from the corresponding energy distribution and producing images for each energy, before combining these to obtain an average. In these cases, the equations 4.23-4.28 were the core of the simulation to be differentiated (see caption).

The upper part of fig. 4.6 extends this application to the practically important case in which we have information from multiple films, each with different energy response. Our use of AD allows us to trivially incorporate this new information, through additions to the error function for each separate film. This is a case where previous approaches struggle to find a solution that can reflect these multiple sources of information on the fields. Beyond the improved integrated field estimates constructed here, the use of AD also allows for more general tomography problems to be treated with little extra setup – by considering multiple screens at multiple angles and positions (where experimental constraints al-




Figure 4.4: The result of the application of a gradient-based optimiser to single-film, single-energy proton radiography. Upper: the image extracted following gradient-driven optimisation of the potential function. Centre: the target image corresponding to the input artificial potential. Lower: the square errors in image intensity across the film. We see only very small error in the reconstruction of the target image.




Figure 4.5: The error distributions in images produced by gradient-descent for a thermal distribution of protons, as compared to the exact analytic result for a monoenergetic beam. We see that we are able to significantly reduce errors through a first-principles gradient descent approach, in which the beam energy distribution is exactly accounted for. Upper: log-square errors in the image produced by the exact monoenergetic solution. Lower: log-square errors in the image produced by gradient-descent for a thermal distribution of protons. We see reduction in the mean square error of approximately two orders of magnitude. Sampling of input energies was performed with a simple Metropolis Hastings algorithm, with equilibrium distribution set equal to the product of the proton source energy distribution and the energy-dependent response of the film. Sampling the product here removes the need to sample the energy distribution and energy response function separately (thereby saving computation time, as energies for which the film does not respond are not sampled and propagated). We are able to substantially improve upon simple methods due to the generality of AD – the difficulty in constructing a reverse algorithm is reduced to that of constructing forward sampling.




Figure 4.6: Extracted distributions over the mean squared magnetic field of the inverted system, for Gaussian error bars on the image pixels with standard deviation equal to 1% of the mean target image intensity. In each case the Hessian of the log-likelihood was calculated at the modal field point, and global preconditioning applied as per section 4.3.4 (alongside the early estimation step of that section for initialisation of walker positions). Upper: the mean x (left) and y (right) components of the path-integrated magnetic fields for the target image in dimensionless units (d.u.). Centre: the sampled variances of the field, exploiting information from one film. Lower: the sampled variances of the field, exploiting information from three independent films.

low this), a general-purpose gradient descent method could be applied in a near-identical manner to obtain field properties in multiple dimensions.

Beyond allowing access to more accurate solutions that are more faithful to the experimental setup, though, our use of AD enables accelerated MCMC application via MALA and SGLD. The central and lower parts of fig. 4.6 illustrate the uncertainties obtained via application of SGLD in combination with the ancillary methods developed in this chapter (details in caption). While – at least on a theoretical level – we cannot dismiss the possibility of viable unexplored regions of field space, these sampled variances act at worst as lower bounds on field uncertainty.

Where a direct inverse solution is calculable – such as in the monoenergetic and single screen case – it is also possible to apply a direct method for sampling. This direct method samples the posterior by adding noise to the target image – distributed according to the error distribution on the image measurements – by inverting the image to a field, and then storing this field as a sample of the posterior. Ref. [109] discusses the direct method in more detail, for the case in which the input intensity profile is known only in distribution. This approach offers an alternative, reliable way to obtain (uncorrelated) field samples when fast, direct inversion is accessible.

4.4.3 General Field Inversion

The advantages of the application of AD to this problem go beyond even the generalisations highlighted above. Rather than relying on the mapping equations 4.23-4.28, we can allow our simulation to go completely from first principles, propagating the protons using Molecular Dynamics directly. In doing so, any issues due to strong fields causing beam overlap are naturally accounted for, extending the range of experimental conditions we can analyse yet further. While, as discussed

in ref. [108], there is not a unique field solution in these cases, we can in principle employ `SGLD` to find the corresponding solution distributions (the volume of potential solutions remains to be evaluated however, and a bound on this would be helpful to evaluate convergence). This approach really uses AD to its advantage, highlighting a class of problems in which it is possible to move the burden of constructing direct analytic inversion techniques to AD software.

We treat this case only as a proof of concept, due to its much higher computational demands. We focus on a relatively simple system model with coarse-grained fields. The differentiable simulation used is more involved than the simple particle dynamics or linearised deflection cases. Particles are projected through the fields with the usual dynamical `MD` updates seen in standard simulations for a fixed time period, for sufficient time that all protons have left the region of the field. The protons are then projected along their linear paths onto the screen. Following this, kernel density estimation (`KDE`) is used to approximate the density created on the screen by the protons – put simply, the density is approximated by a sum of small-width Gaussian blobs placed at each impact point. Initial conditions are set based on the known properties of the proton source; each iteration samples from the known input density of protons. Fig. 4.7 illustrates this approach and the corresponding results.

As a final aside regarding this analysis, we recall that, in order to move beyond optimisation and into `MCMC` sampling, we must ensure that the error estimate achieved by comparing the kernel density estimate and the true image is unbiased (or at least has sufficiently small bias that it can be readily ignored). As it stands, the error becomes unbiased only as the kernel width tends to zero, and the number of simulated trajectories to infinity. However, we can again turn to a debiasing technique to (partially) resolve this issue, as in the previous section.

4 Inferring Static Structure: Inverse Problem Approaches for Plasmas

By constructing a truncated sequence of kernel density estimate images, with increasing trajectory count and decreasing kernel width, a balance can be reached between estimator variance and bias.




Figure 4.7: A toy example of first principles inversion of a (synthetic) proton radiography film, via explicit modelling of proton dynamics. The bottom two images compare the density produced by our inverted field (left) and the density of the trial synthetic image (right). Protons are initially produced with uniform spatial density (top-left), and evolved through time via MD (from left to right, and top to bottom). The densities shown in the first nine images are snapshots at equally spaced time steps – calculated using Kernel Density Estimation (see text) – demonstrating how the proton paths converge to the final target image.

4.5 Application: Particulate Systems

4.5.1 Problem Overview

For our second core application of gradient-based optimisation and MCMC, we consider Monte Carlo potential inversion methods. While the proton radiography application above provides a useful analysis tool for quantum plasma experiment, this section will allow us to make more concrete contact with our prior predictive work on static and dynamic structure factors.

We will focus on methods that will allow us to determine sets of inter-particle potentials that generate a given set of configurations of a system or specified correlation functions, along with associated error bars. We have met some basic potential inversion methods in chapter 3 – particularly Inverse Monte Carlo, IMC, and its generalisation to systems with functional potentials. This section can be viewed as a further generalisation of IMC to a more rigorous treatment of the inverse problem (complete with implied parameter distributions).

In a similar vein to our fitting of parameters for Bohmian mechanics in chapter 3, potential inversion techniques are often used for coarse-graining. This process involves obtaining a simplified system – with fewer degrees of freedom – from a more complex one, while retaining the interesting aspects of the system dynamics (see e.g. [110] for a comprehensive overview). Typically, this involves finding effective potentials for slow degrees of freedom, so that fast degrees of freedom can be removed from the model. Coarse-graining is usually treated as an optimisation task, with approaches based on relative entropy minimisation offering general, practical tools[111]. By taking a more thorough Bayesian approach, however, we can establish not only most-likely effective potentials, but sample from a

potential distribution consistent with our knowledge of the system (see e.g. refs. [112, 113, 114]). This opens up rigorous error analysis of effective potentials.

We'll develop and apply two general purpose, practical, and robust inversion methods here.

4.5.2 Configuration Inversion

Suppose we have a system of particles in a box with potential equal to $V(x; W)$, where x is a vector containing all individual particle positions and W is a parameter controlling the potential form. For example, W could be a vector of values dictating the radial pair interaction potential between particles at a set of fixed radial separations (i.e. W could be a discretised pair potential function).

For simplicity, we'll focus on a finite box with equal-length sides (of length L), although this restriction can be relaxed relatively easily, and we'll take the dimension of configuration space to be D . Now suppose we have access to a stream of n values of x , denoted $x_{target,i}$, $1 \leq i \leq n$, on a target system of the same specification with unknown effective potentials between particles – that is, an unknown value of, or distribution over, W . This stream could come directly from another simulation method – perhaps a finely-grained system we want to coarse-grain – or could be generated based on experimental data. In the commonplace case in which we want to match specified radial distribution functions between particle species, we can use Reverse Monte Carlo[115] to trivially generate such a stream (under the widely employed assumption that pair interactions dominate the potential). We can apply SGLD to obtain a distribution over W – and hence over the effective potential – that reproduces the target system stream.

4 Inferring Static Structure: Inverse Problem Approaches for Plasmas

In general, the likelihood of W can be written as:

$$\pi(W) = \prod_i \exp(-\beta V(x_{target,i}; W)) / Z(W), \quad (4.29)$$

where $Z(W)$ is the partition function at W , given by:

$$Z(W) = \int \exp(-\beta V(x; W)) dx. \quad (4.30)$$

We can then write the log-likelihood as:

$$\log \pi(W) = - \sum_i^n \beta V(x_{target,i}; W) - n \log Z(W), \quad (4.31)$$

Provided we can obtain an unbiased estimator of this quantity – or, more accurately, an unbiased estimator of its gradient – we can then apply the standard optimisation and sampling procedures outlined earlier in this chapter.

It is worth noting, at this point, that the problem of constructing an unbiased estimator in this case is illustrative of one of the particular strengths of analysis via SGLD: While an estimator of the gradient can be constructed relatively straightforwardly, the log-likelihood itself is far less accessible. The appearance of the partition function in the log-likelihood classifies the problem as ‘doubly-intractable’ (discussed, for example, in [99]): the dependence of Z on W means that – unlike in the standard case we are used to with MCMC – partition function evaluations are not naturally avoided. These problems appear in a variety of contexts in statistical applications (exponential random graph models and autologistic models for example[99]). Various methods have been developed to cope with double-intractability, most notably auxiliary variable methods. These variants of conventional Monte Carlo methods allow contributions from Z appearing in the ac-

cept/reject step stage to cancel out, resulting in a tractable algorithm[87]. In applying SGLD, however, we need only an estimate of the gradient of the log-likelihood, which simplifies matters considerably – we can turn to the standard methods discussed and developed in section 4.3.1.

We start by writing the gradient of the log-likelihood as:

$$\nabla \log \pi(W) = n\beta E_W \left[\frac{\partial V(x; W)}{\partial W} \right] - \sum_i^n \beta \frac{\partial V(x_{target,i}; W)}{\partial W}, \quad (4.32)$$

where $E_W[\cdot]$ denotes that we take the expectation over the canonical distribution in x generated by the potential with parameter W . The derivative of V with respect to W can be quickly found through AD as previously (although the derivative is often analytically calculable, as an alternative). As for the new expectation term, we can start by sampling particle configurations with Classical Monte Carlo – with the potential implied by the current value of W – and supply these configurations to the potential function.

To partially account for the the starting-point bias introduced when we use only a finite number of Monte Carlo steps, we can employ a burn-in period, and can also consider a (finite sequence) Russian Roulette approach. Labelling the configurations emerging from the CMC sampler as $x_k, k = 1, 2, 3, \dots$, we can write

$$E_W^n[g] = \frac{1}{M(n)} \sum_{k=1}^{M(n)} g(x_k; W), \quad (4.33)$$

for an increasing function M . We then have the property:

$$E_W[g] = \lim_{n \rightarrow \infty} E_W^n[g]. \quad (4.34)$$

Setting $g = \partial V / \partial W$ and using the delta averaging approach in eq. 4.12, we can arrive at the (approximately) unbiased estimator required.

This approach allows us to naturally incorporate all the known constraints on the system potentials as priors, and allows input streams from multiple temperatures/densities etc.. Iterative Boltzmann Inversion[116], a simplified potential inversion scheme using inexact gradients, has been shown to arrive at more reliable potentials through additional constraints in a similar manner via multiple input streams[117] – the approach here can be seen as a formalisation of this idea, allowing incorporation of an arbitrary number of known constraints and input configurations, and thereby allowing the posterior distribution of potentials to be sampled. Uncertainties in the data stream – in the form of a limited number of configurations, or variance within those configurations – manifest themselves as uncertainties in the derived potentials. In the case where we have access to an infinite stream of configurations from an exact representation of the system, this method essentially reduces to optimisation.

To demonstrate the practicality of the method developed above, we consider a collection of particles in a three-dimensional box, with model Gaussian inter-particle potentials. We generate configurations from this potential using a CMC sampler, and show that we are able to accurately reconstruct the potential from these configurations (fig. 4.8). It is also straightforward to generalise this method to the case where we have not only a set of configurations, but a set of particle paths over time that we wish to match – we will revisit this idea in the next chapter.

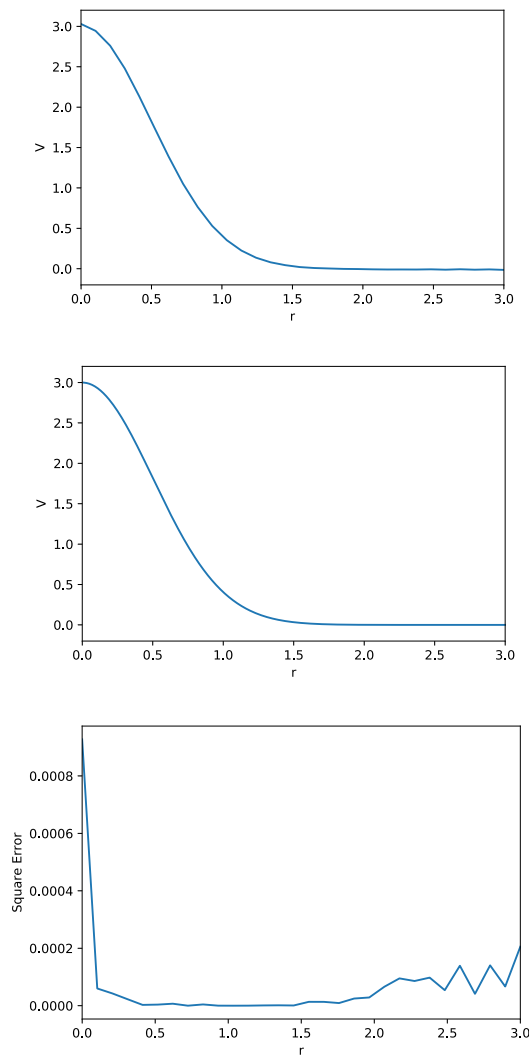


Figure 4.8: Inverted potential results for the first algorithm (configuration inversion), showing accurate reconstruction of particle potentials from a continuous stream of particle configurations. The upper image shows the reconstructed potential; the central image shows the target potential; and the lower image shows the square errors in the reconstructed potential as a function of r . The system considered was a set of particles with - using Hartree atomic units throughout - pair potentials $V(r) = A \exp(-Br^2)$, with $A = 3$ and $B = 2$. The density of particles was set to 0.5, in a box of length 11 (corresponding to 665 particles), and temperature $T = 1$. Periodic boundary conditions were assumed.

4.5.3 General Error Function Inversion

Suppose we have a set of experimental measurements of the system, $\{F_{exp}^i\}$, such that we can write a log-likelihood function as:

$$\log \pi(W) = L(\{F_{exp}^i\}, \{F^i(W)\}) \quad (4.35)$$

where $\{F^i(W)\}$ are the set of equivalent measurements when the system weights are equal to W . We assume $F^i(W)$ can be an average over particle configurations for the given W , i.e.:

$$F^i(W) = \frac{\int \exp(-\beta V(x; W)) f^i(x) dx}{\int \exp(-\beta V(x; W)) dx} = E_W[f^i(x)]. \quad (4.36)$$

Then we can write:

$$\nabla \log \pi(W) = \frac{\partial L}{\partial W} = \sum_i \frac{\partial L}{\partial F^i(W)} \frac{\partial F^i(W)}{\partial W}, \quad (4.37)$$

and:

$$\frac{\partial L}{\partial W} = -\beta \sum_i \frac{\partial L}{\partial F^i} \left(E_W \left[\frac{\partial V}{\partial W} f^i \right] - E_W [f^i] E_W \left[\frac{\partial V}{\partial W} \right] \right), \quad (4.38)$$

which is just a generalisation of the result we have exploited via IMC methods (IMC relies on a special case of eq. 4.38 to construct updates to W via Newton's method – the expression used here is exploited by a more general form of IMC, Newton Inversion[118]). Now, assume we can construct an approximate configuration sampler – i.e. a sampler in x space, producing a sequence of values $x_k, k = 1, 2, 3, \dots$ – for fixed W . Defining $E_W^n[g]$ as per equation 4.33, we assume the property $E_W[g] =$

$\lim_{n \rightarrow \infty} E_W^n[g]$ (which is straightforward to achieve with a CMC sampler, as previously).

We can then construct the quantity:

$$\frac{\partial L^n}{\partial W} = -\beta \sum_i \frac{\partial L}{\partial F^i} \left(E_W^n \left[\frac{\partial V}{\partial W} f^i \right] - E_W^n [f^i] E_W^n \left[\frac{\partial V}{\partial W} \right] \right), \quad (4.39)$$

in which appearances of F^i (in $\partial L / \partial F^i$) are evaluated as $E_W^n[f^i(x)]$. This quantity inherits the property:

$$\frac{\partial L}{\partial W} = \lim_{n \rightarrow \infty} \frac{\partial L^n}{\partial W}. \quad (4.40)$$

In the cases we will consider here, $\partial L / \partial F^i$ will yield a simple polynomial of expectations of functions of f^i ; as such, we can fall back on our cumulant approach to debiasing in direct combination with eq. 4.38. In the more general case, however, we can apply a (finite sequence) Russian Roulette technique to obtain an approximately unbiased estimator, $\partial \bar{L} / \partial W$, via 4.39. Applying the dynamics in eq. 6 (SGLD updates), with $[\nabla \log \pi(W)]_{est.} = \partial \bar{L} / \partial W$, we then sample the posterior distribution over W .

In order to validate this method, we again consider an application to a system of particles with Gaussian interaction potentials. We assume we have knowledge of the radial distribution function, up to Gaussian error bars – fig. 4.9 shows the resulting calculations of potential distributions, assuming a potential constrained to a parametrised form.

Bayesian methods for construction of potentials, in a similar manner, are prominent in the biophysics literature (see e.g. ref. [110]); however, the focus of these methods tends to centre on the optimisation problem, as opposed to extracting the posterior distribution over potential forms. While these approaches are evidently practical and useful, they cannot provide insight into uncertainties in the system.

4 Inferring Static Structure: Inverse Problem Approaches for Plasmas

Going beyond optimisation, some approximate posterior distribution inference has been demonstrated in ref. [114] – however, while the approximate approach used therein has the (important) advantage of speed, it assumes a simple form for the posterior distribution about the maximum likelihood. This assumption can be surprisingly inadequate for complex systems[86]. The present method is then a next step from this, attempting in-principle exact inference (up to residual bias in our estimators) with properly debiased Monte Carlo estimates. For the cost of additional computation time, the use of SGLD ensures more robust error bars.

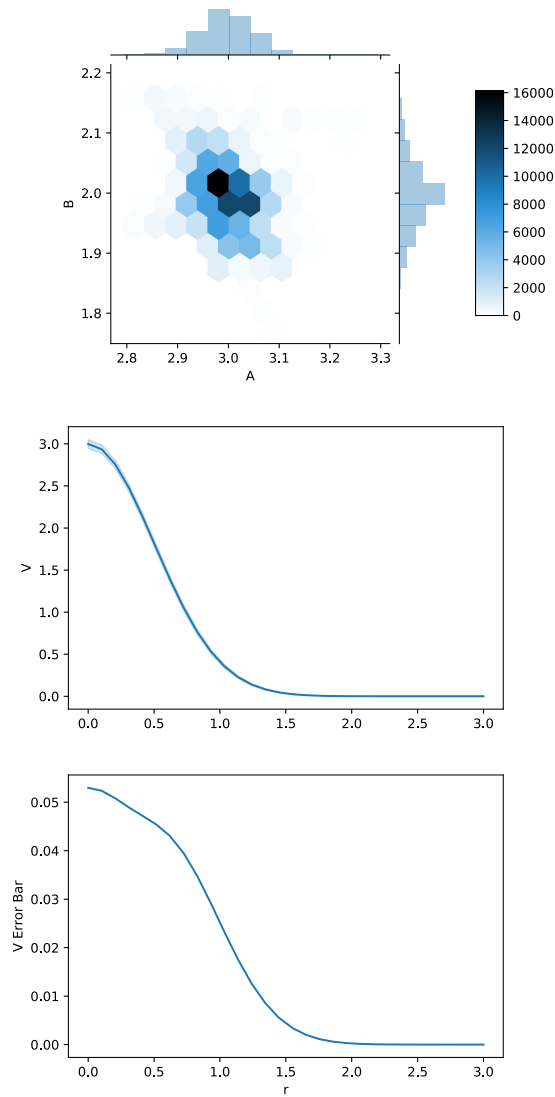


Figure 4.9: Sampling results for general error function inversion. Upper: the sampled distribution of parameters A and B - obtained via *SGLD* - corresponding to the target potential (as outlined in the caption for fig. 4.8) with Gaussian error bars. These results were inferred from an input of 15 equidistant radial distribution measurements up to $r = 3$ (taken from a separate simulation with $A = 3$ and $B = 2$), with fixed Gaussian error bars with variance $1/2$, under constraint of the potential form to $V(r) = A \exp(-Br^2)$. Centre: the corresponding inverted potential and error bars. Lower: a closer view of the error bars on the inverted potential as a function of r .

5

Inferring Dynamic Structure: Analysis of DSF Moments via SGMCMC

5.1 Introduction

The intention of the work in this chapter is to round out our analysis of wDM . In chapter 3, we developed methods for the prediction of static and dynamic properties of wDM ; in chapter 4, we developed methods for inference of static properties of wDM ; this chapter now completes our analysis by developing methods for inference of dynamic wDM properties. To this end, this chapter discusses the process of efficiently calculating moments of the DSF from static particle configurations, constructing DSFs from these calculations, and, ultimately, inverting dynamic structure measurements to constrain system properties.

In section 2.2.2, we briefly discussed the moments of the linear response function – or, equivalently, the moments of the DSF – and their physical meaning. We also alluded to the possibility of a more thorough analysis of the DSF through its moments – in particular, the feasibility of constructing complete DSF estimates from a finite number of moment values. It is this form of moment analysis that will allow us to extend our inversion techniques from static to dynamic parameters.

To begin, we consider moment properties in the traditional context of reversible dynamics, outlining the physical properties of individual DSF moments. We then determine the theoretical effects of noise on the DSF 's structure through stochastic calculus methods and the theory of memory functions. Finally, by combining this theory with the MCMC methods discussed in the previous chapter, we infer dynamic properties of an example Yukawa+SRR system with Langevin noise.

5.2 Moment Methods and Stochastic Calculus

5.2.1 Moments Under Smooth Reversible Dynamics

The central objects that we will consider in this chapter are the frequency moments of the DSF, M_n . In order to use these quantities in dynamic parameter inference schemes, we must first determine their makeup, and pathways to calculating them numerically. As a first step towards doing so, we can consider a low order expansion of its time/frequency Fourier transform, the Intermediate Scattering Function (ISF), given by:

$$F(\mathbf{q}, t) = \frac{1}{N} \langle \rho(\mathbf{q}, t) \rho(-\mathbf{q}, 0) \rangle, \quad (5.1)$$

where N is the number of particles and ρ the particle density function. Defining the DSF moments via:

$$M_n(\mathbf{q}) = \int_{-\infty}^{\infty} \omega^n S(\mathbf{q}, \omega) d\omega, \quad (5.2)$$

the Fourier transform relationship allows us to write the short time expansion of F as:

$$F(\mathbf{q}, t) = M_0 - M_2 \frac{t^2}{2!} + M_4 \frac{t^4}{4!} - \dots \quad (5.3)$$

Here we have initially limited ourselves to smooth, time-reversible dynamics, for which odd terms in the expansion are zero (this follows from the fact that the DSF is real, even in frequency, and infinitely differentiable under these conditions). Obtaining expressions for individual moments is then a matter of constructing a Taylor expansion in time of \tilde{F} as given in eq. 5.1, and comparing term by term with eq. 5.3.

For M_0 , we can just use eq. 5.2 to immediately arrive at:

$$M_0 = \int_{-\infty}^{\infty} S(\mathbf{q}, \omega) d\omega = S(\mathbf{q}). \quad (5.4)$$

Assuming we can treat our system as a collection of semi-classical point particles, we can calculate higher moments by writing \tilde{F} explicitly as:

$$F(q, t) = \frac{1}{N} E \left[\sum_{\alpha, \gamma}^N e^{i\mathbf{q} \cdot \mathbf{x}_\alpha(t)} \cdot e^{-i\mathbf{q} \cdot \mathbf{x}_\gamma(0)} \right], \quad (5.5)$$

where $x_\alpha(t)$ is the position of particle α at time t , and repeatedly differentiating with respect to t . To arrive at useful expressions, we can also apply standard tricks for manipulating the results (see e.g. ref. [119]). In particular, we can apply the Yvon theorem:

$$E \left[G(\mathbf{x}) \frac{\partial U}{\partial \mathbf{x}} \right] = k_B T E \left[\frac{\partial G(\mathbf{x})}{\partial \mathbf{x}} \right], \quad (5.6)$$

where $G(\mathbf{x})$ is a function of the particle coordinates, and U is the energy of the system.

Carrying out these derivatives, we obtain expressions for M_2, M_4 :

$$M_2 = \omega_0^2, \quad (5.7)$$

$$M_4 = 3\omega_0^4 + \omega_0^2 \Omega_0^2 - \omega_0^2 \Omega_L^2, \quad (5.8)$$

where

$$\omega_0^2 = q^2 k_B T / m, \quad (5.9)$$

$$\Omega_0^2 = \frac{\rho}{3m} \int d\mathbf{r} g(\mathbf{r}) \nabla^2 \phi, \quad (5.10)$$

$$\Omega_L^2 = \frac{\rho}{m} \int d\mathbf{r} \exp(-i\mathbf{q} \cdot \mathbf{z}) \partial_z^2 \phi g(\mathbf{r}), \quad (5.11)$$

and ϕ is the pairwise interaction potential. Higher order moments can be constructed by the same procedure, although they grow both in complexity and the degree of particle correlations required (the sixth moment is a functional of the three-particle correlation function $g_3(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$, the eighth requires g_4 , and so on).

Equipped with these expressions, we can map static configurations of particles onto dynamical models via an appropriate ansatz. Before doing so, however, we will generalise these results to systems driven by random forces; to accomplish this, we turn to stochastic calculus methods.

5.2.2 Stochastic Calculus

In elementary terms, stochastic calculus is based on a rigorous mathematical description of noise. Throughout this chapter we will deal with stochastic processes; very loosely, these processes can be viewed as functions of both continuous time t and a random variable ω drawn from some sample space Ω (see e.g. [120] for a rigorous definition). The central object of stochastic calculus – at least for our purposes – is the Wiener process, B_t (we will use B rather than the more typical W in order to distinguish the Wiener process from our system parameters W). This process is defined as the unique continuous-time stochastic process with the properties:

- $B_0 = 0$
- B has independent increments; $B_{t+s} - B_t$, $s > 0$ is independent of all values of B_u for $u < t$
- B_t is continuous in time
- $B_{t+s} - B_t$ is normally distributed with mean 0 and variance s

Stochastic calculus allows us to explicitly define integrals with respect to (and with integrands including) B . A full derivation of this (Itô) method of integration is omitted here, but can be found in e.g. ref. [120]. Through these integrals, we can obtain analytical properties of systems with equations of motion that can be cast in terms of B . The Gaussian increments of the Wiener process allow it to be interpreted – at a heuristic level – as the time integral of white noise; as such, manipulation of its derivative dB allows us to calculate the effects of white noise on particle statistics.

Turning to our immediate objectives, we can calculate properties of systems subject to a Langevin thermostat. As we discussed in chapter 3, the application of a Langevin thermostat is simply the application of a stochastic differential equation model of particle motion – this will be made explicit in the following section.

More generally, as well, the properties of the Wiener process allow for simple and exact inference of stochastic models of particle motion. The well-defined probability distribution of Wiener process increments makes the assignment of likelihoods to particle trajectories straightforward and natural. We will exploit this fact further into this chapter, in section 5.2.7, as we go on to consider the inverse problem for particle dynamics.

5.2.3 Langevin Moment Modifications

We now detail how stochastic calculus methods can be exploited to determine moment corrections due to Langevin friction. Let's start with a general stochastic system – the equations of motion of this system can be written:

$$d\mathbf{X} = \boldsymbol{\mu}dt + \boldsymbol{\sigma} \cdot d\mathbf{B}, \tag{5.12}$$

where μ and σ are, in general, functions of the realised phase space coordinates \mathbf{X} , and B is an m -dimensional Wiener process (i.e. a vector of m one-dimensional, independent Wiener processes). Throughout we follow the convention that omission of the time argument implies an argument t (i.e. $\mathbf{X} := \mathbf{X}_t$). Defining a diffusion operator L^0 via:

$$L^0 = \sum_{k=1}^d \mu_k \frac{\partial}{\partial x_k} + \frac{1}{2} \sum_{k,l=1}^d \sum_{j=1}^m \sigma_{k,j} \sigma_{l,j} \frac{\partial}{\partial x_k \partial x_l}, \quad (5.13)$$

and a set of operators L^j via:

$$L^j = \sum_{k=1}^d \sigma_{k,j} \frac{\partial}{\partial x_k}, \quad (5.14)$$

recursive application of Itô's formula then yields an expansion for a general function $f(\mathbf{X})$:

$$\begin{aligned} f(\mathbf{X}) &= f(\mathbf{X}_{t_0}) + \int_{t_0}^t L^0 f(\mathbf{X}_s) ds \\ &\quad + \sum_{j=1}^m \int_{t_0}^t L^j f(\mathbf{X}_s) dB_s^j \\ &= f(\mathbf{X}_{t_0}) + \int_{t_0}^t [L^0 f(\mathbf{X}_{t_0}) + \int_{t_0}^s L^0 L^0 f(\mathbf{X}_u) du \\ &\quad + \sum_{j=1}^m \int_{t_0}^s L^j L^0 f(\mathbf{X}_u) dB_u^j] ds \\ &\quad + \sum_{j=1}^m \int_{t_0}^t L^j f(\mathbf{X}_s) dB_s^j = \dots \end{aligned} \quad (5.15)$$

As a consequence, we may write:

$$E^{t,x}[f(\mathbf{X})|\mathbf{x}_0, t_0] = \sum_{i=0}^{\infty} \frac{(L^0)^i f(\mathbf{x})}{i!} \Big|_{\mathbf{x}=\mathbf{x}_0} (t - t_0)^i, \quad (5.16)$$

where $E^{t,x}$ is an expectation over realisations of the Wiener process B (i.e. an expectation over the distribution of values taken by the noise associated with B). Terms involving L^j for $j \neq 0$ vanish under the expectation, as each involves at least one integration with respect to a Wiener process (and thus has zero expectation, as a standard stochastic calculus result[120]). We also define the full thermodynamic expectation via:

$$\begin{aligned} E[f(\mathbf{X})] &= \int e^{-\beta\epsilon(\mathbf{x}_0)} E^{t,x}[f(\mathbf{X})|\mathbf{x}_0, t_0] d\mathbf{x}_0 \\ &= E^{\mathbf{x}_0}[E^{t,x}[f(\mathbf{X})|\mathbf{x}_0, t_0]] \end{aligned} \quad (5.17)$$

where ϵ is the energy (via the tower property of conditional expectations)[121]. In the particular case of a dynamic particle subsystem, we can set:

$$\mu_i = \begin{cases} V_i & i \leq N/2 \\ F_k/m - \alpha V_k & \textit{otherwise} \end{cases}, \quad k = i - N/2, \quad (5.18)$$

which defines the force F , and:

$$\sigma_{i,j} = \begin{cases} \sigma & i = j > N/2 \\ 0 & \textit{otherwise} \end{cases}. \quad (5.19)$$

In doing this, we make an approximation with respect to the introduced noise from collisions with omitted degrees of freedom: namely, that the collisions can be approximated as uniform in space and frequency, with rate dictated by σ . This is an assertion that Langevin dynamics/the Langevin thermostat are adequate descriptions of the dynamics. The constants σ , α are then related by the usual fluc-

tuation dissipation relationship. That is to say, if we relate the quantities via:

$$\alpha = \frac{\sigma^2 m \beta}{2} = \frac{\sigma^2}{2 \langle v_1^2 \rangle}, \quad (5.20)$$

then the equilibrium distribution of the stochastic differential equation 5.12 is the canonical distribution:

$$P(\mathbf{x}) = \frac{\exp[-\beta U(\mathbf{x})]}{Z}. \quad (5.21)$$

The corresponding operator L^0 for this case is given by:

$$L^0 = \mathbf{v} \cdot \frac{\partial}{\partial \mathbf{x}} + (\mathbf{F}/m - \alpha \mathbf{v}) \cdot \frac{\partial}{\partial \mathbf{v}} + \frac{1}{2} \sigma^2 \frac{\partial^2}{\partial \mathbf{v}^2}. \quad (5.22)$$

Throughout we will focus on this standard description of Langevin dynamics, but we note that generalising to arbitrary Molecular Dynamics thermostats is straightforward, and the analysis proceeds more or less identically. Applying a new thermostat model (in line with those discussed in chapter 3 for Bohmian trajectories) just requires the appropriate L^0 operator to be constructed via eq. 5.13.

We are interested in the short time expansion of the scattering function, $F(q, t)$, which requires us to consider particles individually. To do so, we define vectors in the full configuration space as functions of the three-dimensional \mathbf{q} via:

$$\mathbf{q}_\alpha(\mathbf{q}) \cdot \mathbf{x} = \mathbf{q} \cdot \mathbf{x}_\alpha, \quad (5.23)$$

where \mathbf{x}_α are the three-dimensional positional coordinates of particle α . We then have:

$$F(q, t) = \frac{1}{N} E \left[\sum_{\alpha, \gamma} e^{i \mathbf{q}_\alpha \cdot \mathbf{x}} \cdot e^{-i \mathbf{q}_\gamma \cdot \mathbf{x}_0} \right]. \quad (5.24)$$

The moments appearing in the expansion of the intermediate scattering function are then given by:

$$M_j = E^{x_0} \left[\frac{(L^0)^j}{N} \sum_{\alpha, \gamma} e^{iq_\alpha \cdot x} \cdot e^{-iq_\gamma \cdot x_0} \Big|_{x=x_0} \right]. \quad (5.25)$$

This can be related to the non-stochastic (Born-Oppenheimer without electron collisions) case, via:

$$M_j^{BO} = M_j|_{\alpha=\sigma=0}, \quad (5.26)$$

allowing the difference between these values to be retroactively added on to simulate collisions.

From the form of L_0 in eq. 5.22 and the general expression for M_j in eq. 5.25, we see each moment can be expressed as a sum of expectations of functionals of the velocity v and (deterministic) force F . We observe at this point that these expectation values are identical for all choices of σ, α satisfying the fluctuation dissipation theorem (eq. 5.20). This follows from the fact that in each case the equilibrium probability distribution is identical (eq. 5.21). As a result, variations in each moment with respect to σ, α are only present where these parameters appear explicitly – functions of the particle configurations themselves are necessarily independent of the noise parameters. Given moment calculation results from one set of these parameters, it is then straightforward to obtain moments corresponding to an alternate set.

Calculations in Three Dimensions

Equipped with the above results, we can now calculate expressions for systems containing many particles in three dimensions. We arrive at:

$$M_0 = M_0^{BO}, \quad (5.27)$$

$$M_1 = 0, \quad (5.28)$$

$$M_2 = M_2^{BO}, \quad (5.29)$$

$$M_3 = -\alpha M_2^{BO}, \quad (5.30)$$

$$M_4 = M_4^{BO} + \alpha^2 M_2^{BO}, \quad (5.31)$$

for the first four terms, where the Born-Oppenheimer contributions are the equivalent quantities under smooth dynamics, as calculated in section 5.2.1.

5.2.4 Practical Sampling of Moments

We now have complete expressions for the early moments of a many-particle system; the utility of these expressions, however, is predicated on our ability to calculate them for model systems. We can do this via Classical Monte Carlo simulation – we can sample particle configurations from the canonical distribution through the usual method, and from each configuration calculate an estimate of each moment. Considering many-body systems interacting via pair potentials – as we have done throughout this thesis – it is in fact straightforward to construct moment estimates (even up to high order) in $O(N^2)$ time, with N the number of particles. That is to say, moments can be estimated within the same time complexity as that required to calculate the force on every particle.

The favourable time complexity follows directly from the form of the potential via a simple inductive argument. For this analysis we are free to focus on the Born-Oppenheimer case for the moments we've considered, as we've seen that it is trivial to build moments for arbitrary Langevin noise from these. We can therefore work directly with time derivatives, without concern about their existence under stochastic dynamics.

We assume the n^{th} velocity derivative for an arbitrary particle in the system has the form (for $n > 1$):

$$\mathbf{v}_\alpha^{(n)} = \sum_{\gamma \neq \alpha} \mathbf{f}_n \left(\mathbf{x}_\alpha - \mathbf{x}_\gamma, \mathbf{v}_\alpha, \dot{\mathbf{v}}_\alpha, \dots, \mathbf{v}_i^{(n-1)}, \mathbf{v}_\gamma, \dot{\mathbf{v}}_\gamma, \dots, \mathbf{v}_\gamma^{(n-1)} \right), \quad (5.32)$$

for some generalised pair interaction function \mathbf{f}_n . The time derivative of this quantity then has the form:

$$\mathbf{v}_\alpha^{(n+1)} = \sum_{\gamma \neq \alpha} \left(\frac{\partial \mathbf{f}_n}{\partial (\mathbf{x}_\alpha - \mathbf{x}_\gamma)} \cdot (\mathbf{v}_\alpha - \mathbf{v}_\gamma) + \sum_{m=0}^{n-1} \left(\frac{\partial \mathbf{f}_n}{\partial \mathbf{v}_\alpha^{(m)}} \cdot \mathbf{v}_\alpha^{m+1} + \frac{\partial \mathbf{f}_n}{\partial \mathbf{v}_\gamma^{(m)}} \cdot \mathbf{v}_\gamma^{m+1} \right) \right), \quad (5.33)$$

which, by inspection, can be written in the same way:

$$\mathbf{v}_\alpha^{(n+1)} = \sum_{\gamma \neq \alpha} \mathbf{f}_{n+1} \left(\mathbf{x}_\alpha - \mathbf{x}_\gamma, \mathbf{v}_\alpha, \dot{\mathbf{v}}_\alpha, \dots, \mathbf{v}_\alpha^{(n)}, \mathbf{v}_\gamma, \dot{\mathbf{v}}_\gamma, \dots, \mathbf{v}_\gamma^{(n)} \right). \quad (5.34)$$

As it is trivial to write $\dot{\mathbf{v}}$ in the form of eq. 5.32, it follows that all velocity derivatives for a particle can be expressed as a linear sum over generalised pair interaction functions, \mathbf{f} , dependent on the positions and lower order velocity derivatives of the interacting particles. Given all lower order derivatives for all particles, it is clear that $\mathbf{v}_\alpha^{(n)}$ can be calculated in $O(N)$ time; calculating $\mathbf{v}^{(n)}$ for all particles is then an $O(N^2)$ operation. Calculating all of the derivatives for all particles up to the n^{th} order can therefore be performed as a sequence of a constant number of

$O(N^2)$ operations, calculating derivatives in order from the lowest to the highest, with the constant factor an increasing function of n .

Once these calculations have been carried out, we can use the resultant quantities to estimate moments. A general moment expression has the form:

$$M_n = \sum_{\alpha, \gamma} m_n (\mathbf{x}_\alpha - \mathbf{x}_\gamma, \mathbf{v}_\alpha, \dot{\mathbf{v}}_\alpha, \dots, \mathbf{v}_\alpha^{(n-1)}), \quad (5.35)$$

for some scalar function m_n , and, given all of the required velocity derivatives and positions, can be calculated in $O(N^2)$ time. As such, the full calculation of all of the moments up to the n^{th} from a single particle configuration is at worst an $O(N^2)$ operation.

5.2.5 Memory Function Methods

Now that we have established a fast and robust path to numerical calculation of moments, we can progress to developing models of the DSF that exploit these results. Our focus here – the memory function method – will allow us to estimate the full DSF from a limited number of moments. In turn, this estimation procedure will open up further possibilities for inference of system properties from the DSF.

To begin building the memory function method, we closely follow the derivation (for the case without Langevin noise) in ref. [25]. Let A, B , be a pair of dynamical system variables, which depend on the coordinates and momenta of the particles in our system. For non-stochastic, time reversible dynamics, the time evolution of functions of this type can be written as:

$$A(t) = \exp(Lt)A(0), \quad (5.36)$$

where L is (up to a factor of i) the standard Liouville operator, with the form:

$$L = \sum_{i=1}^N \left[\frac{\partial H}{\partial p_i} \frac{\partial}{\partial x_i} - \frac{\partial H}{\partial x_i} \frac{\partial}{\partial p_i} \right], \quad (5.37)$$

and x_i, p_i are the coordinates and momenta of the system particles, and H is the Hamiltonian. This operator takes on an identical form to our stochastic evolution operator defined in 5.22, with noise (for the time being) set to zero.

We can define the equilibrium time correlation function between A and B as the expectation:

$$C_{AB}(t', t'') = \langle A(t') B^*(t'') \rangle, \quad (5.38)$$

and can similarly construct the inner product between the two:

$$(A(t'), B(t'')) = \langle A^*(t') B(t'') \rangle. \quad (5.39)$$

It is straightforward to show that this construction satisfies the usual requirements for an inner product, and that the Liouville operator (again, once we account for a factor of i) is Hermitian with respect to it[25].

The memory function method can now be rigorously derived through Zwanzig's projection operator formalism[122]. To simplify our notation, we'll omit time arguments when those arguments are zero (i.e. $A \equiv A(0)$). We start by defining a projection operator, P , in terms of the inner product of eq. 5.39, via:

$$PB(t) = (A, B(t)) (A, A)^{-1} A, \quad (5.40)$$

and a complementary operator Q through:

$$Q = 1 - P. \quad (5.41)$$

The action of P is then to project out the element of its operand along A (and conversely, Q projects out the element that is perpendicular to A). Projecting $A(t)$ along A , we can write:

$$PA(t) = Y(t)A, \quad (5.42)$$

with

$$Y(t) = (A, A(t))(A, A)^{-1} = C_{AA}(t)/C_{AA}(0). \quad (5.43)$$

To calculate the correlation function, then, we can examine the time evolution of $Y(t)$. We can start by taking the Laplace transform of eq. 5.36, giving:

$$(z + L)\tilde{A}(z) \equiv (z + L)(P + Q)\tilde{A}(z) = iA. \quad (5.44)$$

We can then split this equation into perpendicular components by separate application of P and Q . After carrying this out, some simple algebraic manipulation (see e.g. ref. [25]) allows us to write:

$$(-iz - i\Omega)\tilde{Y}(z) + (K, \tilde{R}(z))(A, A)^{-1}\tilde{Y}(z) = 1, \quad (5.45)$$

where we define K via:

$$K = Q\dot{A} = QLA, \quad (5.46)$$

R is a so-called random force (although it does not have the dimensions of a force in general) defined by:

$$\tilde{R}(z) = (iz + QLQ)^{-1}K, \quad (5.47)$$

and we introduce a frequency matrix Ω through:

$$i\Omega = (A, \dot{A})(A, A)^{-1} = \dot{Y}(0). \quad (5.48)$$

Now, defining the memory function M via:

$$M(t) = (R, R(t))(A, A)^{-1}, \quad (5.49)$$

we can rewrite eq. 5.45 as:

$$\tilde{Y}(z) = \left[-iz - i\Omega + \tilde{M}(z) \right]^{-1}, \quad (5.50)$$

which, in the time domain, can be written as:

$$\dot{Y}(t) - i\Omega Y(t) + \int_0^t M(t-s)Y(s)ds = 0. \quad (5.51)$$

We can also obtain a similar result for the orthogonal projection of $A(t)$. From this point, we can further derive recursive dynamical equations for a sequence of memory functions M_n , with $M_1 = M$, each following an equation of motion identical in form to 5.51; we do not pursue this further at this stage, but this technique opens up further paths to approximating the DSF (see, for example, the detailed analysis in ref. [123]).

Generalising eq. 5.51 slightly, we can repeat the same derivation with a vector A of n independent components, A_1, A_2, \dots, A_n in lieu of the single A , arriving at:

$$\dot{\mathbf{Y}}(t) - i\Omega \mathbf{Y}(t) + \int_0^t \mathbf{M}(t-s) \cdot \mathbf{Y}(s)ds = \mathbf{0}, \quad (5.52)$$

with vector quantities analogous to those in the scalar case. Eq. 5.52, and the associated definitions of each of its parts, allow us to directly determine the time evolution of $\mathbf{Y}(t)$ in terms of well-defined quantities.

While the memory function itself is difficult to calculate from first principles, its expansion terms can be directly linked to the expansion terms of $\mathbf{Y}(t)$; it is then

straightforward to constrain ansatzes for M based on the short-time behaviour of Y . Eq. 5.52 thereby allows us to construct descriptive models of the ISF and DSF in terms of a memory function (which, in practice, can be accurately approximated with a much simpler functional form).

The real value of the memory function method becomes clear in application. Considering now the specific case of the ISF \tilde{F} as our correlation quantity, we take the elements of A to be ρ_q , the longitudinal particle current $j_q = \mathbf{q} \cdot \mathbf{j}_q/q$, and a temperature variable T_q , constructed to be perpendicular to the other components via:

$$T_q = e_q - (\rho_q, e_q)(\rho_q, \rho_q)^{-1} \rho_q, \quad (5.53)$$

where e_q is the Fourier energy density, defined for a particulate system as:

$$\sum_i \left(\frac{1}{2} m_i \mathbf{v}_i^2 + \frac{1}{2} \sum_{j \neq i} U(r_{ij}) \right) \exp(-i\mathbf{q} \cdot \mathbf{r}_i), \quad (5.54)$$

in which variables take their standard meanings. From this we can write the static correlation matrix as:

$$(\mathbf{A}, \mathbf{A}) = \begin{pmatrix} NS(k) & 0 & 0 \\ 0 & \frac{Nk_B T}{m} & 0 \\ 0 & 0 & \langle T_q T_{-q} \rangle \end{pmatrix}, \quad (5.55)$$

and the frequency matrix as:

$$-i\Omega = \begin{pmatrix} 0 & -iq & 0 \\ \frac{-iq}{S(q)} \left(\frac{k_B T}{m} \right) & 0 & \frac{\langle j_q T_{-q} \rangle}{\langle T_q T_{-q} \rangle} \\ 0 & -\frac{\langle T_q j_{-q} \rangle}{Nk_B T/m} & 0 \end{pmatrix}. \quad (5.56)$$

We also have that the first component of \mathbf{R} is zero, as \dot{A}_1 is proportional to A_2 ; the memory function matrix is then:

$$\mathbf{M}(q, t) = \begin{pmatrix} 0 & 0 & 0 \\ 0 & M_{22}(q, t) & M_{23}(q, t) \\ 0 & M_{32}(q, t) & M_{33}(q, t) \end{pmatrix}. \quad (5.57)$$

Turning back to eq. 5.50, we can then write the correlation matrix as:

$$\tilde{\mathbf{Y}}(q, z) = \begin{pmatrix} -iz & iq & 0 \\ \frac{iq}{S(q)} \left(\frac{k_B T}{m}\right) & -iz + \tilde{M}_{22}(q, z) & -i\Omega_{23} + \tilde{M}_{23}(q, z) \\ 0 & -i\Omega_{32} + \tilde{M}_{32}(q, z) & -iz + \tilde{M}_{33}(q, z) \end{pmatrix}. \quad (5.58)$$

In the Laplace domain we then have the solution for \tilde{F} , through:

$$\tilde{F}(q, z) = S(q)\tilde{Y}_{11}(q, z) = \frac{S(q)}{-iz + \frac{1}{S(q)} \left(\frac{\omega_0^2}{-iz + \tilde{N}_l(q, z)}\right)}, \quad (5.59)$$

where

$$\tilde{N}_l(q, z) = \tilde{M}_{22}(q, z) - \frac{\Theta(q, z)}{-iz + \tilde{M}_{33}(q, z)}, \quad (5.60)$$

and

$$\Theta(q, z) = \left(\tilde{M}_{23}(q, z) - \frac{\langle \dot{j}_q T_{-q} \rangle}{\langle T_q T_{-q} \rangle}\right) \left(\tilde{M}_{32}(q, z) + \frac{\langle T_q \dot{j}_{-q} \rangle}{Nk_B T/m}\right). \quad (5.61)$$

Finally, in the frequency domain ($z = \omega + i\epsilon$ as $\epsilon \rightarrow 0+$) we have:

$$S(q, \omega) = \frac{1}{\pi} \frac{\omega_0^2 N_l'(q, \omega)}{[\omega^2 - \omega_0^2/S(q) - \omega N_l''(q, \omega)]^2 + [\omega N_l'(q, \omega)]^2}, \quad (5.62)$$

where $N'_l(q, \omega)$ and $N''_l(q, \omega)$ are the real and imaginary parts of $\tilde{N}_l(q, \omega)$ respectively. This fully describes the DSF in terms of static quantities and the memory function.

The explicit physical definition of the frequency matrix now allows to straightforwardly integrate Langevin noise into the model – we can consider an identical procedure to that above, but with the full stochastic evolution operator in eq. 5.22 effecting the time derivative. We simplify the theory to the essentials by considering a two-component \mathbf{A} vector, with components ρ_q and j_q . Assuming our usual Langevin equation of motion, we note that the drag term arising from the action of the full evolution operator introduces a correlation between j_q and its time derivative, via $\langle j_q \dot{j}_{-q} \rangle = -\alpha \langle j_q j_{-q} \rangle$. As a result, we now have $-i\Omega_{22} = \alpha$.

Looking then to eq. 5.58, we see this element update (in the two-component model) is exactly equivalent to making the update $\tilde{M}_{22}(q, z) \rightarrow \tilde{M}_{22}(q, z) + \alpha$; adding noise to the DSF is then just a matter of adding α to $N'_l(q, \omega)$ wherever it appears.

From a phenomenological standpoint, Langevin collisions should introduce an instantaneous decay term to the memory function; this term would then account for randomisation that occurs on a timescale much shorter than that of the characteristic motion of the system. To model this, we could add a term proportional to $\delta(t)$ to $N_l(q, t)$ – clearly, this is equivalent to adding a constant to $N'_l(q, \omega)$, and is consistent with the microscopic result.

Now we are in a position to relate the memory function model back to our full stochastic moment expansion. Our equation of motion for the ISF, $F(q, t)$, can be written:

$$\ddot{F}(q, t) + \alpha \dot{F}(q, t) + \Omega_q^2 F(q, t) + \int_0^t M(q, t - t') \dot{F}(q, t') dt'. \quad (5.63)$$

5 Inferring Dynamic Structure: Analysis of DSF Moments via SGMCMC

where $\Omega_q^2 = \omega_0^2/S(q)$. Repeated differentiation in time then allows us to derive a simple recurrence relationship between the time derivatives of F at $t = 0$:

$$F^{(n)}(0) + \alpha F^{(n-1)}(0) + \Omega_q^2 F^{(n-2)}(0) + \sum_{i=1}^{n-2} M^{(n-2-i)}(0) F^{(i)}(0). \quad (5.64)$$

Terms in α do not appear for $n = 2$; considering terms up to the fourth derivative, we then have:

$$\tilde{F}^{(3,\alpha)} = -\alpha \tilde{F}^{(2,0)}, \quad (5.65)$$

$$\tilde{F}^{(4,\alpha)} = (-M^{(0,\alpha)} - \Omega_q^2 + \alpha^2) \tilde{F}^{(2,0)}, \quad (5.66)$$

where we've adopted the notation that $G^{(l,m)}$ is the l^{th} derivative of G at $t = 0$ and $\alpha = m$. Writing our full stochastic expansion results in an identical form, we have:

$$\tilde{F}_{sto}^{(3,\alpha)} = -\alpha \tilde{F}^{(2,0)}, \quad (5.67)$$

$$\tilde{F}_{sto}^{(4,\alpha)} = \tilde{F}_{sto}^{(4,0)} + \alpha^2 \tilde{F}^{(2,0)}. \quad (5.68)$$

Clearly the third derivative terms are consistent with one another. Now, equating $\tilde{F}_{sto}^{(4,\alpha)} - \tilde{F}_{sto}^{(4,0)}$ with $\tilde{F}^{(4,\alpha)} - \tilde{F}^{(4,0)}$, we see also that setting $M^{(0,\alpha)} = M^{(0,0)}$ ensures consistency up to the fourth derivative in time. As such, the value of M at $t = 0$ is independent of α for an arbitrary system of interacting particles. This result can also be extracted directly from the definition of the memory function in the stochastic two-component case. The random force vector has a single component given by:

$$K_q = QLj_q, \quad (5.69)$$

and M can be written as:

$$M(q, t) = \frac{m}{Nk_B T} (K_q, R_q(t)), \quad (5.70)$$

with

$$R_q(t) = \exp(QLQt)K_q, \quad (5.71)$$

such that

$$M(q, 0) = \frac{m}{Nk_B T} (K_q, K_q). \quad (5.72)$$

The action of the stochastic evolution operator on j_q introduces an additional term (relative to the non-stochastic case) proportional to αj_q ; the action of Q on this additional term, however, is to eliminate it, as it is proportional to one of our chosen A components. The random force vector is then independent of α , and so therefore is the value of M at $t = 0$.

This α independence property leads us to a useful result: given a full memory function for a system (or equivalently its full ISF) and its noise parameter α , we can trivially construct an ISF for the same system under different noise conditions, exact up to the fourth derivative – all that is required is to keep M fixed, and adjust the value of α in eq. 5.63. In particular, we can take results for the ISF calculated with negligible noise via expensive calculation methods (such as DFT-MD with a Nosé-Hoover thermostat), and very quickly construct estimates of the ISF and DSF at arbitrary levels of Langevin noise. Calculations already run under Born-Oppenheimer conditions can then be rapidly corrected to account for Langevin-level collisions without further expensive simulations.

To carry out such a correction, we first must invert the ISF to obtain the memory function. This can be done quite simply by numerically solving for values of M at different times in eq. 5.63 (through e.g. a straightforward generalisation of the

iterative approach in ref. [124] to our slightly more complicated case). Alternatively, we can start with a memory function form from the literature (e.g. one of those in refs [20], [123], [125]), constrained to fit the Born-Oppenheimer moments. Then we can solve for F using this memory function and an alternative value of α , chosen to reflect an estimate of noise effects in the system. Fig. 5.1 demonstrates this procedure and the basics of the relationship between the DSF and α . We can therefore both correct existing simulation data and existing semi-analytical memory function approaches for noise.

Beyond this low-order correction, we can also consider higher moments. Although tedious, it is merely a mechanical task to calculate higher derivatives of F in time through repeated action of the evolution operator. The resulting derivative values can then be compared with derivative values from eq. 5.64, and the memory function can be constrained to have the correct α dependence up to arbitrary order.

5.2.6 Setting α Theoretically

Having established the link between friction and the DSF in detail (and methods for manipulating the DSF in Langevin parameter space), we now consider methods for establishing values of α for use in predictive modelling. In order to do this, we have two options:

1. Derive the characteristics of the noise theoretically, via (for example) linear response theory and properties of the electronic stopping power
2. Perform a full calculation using (for example) non-adiabatic Time-Dependent Density Functional Theory for a sufficient time to extract effective noise properties




Figure 5.1: Approximate injection of Langevin noise into Born-Oppenheimer DSF calculations. Beginning with a Born-Oppenheimer estimate of $F(q_0, t)$, we can extract an associated memory function at $\alpha = 0$. With this memory function, F can be approximately recalculated for the target system and arbitrary α , through numerical solution of eq. 5.63. Results of this procedure for the DSF are shown for a synthetic F input; qualitatively, the variation of the DSF with increasing α is consistent with that explicitly calculated in ref. [15] (decaying frequency and relative amplitude of the acoustic side peaks). For basic demonstration purposes, the (toy) memory function in this case is taken as an exponentially decaying cosine function, with arbitrary values for ω_0 , q_0 , and arbitrary units for α .

The prior approach is far simpler to carry out, and we consider it first.

The electronic stopping power is the rate of energy loss (per unit distance) experienced by a projectile ion traversing a material due to its interaction with the electrons therein. It can be written as $S = -dE/dx$, where E is the ion energy and x its position[126]. By equating the drag associated with α in the Langevin equations of motion with the drag associated with the stopping power (at low velocities, for which linear drag in v dominates), we can attempt to exploit standard calculations of S to obtain α .

Of course, this approach's accuracy is tied to that of the model used for S - we adopt a fairly straightforward point-charge model here for demonstration purposes/speed of calculation. The basic expressions for S and the RPA linear response used below can be found in refs. [126, 18].

The stopping power for an ion with charge distribution ρ_{ion} and velocity v can be written:

$$S(v) = \frac{2}{\pi v^2} \int_0^\infty \frac{dk}{k} |\rho_{ion}(k)|^2 \int_0^{kv} d\omega \omega \operatorname{Im} \left[\frac{-1}{\epsilon(k, \omega)} \right], \quad (5.73)$$

where $\epsilon(k, \omega)$ is the dielectric function associated with the present (valence) electrons. For a point charge, this reduces to:

$$S(v) = \frac{2Z^2}{\pi v^2} \int_0^\infty \frac{dk}{k} \int_0^{kv} d\omega \omega \operatorname{Im} \left[\frac{-1}{\epsilon(k, \omega)} \right]. \quad (5.74)$$

To simplify matters, we divide ϵ into real and imaginary parts through $\epsilon(k, \omega) = \epsilon_1 + i\epsilon_2$. At zero frequency, $\epsilon_2 = 0$. The Born-Oppenheimer approximation, by neglecting the electron dynamics, then essentially fixes ϵ_2 , and hence the stopping power, to 0. Fixing the Langevin parameter by expanding beyond the static limit

(i.e. by expanding ϵ_2 to linear order in ω) can therefore be seen as a first order correction to this.

Defining

$$\kappa(k) = \frac{d\epsilon_2}{d\omega}_{\omega=0}, \quad (5.75)$$

at low velocity the second integral becomes:

$$\int_0^{kv} d\omega \omega \operatorname{Im} \left[\frac{-1}{\epsilon(k, \omega)} \right] = \int_0^{kv} d\omega \omega^2 \frac{\kappa(k)}{\epsilon_1(k, 0)^2} = \frac{k^3}{3} \frac{\kappa(k)}{\epsilon_1(k, 0)^2} v^3. \quad (5.76)$$

The initial gradient of S is given by:

$$\frac{dS}{dv}_{v=0} = \frac{2}{3\pi} \int_0^\infty dk |\rho_{ion}(k)|^2 \frac{k^2 \kappa(k)}{\epsilon_1(k, 0)^2}. \quad (5.77)$$

Also, as

$$S = -\frac{dE}{dx} = -mv \frac{dv}{dx} = -m \frac{dv}{dt} = +\alpha v, \quad (5.78)$$

we have

$$\alpha = \frac{2}{3\pi} \int_0^\infty dk |\rho_{ion}(k)|^2 \frac{k^2 \kappa(k)}{\epsilon_1(k, 0)^2}, \quad (5.79)$$

or for a point charge approximation,

$$\alpha = \frac{2Z^2}{3\pi} \int_0^\infty dk \frac{k^2 \kappa(k)}{\epsilon_1(k, 0)^2}. \quad (5.80)$$

At the level of RPA (although note static local field corrections are straightforward to include), we have:

$$\epsilon(k, \omega) = 1 - v_q \chi_0(k, \omega), \quad (5.81)$$

with v_q the Coulomb potential, and χ_0 the Lindhard function. The real and imaginary parts need to be treated separately. The real part of the Lindhard function is

given by:

$$\frac{\text{Re } \chi_0(k, \omega)}{N(0)} = -\frac{1}{2} \int_0^\infty dx \frac{F(x, T)}{\bar{k}} \left[\ln \left(\frac{x - \nu_-}{x + \nu_-} \right) - \ln \left(\frac{x - \nu_+}{x + \nu_+} \right) \right], \quad (5.82)$$

where

$$F(x, T) = \frac{x}{e^{\beta(x^2 \epsilon_F - \mu)} + 1}, \quad (5.83)$$

$$\bar{k} = \frac{k}{k_F}, \quad (5.84)$$

$$\nu_\pm = \frac{\omega}{k v_F} \pm \frac{k}{2k_F}, \quad (5.85)$$

and the imaginary part is:

$$\frac{\text{Im } \chi_0(k, \omega)}{N(0)} = -\pi \left[\frac{\omega}{v_F k} + \frac{k_B T}{\hbar v_F k} \ln \left(\frac{1 + \exp(\beta[\nu_-^2 \epsilon_F - \mu])}{1 + \exp(\beta[\nu_+^2 \epsilon_F - \mu])} \right) \right]. \quad (5.86)$$

Differentiating this, we arrive at:

$$\frac{\kappa(k)}{N(0)} = -\pi \left[\frac{1}{v_F k} - \frac{1}{\hbar v_F^2 k^2} \left(\frac{4\nu_0 \epsilon_F \exp(\beta[\nu_0^2 \epsilon_F - \mu])}{1 + \exp(\beta[\nu_0^2 \epsilon_F - \mu])} \right) \right], \quad (5.87)$$

where $\nu_0 = k/2k_F$. The static part of eq. 5.82 can also be simplified to obtain:

$$\frac{\text{Re } \chi_0(k, \omega)}{N(0)} = - \int_0^\infty dx \frac{F(x, T)}{\bar{k}} \ln \left(\frac{x + \nu_0}{x - \nu_0} \right). \quad (5.88)$$

The values entering into eq. 5.80 are then completely determined, and α can be calculated by carrying out the integrals involved numerically. We have:

$$\epsilon_1(k, \omega) = \epsilon_\infty + \frac{4e^2 m_e^2}{\pi^{1/2} k^3 \hbar^4 \beta} e^{\beta\mu} [D(y_1) + D(y_2)], \quad (5.89)$$

where ϵ_∞ is a background dielectric constant contribution due to inner electrons in the ion core (which we estimate as 1), $D(y) = \exp(-y^2) \int_0^y \exp(x^2) dx$ is Dawson's

integral,

$$y_1^2 = \frac{\beta \hbar^2}{2}, \quad (5.90)$$

and

$$k_{DH}^2 = \frac{4\pi e^2 n \beta}{\epsilon_\infty}. \quad (5.91)$$

Setting α in this way, our memory function model (alongside static moment calculations in the Born-Oppenheimer approximation) becomes a fully predictive model of the DSF.

5.2.7 Inferring α from Trajectory Data

To go beyond the simple estimate of α established above, we can take a first principles approach with TDDFT. In doing so, we can achieve ‘the best of both worlds’ with respect to DFT calculations of the DSF. TDDFT, while too expensive for full resolution of the ion-ion DSF, can be used to establish effective non-adiabatic properties of the system. A standard adiabatic technique (e.g. Born-Oppenheimer DFT-MD) can then be used in conjunction with a Langevin thermostat with parameters calibrated by TDDFT. The combined result is fast adiabatic resolution of the DSF, with accurate non-adiabatic corrections.

The basis of this approach is the natural probabilistic interpretation of particle motion under a stochastic differential equation model (as briefly discussed in section 5.2.2). Initially, we consider the evolution of the particles in our system over a single time step of length Δt . Under our simple Langevin model, and to first order in Δt under a basic Euler-Maruyama update scheme[120], we would update system velocities via:

$$\mathbf{v}_{t+\Delta t}^{EM} = \mathbf{v}_t + [\mathbf{F}(\mathbf{x}_t)/m - \alpha \mathbf{v}_t] \Delta t + \sigma \Delta \mathbf{B}, \quad (5.92)$$

with ΔB the increment in a $3N$ -dimensional Wiener process over time Δt . We can, then, write the implied probability distribution of $\mathbf{v}_{t+\Delta t}$ via:

$$\mathbf{v}_{t+\Delta t}^{EM} \sim \mathbf{N}(\mathbf{A}(t), C\mathbf{1}), \quad (5.93)$$

$$\mathbf{A}(t) = \mathbf{v}_t + [\mathbf{F}(\mathbf{x}_t)/m - \alpha\mathbf{v}_t] \Delta t, \quad (5.94)$$

$$C = \frac{2\alpha k_B T}{m} \Delta t, \quad (5.95)$$

where N is a $3N$ -dimensional normal distribution, $\mathbf{1}$ the identity matrix, and we have used the fluctuation dissipation relationship between σ and α . Given a sequence of positions/velocities for the system from TDDFT, the log-likelihood of a given force function \mathbf{F} and Langevin friction α is given by:

$$-\frac{1}{2} \sum_t \left(\frac{[\mathbf{v}_{t+\Delta t} - \mathbf{A}(t)]^2}{C} + 3N \log(C) \right). \quad (5.96)$$

It is then just a matter of applying an MCMC method to obtain an implied distribution over α and any force parameters. As the static correlations of the system are independent of α , it is typically a reasonable approximation to first determine an optimal potential/force first (using the methods outlined in the preceding chapter), and then to carry out one-dimensional sampling of α conditional on this potential/force - this is trivial with the methods detailed in the previous chapter.

As a final note on this kind of parameter inference, the procedure we have outlined here can be readily extended to more sophisticated models (which are appropriate when the trajectory data set is large enough to justify them). Noise models built on the generalised Langevin Equation[25], for example, are a natural exten-

sion, with equation of motion for particle γ given by:

$$m\dot{\mathbf{v}}_\gamma = \mathbf{F}_\gamma - \int_0^t dt' \sum_\epsilon \alpha_{\gamma\epsilon}(t, t') \mathbf{v}_\epsilon(t') + \delta\mathbf{F}_\gamma(t), \quad (5.97)$$

with corresponding fluctuation dissipation relationship

$$\langle \delta\mathbf{F}_\gamma(t) \delta\mathbf{F}_\epsilon(t') \rangle = k_B T \alpha_{\gamma\epsilon}(t, t'). \quad (5.98)$$

Kernel models of this type go beyond Markovian approximations to the dynamics by inferring electron effects based on the full history of the ion trajectories. The components $\alpha_{\gamma\epsilon}$ can be sampled in an identical fashion to α in the simple Langevin model, although a suitable prior may be necessary to constrain component values to a small range for practical reasons (in particular, constraints on $\alpha_{\gamma\epsilon}$ are typically necessary to ensure a practical form for a probability distribution over random forces that is consistent with eq. 5.98 – see e.g. ref [127]).

5.3 Inference of Dynamic Properties via Moments

5.3.1 Static Sampling of Moment Gradients

The theoretical elements developed above have established an insight into the effects of noise on the ISF; this insight, in turn, has allowed us to extend well-known semi-analytic models of the DSF, and adjust simulation outputs for noise as a post-processing step. To round out this chapter, we will develop the remaining theory required for dynamic property inference using integrated moments of the DSF. This will allow us to demonstrate the applicability of our theoretical work to experimental inference via a synthetic test case.

In section 5.2.7 we considered parameter inference using particle trajectory data; clearly, however, the procedure in that section is limited when it comes to inference from experimental data. We must work with experimentally measurable properties, and cannot rely on full step-by-step data on particle trajectories – we therefore turn to the DSF and its integrated moments.

In order to apply SGMCMC methods to dynamic correlation properties, then, we must be able to sample unbiased gradients – with respect to system parameters – of errors in DSF moments relative to experimental data. We have established in section 5.2.4 that moments themselves can be sampled in $O(N^2)$ time from a Classical Monte Carlo simulation configuration; we now examine various derivatives of these moments with respect to system parameters.

For demonstration purposes, the posterior distribution over α is our primary target. As static correlations of the system are entirely independent of α , calculating its distribution is not possible solely through the techniques developed in section 4.5 – we require a generalisation using the moment theory developed in this chapter.

The derived equations 5.27-5.31 dictate the relationship between moments of a Langevin system and those of a Born-Oppenheimer system under the same conditions. In section 5.2.3 we also argued that, as expectations over coordinates and velocities are independent of α , moments vary with α only where α or σ appear explicitly. Up to fourth order, then, the non-vanishing derivatives we require are simply:

$$\frac{\partial M_3}{\partial \alpha} = -M_2^{BO} = -\omega_0^2, \quad (5.99)$$

$$\frac{\partial M_4}{\partial \alpha} = 2\alpha M_2^{BO} = 2\alpha\omega_0^2. \quad (5.100)$$

As these quantities are linearly proportional to lower moments, it is immediately clear that they too can be calculated in $O(N^2)$ time.

Gradients with respect to parameters defining inter-particle potentials can be found as previously (via the general error inversion technique of section 4.5). We can also consider variation with respect to other system quantities, such as the temperature, with little extra effort. For a general property $F = E_W[f]$ that is an average over coordinate configurations - Born-Oppenheimer moments and static structure factors alike - we can write:

$$\frac{\partial F}{\partial \beta} = E_W[f] E_W[V] - E_W[Vf], \quad (5.101)$$

where V is the total potential energy of the system. Equipped with these gradients, we can move on to sampling.

5.3.2 Application: Yukawa+SRR with Langevin Noise Model

We now delineate a complete implementation of the numerical techniques discussed in this chapter, in order to obtain detailed insight into a particulate system

under WDM-like conditions. Combining the various elements of preceding sections, we can describe a general inference procedure as follows:

- Sample inter-ion potentials and dynamic parameters using the static structure factor and dynamic structure factor moment expansions, and an SGMCMC sampler
- Extract the target property from each of these samples; this may be done directly if possible, or else simulations can be created for each sampled parameter set, and from each of these a target property calculated
- Infer a posterior distribution over the target property from the samples

As is typical of experimental input data, we will assume knowledge of the static structure factor and dynamic structure factor moments (and associated error bars on these quantities). As higher moments of the DSF have increasingly large error bars, we restrict our analysis to the first few moments only – measurements with large error bars do not act as significant constraints on our model systems, as they represent very small contributions to the likelihood function. The used experimental inputs are illustrated in table 5.1.

For our example case, we work directly with DSF moment values and error bars from (synthetic) experiment, which we assume would be accessible from integrals over the DSF extracted from real experimental data. As a side point, though, we note that the resolution of the DSF may be so poor for some experimental data that these integrals are no longer meaningful. In these cases, we can apply the same form of analysis with an extended forward model – rather than working with moments directly, we can use full Molecular Dynamics to determine the full DSF for comparison, or employ an ansatz for the memory function. The cost of this is increased computational expense (for MD), or reduced inference accuracy due to an

Measurement	Value	Error Bar
M_2	1	3%
M_4	2.083	3%

Table 5.1: The assumed experimental input data and associated error bar standard deviations (assumed Gaussian) for our demonstration of dynamic property inference.

additional forward model approximation (the choice of ansatz), but the procedure is much the same.

We model the system as an ensemble of Yukawa+SRR representative systems, with Langevin noise representing electron-ion dynamic interaction, at fixed density. We further assume a fixed potential function for simplicity (which in general can be inferred via the methods of the previous chapter). Our primary aim is to find temperature and stochastic parameter distributions implied by our experimental data; fig. 5.2 displays the marginalised posterior distributions over these parameters of the model, as determined by SGLD in combination with an approximately unbiased gradient estimator (extracted with Classical Monte Carlo – see caption). These distributions provide an immediate insight into the system and its properties, and lay the groundwork for possible extensions, discussed below.

5.3.3 Extension to Arbitrary Properties: a Brief Outline

In addition to directly sampled properties (i.e. temperature, potential parameters, α , etc.), comprehensive system properties implied by our data can be straightforwardly obtained using the posterior distributions we have derived. Selecting points at random from our temperature, potential and α distributions, we can in principle construct a range of Molecular Dynamics simulations with properties that are consistent with the data (under the assumption that the Yukawa+SRR with Langevin dynamics model is sufficient to represent the system for well-chosen pa-




Figure 5.2: The sampled posterior distribution over the system (inverse) temperature β and the friction parameter α . The sampled distribution reflects the (approximate) relative probability of each combination of temperature and α of our Yukawa+SRR system, as implied by the measured moment values and error bars in table 5.1. The likelihood gradients required for each step of our sampling method (SGLD) were calculated through CMC (using the moment expressions derived in this chapter).

rameters). The distributions over completely arbitrary, dynamical system properties can then be sampled (via extraction of these properties to be sampled from the individual MD results).

5.3.4 Extension to Mixed States: a Brief Outline

Suppose our focus is on extracting, say, density and temperature from experimental data. In this case, we can extend our analysis of experimental scattering data (via the DSF) to mixed states: It is in fact straightforward to account for scattering from a range of different temperatures and densities. We make the assumption that the total observed scattering signal is a weighted superposition of scattering signals from individual density/temperature (ρ/T) states, i.e.:

$$S(q, \omega) = \sum_{\rho} \sum_T S_{\rho,T}(q, \omega). \quad (5.102)$$

We then extend our our definition of the system weights W : rather than containing density and temperature alone, W becomes a two-dimensional grid of weights in density-temperature space, each element dictating the quantity of material at a given density and temperature. The range of temperatures and densities over the grid considered can be selected to encompass all feasible experimental conditions. Sampling can then be performed as previously, the forward model now a linear sum over signals for each possible density-temperature pair; the signals at each set of physical parameters can be calculated prior to sampling.

The practicality of this approach is limited by the need to calculate scattering signals over a large grid of feasible parameters (and the much larger dimension of the sample space). This can be ameliorated to some extent through creating unbiased statistical estimates of the signal error for a given W . To do so, we can draw a

finite number of samples from the distribution over density and temperature implied by W , and combine them into an unbiased estimator of the total error in the signal. This is completely analogous to the single density-temperature case, but with samples drawn using densities and temperatures sampled from W .

Further simplification is also possible if some constraints can be set on the temperature and density profiles of the system. Assuming, for example, that Gaussian distributions of densities and temperatures provide good models of the system under scrutiny, we can parametrise W with only a mean and covariance in temperature-density space.

In either case, our samples of W can allow us to establish posterior distributions over mixture components, when treatment of a system as a single temperature/density is not appropriate; this opens an avenue for further development and application to experimental systems.

6

Future Directions and Summary

Based on the experience of my life, which I have not exactly hit out of the park, I tend to agree with that thing about, If it's not broke, don't fix it. And would go even further to: Even if it is broke, leave it alone, you'll probably make it worse.

My Chivalric Fiasco
George Saunders

6.1 Introduction

The various methods developed throughout chapters 3, 4, and 5 have allowed us to establish an end-to-end procedure for calculating quantum plasma properties – both predictively and by exploiting experimental measurements. However, there are multiple routes to extending these methods that are worth exploring in the future, for both broadening their realm of applicability and improving their computational efficiency.

This chapter considers multiple possible extensions to the core work of this thesis. In doing so, it showcases several algorithmic extensions – while the theory of these extensions is developed and may prove a useful contribution in its own right, complete implementation and validation is left to future research.

The chapter begins with an outline of possible future directions for the Bohmian trajectory method of chapter 3, and follows this up with avenues for extending the MCMC work of chapters 4 and 5. Finally, the overall contribution of the work in this thesis is summarised.

6.2 Extending (or Overhauling) the Bohm Approach

6.2.1 Configuration Monte Carlo

First, let's consider a direct addition to the overall algorithm of chapter 3. As an alternative to our search for a λ factor arising from the linearisation of the Bohm potential in section 3.3.4, we may consider a Monte Carlo scheme to sample valid sets $\{g_\alpha(r)\}$ of pair correlations with correct (Boltzmann) probability. This has the potential to remove the need for simultaneous DFT simulation, at a cost of a more computationally intensive Monte Carlo approach. While this approach will require significantly more work before it can be robustly implemented, we present an outline of its core components here.

The configuration potential energy can be expressed in terms of $\{g_\alpha(r)\}$ and the system volume Ω under the previously used pair potential approximation:

$$E[\{g_\alpha\}] = \Omega \sum_{i,j} \pi \rho_i \rho_j \int r^2 u_{ij}(r|g_{ij}) g_{ij}(r) dr. \quad (6.1)$$

For a Monte Carlo step in which we effect the movement of a single particle, provided the r domain is sufficiently large that correlations can be neglected at r_{max} , the change in energy becomes:

$$\Delta E = \Omega \sum_{i,j} \pi \rho_i \rho_j \int_0^{r_{max}} r^2 \Delta U_{ij}(r) dr, \quad (6.2)$$

where $U_{ij}(r) = u_{ij}(r|g_{ij})g_{ij}(r)$. Given the species index of the moved particle is k , the update corresponds to a change in g_{ij} through:

$$g_{ij} \rightarrow g_{ij} + \frac{\Delta G_{ij}}{N_k}, \quad (6.3)$$

6 Future Directions and Summary

where $\Delta G_{ij}(r)$ is the change in pair correlation calculated relative to the moved particle only (and is zero when $k \notin \{i, j\}$). ΔU_{ij} can then be expanded quite generally in terms of ΔG_{ij} and its derivatives, to obtain a volume-independent (and valid in the thermodynamic limit) expression for ΔE . For computational purposes, however, we simulate a finite volume, so that Ω is well-defined. Hence, we use directly eq. 6.1 using g_{ij} both before and after application of the map 6.3.

This non-linear Monte Carlo approach is then similar to that in the linearised case, but rather than evaluating the single particle energy change associated with each move, the energy functional change of eq. 6.2 is used directly, or else the energy is fully re-evaluated for all particles following the potential change implied by map 6.3. For a given temperature, we are able to sample the non-linear pair density distribution:

$$P[\{g_\alpha\}] = \frac{G_\alpha \exp(-\beta E[\{g_\alpha\}])}{\sum_{\alpha'} G_{\alpha'} \exp(-\beta E[\{g_{\alpha'}\}])}, \quad (6.4)$$

where the summation is over all pair correlations realizable by the system, with degeneracy factor G_α taken care of as a consequence of sampling from a particulate system.

The principle approximation in this approach is then the structure of the energy functional as described by eq. 6.1. The distribution of eq. 6.4 exactly replicates that of the real quantum system in the case that the functional E is such that $P[\{g_\alpha\}]$ becomes equivalent to:

$$\begin{cases} \frac{G_{\alpha,exact} \exp(-\beta E_{exact}[\{g_\alpha\}])}{\sum_{\gamma} G_{\gamma,exact} \exp(-\beta E_{exact}[\{g_\gamma\}])} & \{g_\alpha\} \in \{\{g_\gamma\}\} \\ 0 & otherwise \end{cases}, \quad (6.5)$$

where the summation over γ is over the true eigenstates of the system, inclusive of symmetry considerations. Again, we use the addition of the effective Pauli potential as described by eq. 3.28 to obtain a first approximation to the correction to the functional E . From application of this non-linear Monte Carlo simulation we can determine the mean thermal pair correlations and associated averages, including effective pair potentials for use in subsequent dynamic simulation, as previously obtained from the linearised theory.

6.2.2 Conditional Wavefunction Sampling

While we have discussed multiple classical trajectory models of wDM, this thesis has not fully explored the possibility of combining classical trajectories with associated quantum properties. Rather than simply extending our Bohmian method, then, we could consider a more substantial reworking of it, through augmenting a trajectory of each simulation species with a conditional wavefunction (as discussed in section 3.2.1). This provides a computationally inexpensive enhancement to the structural pair potential approach, and offers an interesting route for building on Classical Map methods.

The evolution of the conditional wavefunction of a single particle of each species ('tracker' particles) can be determined by exploiting the equation of motion 3.9, with the corresponding particles having their trajectory dynamics dictated by the fixed phase gradient to velocity relationship. The remaining particles in the system ('background' particles) can then be approximately treated using effective pair potentials, with input pair correlations for constructing these potentials extracted from the wavefunction-augmented particles.

Consider first the background particles, which are treated with classical potentials. As a first approximation, we may use standard Quantum Statistical Po-

6 Future Directions and Summary

tentials for the pair interactions, as outlined in section 2.3.7. Then, in a similar procedure to that in section 3.4.2, we may iteratively improve the used potentials after a period of simulation. To do so, we can simply examine the pair correlations as measured relative to the tracker particles only, and then invert the Hypernetted Chain equations to obtain effective statistical potentials for use between background particles. As such, the interactions between background particles can be improved self-consistently until correlations between background particles become equal to those experienced by the tracker particles.

We then consider the conditional wavefunction potential terms G_a , J_a , as defined by equations 3.7, 3.8, in order to determine the equations of motion for the tracker particles. To simplify these terms we first insert a zero-order Taylor expansion approximation for the background particle velocity fields. That is, we take $\partial S/\partial x_b(x_a, x_b, t) \approx \partial S/\partial x_b(x_a, x_b[t], t) = mv_b(t)$; this yields immediately $J_a \approx 0$. This follows along fairly similar lines to the original successful approach of Oriols[54], wherein the potential terms themselves are evaluated at this point. The remaining G_a contribution then becomes:

$$G_a = U_b + \sum_{k \neq a} \left\{ V_{B,k} - \frac{1}{2} m_k v_k^2 \right\}, \quad (6.6)$$

where the velocity-dependent term can further be neglected, as it produces only a global phase.

It remains to approximate the Bohm potential term $V_{B,k}$. In order to do this, we can consider again the Quantum Hamilton-Jacobi equation, eq. 3.3. From this, it is trivial to see that the accelerating ‘force’ on each quantum trajectory is due to the (negative) gradient of the real potential U and Bohm potential only. As such, we can identify an approximation for the Bohm potential as simply the applied statistical potential for the background particles (i.e. the effective potential used can

be properly treated as the ‘real’ energy entering into G_a). In order to perform this substitution when employing statistical pair potentials, we use the corresponding equality

$$\frac{1}{2} \sum_i \sum_{j \neq i} V_{QSP}(x_i, x_j) = \sum_i V_{B,i}. \quad (6.7)$$

Note that there is an ambiguity here in what proportion of the energy associated with each pair interaction to allocate to the Bohm potential of each particle; remaining consistent with the proportions used in the structural pair potential closure approach employed in sections 3.4.2, 3.3.1, implicit in eq. 3.27, the allocation to each particle in a pair interaction can be weighted as proportional to the inverse masses. Then, for $V_{B,i}$, we have:

$$V_{B,i} = \sum_{j \neq i} \frac{m_i^{-1}}{m_i^{-1} + m_j^{-1}} V_{QSP}(x_i, x_j). \quad (6.8)$$

The system is then closed, up to iterative self-consistent improvement of the effective background pair potentials; however, exchange effects require additional consideration. These effects can be treated with an effective Hamiltonian approach, analogous to the effective Pauli repulsion potential used in section 3.3.2 when treating the particles (semi-)classically. For the non-interacting electron gas, we have known N-particle correlation ρ_N . Given the positions of the other $N - 1$ particles x_b , a conditional wavefunction of a single electron in the non-interacting gas with correct probability distribution can be written as:

$$\psi_a^0(x_a, t) = A \sqrt{\rho_N(x_a, x_b[t])}, \quad (6.9)$$

with A a normalisation factor. This can be inverted using eq. 3.9 to give an effective Hamiltonian contribution of:

$$H_P = \frac{\hbar^2}{2m} \frac{\partial^2 \psi_a^0}{\psi_a^0}. \quad (6.10)$$

This can be viewed as the zero-temperature quantum form of the approximation seen in the classical potential of eq. 3.28.

6.2.3 Other Possibilities

While not as developed as the algorithms of the preceding sections, there are several other research paths worth considering for extending the Bohmian approach. In particular, an extension of the approach to non-equilibrium dynamics would open up analysis of a far greater range of quantum plasma systems. While a complete, robust treatment will require further development and validation, a first step towards treatment of non-equilibrium system could be constructed by ‘stitching together’ equilibrium results at different temperatures and densities.

By determining Bohmian potentials for a range of different conditions, it would be feasible to treat a non-equilibrium system – approximately – with a combination of potentials applicable to different densities and temperatures. A full simulation could then be constructed with Bohm potentials varying with position, according to local temperature and density conditions. Determining derivatives of the Bohm potential with respect to density and temperature, too, via Automatic Differentiation, could simplify the process of determining (equilibrium) Bohm potentials across a range of parameters (not to mention the possibility of accelerating the Bohm method via gradient-based optimisation).

Alongside this is the possibility of alternative thermostats. We have developed approximate Fermi thermostats to emulate electronic behaviour, but these are not the only reasonable options for non-classical equations of motion. It would be well worth exploring the effect of other modified thermostats for quantum trajectories, to more formally address the problem of generalising Classical Map methods to include dynamic modifications (including the possibility of non-Markovian equations of motion along the lines of eq. 5.97, and the possibility of treating the coupling between momentum and real space implied by the full, exact Fermi distribution). In addition, we could consider more sophisticated target momentum distributions (beyond those of a non-interacting electron gas) through approximate DFT calculations of electronic states.

If we are willing to consider a more substantial overhaul of the Bohmian approach, we could also consider broadening the scope of the method's theoretical basis. Thus far, we've relied on an approximate coupling of the system to a heat bath; this requirement emerged as a result of our starting point – a non-thermal, pure state. It is worth asking the question: Is it possible to begin with a fully thermal description of Bohmian mechanics, and progress from there? For very simple systems, the answer is certainly yes. Considering the density matrix of a one-dimensional system, for example, we can follow ref. [128] to establish dynamical equations for trajectories in two dimensions (i.e. in the space of the density matrix). This will yield an analogue of the Bohmian mechanical equations, this time for thermal systems – we can then attempt to construct practical three-dimensional equations of motion using this formalism as an approximate basis.

Beyond these modifications, we can also consider acceleration through emulation of key steps of our simulations. We have already alluded to machine learning emulation of likelihood calculations for the inverse problem, but we can also con-

sider emulation of individual steps in our forward calculations. This could take the form of replacing force calculations in our MD step with fast approximate methods, for example; where neural networks are employed, we can account for statistical uncertainty in our emulated simulation step via Bayesian techniques[129]. The full range of options here is beyond the scope of this discussion, but emulation techniques in general are well worth exploring for (further) accelerating the simulations we have discussed.

Finally, we can consider semi-classical models more generally, via use of the inverse algorithms we have already developed. Suppose we have access to static DFT simulation data for an electron-ion system (or, indeed, static data from any of the simulation schemes we've examined). Assuming both the electron-ion and ion-ion static correlations can be extracted with reasonable accuracy, we may exploit the Ornstein-Zernike equation and Hypernetted Chain closure to 'inject' dynamic, semi-classical electrons. This is simply a matter of using the generalised IMC methods developed in chapter 4 with fixed knowledge of electron-ion and ion-ion correlations and the ion-ion potential. Through this extension, effective (classical) ion-electron and electron-electron potentials could be determined for use in Molecular Dynamics. The implementation and validation of this technique could prove a nice (and very simple to implement) baseline for further development of predictive semi-classical simulations.

6.3 Extending the Inverse Problem Approach

6.3.1 Further Exploiting Connections to MD

Clearly, there are strong links between the MD and SGMCMC methods we've discussed – for the most part, SGMCMC has been treated as an extension of MD to parameter space. As such, there are tools that have been developed for MD that it may be worth applying to SGMCMC in future work that we have not fully explored here. Alongside use of correlation metrics that are typically applied to MD, we also have the option to exploit different thermostats and equations of motion (e.g. equations of motion incorporating inter-walker potentials) for SGMCMC – this is discussed in general terms in ref. [95], and is worth researching further for the contexts we have considered.

6.3.2 Underlying Estimator Improvements

In chapters 4 and 5 we have discussed Russian Roulette techniques for debiasing estimators. Again, as with the thermostats we developed in chapter 3, we have considered just one option from many possibilities, and it would be worth examining other options with an eye to minimising estimator variance (e.g. the examples developed in ref. [130]). Alternative debiasing methods will also, in turn, open up new possibilities for multifidelity composition methods to extend the work of section 4.3.2.

6.3.3 Other Possibilities

Chapter 5 considered a general purpose method for DSF inversion via its moments; while we demonstrated application of this approach to a Yukawa+SRR model, it

would also be interesting to consider application to chapter 3's Bohmian dynamics model. To do so would be a fairly straightforward extension of the theory already developed, requiring only an adaptation of the evolution operator (used to construct moments) to account for the non-classical electron thermostat.

Applying a moment approach to the Bohmian model would be particularly valuable due to the model's explicit inclusion of electrons. Whereas MD must deal with the separation of electronic and ionic timescales, extracting dynamical property moments via CMC avoids this issue by marginalising over particle velocities – as there is no need to directly model the dynamics, there is no need to decrease a corresponding simulation time step to accommodate the small electron mass.

Finally, we recall that at the tail end of chapter 5 we explained two key extensions for future work – namely applying the methods we have developed to arbitrary system properties and mixed states. These two proposed extensions would not require a substantial degree of additional setup, and could be validated fairly quickly with simple additions to our existing sampling pipeline. In particular, extension to arbitrary properties can be carried out using existing sampling data with little additional statistical work, and – depending on the properties explored – would be a worthwhile research option for the near future.

6.4 Summary

Our goal in this thesis has been to develop and validate theoretical and computational methods for analysing quantum plasmas and, in particular, Warm Dense Matter. To accomplish this goal, we have considered analysis from two complementary angles – prediction of WDM properties with novel *ab initio* simulations, and inference of WDM properties through a combination of *ab initio* simulations and Monte Carlo techniques. Our results, in summation, provide a comprehensive toolset for designing and interpreting WDM systems.

In chapter 3, we developed a new approach to predictive WDM simulations. Fundamentally based on Bohm’s theory of quantum mechanics, our approach drew on approximation methods present in Quantum Hydrodynamics and Classical Map techniques in order to incorporate fully dynamic electrons – while still being fast enough to resolve the ion-ion Dynamic Structure Factor. We demonstrated the effectiveness of our approach in multiple WDM test cases, and have highlighted many potential avenues for developing this work further – perhaps most notably, extension to non-equilibrium systems – that are under active investigation.

Chapter 4 took existing Stochastic Gradient Markov Chain Monte Carlo methods, and, with a collection of novel general and domain-specific extensions, applied them to key static inference problems for WDM systems. The chapter developed strategies for inference of electromagnetic fields in plasma systems – and their error bars – alongside effective potential inference for particulate systems.

Finally, chapter 5 developed a methodology for full inference of dynamic properties of WDM from DSF and ISF data. The chapter combined the results of chapters 2, 3 and 4 with a stochastic moment expansion of the DSF, rounding out our complete treatment of both the forward and inverse problems of WDM analysis.

6 *Future Directions and Summary*

While a complete, exact treatment of WDM systems remains far beyond current theoretical and computational capabilities, the methods we have developed should prove to be invaluable tools for future WDM calculations. In the case of the Bohmian method, we have a method that is capable of predicting properties of WDM without neglecting crucial dynamics in the electron-ion interaction; in the case of chapters 4 and 5, we have a range of methods that extend – and make robust – the range of WDM properties that can be extracted from experiment. Alongside these developments, we expect that the extensions suggested in this chapter have the potential to close the gap between simulation and reality yet further.

Bibliography

- [1] B. Larder, D. O. Gericke, S. Richardson, P. Mabey, T. White, and G. Gregori, “Fast nonadiabatic dynamics of many-body quantum systems,” *Science Advances*, vol. 5, 2019.
- [2] G. Gregori, B. Reville, and B. Larder, “Modified Friedmann equations via conformal Bohm-de Broglie gravity,” *The Astrophysical Journal*, vol. 886, p. 50, nov 2019.
- [3] A. I. Morozov, *Introduction to plasma dynamics*. CRC Press, 2012.
- [4] F. Haas, “An introduction to quantum plasmas,” *Brazilian Journal of Physics*, vol. 41, pp. 349–363, Dec 2011.
- [5] T. Guillot, “Interiors of Giant Planets Inside and Outside the Solar System,” *Science*, vol. 286, no. 5437, pp. 72–77, 1999.
- [6] G. I. Kerley, “Equation of state and phase diagram of dense hydrogen,” *Physics of the Earth and Planetary Interiors*, vol. 6, no. 1-3, pp. 78–82, 1972.
- [7] R. Ernstorfer, M. Harb, C. T. Hebeisen, G. Sciaini, T. Dartigalongue, and R. J. D. Miller, “The Formation of Warm Dense Matter: Experimental Evidence for Electronic Bond Hardening in Gold,” *Science*, vol. 323, no. 5917, pp. 1033–1037, 2009.
- [8] S. H. Glenzer, B. J. MacGowan, P. Michel, N. B. Meezan, L. J. Suter, S. N. Dixit, J. L. Kline, G. A. Kyrala, D. K. Bradley, D. A. Callahan, E. L. Dewald, L. Divol, E. Dzenitis, M. J. Edwards, A. V. Hamza, C. A. Haynam, D. E. Hinkel, D. H. Kalantar, J. D. Kilkenny, O. L. Landen, J. D. Lindl, S. LePape, J. D. Moody, A. Nikroo, T. Parham, M. B. Schneider, R. P. J. Town, P. Wegner, K. Widmann, P. Whitman, B. K. F. Young, B. Van Wonterghem, L. J. Ather-ton, and E. I. Moses, “Symmetric Inertial Confinement Fusion Implosions at Ultra-High Laser Energies,” *Science*, vol. 327, no. 5970, pp. 1228–1231, 2010.
- [9] S. X. Hu, V. N. Goncharov, T. R. Boehly, R. L. McCrory, S. Skupsky, L. A. Collins, J. D. Kress, and B. Militzer, “Impact of first-principles properties of

- deuterium-tritium on inertial confinement fusion target designs,” *Physics of Plasmas*, vol. 22, no. 5, 2015.
- [10] S. Hu, L. Collins, T. Boehly, Y. Ding, P. B. Radha, V. Goncharov, V. Karasiev, G. Collins, S. Regan, and M. Campbell, “A review on ab initio studies of static, transport, and optical properties of polystyrene under extreme conditions for inertial confinement fusion applications,” *Physics of Plasmas*, vol. 25, 05 2018.
- [11] E. García Saiz, G. Gregori, D. O. Gericke, J. Vorberger, B. Barbrel, R. J. Clarke, R. R. Freeman, S. H. Glenzer, F. Y. Khattak, M. Koenig, O. L. Landen, D. Neely, P. Neumayer, M. M. Notley, A. Pelka, D. Price, M. Roth, M. Schollmeier, C. Spindloe, R. L. Weber, L. van Woerkom, K. Wünsch, and D. Riley, “Probing warm dense lithium by inelastic X-ray scattering,” *Nature Physics*, vol. 4, no. 12, pp. 940–944, 2008.
- [12] S. H. Glenzer and R. Redmer, “X-ray Thomson scattering in high energy density plasmas,” *Reviews of Modern Physics*, vol. 81, no. 4, pp. 1625–1663, 2009.
- [13] L. B. Fletcher, H. J. Lee, T. Döppner, E. Galtier, B. Nagler, P. Heimann, C. Fortmann, S. LePape, T. Ma, M. Millot, A. Pak, D. Turnbull, D. A. Chapman, D. O. Gericke, J. Vorberger, T. White, G. Gregori, M. Wei, B. Barbrel, R. W. Falcone, C.-C. Kao, H. Nuhn, J. Welch, U. Zastra, P. Neumayer, J. B. Hastings, and S. H. Glenzer, “Ultrabright X-ray laser scattering for dynamic warm dense matter physics,” *Nature Photonics*, vol. 9, no. 4, pp. 274–279, 2015.
- [14] T. G. White, S. Richardson, B. J. B. Crowley, L. K. Pattison, J. W. O. Harris, and G. Gregori, “Orbital-free density-functional theory simulations of the dynamic structure factor of warm dense aluminum,” *Physical Review Letters*, vol. 111, no. 17, 2013.
- [15] P. Mabey, S. Richardson, T. G. White, L. B. Fletcher, S. H. Glenzer, N. J. Hartley, J. Vorberger, D. O. Gericke, and G. Gregori, “A strong diffusive ion mode in dense ionized matter predicted by Langevin dynamics,” *Nature Communications*, vol. 8, p. 14125, 2017.
- [16] J. Chihara, “Interaction of photons with plasmas and liquid metals: photoabsorption and scattering,” *J. Phys.: Condens. Matter*, vol. 12, p. 231, 2000.
- [17] K. Sturm, “Dynamic Structure Factor: An Introduction,” *Zeitschrift für Naturforschung A*, vol. 48, pp. 233–242, 1993.
- [18] G. F. Giuliani and G. Vignale, *Quantum Theory of the Electron Liquid*. Cambridge University Press, 2005.

- [19] C. R. D. Brown, D. O. Gericke, M. Cammarata, B. I. Cho, T. Döppner, K. Engelhorn, E. Förster, C. Fortmann, D. Fritz, E. Galtier, S. H. Glenzer, M. Harmand, P. Heimann, N. L. Kugland, D. Q. Lamb, H. J. Lee, R. W. Lee, H. Lemke, M. Makita, A. Moinard, C. D. Murphy, B. Nagler, P. Neumayer, K.-U. Plagemann, R. Redmer, D. Riley, F. B. Rosmej, P. Sperling, S. Toleikis, S. M. Vinko, J. Vorberger, S. White, T. G. White, K. Wünsch, U. Zastra, D. Zhu, T. Tschentscher, and G. Gregori, “Evidence for a glassy state in strongly driven carbon,” *Scientific Reports*, vol. 4, pp. 1–6, 2014.
- [20] G. Gregori and D. O. Gericke, “Low frequency structural dynamics of warm dense matter,” *Physics of Plasmas*, vol. 16, no. 5, 2009.
- [21] D. Bohm, F. D. Filosoja, C. Letras, U. D. S. Paulo, and S. Paulo, “A Collective Description of-Electron Interactions: III. Coulomb Interactions,” vol. 92, no. 3, pp. 609–625, 1953.
- [22] A. Schindlmayr and R. W. Godby, “Systematic vertex corrections through iterative solution of hedin’s equations beyond the gw approximation,” *Physical review letters*, vol. 80, no. 8, p. 1702, 1998.
- [23] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids*. Oxford University Press, 2017.
- [24] W. K. Hastings, “Monte carlo sampling methods using markov chains and their applications,” *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.
- [25] J. Hansen and I. R. McDonald, *Theory of Simple Liquids. With Applications to Soft Matter*, vol. 10. Academic Press, 2013.
- [26] J. Dufty and S. Dutta, “Classical representation of a quantum system at equilibrium: Theory,” *Physical Review E*, vol. 87, no. 3, p. 032101, 2013.
- [27] V. Schwarz, T. Bornath, W.-D. Kraeft, S. Glenzer, A. Höll, and R. Redmer, “Hypernetted chain calculations for two-component plasmas,” *Contributions to Plasma Physics*, vol. 47, no. 4-5, pp. 324–330, 2007.
- [28] T. Pang, *An Introduction to Quantum Monte Carlo Methods*. 2053-2571, Morgan Claypool Publishers, 2016.
- [29] R. Fantoni, “Radial distribution function in a diffusion Monte Carlo simulation of a Fermion fluid between the ideal gas and the Jellium model,” *European Physical Journal B*, vol. 86, no. 6, 2013.
- [30] J. B. Anderson, “Quantum chemistry by random walk,” *The Journal of Chemical Physics*, vol. 65, no. 10, pp. 4121–4127, 1976.
- [31] D. M. Ceperley and B. J. Alder, “Ground state of the electron gas by a stochastic method,” *Physical Review Letters*, vol. 45, no. 7, pp. 566–569, 1980.

Bibliography

- [32] E. W. Brown, B. K. Clark, J. L. Dubois, and D. M. Ceperley, "Path-integral Monte Carlo simulation of the warm dense homogeneous electron gas," *Physical Review Letters*, vol. 110, no. 14, pp. 1–5, 2013.
- [33] D. M. Ceperley, "Path integrals in the theory of condensed helium," *Reviews of Modern Physics*, vol. 67, no. 2, pp. 279–355, 1995.
- [34] T. Dornheim, S. Groth, J. Vorberger, and M. Bonitz, "Ab initio path integral monte carlo results for the dynamic structure factor of correlated electrons: From the electron liquid to warm dense matter," *Phys. Rev. Lett.*, vol. 121, p. 255001, Dec 2018.
- [35] S. Zhang, K. Driver, F. Soubiran, and B. Militzer, "Path Integral Monte Carlo Simulations of Warm Dense Sodium," *High Energy Density Physics*, vol. 21, pp. 19–22, 2016.
- [36] B. B. Witte, M. Shihab, S. H. Glenzer, and R. Redmer, "Ab initio simulations of the dynamic ion structure factor of warm dense lithium," *Physical Review B - Condensed Matter and Materials Physics*, vol. 95, no. 14, pp. 1–9, 2017.
- [37] M. W. C. Dharma-wardana, "Static and dynamic conductivity of warm dense matter within a density-functional approach: Application to aluminum and gold," *Phys. Rev. E*, vol. 73, p. 036401, Mar 2006.
- [38] R. Bauernschmitt and R. Ahlrichs, "Treatment of electronic excitations within the adiabatic approximation of time dependent density functional theory.," *Chemical Physics Letters*, vol. 256, no. 4-5, pp. 454–464, 1996.
- [39] K. Burke, "Perspective on density functional theory," *Journal of Chemical Physics*, vol. 136, no. 15, 2012.
- [40] T. Ziegler, "Approximate Density Functional Theory as a Practical Tool in Molecular Energetics and Dynamics," *Chemical Reviews*, vol. 91, no. 5, pp. 651–667, 1991.
- [41] R. O. Jones, "Density functional theory: Its origins, rise to prominence, and future," *Rev. Mod. Phys.*, vol. 87, pp. 897–923, Aug 2015.
- [42] P. Hohenberg and W. Kohn, "Inhomogeneous electron gas," *Phys. Rev.*, vol. 136, pp. B864–B871, Nov 1964.
- [43] W. Kohn and L. J. Sham, "Self-consistent equations including exchange and correlation effects," *Physical Review*, vol. 140, no. 4A, 1965.
- [44] V. L. Lignères and E. A. Carter, *An Introduction to Orbital-Free Density Functional Theory*, pp. 137–148. Dordrecht: Springer Netherlands, 2005.

- [45] M. A. L. Marques and E. K. U. Gross, "Time-dependent density functional theory," *Annual review of physical chemistry*, vol. 55, no. 1, pp. 427–455, 2004.
- [46] X. Andrade, D. Strubbe, U. De Giovannini, A. H. Larsen, M. J. Oliveira, J. Alberdi-Rodriguez, A. Varas, I. Theophilou, N. Helbig, M. J. Verstraete, *et al.*, "Real-space grids and the octopus code as tools for the development of new simulation approaches for electronic systems," *Physical Chemistry Chemical Physics*, vol. 17, no. 47, pp. 31371–31396, 2015.
- [47] A. D. Baczewski, L. Shulenburger, M. P. Desjarlais, S. B. Hansen, and R. J. Magyar, "X-ray Thomson Scattering in Warm Dense Matter without the Chihara Decomposition," *Physical Review Letters*, vol. 116, no. 11, pp. 1–6, 2016.
- [48] L. Harbour, M. Dharma-Wardana, D. D. Klug, and L. J. Lewis, "Pair potentials for warm dense matter and their application to x-ray thomson scattering in aluminum and beryllium," *Physical Review E*, vol. 94, no. 5, p. 053211, 2016.
- [49] J. Vorberger, Z. Donko, I. Tkachenko, and D. Gericke, "Dynamic ion structure factor of warm dense matter," *Physical review letters*, vol. 109, no. 22, p. 225001, 2012.
- [50] A. V. Filinov, M. Bonitz, and W. Ebeling, "Improved Kelbg potential for correlated Coulomb systems," *Journal of Physics A: Mathematical and General*, vol. 36, no. 22, pp. 5957–5962, 2003.
- [51] M. Dharma-wardana, "Hot-electron relaxation in dense two-temperature hydrogen," *arXiv preprint arXiv:1705.10221*, 2017.
- [52] A. P. Lyubartsev, M. Karttunen, I. Vattulainen, and A. Laaksonen, "On coarse-graining by the inverse monte carlo method: Dissipative particle dynamics simulations made to a precise tool in soft matter modeling," *Soft Materials*, vol. 1, no. 1, pp. 121–137, 2002.
- [53] D. Bohm, "A suggested interpretation of the quantum theory in terms of "hidden" variables. I," *Physical Review*, vol. 85, no. 2, pp. 166–179, 1952.
- [54] X. Oriols, "Quantum-trajectory approach to time-dependent transport in mesoscopic systems with electron-electron interactions," *Physical Review Letters*, vol. 98, no. 6, pp. 1–4, 2007.
- [55] G. Manfredi, "How to model quantum plasmas," *Fields Inst. Commun*, vol. 46, pp. 263–287, 2005.
- [56] A. P. Lyubartsev and A. Laaksonen, "Calculation of effective interaction potentials from radial distribution functions: A reverse monte carlo approach," *Physical Review E*, vol. 52, no. 4, p. 3730, 1995.

Bibliography

- [57] F. Lado, "Effective potential description of the quantum ideal gases," *The Journal of Chemical Physics*, vol. 47, no. 12, pp. 5369–5375, 1967.
- [58] M. Dharma-Wardana and F. Perrot, "Simple classical mapping of the spin-polarized quantum electron gas: Distribution functions and local-field corrections," *Physical Review Letters*, vol. 84, no. 5, p. 959, 2000.
- [59] M. W. C. Dharma-wardana, "Pair-distribution functions of two-temperature two-mass systems: Comparison of molecular dynamics, classical-map hypernetted chain, quantum Monte Carlo, and Kohn-Sham calculations for dense hydrogen," *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 2008.
- [60] M. W. C. Dharma-Wardana, "The classical-map hyper-netted-chain (CHNC) method and associated novel density-functional techniques for warm dense matter," *International Journal of Quantum Chemistry*, vol. 112, no. 1, pp. 53–64, 2012.
- [61] S. Dutta and J. Dufty, "Classical representation of a quantum system at equilibrium: Applications," *Physical Review E*, vol. 87, no. 3, p. 032102, 2013.
- [62] S. Dutta and J. Dufty, "Uniform electron gas at warm, dense matter conditions," *EPL (Europhysics Letters)*, vol. 102, no. 6, p. 67005, 2013.
- [63] N. Troullier and J. L. Martins, "Efficient pseudopotentials for plane-wave calculations. II. Operators for fast iterative diagonalization," *Physical Review B*, vol. 43, no. 11, pp. 8861–8869, 1991.
- [64] R. Benguria, H. Brezis, and E. H. Lieb, "The Thomas-Fermi-von Weizsacker theory of atoms and molecules," *Communications in Mathematical Physics*, vol. 79, pp. 167–180, 1981.
- [65] D. Michta, F. Graziani, and M. Bonitz, "Quantum Hydrodynamics for Plasmas - A Thomas-Fermi Theory Perspective," *Contributions to Plasma Physics*, vol. 55, no. 6, pp. 437–443, 2015.
- [66] S. Nosé, "A molecular dynamics method for simulations in the canonical ensemble," *Molecular Physics*, vol. 52, no. 2, pp. 255–268, 1984.
- [67] J. R. Bingle, N. Chelikowsky, "Langevin molecular dynamics with quantum forces: Application to silicon clusters," *Physical Review B*, vol. 50, no. 16, pp. 764–770, 1994.
- [68] G. A. Pavliotis, *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations*, vol. 60. Springer, 2014.
- [69] G. J. Martyna and M. L. Klein, "Nose-Hoover chains: the canonical ensemble via continuous dynamics," *Physical review letters*, vol. 97, no. 4, p. 2643, 1992.

- [70] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. S. Seljebotn, and K. Smith, "Cython: The best of both worlds," *Computing in Science and Engineering*, vol. 13, no. 2, pp. 31–39, 2011.
- [71] J. P. Mithen, *Molecular dynamics simulations of the equilibrium dynamics of non-ideal plasmas*. PhD thesis, 2012.
- [72] S. Plimpton, "Fast Parallel Algorithms for Short-Range Molecular Dynamics," *Journal of Computational Physics*, vol. 117, no. June 1994, pp. 1–19, 1995.
- [73] X. Gonze, B. Amadon, P. M. Anglade, J. M. Beuken, F. Bottin, P. Boulanger, F. Bruneval, D. Caliste, R. Caracas, M. Cote, T. Deutsch, L. Genovese, P. Ghosez, M. Giantomassi, S. Goedecker, D. R. Hamann, P. Hermet, F. Jollet, G. Jomard, S. Leroux, M. Mancini, S. Mazevet, M. J. T. Oliveira, G. Onida, Y. Pouillon, T. Rangel, G. M. Rignanese, D. Sangalli, R. Shaltaf, M. Torrent, M. J. Verstraete, G. Zerah, and J. W. Zwanziger, "ABINIT: First-principles approach to material and nanosystem properties," *Computer Physics Communications*, vol. 180, no. 12, pp. 2582–2615, 2009.
- [74] M. Motta, D. M. Ceperley, G. K. L. Chan, J. A. Gomez, E. Gull, S. Guo, C. A. Jiménez-Hoyos, T. N. Lan, J. Li, F. Ma, A. J. Millis, N. V. Prokof'ev, U. Ray, G. E. Scuseria, S. Sorella, E. M. Stoudenmire, Q. Sun, I. S. Tupitsyn, S. R. White, D. Zgid, and S. Zhang, "Towards the solution of the many-electron problem in real materials: Equation of state of the hydrogen chain with state-of-the-art many-body methods," *Physical Review X*, vol. 7, no. 3, pp. 1–28, 2017.
- [75] G. Kresse and J. Hafner, "Ab initio molecular dynamics for liquid metals," *Physical Review B*, vol. 47, no. 1, pp. 558–561, 1993.
- [76] G. Kresse and J. Hafner, "Ab initio molecular-dynamics simulation of the liquid-metal-amorphous-semiconductor transition in germanium," *Physical Review B*, vol. 49, no. 20, pp. 14251–14269, 1994.
- [77] G. Kresse and J. Furthmüller, "Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set," *Physical review B*, vol. 54, no. 16, p. 11169, 1996.
- [78] G. Kresse and J. Furthmüller, "Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set," *Computational Materials Science*, vol. 6, pp. 15–50, 1996.
- [79] J. Enkovaara, C. Rostgaard, J. J. Mortensen, J. Chen, M. Duřak, L. Ferrighi, J. Gavnholt, C. Glinsvad, V. Haikola, H. A. Hansen, H. H. Kristoffersen, M. Kuisma, A. H. Larsen, L. Lehtovaara, M. Ljungberg, O. Lopez-Acevedo, P. G. Moses, J. Ojanen, T. Olsen, V. Petzold, N. A. Romero, J. Stausholm-Møller, M. Strange, G. A. Tritsarlis, M. Vanin, M. Walter, B. Hammer, H. Häkkinen, G. K. Madsen, R. M. Nieminen, J. K. Nørskov, M. Puska, T. T.

Bibliography

- Rantala, J. Schiøtz, K. S. Thygesen, and K. W. Jacobsen, “Electronic structure calculations with GPAW: A real-space implementation of the projector augmented-wave method,” *Journal of Physics Condensed Matter*, vol. 22, no. 25, 2010.
- [80] J. J. Mortensen, L. B. Hansen, and K. W. Jacobsen, “Real-space grid implementation of the projector augmented wave method,” *Physical Review B - Condensed Matter and Materials Physics*, vol. 71, no. 3, pp. 1–11, 2005.
- [81] A. Larsen, J. Mortensen, J. Blomqvist, and K. Jacobsen, “The atomic simulation environment—a Python library for working with atoms,” *Journal of Physics: Condensed Matter*, vol. 29, no. 27, p. 273002, 2017.
- [82] J. P. Perdew, K. Burke, and M. Ernzerhof, “Generalized Gradient Approximation Made Simple,” *Physical Review Letters*, vol. 77, no. 18, pp. 3865–3868, 1996.
- [83] K. A. Dill and J. L. MacCallum, “The protein-folding problem, 50 years on,” *Science*, vol. 338, no. 6110, pp. 1042–1046, 2012.
- [84] N. Poccia, M. Fratini, A. Ricci, G. Campi, L. Barba, A. Vittorini-Orgeas, G. Bianconi, G. Aeppli, and A. Bianconi, “Evolution and control of oxygen order in a cuprate superconductor,” *Nature materials*, vol. 10, no. 10, pp. 733–736, 2011.
- [85] M. Suga, F. Akita, M. Sugahara, M. Kubo, Y. Nakajima, T. Nakane, K. Yamashita, Y. Umena, M. Nakabayashi, T. Yamane, *et al.*, “Light-induced structural changes and the site of o= o bond formation in psii caught by xfel,” *Nature*, vol. 543, no. 7643, pp. 131–135, 2017.
- [86] M. F. Kasim, T. P. Galligan, J. Topp-Mugglestone, G. Gregori, and S. M. Vinko, “Inverse problem instabilities in large-scale modeling of matter in extreme conditions,” *Physics of Plasmas*, vol. 26, no. 11, p. 112706, 2019.
- [87] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, *Handbook of Markov Chain Monte Carlo*. Chapman Hall/CRC, 01 2011.
- [88] H. Haario, M. Laine, A. Mira, and E. Saksman, “Dram: efficient adaptive mcmc,” *Statistics and Computing*, vol. 16, pp. 339–354, 12 2006.
- [89] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, “emcee: the mcmc hammer,” *Publications of the Astronomical Society of the Pacific*, vol. 125, no. 925, p. 306, 2013.
- [90] G. O. Roberts and J. S. Rosenthal, “Optimal scaling of discrete approximations to langevin diffusions,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 60, no. 1, pp. 255–268, 1998.

- [91] M. Barborini, S. Sorella, and L. Guidoni, “Structural optimization by quantum monte carlo: investigating the low-lying excited states of ethylene,” *Journal of chemical theory and computation*, vol. 8, no. 4, pp. 1260–1269, 2012.
- [92] U. Ekstrom, L. Visscher, R. Bast, A. Thorvaldsen, and K. Ruud, “Arbitrary-order density functional response theory from automatic differentiation,” *Journal of Chemical Theory and Computation*, vol. 6, p. 1971, 2010.
- [93] T. Tamayo-Mendoza, C. Kreisbeck, R. Lindh, and A. Aspuru-Guzik, “Automatic differentiation in quantum chemistry with applications to fully variational hartree-fock,” *ACS Central Science*, vol. 4, pp. 559 – 566, 2018.
- [94] T. Isaac, N. Petra, G. Stadler, and O. Ghattas, “Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the antarctic ice sheet,” *Journal of Computational Physics*, vol. 296, pp. 348 – 368, 2015.
- [95] Y. W. Teh, A. H. Thiery, and S. Vollmer, “Consistency and fluctuations for stochastic gradient langevin dynamics,” *J. Mach. Learn. Res.*, vol. 17, pp. 7:1–7:33, 2016.
- [96] Y. Ma, Y.-A. Ma, T. Chen, and E. B. Fox, “A complete recipe for stochastic gradient mcmc,” in *NIPS*, 2015.
- [97] R. A. Fisher, “Moments and product moments of sampling distributions,” *Proceedings of the London Mathematical Society*, vol. 2, no. 1, pp. 199–238, 1930.
- [98] D. McLeish, “A general method for debiasing a monte carlo estimator,” *Monte Carlo Methods and Applications*, vol. 17, pp. 301 – 315, 2011.
- [99] A. Lyne, M. A. Girolami, Y. Atchadé, H. Strathmann, and D. Simpson, “On russian roulette estimates for bayesian inference with doubly-intractable likelihoods,” *Statistical Science*, vol. 30, pp. 443–467, 2015.
- [100] B. Peherstorfer, K. Willcox, and M. D. Gunzburger, “Survey of multifidelity methods in uncertainty propagation, inference, and optimization,” *SIAM Review*, vol. 60, pp. 550–591, 2018.
- [101] M. Welling and Y. W. Teh, “Bayesian learning via stochastic gradient langevin dynamics,” in *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML’11*, (Madison, WI, USA), p. 681–688, Omnipress, 2011.
- [102] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [103] M. F. Kasim, D. Watson-Parris, L. Deaconu, S. Oliver, P. Hatfield, D. H. Froula, G. Gregori, M. Jarvis, S. Khatiwala, J. Korenaga, J. Topp-Mugglestone, E. Viezzer, and S. M. Vinko, “Up to two billion times acceleration of scientific simulations with deep neural architecture search,” 2020.

Bibliography

- [104] J. Baker, P. Fearnhead, E. B. Fox, and C. Nemeth, “Control variates for stochastic gradient mcmc,” *Statistics and Computing*, vol. 29, no. 3, pp. 599–615, 2019.
- [105] K. A. Dubey, S. J. Reddi, S. Williamson, B. Póczos, A. J. Smola, and E. P. Xing, “Variance reduction in stochastic gradient langevin dynamics,” *Advances in neural information processing systems*, vol. 29, pp. 1154–1162, 2016.
- [106] H. Jeffreys, “An invariant form for the prior probability in estimation problems,” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 186, no. 1007, pp. 453–461, 1946.
- [107] N. Kugland, D. Ryutov, C. Plechaty, J. Ross, and H.-S. Park, “Invited article: Relation between electric and magnetic field structures and their proton-beam images,” *Review of Scientific Instruments*, vol. 83, no. 10, p. 101301, 2012.
- [108] A. F. A. Bott, C. Graziani, P. Tzeferacos, T. G. White, D. Q. Lamb, G. Gregori, and A. A. Schekochihin, “Proton imaging of stochastic magnetic fields,” *Journal of Plasma Physics*, vol. 83, no. 6, p. 905830614, 2017.
- [109] M. F. Kasim, A. F. A. Bott, P. Tzeferacos, D. Q. Lamb, G. Gregori, and S. M. Vinko, “Retrieving fields from proton radiography without source profiles,” *Phys. Rev. E*, vol. 100, p. 033208, Sep 2019.
- [110] M. G. Saunders and G. A. Voth, “Coarse-graining methods for computational biology,” *Annual Review of Biophysics*, vol. 42, no. 1, pp. 73–93, 2013. PMID: 23451897.
- [111] I. Billionis and N. Zabaras, “A stochastic optimization approach to coarse-graining using a relative-entropy framework,” *The Journal of Chemical Physics*, vol. 138, no. 4, p. 044313, 2013.
- [112] P. Liu, Q. Shi, H. Daumé III, and G. A. Voth, “A bayesian statistics approach to multiscale coarse graining,” *The Journal of Chemical Physics*, vol. 129, no. 21, p. 214114, 2008.
- [113] K. Farrell, J. T. Oden, and D. Faghihi, “A bayesian framework for adaptive selection, calibration, and validation of coarse-grained models of atomistic systems,” *Journal of Computational Physics*, vol. 295, pp. 189 – 208, 2015.
- [114] M. Schöberl, N. Zabaras, and P.-S. Koutsourelakis, “Predictive coarse-graining,” *J. Comput. Phys.*, vol. 333, pp. 49–77, 2017.
- [115] R. L. McGreevy and L. Pusztai, “Reverse monte carlo simulation: A new technique for the determination of disordered structures,” *Molecular Simulation*, vol. 1, no. 6, pp. 359–367, 1988.

- [116] D. Reith, M. Pütz, and F. Müller-Plathe, “Deriving effective mesoscale potentials from atomistic simulations,” *Journal of Computational Chemistry*, vol. 24, no. 13, pp. 1624–1636, 2003.
- [117] T. C. Moore, C. R. Iacovella, and C. McCabe, “Derivation of coarse-grained potentials via multistate iterative boltzmann inversion,” *The Journal of Chemical Physics*, vol. 140, no. 22, p. 224104, 2014.
- [118] A. Lyubartsev, A. Mirzoev, L. Chen, and A. Laaksonen, “Systematic coarse-graining of molecular models by the newton inversion method,” *Faraday discussions*, vol. 144, pp. 43–56; discussion 93, 01 2010.
- [119] R. Bansal and K. N. Pathak, “Sum rules and atomic correlations in classical liquids,” *Physical Review A*, vol. 9, no. 6, pp. 2773–2782, 1974.
- [120] B. Oksendal, *Stochastic Differential Equations (3rd Ed.): An Introduction with Applications*. Berlin, Heidelberg: Springer-Verlag, 1992.
- [121] D. Williams, *Probability with Martingales*. Cambridge mathematical textbooks, Cambridge University Press, 1991.
- [122] R. Zwanzig, “Memory effects in irreversible thermodynamics,” *Phys. Rev.*, vol. 124, pp. 983–992, Nov 1961.
- [123] A. V. Mokshin, “Relaxation Processes in Many Particle Systems – Recurrence Relations Approach,” *The interdisciplinary journal of Discontinuity, Non-linearity, and Complexity*, vol. 2, no. 1, pp. 43–56, 2013.
- [124] R. D. Mountain, “Single particle motions in liquids: Qualitative features of memory functions,” *Journal of research of the National Bureau of Standards. Section A, Physics and chemistry*, vol. 78A 3, pp. 413–420, 1974.
- [125] A. V. Mokshin and B. N. Galimzyanov, “Self-consistent description of local density dynamics in simple liquids. the case of molten lithium,” *Journal of Physics: Condensed Matter*, vol. 30, no. 8, p. 085102, 2018.
- [126] N. R. Arista, “Energy loss of ions in solids: Non-linear calculations for slow and swift ions,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, vol. 195, no. 1, pp. 91 – 105, 2002.
- [127] G. Jung, M. Hanke, and F. Schmid, “Generalized langevin dynamics: construction and numerical integration of non-markovian particle-based models,” *Soft Matter*, vol. 14, pp. 9368–9382, 2018.
- [128] J. B. Maddox and E. R. Bittner, “Quantum dissipation in unbounded systems,” *Phys. Rev. E*, vol. 65, p. 026143, Jan 2002.

Bibliography

- [129] A. G. Wilson and P. Izmailov, “Bayesian deep learning and a probabilistic perspective of generalization,” *ArXiv*, vol. abs/2002.08791, 2020.
- [130] A. Beatson and R. P. Adams, “Efficient optimization of loops and limits with randomized telescoping sums,” *ArXiv*, vol. abs/1905.07006, 2019.