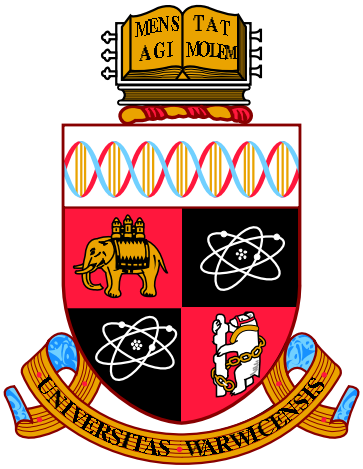

Uncertainty Quantification for Classical Effective Potentials

by

Sarah Longbottom

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

School of Engineering

November 2019

Contents

List of Tables

List of Figures

Acknowledgments

There are many people I would like to thank for their support and patience throughout the course of this work. The first of whom is my supervisor Dr. Peter Brommer, whose guidance and knowledge I am incredibly grateful for. Not only have I been shown a great deal of patience whilst trying (always long-windedly) to explain my thinking, but also a greatly appreciated, empathetic understanding of personal challenges along the way.

Throughout the course of four years within the Engineering department, I'm thankful to have shared my experience with Punit, Sam, Eleanor, Matt, Rhys, Natalie, Joe, and Alfie, among others, whose CDWM evenings, fresh-air coffees, and Lunch^{TM} breaks have kept me sane, and made even the most frustrating days bearable. To Sam, Andrea and Ross, thank you for the many relaxed (provided Andrea isn't navigating) hiking weekends, I can't wait for the next one. I acknowledge financial support from EPSRC during the course of my PhD and for multiple workshops and conferences. The computational resources were provided by the Centre for Scientific Computing at Warwick.

I'd like to thank my parents, Anne-Marie and Eric, for all of their unrelenting support throughout my studies and the stresses which turned up alongside. I appreciate to no end the interest you have shown in my work, and your willingness to listen and provide reassurance for the umpteenth time. Mum, you've instilled in me a curiosity about the world which has brought me to where I am. Dad, I hope that our many late-night debates and extensive discussions about anything and everything, which I love, have prepared me to defend this work as well as I hope to.

To my parents-in-law Julie and Simon, I am so grateful for your support and patience with the both of us throughout the last few years. Escaping to Bedford for a gin and tonic, broad beans from the allotment, and long chats over dinner, is something I always look forward to.

Brooke, I could not have done this without you. Your patience has exceeded what I thought possible, and your love of science has helped me to find mine when it gets lost. I am so proud to finish this PhD as your wife, I love you endlessly. With your support, and plenty of cuddles from our Burns supper, I feel fully prepared for whatever challenge is next.

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree. The work presented (including data generated and data analysis) was carried out by the author.

Parts of this thesis have been published by the author and are detailed in the publication

- **Longbottom, S.**; Brommer, P. Uncertainty Quantification for Classical Effective Potentials: an Extension to *potfit*. *Modelling and Simulation in Materials Science and Engineering* 2019, **27** 044001.

Abstract

Effective potentials are an essential ingredient of classical molecular dynamics simulations. Little is understood of the errors incurred in representing the complex energy landscape of an atomic configuration by an effective potential containing considerably fewer parameters. This thesis details the introduction of an uncertainty quantification framework into the potential fitting process within the *potfit* force matching code. The probabilistic sloppy model method has been implemented within *potfit* as a means to quantify the uncertainties in analytic potential parameters, and in subsequent quantities measured using the fitted potential. Uncertainties in the effective potential are propagated through molecular dynamics simulations to obtain uncertainties in quantities of interest, which are a measure of the confidence in the model predictions.

The implementation has been designed to fit flexibly within the existing *potfit* workflow, and is generalised to work with any potential model or material. The uncertainty quantification software contains a variety of controllable parameters, which provide the user with diagnostic capabilities to understand the nature of the fitting landscape defined by their potential model and reference data. The implementation is available for use by the materials modelling community as part of the open source *potfit* software.

The uncertainty quantification technique is demonstrated using three potentials for nickel: two simple pair potentials, Lennard-Jones and Morse, and a local density dependent EAM potential. A sloppy model fit to *ab initio* reference data is constructed for each potential to calculate the uncertainties in lattice constants, elastic constants and thermal expansion. These can be used to show the unsuitability of pair potentials for nickel. In contrast, with EAM we observe a decreased uncertainty in the model predictions. This shows that our method can capture the effects of the error incurred in the potential generation process without resorting to comparison with experiment or *ab inito* calculations, which is an essential part to assess the predictive power of molecular dynamics simulations.

Abbreviations

DFT	**D**ensity **F**unctional **T**heory
MD	**M**olecular **Dy**namics
MC	**M**onte **C**arlo
FM	**F**orce **M**atching
XC	E**x**change-**C**orrelation (functional)
UQ	**U**ncertainty **Q**uantification
KIM	**K**nowledgebase of **I**nteratomic **M**odels
QoI	**Q**uantities **o**f **I**nterest
LDA	**L**ocal **D**ensity **A**pproximation
PBE	**P**erdew, **B**urke and **E**rnzerhof (functional)
pbc	**p**eriodic **b**oundary **c**onditions
MP	**M**onkhurst-**P**ack (grid)
PAW	**P**rojector **A**ugmented **W**ave (method)
GAP	**G**aussian **a**pproximation **p**otentials
PES	**P**otential **E**nergy **S**urface
LJ	**L**ennard-**J**ones
EAM	**E**mbedded **A**tom **M**ethod
k-ART	**K**inetic **A**ctivation-**R**elaxation **T**echnique
AIRSS	***A**b **i**nitio* **R**andom **S**tructure **S**earching
PCA	**P**rinciple **C**omponent **A**nalysis
MCMC	**M**arkov **C**hain **M**onte **C**arlo
ASE	**A**tomic **S**imulation **E**nvironment

LAMMPS	**L**arge-scale **A**tomic/**M**olecular **M**assively **P**arallel **S**imulator
LAPACK	**L**inear **A**lgebra **P**ackage
EVD	**E**igenvalue **D**ecomposition
SVD	**S**ingular **V**alue **D**ecomposition
RDF	**R**adial **D**istribution **F**unction
EOS	**E**quation **o**f **S**tate
IQR	**I**nter-**q**uartile **R**ange
GPR	**G**aussian **P**ocess **R**egression

1

Introduction

"Truth is much too complicated to allow anything but approximations."

– John von Neumann

Materials modelling is an important tool used to predict and understand physical phenomena and to verify the underlying theories through comparison with experiment. Such simulations also allow us to predict new properties and structures which can then be reproduced experimentally. In order to facilitate the modelling, interatomic potentials have long been used to circumvent the limitations on simulation timescales and system size of costly first principles calculations, by eliminating electronic degrees of freedom and specifying the energy dependence solely upon the atomic positions. This functional representation is a key constituent of molecular dynamics (MD) simulations, where the quality of the output relies predominantly on the potential employed. Currently, the systematic error incurred in using an interatomic potential is generally unknown, as is the resulting effect on the quantities of interest it is used to predict, therein forming the motivation for this work.

Figure 1.1 illustrates the methods used to access simulations on a given timescale

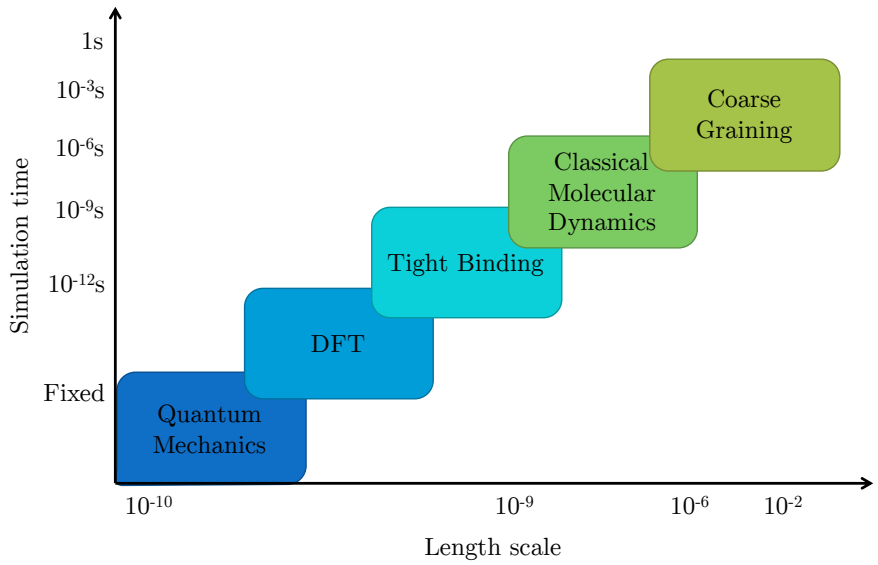

Figure 1.1: The material modelling methods used to simulate increasing system sizes for longer timescales.

and order of system size. Moving through the methods in figure 1.1 from left to right, each subsequent approach contains a greater level of approximation to the true physics of a system than the method preceding it. Ideally, all simulations would involve the exact solving of the Schrödinger equation, for all electrons in the system, with minimal approximations (i.e. an all electron, all nuclei, fully interacting treatment). As the computation of a fully quantum mechanical treatment scales unfavourably for more than a handful of atoms, approximations need to be made. For systems on the order of a few hundred atoms, electronic structure methods (i.e. tight binding model and density functional theory (DFT)) are used to obtain the electronic wavefunctions at a given moment in time. With the advancements in computing power, such *ab initio* methods can even be used to model systems with great accuracy for a short period in what is known as *ab initio*-molecular dynamics. For larger systems containing thousands, to at most trillions of atoms, molecular mechanics methods, such as classical MD and Monte Carlo (MC) methods, are used to explore the behaviour of systems over time. These methods rely predominantly on interatomic potentials to approximate the forces between atoms, which can take a variety of forms. The first potentials used intuitive functional forms, fitted to experimental data; however, new poten-

tials are frequently fitted to *ab-initio* data such as atomic forces, energies and stresses using a method known as force matching (FM).[1,2] The force matching method is the approach utilised by the *potfit* code,[3–6] within which this work is implemented. A recent advancement in potential development has been to use machine learning to construct potentials,[7] examples of which are Gaussian approximation potentials (GAP),[8–11] which use Gaussian process regression to fit to *ab initio* data, as well as neural-network based potentials, which follow the approach outlined by Behler and Parrinello.[12–16] Additionally, there exist hybrid methods, which combine molecular mechanics simulations with the accuracy of quantum mechanical approaches (so called QM-MM methods), within which regions of the classical molecular mechanics system are subjected to a quantum mechanical treatment.[17–19] As system sizes and simulation timescales increase further, coarse graining is a common technique whereby atoms or molecules are grouped into blocks, in order to reduce the computational complexity.[20]

As simulations increase in scale, so do the number of approximations required in order to ensure they are computationally feasible. Due to the stacking of approximations through the scales, the amount of unquantified error in simulations could quickly result in erroneous predictions of new phenomena, i.e. simulation artefacts, in the worst case. Therefore, it is of great interest to the modelling community that the uncertainties incurred as a result of specific approximations within a method are quantified, as well as those which were inherently introduced in any of the smaller-scale, more precise methods they are built upon.

In a brief review of the uncertainties quantified at each scale to date, we begin with systems which are treated quantum mechanically using *ab initio* electronic structure methods such as the tight binding model and DFT. As the majority of interesting physical phenomena require at least a few hundreds of atoms to simulate, the accuracy of a truly quantum mechanical treatment quickly becomes disproportionately computationally intensive. Therefore, the majority of simulation methods at this scale inherently contain approximations to the purest mathematical treatment. Some of these approximations are largely controllable (e.g. k-point density, planewave basis cutoff energy and pseudopotential representation) and can be used to improve the quality of a simulation at the expense of speed and

memory, whereas other are not (e.g. the exchange-correlation (XC) approximation). Work has been done to begin quantifying the errors introduced through the approximations made in first principles simulation methods, most notably of the error due to the treatment of the electron exchange-correlation term within DFT.[21,22] It has also been reported that the choice of pseudopotential used in DFT, in order to reduce the number of required electronic calculations, can introduce error depending on the functional representation.[23] There are also ongoing efforts to quantify the systematic error in predictions amongst available DFT software implementations.[24]

The approximation which bridges from the scale of *ab initio* methods to atomistic simulations is interatomic potentials. These potentials are implemented to reduce the number of parameters used to model interatomic forces between atoms, by reformulating the potential energy as an analytic function dependent on the atomic separation. Whilst these classical effective potentials permit access to a new order of simulation scales, the error introduced by the choice of potential model can be difficult to quantify.[25,26] Within the classical simulation methods there has been a significant recent effort in the literature to quantify uncertainty using a Bayesian framework, for a variety of interatomic models and force fields.[27–29] Similarly, work has been done to define a framework for uncertainty quantification (UQ) of machine-learning models.[30] There has also been research concerning the quantification of uncertainty due to the potential fitting reference set,[31] as well as a proposed framework to efficiently propagate parameter uncertainties to molecular dynamics outputs.[32] More specifically, quantification of parameter uncertainty for single potentials has been undertaken in a handful of cases.[33–36] This work adds to the growing body of uncertainty quantification in potential development and application by providing an open source implementation of the potential ensemble method[27,37] for use in future potential development projects.

The potential ensemble method underlying the *potfit* UQ implementation was chosen due to its efficient sampling of n-dimensional spaces. Within the potential fitting process, a fitting landscape with dimensionality equal to the number of potential parameters being fitted, is explored. Within such landscapes, the most suitable potential is usually believed to correspond to the "lowest" point in the space, i.e. there is least difference between forces

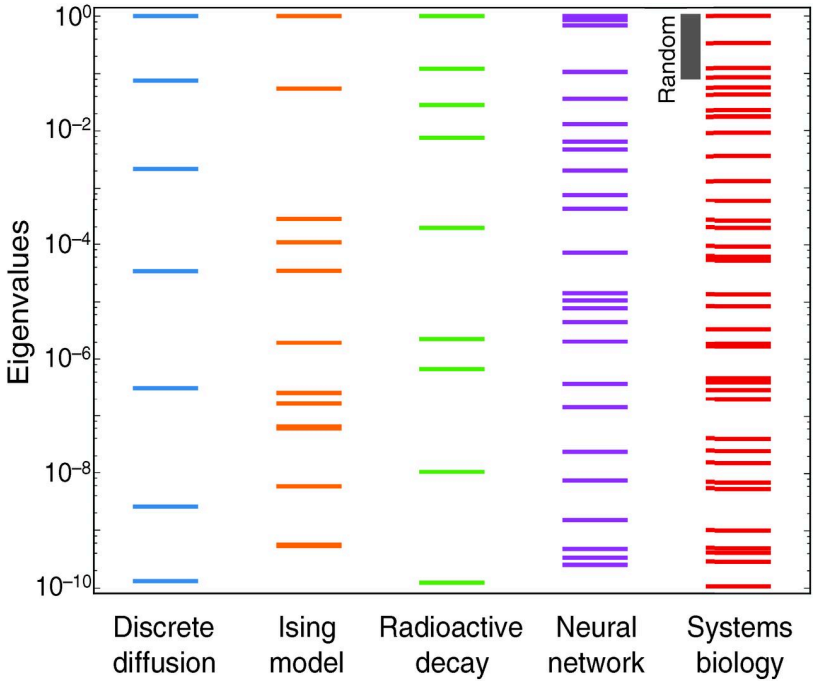

Figure 1.2: The spread of eigenvalues of the principle axes of curvature in spaces defined within a variety of sloppy models. From Machta, B. *et al.* Science, **342** 6158 (2013).[38] Reprinted with permission from AAAS.

predicted by this potential and those given in the reference data. This follows from the assumption that the reference data used in the fitting process is the most suitable set given the desired purpose of the eventual potential. As this may not be true, it could follow that alternative potentials within a range of the landscape minimum are better suited to predicting the desired quantities. Therefore, sampling alternative potentials from around the fitting space minimum could provide insight into the error within the system.

It has been observed that for some potential models, the order of landscape curvature around the minimum varies significantly along each of the principle axes of curvature.[6,27] Figure 1.2 illustrates some examples of similar model spaces which also exhibit this range of curvatures, and which are referred to as sloppy models.[38] In order to efficiently sample from such spaces, a Monte Carlo method with an adaptive step size dependent on the curvature was proposed, and formulated to work within the potential fitting framework by Frederiksen *et al.*[27] This potential ensemble method was chosen as a basis for the *potfit* uncertainty quantification implementation, as it provides a generalised mathematical framework which

can be used to efficiently sample any fitting space defined by a potential model + reference dataset. Additionally, as *potfit* has recently been updated to work within the OpenKIM framework,[39–41] this has now enabled users to fit any and all potential models which have not been implemented directly within *potfit* but are provided by the OpenKIM potential repository, which the potential ensemble method can also handle.

1.1 Thesis Outline

The aim of this thesis is to outline and demonstrate how the newly implemented uncertainty quantification workflow within *potfit* is used to fit interatomic potentials for MD simulations with a known uncertainty in their parameters, and demonstrate how the knowledge of this error can then be used to quantify the resulting uncertainties in quantities measured using MD. The basis of the *potfit* UQ implementation uses the potential ensemble method[27] to represent the error in analytic potential parameters. As discussed in depth by Pernot and Cailliez,[25,26] in our implementation of the ensemble method, we assume that the system uncertainty is dominated by potential model inadequacy rather than reference data uncertainty. This means that we assume the predominant error in the ability of a fitted potential to reproduce reality, is due to the choice (and limitations) of using a particular analytic function to model the interatomic forces, rather than the error inherent in the *ab initio* reference data, resulting from the approximations within DFT simulations, with which the reference set is created. Methods to minimise the reference data uncertainty are discussed within the relevant sections of this work. A demonstration of the UQ implementation for three increasingly complex potential models of nickel are introduced. Additionally the propagation of potential parameter error is demonstrated, whereby the incurred uncertainties in both equilibrium and non-equilibrium quantities of interest (QoI) are measured, namely the equilibrium lattice constant, elastic constants C_{11}, C_{12} and C_{44}, and the thermal expansion coefficient at 300 K.

In **Chapter 2** the relevant background theory is introduced. The chapter begins by introducing the theory underpinning *ab initio* and classical molecular dynamics techniques commonly used in atomistic modelling, including formalised explanations of DFT and MD.

The discussion of molecular dynamics leads into the formal introduction of interatomic potentials, and the fitting (cost) landscape which is sampled from during the potential fitting process. Finally, the potential ensemble method is outlined, and a basic background about the elastic constants and thermal expansion coefficient introduced.

Chapter 3 pertains to the implementation of the ensemble method within *potfit*. A detailed discussion of design considerations and technical aspects of the coding are documented, as well as a summary of the uncertainty quantification workflow. Included within the chapter are a series of discussions and results illustrating cases of interest which have been encountered to date when using the ensemble method. As a result of these observations, additional functionality has been implemented within *potfit*, and introduced within the chapter, which enables the user to better interrogate and understand the nuances of the ensemble generation process.

The results of the three potentials fitted for nickel are detailed in **Chapter 4**. To begin, the reference data used in the fitting process is discussed, as well as the limitations and considerations of the three analytical potential models chosen. A comprehensive walkthrough of the entire ensemble method process is presented, providing a significant amount of supplementary detail to the key results already published by the author.[42]

Finally, **Chapter 5** concludes the thesis with an evaluation of the uncertainty quantification implementation, and discussion of its possible limitations, as well as suggestions for further improvements and related avenues for further research.

2

Theoretical Background

"Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful."

– George E. P. Box

The modelling of materials is a multifaceted undertaking, requiring the careful weaving together of theoretical methods to create a complete understanding of the physical properties of a material. Robust bridges between methods, which describe behaviour on adjoining scales, are imperative in order to simulate properties with the greatest accessible accuracy within the constraints of computational cost. Of importance to this work, are the scales concerning atomistic phenomena, where quantum mechanics based methods meet classical atomistic descriptions. The timescales of such simulations are typically on the order of nanoseconds at most, with the largest accessible systems on the order of nanometers in size.

Models whose underlying theory is based in quantum mechanics, also known as *ab initio* methods, provide accurate atomic energies, forces and system stress information.

To simulate time-dependent quantities, we cross a multiscale bridge into classical models, where thermodynamical properties are readily accessible, at the expense of highly accurate descriptions of the electronic and interatomic interactions.

In this chapter, the necessary theoretical background to conduct the fitting and uncertainty quantification of interatomic potentials is outlined. Initially, in section 2.1, the focus is confined to quantum mechanical, *ab initio* methods. The Schrödinger equation is first introduced in section 2.1.1, before discussing a key development in *ab initio* simulation: the Born Oppenheimer approximation. In section 2.1.3 the Hartree-Fock formalism is introduced, before moving onto one of two simulation techniques used as part of this work, density functional theory (DFT), in section 2.1.4. Section 2.2 moves on to discuss molecular dynamics (MD), the second simulation technique of vital importance to this work. A key constituent in classical molecular dynamics are the effective potentials after which this work is titled. Section 2.3 is dedicated to the discussion of effective potentials, and the potential fitting process. In sections 2.3.1 and 2.3.2 a few of the effective potential models are outlined, before moving onto the fitting process in section 2.3.3. The landscape sampled during the potential fitting process warrants its own discussion in 2.3.5. Within section 2.3.7, the uncertainties introduced in using an effective potential are discussed. The remaining sections are concerned with necessary background theory to quantify the uncertainties in fitting effective potentials. In section 2.4 the sampling used in the uncertainty quantification method is outlined, followed by a detailed discussion of the uncertainty quantification method in section 2.5.

2.1 *Ab Initio* Methods

First principles (*ab initio*) methods are those which provide approximate solutions to the Schrödinger equation. Such methods provide electronic structure information, for systems up to the order of hundreds of atoms, which is obtained from the description of the atoms by a many-body wave function, $\mathbf{\Psi}$.

2.1.1 The Schrödinger Equation

The time-independent (non-relativistic) Schrödinger equation for a system of interacting particles is

$$\hat{H}|\Psi\rangle = E|\Psi\rangle \tag{2.1}$$

where E is the energy of the current system state, and $\hat{H}$ the Hamiltonian describing the kinetic and potential energies of the nuclei and electrons. Known as the Coulomb Hamiltonian, it is composed of electronic and nuclear energy contributions

$$\hat{H} = \hat{T}_n + \hat{T}_e + \hat{V}_{en} + \hat{V}_{ee} + \hat{V}_{nn}. \tag{2.2}$$

The states of the atoms and electrons in the system are described by their wavefunctions, Ψ. Solving the time-independent Schrödinger equation for a given wavefunction gives the energies of the states. In practise, the exact wavefunctions of many-body systems are not known exactly, therefore *ab initio* methods require iterative improvements to an initial guess wavefunction until the ground state energy is found.

2.1.2 The Born Oppenheimer Approximation

One can argue that the timescale over which electrons move is significantly faster than that of the larger, more massive nuclei. As a result, for any description of the nuclei for the system in a given state, it can be assumed that the electrons have almost instantaneously adopted their positions in the relevant eigenstate. Therefore, the wavefunction can be decomposed into electronic and ionic contributions, $\mathbf{\Psi} = \psi_e \otimes \psi_n$. The same approximation can be approached from another angle; the kinetic energy contribution from the relatively massive ions, T_n, is negligible compared to that of the electrons, T_e, where $T_i \propto 1/m_i$. As a result, T_n is neglected from the Hamiltonian, and the system is effectively repositioned to describe the movement of electrons, relative to the positions of stationary nuclei. This is known as the Born-Oppenheimer approximation, and was a crucial development towards the formulation of a computationally tractable quantum mechanical model.

For a system with N_n nuclei and N_e electrons the many-body Hamiltonian can then

be written, in atomic units, as

$$\hat{H} = -\frac{1}{2}\sum_{i}^{N_e}\nabla_i^2 - \sum_{i}^{N_e}\sum_{I}^{N_n}\frac{Z_I}{|\boldsymbol{r}_i - \boldsymbol{R}_I|} + \frac{1}{2}\sum_{i}^{N_e}\sum_{j\neq i}^{N_e}\frac{1}{|\boldsymbol{r}_i - \boldsymbol{r}_j|} + \frac{1}{2}\sum_{I}^{N_n}\sum_{J\neq I}^{N_n}\frac{Z_I Z_J}{|\boldsymbol{R}_I - \boldsymbol{R}_J|} \tag{2.3}$$

where I and J index the nuclei, and i and j are summed over the electrons. The atomic number of the I-th nucleus is denoted Z_I, with position R_I and mass m_I, and similarly for electrons.

Exact solutions to the time-independent Schrödinger equation using the Hamiltonian in equation (2.3), are still incredibly computationally expensive for systems of atoms with more than a handful of electrons. Therefore further approximations are required. Two methods of approach arose: either an approximation to the wave function is applied, as in the Hartree-Fock method, or an approximation to the instantaneous electronic positions, reformulated as an electron density functional is made, resulting in the Kohn-Sham equations, with the latter case forming the backbone of density functional theory.

2.1.3 The Hartree-Fock Equation

The method formulated by Hartree assumes that the electrons in the many-body wavefunction, under the Born-Oppenheimer approximation, can be written as a product of individual electronic wavefunctions,

$$\boldsymbol{\Psi}(\boldsymbol{r}_1, \boldsymbol{r}_2, ..., \boldsymbol{r}_N) = \psi_1(\boldsymbol{r}_1) \otimes \psi_2(\boldsymbol{r}_2) \otimes ... \otimes \psi_N(\boldsymbol{r}_N), \tag{2.4}$$

which interact through a spherically symmetric mean-field Coulomb potential. The electron-electron contribution to the Hamiltonian in equation (2.3), is then redefined in terms of a continuous, smoothed-out charge density, $\rho(\boldsymbol{r}) = \sum_i |\psi(\boldsymbol{r})|^2$, with energy in atomic units

$$\hat{V}_{ee}(\boldsymbol{r}) = -\int \frac{\rho(\boldsymbol{s})}{|\boldsymbol{r} - \boldsymbol{s}|} d\boldsymbol{s}. \tag{2.5}$$

The Schrödinger equation can then be self-consistently solved to find the minimum energy solutions. This approach however, fails to account for the Pauli exclusion principle, whereby the many-body wavefuction is required to be anti-symmetric under the exchange of any two

electron positions,

$$\boldsymbol{\Psi}(\boldsymbol{r}_1, \boldsymbol{r}_2, ..., \boldsymbol{r}_N) = -\boldsymbol{\Psi}(\boldsymbol{r}_2, \boldsymbol{r}_1, ..., \boldsymbol{r}_N). \tag{2.6}$$

In 1930, Fock[43] remedied the particle spin issue by describing the solutions to the Hartree many-body wavefunction as a Slater determinant of individual orbitals. The inclusion of the Slater determinant, ensured that anti-symmetric combinations of electronic orbitals are chosen, so as to correctly encode electron spin:

$$\boldsymbol{\Psi}(\boldsymbol{r}_1, \boldsymbol{r}_2, ..., \boldsymbol{r}_N) = \frac{1}{\sqrt{N!}} = \begin{vmatrix} \psi_1(\boldsymbol{r}_1) & \psi_1(\boldsymbol{r}_2) & \cdots & \psi_1(\boldsymbol{r}_N) \\ \psi_2(\boldsymbol{r}_1) & \psi_2(\boldsymbol{r}_2) & \cdots & \psi_2(\boldsymbol{r}_N) \\ \vdots & \vdots & \ddots & \vdots \\ \psi_N(\boldsymbol{r}_1) & \psi_N(\boldsymbol{r}_2) & \cdots & \psi_N(\boldsymbol{r}_N) \end{vmatrix} \tag{2.7}$$

This self-consistent solution to the Schrödinger equation, with a mean-field description of the electronic interactions and the many-body electronic wavefunction, is known as the Hartree-Fock method. A limitation of this method, is that the electron-electron potential does not account for electron exchange or correlation, the latter of which is essential for consistency with the time-independent Schrödinger equation. Therefore, the accuracy with which the Hartree-Fock method can predict the wavefunction energy is known as the Hartree-Fock limit.

2.1.4 Density Functional Theory

The alternative approximation proposed to the time-independent Schrödinger equation solution, is the density functional theory formulation. This method shifts the focus from a product of electronic wavefunctions to an electron density-centred approach. Within this model, the solution to the electronic orbital wavefunctions are found iteratively. In DFT the system of atoms are described by an analogue system comprising of non-interacting electrons exposed to a local effective potential, V_{KS}, which captures the electron-electron interactions. This allows for the decoupling of the many-body electron wavefunction into its individual orbital representations, $\varphi_i(\boldsymbol{r})$. The system is again solved self-consistently, to

find the correct electronic orbitals which reproduce the system density, $n(\boldsymbol{r})$.

The Hohenberg-Kohn Theorems

In 1964, the underlying motivation for the density functional theory approach was first formalised as a pair a theorems by Hohenberg and Kohn.[44] As in the Born-Oppenheimer approximation, an inhomogeneous system of interacting electrons subject to an external potential, V_{ext}, have an electronic Hamiltonian of the form

$$\hat{H} = -\frac{1}{2}\sum_{i}\nabla_i^2 + \frac{1}{2}\sum_{i}\sum_{j\neq i}\frac{1}{|\boldsymbol{r}_i - \boldsymbol{r}_j|} + V_{ext}. \tag{2.8}$$

The first theorem then states that the (non-degenerate) ground state electronic density, $n(\boldsymbol{r})$, uniquely defines $V_{ext} = \int v_{ext}(\boldsymbol{r})n(\boldsymbol{r})\,d^3r$, up to the addition of a constant ($v_{ext} - v_{ext}^{'} = c$). The significant implication of this statement is that all properties of the system, including descriptions of excited states, can be determined exactly from the ground-state density.

The second theorem then moves to reformulate the electronic components of the Hamiltonian in line with the ground-state density description. Hohenberg and Kohn show in the second theorem, that the (non-degenerate) ground state energy of the system can be formulated as the sum of an external potential with the addition of a density dependent functional, $F[n]$:

$$E[n] = F[n] + \int v_{ext}(\boldsymbol{r})n(\boldsymbol{r})\,d^3r, \tag{2.9}$$

where $F[n]$ is a universal functional of any electronic system of unknown form. However, if the functional form of $F[n]$ is simple enough, the many-body problem has been reduced to a minimisation of a three-dimensional functional, resulting in a significant reduction in computational cost when juxtaposed with the $3N$-variable Hartree-Fock approach.

Later, in 1979, Levy[45] circumvented the assumption of a non-degenerate ground state in the Hohenberg-Kohn theorems. He did so through the introduction of a constrained search formalism, which ensures that the search to minimise the universal energy functional conserves the direct relationship between the Hamiltonian and the system wavefunction.

The Kohn-Sham Equations

Working from the Hohenberg-Kohn theorems, Kohn and Sham[46] found that a result of re-basing the density-dependent description to a fictitious system of non-interacting "electrons" meant that the universal functional could be separated into three parts,

$$F[n] = T[n] + \frac{1}{2}\int \frac{n(\boldsymbol{r})n(\boldsymbol{s})}{|\boldsymbol{r}-\boldsymbol{s}|}\, d^3r d^3s + E_{XC}[n]. \tag{2.10}$$

Here, the kinetic contribution, $T[n]$, comes from the non-interacting system of electrons, the second term is the Hartree energy as in equation (2.5), and finally, $E_{XC}[n]$ accounts for the energy resulting from electron exchange and correlation effects. It is important here to highlight that the kinetic energy is no longer that for a system of interacting particles, and is part of the change in system description unique to the Kohn-Sham formulation.

By constraining the total number of electrons to be constant, as well as grouping the potential energy contributions for non-interacting particles in an external potential, the set of individual electronic Schrödinger equations can be written as

$$\Big[-\frac{1}{2}\nabla^2 + V_{KS}(\boldsymbol{r})\Big]\varphi_i(\boldsymbol{r}) = \epsilon_i\varphi_i(\boldsymbol{r}). \tag{2.11}$$

These are known as the Kohn-Sham equations which describe individual electronic orbitals, $\varphi_i(\boldsymbol{r})$, using the Kohn-Sham potential

$$V_{KS}(\boldsymbol{r}) = \underbrace{\int \frac{n(\boldsymbol{s})}{|\boldsymbol{r}-\boldsymbol{s}|}\, d^3s}_{V_H} + V_{XC}(\boldsymbol{r}) + V_{ext}(\boldsymbol{r}). \tag{2.12}$$

The system of Kohn-Sham equations can now be self-consistently solved to converge to the ground state density, which is constructed by summing over $N/2$ particle states,

$$n(\boldsymbol{r}) = 2\sum_{i=1}^{N/2} |\varphi_i(\boldsymbol{r})|^2. \tag{2.13}$$

The factor of two in the summation of the density above accounts for the electron degeneracy due to spin. In the Kohn-Sham formulation, the only remaining component unaccounted

for up to this point relates to the exchange and correlations of electrons, V_{XC}.

The Exchange-Correlation (XC) Energy

The exchange-correlation energy is the only term in DFT whose functional form is indirectly approximated, as it is likely too complex to be represented analytically. The exchange-correlation contribution is formulated as a functional of the density, $n(\boldsymbol{r})$,

$$V_{XC}(\boldsymbol{r}) = \frac{\delta E_{XC}[n(\boldsymbol{r})]}{\delta n(\boldsymbol{r})}. \tag{2.14}$$

Two popular parametrisations of the exchange-correlation approximation are: the local density approximation (LDA)[46] and the Perdew, Burke and Ernzerhof (PBE), generalised gradient, functional.[47] The LDA functional is solely dependent on the local electron density, separating the exchange and correlation energies and treating them independently,

$$E_{XC}^{LDA}[n] = \int n(\boldsymbol{r})\varepsilon_{XC}(n(\boldsymbol{r}))\, dr. \tag{2.15}$$

However, as the LDA approximation is designed to describe a homogeneous electron gas, the functional can struggle with systems with rapidly changing electron densities. The PBE approach, is one of a class of generalised gradient approximations, which builds upon the LDA functional to include the rate of change of density throughout the material,

$$E_{XC}^{PBE}[n] = \int n(\boldsymbol{r})\varepsilon_{XC}(n(\boldsymbol{r}), \nabla n(\boldsymbol{r}))\, dr. \tag{2.16}$$

This inclusion of the gradient as a component of the functional means that the PBE approximation prefers an inhomogeneous electron density compared to the LDA.[48] This also results in improved descriptions of total energies, structural differences and atomisation energies,[49] and so the PBE functional is employed in all DFT calculations throughout this work.

Bloch's Theorem

Up to this point, the many-body Schrödinger equation in (2.1) has been reformulated using the Hohenberg-Kohn theorems, to describe a series of independent Kohn-Sham particles,

as in equation (2.11). Finding the ground-state energy of this system then becomes a case of solving the Kohn-Sham equations for enough states to account for all electrons in the system. Clearly, for simulations of solid state materials with more than a few atoms, the number of required electronic states quickly adds up, even in systems with degeneracy in the election energies.

Bloch's theorem uses the periodicity of atomic positions to simplify the wavefunctions of a repeating atomic motif. It is due to this design choice that many planewave DFT codes use periodic boundary conditions (pbc). Following the decision that the system of atoms have translational symmetry in all three directions, it is implied that the electrons subjected to the nuclear potential are also periodic. Therefore, for a repeating cell with lattice vector $\boldsymbol{a}$, the periodic potential satisfies $V(\boldsymbol{r}) = V(\boldsymbol{r} + \boldsymbol{a})$. Naturally, this also implies that the same is true for the density, $n(\boldsymbol{r}) = n(\boldsymbol{r} + \boldsymbol{a})$.

Finally, Bloch's theorem states that in a periodic system, solutions to the wavefunction can be written as a 'quasi-periodic' product. The wavefunctions consist of a repeating cell part, $u_j(\boldsymbol{r}) = u_j(\boldsymbol{r} + \boldsymbol{a})$, and an arbitrary phase with wavevector $\boldsymbol{k}$, which is restricted to the first Brillouin zone:

$$\begin{aligned} \varphi_j(\boldsymbol{r}) &= e^{i\boldsymbol{k}\cdot\boldsymbol{r}} u_j(\boldsymbol{r}), \\ \varphi_j(\boldsymbol{r} + \boldsymbol{a}) &= e^{i\boldsymbol{k}\cdot(\boldsymbol{r}+\boldsymbol{a})} u_j(\boldsymbol{r} + \boldsymbol{a}) = e^{i\boldsymbol{k}\cdot\boldsymbol{a}} \varphi_j(\boldsymbol{a}). \end{aligned} \tag{2.17}$$

Planewave Basis Sets

As $u_j(\boldsymbol{r})$ is periodic, it can be expanded in terms of a Fourier series with reciprocal lattice vectors $\boldsymbol{G}$ (with $\boldsymbol{G} \cdot \boldsymbol{a} = 2\pi m$, $m \in \mathbb{Z}$) and complex coefficients $c_{j,\boldsymbol{G}}$. This means that the electronic wavefunctions can be described by a basis of planewaves:

$$\begin{aligned} u_j(\boldsymbol{r}) &= \sum_{\boldsymbol{G}} c_{j,\boldsymbol{G}} e^{i\boldsymbol{G}\cdot\boldsymbol{r}}, \\ \varphi_j(\boldsymbol{r}) &= e^{i\boldsymbol{k}\cdot\boldsymbol{r}} \sum_{\boldsymbol{G}} c_{j,\boldsymbol{G}} e^{i\boldsymbol{G}\cdot\boldsymbol{r}} = \sum_{\boldsymbol{G}} c_{j,\boldsymbol{G}} e^{i(\boldsymbol{k}+\boldsymbol{G})\cdot\boldsymbol{r}}. \end{aligned} \tag{2.18}$$

This planewave basis can be inserted into the Kohn-Sham equations in 2.11 to obtain a description of the potentials in terms of Fourier coefficients $c_{j,\boldsymbol{k}+\boldsymbol{G}}$,

$$\sum_{\boldsymbol{G}'}\left[\frac{1}{2}|\boldsymbol{k}+\boldsymbol{G}|^2\delta_{\boldsymbol{G}\boldsymbol{G}'}+V_H(\boldsymbol{G}-\boldsymbol{G}')+V_{XC}(\boldsymbol{G}-\boldsymbol{G}')+V_{\text{ext}}(\boldsymbol{G}-\boldsymbol{G}')\right]c_{j,\boldsymbol{k}+\boldsymbol{G}'}=\epsilon_j c_{j,\boldsymbol{k}+\boldsymbol{G}} \quad (2.19)$$

where $\delta_{\boldsymbol{G}\boldsymbol{G}'}$ is the Dirac delta function.

For large wavenumbers $\boldsymbol{G}$, the magnitude of the corresponding coefficients $c_{j,\boldsymbol{G}}$ decreases. Therefore, as the weighting of these wavevectors become negligible, it makes sense to exclude them from the calculations in order to create a finite basis set. As the square of the wavenumber is directly related to the kinetic energy of the state, the exclusion of high wavenumber contributions is enforced by setting a cut-off energy for the planewave basis set,

$$E_{\text{cutoff}}=\frac{\hbar^2}{2m_{\text{e}}}|\boldsymbol{k}+\boldsymbol{G}|^2. \quad (2.20)$$

In every planewave DFT simulation the cutoff energy for the planewave basis is found by trialling a variety of energies until the results of the calculation converge to within a desired tolerance. Similarly to converging the planewave cutoff energy, the k-points at which the Brillouin zone is sampled can also be converged in order to optimise the calculations. A Monkhurst-Pack (MP) grid[50] is the method of choice used within this work when selecting the relevant k-points. The MP k-points are evenly spaced in reciprocal space, where due to the inverse relationship with cell size, larger simulation cells require fewer k-points to compute the planewave solutions. In order to simulate complex morphologies such as surfaces, the simulation cell can be defined in such a way that it includes a large enough vacuum so that periodic images of the slab surfaces do not interact.

Pseudopotentials

To further reduce the computational requirements of DFT simulations, therefore allowing for the computation of more atoms, an approximation to the wavefunction is frequently used. A pseudopotential is used to approximate the components of the electric potential due to the interactions of atomic nuclei (with associated core electrons) and valence electrons. Close to the nuclei, valence electron wavefunctions oscillate rapidly in order to maintain

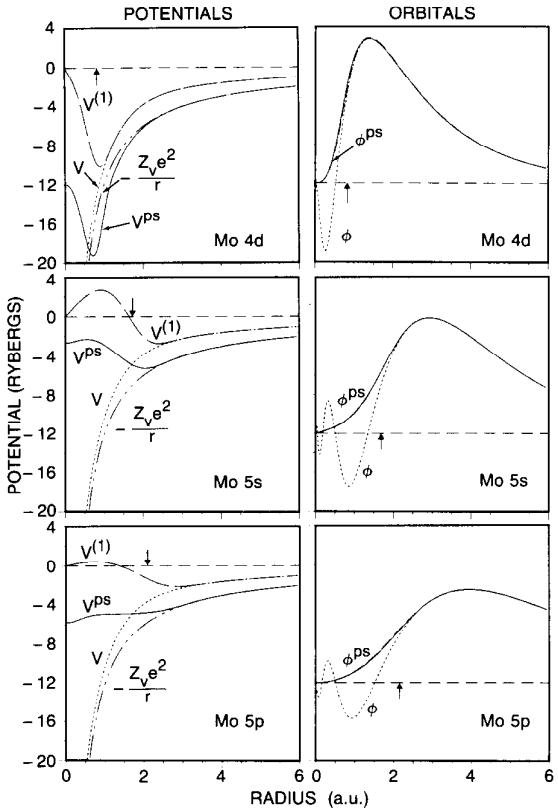

Figure 2.1: The pseudopotential V^{ps} and respective wavefunctions ϕ^{ps} for the 4d, 5s and 5p valence states in molybdenum. Vertical arrows denote the core cutoff distance r_C, and $V^{(1)}$ show the initial pseudopotential trials before convergence. The all-electron, plane wave potential is denoted V, with the the limiting form illustrated alongside. Reprinted from Computer Physics Reports, **9**, Pickett, W. Pseudopotential methods in condensed matter applications, p103., Copyright 1989, with permission from Elsevier.[51]

orthogonality to core electron wavefunctions. Furthermore, in order to strongly bind the core electrons to the nucleus, the Coulomb potential diverges with $1/r$ as $r \to 0$. As the majority of material simulations are interested in electronic, mechanical and chemical properties which are predominantly affected by the valence electrons, replacing the core election wavefunctions and Coulomb potential with smoother functional forms is often a worthwhile trade in precision for the gain in computational speed-up. Due to the strong attractive Coulomb interactions, core electrons reside close to the nucleus, and so their only significant contribution to system is the screening of the nuclei from valence electrons. As long as the screening effect is replicated by the pseudopotential, encapsulating the core electron wavefunctions in a generalised nuclei-core ion description has little effect on the

accuracy of simulations.

Figure 2.1 illustrates a Coulomb potential and wavefunction close to the nuclei, for three valence states in molybdenum. Outwith a chosen cutoff radius, r_C, the pseudopotential and corresponding potential provide a good approximation to the all electron solutions. As the smoother wavefunctions of pseudopotentials oscillate less within the ion core region, the orthogonal valence wavefunction also have fewer oscillations. Therefore a lower planewave cutoff energy is required when using pseudopotentials, yielding a significant computational speedup. Libraries of pre-computed pseudopotentials are now common inclusions in DFT codes,[52] and frequently fall into three of the most popular categories: norm-conserving pseudopotentials,[53] ultrasoft pseudopotentials[54,55] and the projector augmented wave method (PAW) pseudopotentials.[56,57] For further discussions of pseudopotentials the reader is directed to the textbook: "Electronic Structure: Basic Theory and Practical Methods".[58]

The Hellmann-Feynman Theorem

In order to find the ground state of a system, one has to understand how the electronic configuration and atoms in the system should move in order to decrease the total system energy. The Hellmann-Feynman Theorem[59,60] relates the derivatives of the total energy to the expectation value of the derivative of the Hamiltonian for a given parameter, if the wavefunction is an eigenfunction of the Hamiltonian. If this parameter is chosen to be the ion positions in the Kohn-Sham treatment, the forces on these ions can be calculated. Therefore, for the Kohn-Sham wavefunctions in equation (2.19), the forces acting on the ions are be described by the Hellmann-Feynman Theorem by taking the derivatives with respect to the ion positions, $\boldsymbol{R}_I$ as follows

$$F_I = -\frac{\partial E_{KS}}{\partial \boldsymbol{R}_I} = -\sum_i \langle \varphi_i | \frac{\partial H_{KS}}{\partial \boldsymbol{R}_I} | \varphi_i \rangle. \tag{2.21}$$

Geometry Optimization

To undertake a geometry optimization, where the position of atoms in a system are allowed to relax into the most favourable configuration, there are a series of steps which must be followed. Before atoms can be moved, the directions of movement which would lower the

total system energy must first be understood. As the Hellmann-Feynman theorem states, these forces can be calculated for each ion from the derivative of the energy. This theorem only holds if the electronic orbitals are eigenfunctions of the Hamiltonian, and so this is what must first be established. Therefore, the geometry optimization procedure is as follows:

1. Using the variational principle, self-consistently find the correct Kohn-Sham electronic wavefunctions associated to each ion in the system. Utilizing the method proposed by Car and Parrinelllo,[61] the electronic wavefunctions can be treated as dynamical variables, resulting in a significant computational speed up compared to the traditional matrix-diagonalisation method.

2. Use the Hellmann-Feynman Theorem to calculate the incident forces on each ion, and therefore the direction of movement for each which minimizes the total system energy.

3. Move the ions by a specified amount in the identified directions. This positional change will then require recalculation the Kohn-Sham wavefunctions.

4. Repeat steps 1-3 until the Hellmann-Feynman forces are within a given tolerance.

When the forces and energies of any system are required, without completing the geometry optimization process, this is often referred to as a singlepoint DFT calculation. In this case steps 1 and 2 in the routine above are followed to obtain the energies and forces without continuing on to moving the atoms.

2.2 Classical Molecular Dynamics

To access thermodynamical material properties, classical mechanics is used to model interatomic interactions and is known as molecular dynamics (MD). Molecular dynamics is both computationally cost effective and efficient, when compared to any attempt to model such quantities with a quantum description. The basis of MD relies upon integrating Newton's equations of motion numerically for a collection of atoms (or molecules if the system is coarse grained).

The Hamiltonian for a closed system can be separated into two constituent parts: a kinetic energy contribution relating to the deterministic momentum (i.e. mass and veloc-

ity) of each atom, and a potential energy describing the interactions of an atom with its local atomic neighbourhood. Molecular dynamics simulations are largely independent of the potential energy description, which can be provided either from *ab initio* calculations or from a classical model. In this work the potential energy is modelled by a classical effective (interatomic) potential, which is either an analytic or tabulated function.

In essence, there are three components to molecular dynamics: the first is the algorithm used to evolve the atoms over time, which in this case we consider the Velocity-Verlet algorithm, the second is the choice of thermodynamic ensemble being simulated as defined by the thermostat (or absence of one), and third is the description of the interatomic forces, as handled by the interatomic potential model.

2.2.1 The Verlet Algorithm

By considering the Taylor series expansion of an atom with position, $\boldsymbol{r}$, at a time $t \pm \Delta t$, a description for the time evolution of the system can be derived.

$$\boldsymbol{r}(t + \Delta t) = \boldsymbol{r}(t) + \frac{d\boldsymbol{r}(t)}{dt}\Delta t + \frac{d^2\boldsymbol{r}(t)}{dt^2}\frac{\Delta t^2}{2!} + \frac{d^3\boldsymbol{r}(t)}{dt^3}\frac{\Delta t^3}{3!} + ... \tag{2.22}$$

and similarly

$$\boldsymbol{r}(t - \Delta t) = \boldsymbol{r}(t) - \frac{d\boldsymbol{r}(t)}{dt}\Delta t + \frac{d^2\boldsymbol{r}(t)}{dt^2}\frac{\Delta t^2}{2!} - \frac{d^3\boldsymbol{r}(t)}{dt^3}\frac{\Delta t^3}{3!} + ... \tag{2.23}$$

Adding the expansions in equation (2.22) and equation (2.22) together, and rearranging in terms of $t + \Delta t$ gives

$$\boldsymbol{r}(t + \Delta t) = 2\boldsymbol{r}(t) - \boldsymbol{r}(t - \Delta t) + \frac{d^2\boldsymbol{r}(t)}{dt^2}\Delta t^2 \tag{2.24}$$

for terms up to second order. It is then assumed that the acceleration on an atom is a result of the forces from neighbouring atoms. This force is described as the rate of change (derivative) of an interatomic potential energy function, $V(\boldsymbol{r})$. From Newton's Second Law

of Motion, the force on a particle, i, with mass, m is then

$$\boldsymbol{F}_i = m_i \frac{d^2 \boldsymbol{r}_i}{dt^2} = -\frac{dV(\boldsymbol{r}_i)}{d\boldsymbol{r}_i}. \tag{2.25}$$

By replacing the second order term in equation (2.24) with the interatomic potential description in equation (2.25), an easily-implementable, discretised description is obtained for an atom moving over time,

$$\boldsymbol{r}(t + \Delta t) = 2\boldsymbol{r}(t) - \boldsymbol{r}(t - \Delta t) + \frac{\Delta t^2}{m}\boldsymbol{F}(t). \tag{2.26}$$

This equation is known as the Verlet algorithm, and forms the backbone of MD simulations. To calculate the atomic velocities, a central finite difference approach of subtracting the atomic positions at a displacement of $\pm\Delta t$ can be used to approximate the derivative. However, this approach is liable to potential rounding errors arising from the subtracting of atomic positions. This issue can be mitigated by the explicitly evaluating the velocities at the same point in time as the positions. This is a reformulation is known as the velocity-Verlet algorithm, and results in a description of velocities accurate to third order (an accuracy equivalent to that of the the atomic positions)

$$\boldsymbol{r}(t + \Delta t) = \boldsymbol{r}(t) + \dot{\boldsymbol{r}}\Delta t + \frac{1}{2}\frac{\Delta t^2}{m}\boldsymbol{F}(t), \tag{2.27}$$

$$\dot{\boldsymbol{r}}(t + \Delta t) = \dot{\boldsymbol{r}}(t) + \frac{1}{2}\frac{\Delta t}{m}\Big[\boldsymbol{F}(t) + \boldsymbol{F}(t + \Delta t)\Big]. \tag{2.28}$$

2.2.2 Thermodynamic Ensembles

The velocity-Verlet algorithms define a closed system (fixed number of particles, N, and simulation cell volume, V) within which the total energy (E) is conserved. This adiabatic construction defines the NVE ensemble, otherwise known as the microcanonical ensemble. Following from the ergodic hypothesis, which states that each microstate of a constant energy system is accessible with an equal probability, this means that the time-averaged NVE samples of accessible system states can be used to compute macroscopic thermodynamic properties. From the velocity-Verlet algorithms in 2.27 and 2.28, the initial force on an atom

is described by the potential energy due to neighbours, and is described by an interatomic potential, $\boldsymbol{F} = -\nabla V$. The atomic velocities, which introduce kinetic energy contributions into the system, can be initialised using a Maxwell-Boltzman distribution, and scaled to represent a temperature, as defined by the equipartition theorem

$$T = \frac{1}{3Nk_B} \sum_{i=1}^{N} \frac{|\boldsymbol{p}_i|^2}{m_i} \tag{2.29}$$

where $\boldsymbol{p}_i = m_i \dot{\boldsymbol{r}}_i$ is the momentum of atom i, and k_B the Boltzmann constant. With this velocity initialisation, it is necessary to run the simulation until the distribution of kinetic and potential energies equilibrate before extracting results.

Due to choice of time-step, system size or morphology, it is possible for the instantaneous temperature in equation (2.29) to drift over the course of a simulation. Therefore, to regulate the system temperature and realistically simulate macroscopic quantities, a thermostat can be used to manipulate the exchange of energy within the system. This is known as the canonical ensemble (NVT), in which the temperature is kept constant and the total energy able to fluctuate. For large enough systems, a canonical ensemble becomes equivalent to the microcanonical ensemble. In an NVE ensemble, each microstate is equally probable due to the conservation of energy (with temperature allowed to fluctuate), whereas within the NVT ensemble, the probability assigned to each microstate is dependent on the energy of the state

$$P = \frac{1}{Z} e^{-E/k_B T} \quad \text{with} \quad Z = \sum_i e^{-\varepsilon_i / k_B T} \tag{2.30}$$

where Z is the canonical partition function dependent on the Helmholtz free energy. This is a sum of all microstates, with energy ε_i, accessible on the same contour of the potential energy surface (i.e. the system at constant temperature in an NVE ensemble).

To regulate the temperature in a canonical ensemble, the thermostat is required to periodically adjust the kinetic energy of the system by acting upon the atom velocities. There are four main approaches of thermostat which will be briefly touched upon, however for further detail the reader is directed to "Understanding Molecular Simulation: From Algorithms to Applications" by Frenkel and Smit.[62]

From the discussion of initialising atoms with a velocity distribution within the mi-

crocanonical ensemble, it follows that the first approach to temperature control would be to periodically rescale the velocities to sustain the desired temperature. Scaling the atomic velocities to maintain the desired result is referred to as a strong coupling method, one disadvantage of which is that a simulation can easily become destabilised by the instantaneous shifting an atom velocities between disjointed points in phase space (i.e. instantaneous changes to velocities are unrealistic). Therefore, velocity rescaling is more usefully used as a tool to quickly initialise a system with a given temperature, prior to a burn-in (energy equilibration) period, before applying a less intrusive thermostat.

The second type of thermostats are categorised as weak-coupling methods, under which the atom velocities are adjusted towards the direction of the desired temperature, as opposed to being instantaneously shifted. The Berendsen thermostat,[63] is an example of a weak-coupling method, under which the velocities are slowly rescaled to maintain the desired temperature. One major disadvantage of this approach however, is that a Maxwell-Boltzman velocity distribution is not guaranteed, which could arise in certain atom velocities being continuously scaled more than others, resulting in what as known as the "flying ice-cube" effect.[64]

The third approach to thermostatting uses stochastic interactions with the atoms to perturb the velocities. With the Andersen thermostat,[65] these interactions take the form of random "collisions" between an atom and the thermostat, with the resulting velocities are drawn from a Maxwell-Boltzman distribution. As with the thermostatting methods mentioned previously, the Andersen thermostat also risks introducing instabilities in complex morphologies due to potentially drastic velocity perturbations which could unrealistically disrupt simulation of a desired structure. Compared to the Andersen approach, the Langevin thermostat[66] is another stochastic method which is less prone to such instabilities. With a Langevin thermostat, the system is heated through the application of random forces, which are balanced by the removal (i.e. cooling) of the system through a frictional damping term. The additional stability in Langevin dynamics simulations is due to the introduction of this damping force, which is applied randomly to particles and is related to the random 'heating' forces by a dispersion relation. The only disadvantage of the Langevin approach, is that this damping force can be incident on atoms within any neighbourhood. In the case of an

atom in a vacuum, or far from a surface, this can obviously lead to unrealistic results.

The final thermostat to be introduced involves a direct alteration to equation (2.28) in the Velocity-Verlet algorithms, through the introduction of a friction factor as an additional degree of freedom. Known as the Nosé-Hoover method,[67,68] this approach results in deterministic particle trajectories, due to the additional degree of freedom in the system dynamics, unlike the random velocity updating manifested in a variety of forms by previous thermostats. With the Langevin thermostat, the temperature modulation is akin to interactions between the system and a viscous liquid, whereas in the Nosé-Hoover thermostat the system is coupled to an external heat bath through the additional degree of freedom. The Nosé-Hoover thermostat guarantees a Maxwell-Boltzmann velocity distribution for large systems, however for single particles this is not guaranteed, and so the use of an Andersen thermostat is still preferable for small systems.

In addition to the microcanonical (NVE) and canonical (NVT) ensembles, there exists an isobaric-isothermal ensemble (NPT) which simulates the conditions under which the majority of physical experiments are performed. In order to reproduce experimental results, a combination of the thermostats and barostats, of a similar mathematical construction, are used to maintain the ensemble, with improvements to this simulation technique being continually proposed.[69]

2.3 Interatomic Potentials

The final mathematical construction to be detailed in the molecular dynamics formalism, is the encoding of the atom-atom interactions. As atoms are perturbed by the thermostat, the resulting neighbour forces, due to the change in positions, are described by a classical effective (interatomic) potential, $V(\boldsymbol{r})$. This work is concerned with analytic effective potentials which assume explicit functional forms. Tabulated potentials may also be utilized in molecular dynamics, and for further details on tabulated potentials the reader is referred to the textbook "Interatomic Forces in Condensed Matter" by Finnis.[70]

Interatomic potentials encode the material-specific physics into the system, and so great care should be taken when selecting the potential model for use in an MD simulation.

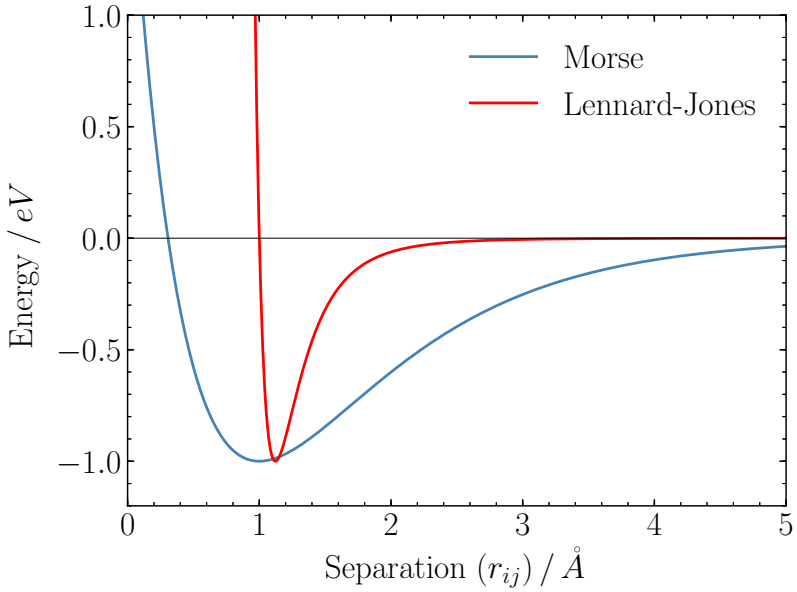

Figure 2.2: A juxtaposition of the Lennard-Jones potential (red) and Morse potential (blue) with all parameters set to unity. The difference in gradient around the minima is due to the additional Morse parameter, a, which controls the gradient of the potential well.

Analytic potential functions have historically been determined based on intuition, or based upon the understanding of the material from experiments. Potential functions were initially designed to reproduce experimental values, such as the cohesive energy, stress/strain curves, expansion coefficients, lattice parameters and elastic constants. The techniques behind the creation of effective potentials have since evolved, to include fitting analytic functions to *ab initio* data, through the use of force-matching,[1] or by machine learning, with the most notable approaches being neural network potentials,[71,72] and Gaussian approximation potentials.[73] Interatomic potentials aim to capture the most important nuances in the interatomic energy landscape. As a result, an effective potential developed for one particular system is unlikely to be transferable to a system with different behaviour. Present day advancements in potential creation improve upon the potential design process, by additionally working to more accurately reproduce quantities and distributions obtained from *ab initio* simulations, such as per-atom forces and defect energies.

2.3.1 Pair Potentials

The first, most simplistic analytic potentials were designed to model liquids, by describing the interatomic forces between nearest neighbour atom pairs. Such functional forms are known as pair potentials, with two examples being the Morse[74] and Lennard-Jones[75] models. Graphical representations of these potential models are illustrated in figure 2.2, each with unit parameter values. For a pair of atoms separated by a distance r, the Lennard-Jones potential is defined as

$$V(r) = 4\varepsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6} \right], \tag{2.31}$$

and the Morse potential has the form

$$V(r) = D_e \left(\left[1 - e^{(-a(r-r_e))} \right]^2 - 1 \right). \tag{2.32}$$

Both potentials describe pairwise interactions, with repulsive behaviour (i.e. negative gradients) for particles with separations less than that of their diameter, and attractive forces (i.e. positive gradients) between nearby particles, up to some cut-off distance. The equilibrium interatomic distances between particles for Lennard-Jones (LJ) and Morse are defined at $\frac{dV(r)}{dr} = 0$, which corresponds to $r_{LJ} = \sqrt[6]{2}\sigma$, and $r_M = r_e$, respectively. The depth of the potential well is given by ε, in LJ, and D_e for Morse. The Morse potential contains an additional system-dependent parameter, a, which inversely controls the width of the potential, i.e. a smaller value of a gives a wider potential. Although initially designed to model liquids, pair potentials can be used to describe metallic systems with some success.[76] However, an inherent limitation of such potentials are their inability to reproduce all three elastic constants simultaneously due to having too few parameters (i.e. degrees of freedom).[77]

2.3.2 The Embedded Atom Method

In the late 1970s and early 1980s, there was a significant effort to construct simpler descriptions of impurities in metals, namely hydrogen. At the time there were two significant approaches to the problem: a local self-consistent cluster-based treatment of the impurity and its nearest neighbours[78] and the "jellium" model, which treated the impurity as an em-

bedding in a homogeneous electron gas (also used as the rationale for the LDA functional in DFT). In 1980, Stott and Zaremba[79] considered the "quasi-atom" approach to simulating embedded impurities, by treating the impurity and its electronic screening cloud as a single unit. From this approach, advances were made in formalising what is now termed the transfer function in the embedded atom method (EAM) potential.

The EAM potential is comprised of three functional parts: a pair potential term, an embedding function, and a transfer function. In 1982, Jens Nørskov proposed a covalent bonding contribution to the embedding term,[80–82] which helped to generalise the description to more complex metals. Finally, the embedded atom method (EAM) as it is now known, utilises a pairwise embedding function, and was finally presented one year later by Daw and Baskes,[83] as a generalisation of the quasi-atom approach used to model the presence of hydrogen in metals.

The functional form of the EAM potential is separated into three functions, a pair term, an embedding and transfer function. The energy of an atom i, surrounded by N neighbours can be written

$$E_i = \frac{1}{2}\sum_{i\neq j}^{\mathrm{N}} \Phi_{\alpha\beta}(r_{ij}) + F_\alpha(n_i) \quad \text{where} \quad n_i = \sum_{j\neq i} \rho_\beta(r_{ij}). \tag{2.33}$$

The pair term, $\Phi(r)$, is a typical pair function, as above (i.e. LJ or Morse), and predominantly describes the symmetric nuclear repulsion. The improved accuracy of the EAM potential in equation (2.33), is a result of the inclusion of a density-dependent embedding term, $F(n)$ and transfer function, $n(\rho(r))$. These functions account for the effect of delocalised electrons in metals, and the resulting non-linear electronic contributions to the local environment of an atom.

A common misconception of the EAM formalism, is that the embedding term describes the energy required to embed an ion core into the local electron density. However, due to the gauge degrees of freedom in the analytic form as shown in equations 2.34, it can be demonstrated that the linear contribution to the electron density can be shared between the pair contribution, and the embedding term. Therefore, a complete description of the system can only be understood when considering the EAM functional form as whole and

not as separable constituent parts.

$$\begin{aligned} \Phi_{\alpha\beta}(r) &\mapsto \Phi_{\alpha\beta}(r) + \lambda_\alpha \rho_\alpha(r) + \lambda_\beta \rho_\beta(r) \\ F(n_\alpha) &\mapsto F(n_\alpha) - \lambda n_\alpha. \end{aligned} \tag{2.34}$$

The EAM functional form in equation 2.33, and description of gauges in 2.34, are generalised to describe interactions between different elements. Here, α, denotes the "embedded" atom type of interest, with β describing the effect of interactions with potentially different neighbouring elements. Within this work, only single element materials have been modelled thus far, and so the subscripts α and β can be dropped.

2.3.3 The Force Matching Method

Interatomic potentials are an invaluable tool in materials modelling, as they reduce the complexity of the problem by reducing dimensionality in the configuration space. DFT provides the forces on each atom in a material for a given morphology and stoichiometry. In the force-matching method,[1] an interatomic potential is fitted to model a given system, by matching the derivative of an initial guess potential model to the array of *ab initio* forces for various instances of the system (i.e. geometries, structures, thermodynamic phases, surfaces). The parameters of an interatomic potential model are then adjusted to optimally reproduce the forces, energies and stresses, which have typically been obtained from DFT calculations. There are a variety of open source codes available which implement the force matching method, in this work the *potfit*[3–6] implementation is used. In *potfit*, the potential parameters, $\boldsymbol{\theta} = \{\theta_1, ..., \theta_{\mathrm{N}}\}$, either belong to an analytic potential, or are the values of the tabulated potential function at sampling points.

In order to fit a potential by force matching, the deviation of the interatomic model from the reference data is quantified by the least squares N-dimensional cost function, $C(\boldsymbol{\theta})$, with

$$C(\boldsymbol{\theta}) = \sum_{k=1}^{\mathrm{M}} a_k \left(F_k(\boldsymbol{\theta}) - F_k^0 \right)^2 + \sum_{r=1}^{\mathrm{N_C}} b_r \left(A_r(\boldsymbol{\theta}) - A_r^0 \right)^2. \tag{2.35}$$

where N, is the number of degrees of freedom (i.e. parameters, $\boldsymbol{\theta}$) in the potential model. The first sum, runs over all M force components of the reference configurations, and where F_k^0, is

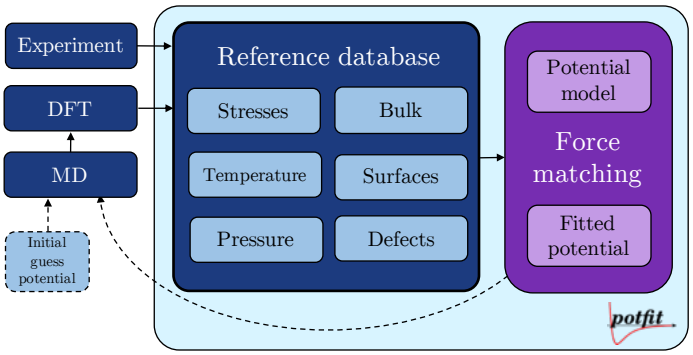

Figure 2.3: The potential fitting workflow, based on force matching, as implemented in *potfit*.

the set of individual atomic forces, each with a weighting, given by a_k. Here, $F_k(\boldsymbol{\theta})$ represents the corresponding set of forces from the potential, for each atom in the configuration. The fit can be enhanced with additional information about the target system energies and stresses (and optionally other quantities), represented by A_r^0. These additional quantities are most often obtained through first principles calculations, and can be given weights b_r, depending on the importance of the accuracy in their descriptions, as given by the potential, $A_r(\boldsymbol{\theta})$. The "best fit" potential parameters $\boldsymbol{\theta}^*$, are those which give the minimum cost value, $C(\boldsymbol{\theta}^*)$.

The weightings a_k (associated with forces) and b_r (energies or stresses) are chosen by inspecting the error sum initially associated to each fitted quantity, and then tuning the weights to improve upon the prediction of a particular quantity at expense of the others. This then allows for finer control over the predictive capabilities of the fitted potentials, by encouraging greater precision in the quantities most closely related to its intended use.

Figure 2.3, illustrates the *potfit* potential fitting workflow, with the light-blue box indicating sections of the fitting process handled within the *potfit* software. On the left of the illustration are three boxes outlining sources of the atomistic forces, energies and stresses, which constitute the reference data. For reference data generated using molecular dynamics, an initial guess potential must of course be used, and so, as indicated by the curving dashed line, it may be favourable to iteratively supplement such reference snapshots with those produced using subsequently fitted potentials. This iterative approach can effectively be circumvented however, by performing singlepoint DFT calculations on each of the MD reference snapshots, to obtain more accurate forces, energies and stresses, independently of the interatomic potential being fitted. In this way, the MD portion of the reference data

generation is only required to sufficiently perturb the atomic positions by enough to create a spread of interatomic forces, with which an interatomic potential can be fit. As a range of forces are required to fit a potential model, it is important when building the reference data to include an array of morphologies and environments. This is indicated by the variety of subsections within the reference database in figure 2.3. The purple boxes within *potfit*, are an incredibly simplistic representation of the computational components of the program at which the force matching is performed.

Whilst sufficient variety in the reference data is important in order to provide a greater quantity of distinct data-points to which the potential is fitted, it is imperative that each reference snapshot is relevant. This means that the reference snapshots chosen should be representative of the system the potential will be used to simulate, i.e. a potential fitted to predict quantities at room temperature should not contain obscenely high temperature reference snapshots in the fitting data.

Until this point, the force matching method provides us with a way to quantify the fit quality, of an interatomic potential,to the reference data, through the cost value. However, it remains unclear how one would proceed in finding which potential parameter values would yield an improved potential fit (i.e. lower cost value). The following subsections outline the methods of space searching (local and global optimisers), which can be used to find improved potential parameters. However, before proceeding to this discussion, the next subsection summarises and describes the nature, and basic construction, of the variety of landscapes (most importantly the cost landscape), encountered within the potential fitting process.

2.3.4 Landscapes of the Potential Fitting Process

It is beneficial to clarify the coordinates of each mathematical space encountered throughout the potential fitting process. Three such spaces have been indirectly encountered up to this point: configuration space, phase space and the cost landscape.

For a fixed number of atoms, n, configuration space describes all possible atomic positions of the system. Therefore, as each atom has three Cartesian coordinates defining its position, $\boldsymbol{r}_i = (x_i, y_i, z_i)$, the dimensionality of configuration space is $\mathbb{R}^{3n}$. The coordinates of a single point in configuration space, contain the positions of all atoms in the system at a

single instant. Moving in configuration space is equivalent to changing the system structure by moving atoms. As the Hamiltonian of a system can be separated into kinetic and potential energy terms, where the potential energy is defined solely by the atomic positions, each coordinate in configuration space has an associated potential energy. Therefore, moving through configuration space, becomes a traverse across a potential energy surface (PES).

Similarly, points in phase space are defined by the atomic positions, additionally with the components of momentum, $\boldsymbol{p_i} = (p_{x_i}, p_{y_i}, p_{z_i})$. Therefore, as phase space contains twice as many coordinates defining each point, its dimensionality is $\mathbb{R}^{6n}$, where n, is the total number of atoms. As kinetic energy contributions are a function of the momenta, in addition to the potential energy configuration space information, each coordinate in phase space can be evaluated to obtain the total system energy. In the discussion of thermodynamic ensembles, this concept is briefly mentioned in equation (2.30), where the probability of a system occupying a given microstate is related to the sum of all phase space coordinates with the same energy (i.e. energy isosurface). Phase space is also explored throughout a molecular dynamics simulation, through the time evolution of the system, where paths in phase space (across the surface of microstate total energies) are forged.

Finally, the cost space introduced in the potential fitting process is fundamentally different in nature to the two above, and upon reflection should perhaps more appropriately be referred to as the cost surface (or landscape). The construction of cost space begins with the definition of the reference data, which may contain as many as m reference configurations. These datasets provide m discrete sample of points in configuration space (positions), along with their first derivatives (forces as $F = -\nabla E$) and the potential energy of each point (total configuration energy), as well as the stress components associated with the coordinate (which links the positions and forces, i.e. for uniaxial normal stress, $\sigma = F/A$). The potential model being fitted, defines a potential model space with dimensionality equal to the number of fitted potential parameters, $\boldsymbol{\theta} = \{\theta_i\}_{i=1}^{\mathrm{N}}$. Each point in this potential parameter space, is a set of parameters corresponding to a different interatomic potential, which are each potential energy functionals of atomic separation, r. The cost landscape is of the same dimensionality and coordinate system as the potential parameter space. However, the evaluation of each point is a cost value, defined by its fit to the reference data. Moving through the cost

landscape is analogous to changing the potential parameters, θ_i. The space explored within the potential fitting process, is then essentially the potential parameter space, relative to the reference data. In order to choose directions within the cost landscape with decreasing cost values, generalised minimisation algorithms, designed for any type of space searching of this nature, are required.

2.3.5 Global Minimisers and Local Minimisers

An understanding of the cost landscape defined by the potential fitting process can quickly become unintuitive, as in the EAM model, where spaces of 10 dimensions (i.e. potential parameters), or greater, are explored. As the potential which most accurately reproduces the reference data has the lowest cost in the landscape, intelligent techniques to efficiently and effectively navigate the cost space must be employed.

Herein lies an open question in multidimensional space searching, how does one know when the global minimum of such a space has been found? Exhaustive mapping of the space is computationally infeasible, and therefore the development of new algorithms to efficiently search for a minima is an area of ongoing research within the computer science community.[84,85] Within *potfit*, the current best-practice approach is to first invoke a "global" minimiser, to find the region of space most likely to contain the minimum, followed by a "local" minimiser, to descend to the lowest point in the minimum basin. Two popular global minimisation algorithms which are implemented in *potfit* are simulated annealing,[86] (used in this work) and differential evolution.[87]

Simulated annealing follows from the process of annealing in metals; where, with a near molten metal, the temperature is gradually reduced or quenched, allowing the atoms to migrate into their optimal crystalline positions. The simulated annealing heuristic follows similarly, where the space is searched and the notion of "temperature" introduced as a constraint on the size of steps taken whilst searching. Downhill steps in the landscape are accepted outright, whereas those which increase the cost are accepted with a probability relating to the temperature, $P_{\text{acc}}(\Delta C) = \exp(-\Delta C/T)$. In this way, the size of the searching steps are gradually reduced, i.e. temperature decreases, and the algorithm settles into what is believed to be the deepest basin. As will be discussed in greater detail in section 2.4,

simulated annealing is an example of a Monte Carlo method, which stem from the work of Metropolis *et al.*[88]

Upon isolating the region likely containing the landscape minimum, a local minimisation algorithm is then used to descend to the bottom. The local minimiser employed in *potfit* is Powell's method,[89] which only requires a single numerical calculation of the basin gradient at initialisation, unlike other conjugate gradient methods, which require an explicit gradient calculation at each step.[90]

It is not always the global minimum which is of greatest interest to researchers, particularly in chemistry, where for example the aim could instead be to model chemical reactions.[91] The saddle-points of a landscape, known as reaction pathways, map out the preferred method of transition from one state to the next. The dimer method,[92] can be used to probe transition states and reaction pathways. The method relies on eigenvalue decomposition to investigate the second order moments of the eigenvalues, obtained from the Hessian, coupled with kinetic Monte Carlo moves, to evolve the system over large timescales. In the uncertainty quantification method implemented in this work, eigenvalue decomposition of the Hessian is also used, and is explained in section 2.5. The kinetic activation-relaxation technique (k-ART) is another method which can be used to describe transitions in condensed matter.[93,94] The k-ART technique has, for example, successfully described movement of interstitials between favourable states in metals.[93]

The final type of configuration space exploration of interest, is concernred with finding new materials and morphologies. *Ab initio* random structure searching (AIRSS),[95] is based upon throwing random concentrations of atoms into a hypothetical box and relaxing them, using DFT, to the most favourable energetic configuration. As a result of AIRSS simulations performed at high pressure, new crystalline morphologies, believed to describe those found inside stars, have been discovered.[96] As with gradient based methods, such as Powell's, the AIRSS dartboard approach relies on the assumption that the deepest potential basins are also the widest.

2.3.6 Reference Data

The aim of a good effective potential, is to match the relevant area of the configuration space as closely as possible, at the sampling points provided in the reference data. The testing set then contains previously unseen points in configuration space. It is assumed that the interpolation of the potential at the new training points will match their true configuration space values sufficiently closely, i.e. to within the fitting error. A common issue in the building of potentials, can be under-, or over-fitting, of the potential.

Under-fitted potentials are the results of sparse training sets. Sparse sets typically contain too few sampling points in the minima and saddle-points of the configuration space. A configuration space minimum typically describes the most energetically favourable arrangements of atoms, such as the ground states of a system, over a variety of pressures and morphologies. Saddle points in configuration space however, frequently describe transitions between such energetically favourable states, for example a chemical reaction process or the movement of an electron-hole pair in semiconductors.

The over-fitting of potentials can be pictorially envisioned as a wrinkling of the potential cost space, compared to the configuration space. This is often the result of too many parameters in the analytic potential form. This means that a complex potential is attempting to capture motifs in the energy landscape which do not exist; the potential wrongly interpolates between training points which it has matched exactly. An unnecessarily complex analytic potential model can be identified by principle component analysis (PCA), which illuminates which parameters contribute least to the fit.

A high quality potential should aim to capture the same trends and motifs of a given configuration space, but need not match exactly for all points in the training reference configurations. Interatomic potentials are ultimately predictions on the shapes of configuration space, given what is already known about particular systems. Therefore when building the reference data, there is an assumption that the configurations chosen for the potential fitting, are representative of the areas of configuration space which are to be modelled by the interatomic potential. Were the reference data ill-chosen in this manner, this may result in the most desirable potential for the simulation of quantities of interest not being that with the lowest cost value. The suitability of the reference data used, introduces the first

of multiple sources for error introduction into the system. This specific type of error affects the specificity (i.e. the precision) of the fitted model, and is discussed in greater detail in section 2.3.7.

A second type of reference data error is due to the quality of the fitting configurations used. The quality of reference data, controls the error introduced into the system, from approximations made in the *ab initio* simulations. Reference data should typically contain both equilibrium and non-equilibrium configurations, the latter of which can be created either directly from *ab initio* MD simulations, or use a snapshot of atomic positions from a classical MD simulation, upon which a DFT calculation is performed to obtain more accurate atomic forces, energies and stresses. One important benefit of the inclusion of non-equilibrium reference data, is that it provides a significantly greater amount of information (especially in forces) for the potential fitting than, for example, a crystalline, 0 K, snapshot would. Therefore, the *ab initio* calculations performed on the non-equilibrium data solve the Kohn-Sham equations for the fixed atomic positions provided, and are known as singlepoint calculations. For all DFT calculations, there is a tolerance to which the Kohn-Sham equations are solved which directly affects the precision of the results. Therefore, one of the most important steps when creating a reference dataset is to ensure that the tolerance used for the DFT calculation is small enough that the calculations are consistent with one another, i.e. that they have converged. The tolerance of the convergence can be controlled predominantly in two ways; through the cutoff energy of the planewave basis set, and through the density of the k-point grid. By ensuring the parameters used in the DFT calculations are such that the they have converged to a high enough precision, uncertainties from the reference data can be minimised.

2.3.7 Sources of Uncertainty

The decision to employ an interatomic potential introduces a variety of sources of error into a simulation, which will either affect the accuracy or precision of predictions. Figure 2.4, illustrates the interplay between predictive accuracy and precision. The uncertainty in interatomic potential parameters is quantified in this work, whereby the accuracy of the potential is assessed by the error in fitting to reference data. The model choice and reference

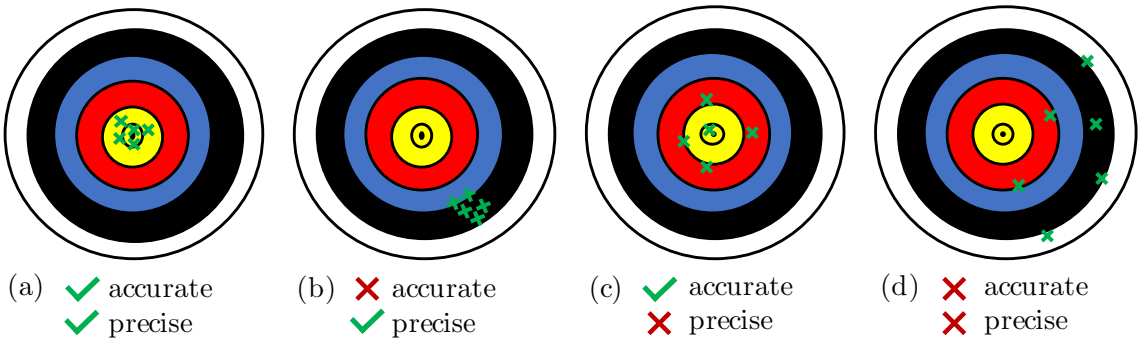

Figure 2.4: A pictorial representation of accuracy vs. precision. Each trial (green cross) represents the result of a simulation, with the centre of the target representing the true value.

data play a role in the precision of predicted quantities of interest which use the potential. The precision of predictions informs on the uncertainty in each value.

Figure 2.4 illustrates the results of four differing levels of accuracy and precision in measurements. It is clear that the aim of potential fitting is to create a fitted potential which is both accurate and precise in its representation of the system being modelled, corresponding to the first subfigure. However, as is the case in figure 2.4b, it is possible to fit a potential which produces consistent predictions for similar systems, although the predictions are systematically incorrect, or off. This can arise in cases where the potential has been fitted to reference data which does not represent the configuration space of the systems it is used to simulate. The evaluation of a potentials ability to simulate configurations which it was not designed for is known as the transferability of a potential.

There are a second set of circumstances with which scenario 2.4b can arise, relating to the choice of potential model being fit. If a potential model is of insufficient complexity to capture the nuances of the system of interest, this can manifest through systematically incorrect predictions of QoI, for configurations similar to those included in the reference data. This is one of the presentations of insufficient model choice.

In figure 2.4c, the results are considered accurate, as they are equally distributed around the true value, however the predictions are lacking in precision, due to their large spread. This case is representative of a fitted potential which fits well to the reference data, i.e. has a small cost value, but is indicative of a reference dataset which is sparely populated with the target configuration spaces being simulated. To improve a potential fit

in this instance, the reference data could be complemented with additional configurations, created using the initial fitted potential, highlighting an accepted best-practise belief that the process of potential fitting should be iterative. There are additional subtle scenarios in which figure 2.4c could occur; one being that the fitted model lacks the complexity required to accurately predict the simulated quantity, and another being that the results are for a robust simulated quantity, which is insensitive to small perturbations of the fitted potential parameters.

The final case shown in figure 2.4d, contains both inaccurate and imprecise results. As in figures 2.4b and 2.4c, such results can arise from a combination of insufficient model choice and unrepresentative reference data. Finally, the scenario in figure 2.4d could also arise if the fitting algorithms for the potential employed were prematurely stopped before converging to the global cost minimum.

Up to this point, the sources of uncertainty discussed have been limited to those introduced by the user in the potential fitting process. There are of course inherent uncertainties associated with employing an effective potential in place of an *ab initio* simulation of the desired quantities. When fitting to *ab initio* reference data, this relative model error, introduced when using a specific potential model, is reported by *potfit* as a total error sum (cost) in the forces, total energy and stresses. When creating a new potential for productions use, the convention is to withhold a portion of the reference data during the fitting phase so that it may be used in testing. As the focus of this work is the development of an uncertainty quantification workflow and not the creation of a new production-grade potential, the entire reference data set is utilized in the fitting process.

In addition to the the model error, the reference data itself carries an inherent error resulting from the quality of DFT simulations used in its production. Convergence tests of the planewave cutoff energy and k-point density are used to select the accuracy of calculations, and give the maximum error in the reference data energies, forces and stresses. Good quality reference data will typically contain errors which are less than 1 % of the associated potential fitting error sums.

When an interatomic potential is used in MD simulations to predict time-averaged quantities of interest, thermal fluctuations impart yet another error into predictions. In

order to isolate the error specifically resulting from the analytic potential, in this work the same seed value is used so that the resulting error for a given timestep is isolated to that from the potential.

An often overlooked source of error arises from the choice of software. Design and implementation choices amongst simulation codes affect the precision of predictions. The Δ-project[24] is an initiative to benchmark a variety of common DFT simulation codes, in order to compare the predictions resulting from implementation choices. As the same software is used consistently in this work, this does not impact the comparison of results.

Finally, within the workflow of this research are additional possible sources of error due to implementation differences of the analytic potential models in fitting and MD codes. This potential for error can be easily addressed by ensuring consistency in force and energy predictions for a test system. Additionally, the OpenKIM API[39–41] provides a framework which integrates with existing software to standardise the implementation of published potential models.

2.4 Monte Carlo Methods

The aim of this work is to implement a method which can quantify the parameter uncertainties, as well as those in predicted quantities of interest, in an effective potential fitted using *potfit*. In order to sample the space defined by the fitting of a potential to the reference data (the cost space) Markov chain Monte Carlo[97] is used. This section introduces the necessary background to Monte Carlo sampling and how it can be applied within the *potfit* force matching framework to build an ensemble of potentials which inform of the parameter uncertainty in a fitted potential.

Monte Carlo (MC) methods randomly sample spaces to build up a numerical approximation of the underlying distribution. This class of algorithms are commonly employed when the landscape of interest is difficult to map directly; either because it is too complex to define analytically, or the exact shape/distribution is not yet known. Unlike more direct methods such as Latin hypercube sampling[98] or Gaussian process regression,[99] MC methods rely on their inherent randomness in order to overcome the 'curse of dimensionality', with

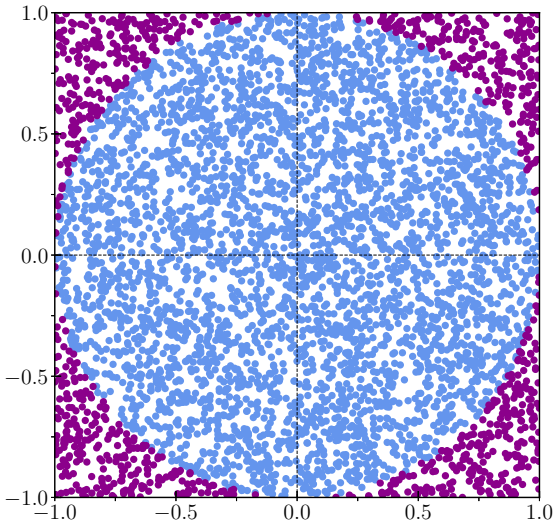

Figure 2.5: A simple example using the Monte Carlo method to approximate the value of π. The coordinates of 5000 points are randomly drawn from a uniform distribution, with those falling within a circle of 0.5 radius in blue.

a convergence rate of $1/\sqrt{N}$ (where N is the number of samples).[100]

Figure 2.5 illustrates simple application of the Monte Carlo method in estimating the value of π. A circle with area πr^2 is enclosed by a square with area $(2r)^2$. Therefore, the value of π can be approximated from the ratio of areas of the circle to square, $\pi r^2/4r^2 = \pi/4$. By following the Monte Carlo method, the coordinates of points are randomly drawn from a uniform distribution, and the ratio of the number of points from inside the circle, to those inside the square (i.e. total number of points drawn), give an approximation of $\pi/4$. In figure 2.5, the 5000 MC samples result in an approximation of $\pi \approx 4 \times N_{\text{circle}}/N_{\text{total}} = 3.1664$. Another example of a MC method, already introduced in section 2.3.5, is simulated annealing. Randomness is used to chose steps, with the range of the perturbation dependent on the gradually reducing temperature (i.e. decreasing range).

2.4.1 Monte Carlo in Materials Simulation

In computational chemistry, Monte Carlo simulations are often instead of molecular dynamics, to understand the reaction pathways of molecules undergoing chemical reactions.[101] MC simulations are also the most common tool employed in the research of protein folding,[102] as well as in the majority materials simulations within which transitions between favourable

system states are not well understood.[103,104] This is because the randomness underpinning MC algorithms make them significantly more successful at simulating the various atomistic states of a reaction, which are often improbable transitions to capture with MD, particularly if the density of the system is low.[101] The minima of the PES explored through MC describe favourable states of the system, which may include the state of the system before and after a reaction,[105] crystallisation of the atoms/molecules involved[106] or adsorption of molecules to a surface,[107] among many other possibilities.[108]

2.4.2 Monte Carlo Methods in Force Matching

As outlined in section 2.3.3, the landscape to be explored in this work, is the cost space, defined when fitting an effective potential to reference data. As the dimensionality of the space is defined by the number of potential fitting parameters, the landscape quickly becomes impossible to visualise fully for potentials with > 3 parameters. Furthermore, as the values of the parameters in the effective potential are the coordinates in the cost space, MC can be used to perturb these parameters in order to move throughout the space.

In order to carry out the uncertainty quantification of potential parameters in cost space, it quickly becomes clear from an assessment of the available information, that a MC method must be used. This is because the cost value (i.e. relative height) of a point in cost space can only be calculated from the corresponding parameter values. Given a point in cost space, it is impossible to chose to move up/downhill in cost without trialling all neighbouring parameter sets in order to find the direction. An alternative way to think about this, is that the force matching equation in equation (2.35) is an injective map, taking the effective potential parameters from parameter space to the reference data-dependent cost space. Therefore, due to the summation of forces ($F(\boldsymbol{\theta}) = dV(\boldsymbol{\theta}, r)/dr$) in equation (2.35), it is impossible to recover the correct potential parameter coordinates from a cost value, as the mapping is not necessarily one-to-one.

2.4.3 Markov Chain Monte Carlo

The example in figure 2.5 introduces the concept of randomness, which underpins all MC methods. In Markov chain Monte Carlo (MCMC), each sample is connected to the one

before it, unlike in figure 2.5 where the samples are drawn independently. Therefore, MCMC algorithms use random walks (i.e. Markov chains) to explore the space. In the case of a cost landscape, MCMC can be used to efficiently sample around the minimum of the space, by making perturbations to the potential parameter values. So in cost space, a Markov chain denoted S, takes steps in parameters such that $S(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}') = C(\boldsymbol{\theta}) \rightarrow C(\boldsymbol{\theta}')$.

Unlike the MC method in figure 2.5, MCMC samples are autocorrelated due to the connectedness of steps in the random walk. Therefore, when using Markov chains to sample an underlying distribution, it is necessary to use a subset of decorrelated samples in the analysis. Furthermore, although not of concern to this work, as explained later in section 2.5.2, most MCMC methods require a "burn-in"' period, within which the Markov chain is allowed to converge to the underlying distribution.

An obvious concern when sampling a local cost space minima would then be that the Markov chain leaves the area of interest, in a manner similar to the transition state searching MC methods. Therefore, to sample from the local cost space distribution of interest, a subset of MCMC methods is used, known as Metropolis-Hastings methods.[88,109]

The Metropolis-Hastings Method

In a Metropolis-Hastings algorithm, the concept of rejection sampling is applied to the Markov chain, in order to encourage sampling from the distribution of interest. This is where a "critic", $P_{\text{acc}}(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}')$, is employed to assign a probability to accepting each proposed step. By introducing the ability to reject steps to the sampling, a new Markov chain, $Q(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}')$, is defined where

$$\begin{aligned} Q(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}') &= S(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}') P_{\text{acc}}(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}') \quad \text{for} \quad \boldsymbol{\theta} \neq \boldsymbol{\theta}', \\ Q(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}) &= S(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}) P_{\text{acc}}(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}) + \sum_{\boldsymbol{\theta} \neq \boldsymbol{\theta}'} S(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}')(1 - P_{\text{acc}}(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}')). \end{aligned} \tag{2.36}$$

As we are interested in sampling the from parameter sets with cost values near the minima of interest (i.e. within the minima basin), it is important to ensure that the newly devised acceptance probability samples from a stationary distribution. In simpler terms, this means

that we check that the probabilities defined by P_{acc}, relate to samples taken from the cost minimum, and for instance, do not actively encourage the Markov chain to leave the minimum basin. The minimum basin distribution, $\pi(\boldsymbol{\theta})$, which the Markov chain samples from, is then said to be stationary if $\pi(\boldsymbol{\theta}') = \pi(\boldsymbol{\theta})Q(\boldsymbol{\theta} \to \boldsymbol{\theta}')$.

It can be shown that the distribution sampled from is stationary if the Markov chain is reversible. This is known as satisfying the detailed balance equation:

$$\pi(\boldsymbol{\theta})Q(\boldsymbol{\theta} \to \boldsymbol{\theta}') = \pi(\boldsymbol{\theta}')Q(\boldsymbol{\theta}' \to \boldsymbol{\theta}). \tag{2.37}$$

It is simple to prove that a Markov chain satisfying the detailed balance condition samples from a stationary distribution, by summing on both sides over all possible parameter sets, to recover the equation for a stationary distribution,

$$\begin{aligned} \sum_{\boldsymbol{\theta}} \pi(\boldsymbol{\theta})Q(\boldsymbol{\theta} \to \boldsymbol{\theta}') &= \sum_{\boldsymbol{\theta}} \pi(\boldsymbol{\theta}')Q(\boldsymbol{\theta}' \to \boldsymbol{\theta}) \\ &= \pi(\boldsymbol{\theta}') \underbrace{\sum_{\boldsymbol{\theta}} Q(\boldsymbol{\theta}' \to \boldsymbol{\theta})}_{=1} = \pi(\boldsymbol{\theta}'). \end{aligned} \tag{2.38}$$

The final step in the construction of the Metropolis-Hastings MCMC method, is to decide upon the form of the acceptance probability, P_{acc}. For the style of MCMC of interest to this work, the acceptance probability works to constrain the searching of the cost landscape to the minima of interest. Therefore, for a Markov chain beginning at the minimum of the cost landscape, the probability of accepting steps which move away from this point (i.e. uphill) should be less than the probability of the reverse (i.e. moving from uphill to downhill). Therefore in this work the acceptance probability distribution is not symmetric. This is due to a combination of the eventual choice of step perturbation function, $S(\boldsymbol{\theta} \to \boldsymbol{\theta}')$, as well as the choice of stationary distribution, $\pi(\boldsymbol{\theta})$, being sampled from. Both functions used in this work are defined and discussed in detail later, in section 2.5, when the uncertainty quantification method is introduced. At present, the only relevance of the asymmetric acceptance probability distribution choice, is that this design choice accounts for the inequality in the following equations. Using the detailed balance equation in 2.37, the choice of acceptance

probability distribution can be formalised as follows

$$\pi(\boldsymbol{\theta})\overbrace{S(\boldsymbol{\theta}\to\boldsymbol{\theta}')P_{\text{acc}}(\boldsymbol{\theta}\to\boldsymbol{\theta}')}^{Q(\boldsymbol{\theta}\to\boldsymbol{\theta}')} = \pi(\boldsymbol{\theta}')\overbrace{S(\boldsymbol{\theta}'\to\boldsymbol{\theta})P_{\text{acc}}(\boldsymbol{\theta}'\to\boldsymbol{\theta})}^{Q(\boldsymbol{\theta}'\to\boldsymbol{\theta})},$$

$$\frac{P_{\text{acc}}(\boldsymbol{\theta}\to\boldsymbol{\theta}')}{P_{\text{acc}}(\boldsymbol{\theta}'\to\boldsymbol{\theta})} = \frac{\pi(\boldsymbol{\theta}')S(\boldsymbol{\theta}\to\boldsymbol{\theta}')}{\pi(\boldsymbol{\theta})S(\boldsymbol{\theta}'\to\boldsymbol{\theta})} \leq 1,$$

$$P_{\text{acc}}(\boldsymbol{\theta}\to\boldsymbol{\theta}') = \min\left\{1, \frac{\pi(\boldsymbol{\theta}')S(\boldsymbol{\theta}\to\boldsymbol{\theta}')}{\pi(\boldsymbol{\theta})S(\boldsymbol{\theta}'\to\boldsymbol{\theta})}\right\}. \tag{2.39}$$

2.4.4 Multidimensional Markov Chain Monte Carlo

In this work, the cost function in (2.35) is being sampled at different points around the cost space minimum. This is the stationary distribution, $\pi(\boldsymbol{\theta})$ from section 2.4.3, which is effectively a Gaussian distribution of the forces predicted by the potential model, $\mathcal{N}(F^0, T^2 I)$. This distribution is centred on the cost minimum force predictions, F^0, with the variance given by a sampling temperature, T, which is contained in the "critic", P_{acc}, and detailed in equation (2.43).

As had been implied throughout the discussion of Monte Carlo methods above, this work samples from cost landscapes of dimensionality equal to the number of potential parameters being fitted. Therefore, when sampling in these potentially high dimensional spaces, Makov chains can suffer from the "curse of dimensionality", which is expressed through long decorrelation times, and high step-rejection rates. Of particular interest to multidimensional MCMC, are a few design choices when creating the sampler which work to circumvent these issues.[110]

The first improvement relates to the size of Markov chain steps. As the degree of basin curvature may differ across the directions (i.e. in sloppy models), taking equally sized steps in all parameter directions is highly inefficient. If steps are too large in a direction of steep curvature, this will lead to high rejection rates and long simulation times. However, in contrast, for steps which are too small in slowly varying (i.e. sloppy) directions, there will be a high acceptance of steps, yet slow exploration of the underlying distribution, which results in poor convergence. Therefore, it follows that a Markov step size which is proportional to the degree of curvature in each direction would most efficiently sample the

space. It is worth noting at this point, that in reality, the principle axes of curvature in high dimensional landscapes may not align with the parameter directions which define the space. Therefore, when adapting the MCMC step size based on the basin curvature, there is often some degree of translation required between curvature directions and parameter directions. This is discussed in more detail in section 2.5. Additionally, as will also be discussed in the following section, one design constraint of a curvature-informed step size, is that the sampled basin is symmetric in curvature along each of the principle axes, which may not be the case. In section 3.7 we have implemented an algorithm which allows the user to interrogate the basin symmetry in their system, before beginning the uncertainty quantification, if desired.

The second, more obvious design consideration, is to ensure that steps are made in all parameter directions at once. Should steps only change a single parameter of the coordinate set at a time, then sampling of the underlying distribution would require significantly more steps to obtain the true covariances between parameters. Finally, as reported by Roberts *et al.*[110] the ideal balance between step size, and rejection rate, is for 23 % of MCMC moves to be accepted. Therefore, having calculated the relative differences in step size between parameters based on the curvature, the whole set of parameter step sizes should be tuned to offer an acceptance rate of 0.23. Having introduced the notion of using Markov chains to sample high dimensional spaces, we can begin the formal discussion of the ensemble method, implemented in this work as a means to quantify the uncertainties encountered in the process of fitting potentials.

2.5 The Potential Ensemble Method for Sloppy Models

As introduced in section 1, the cost landscapes encountered in potential fitting frequently fall into a class of models known as sloppy models.[27,38] The hallmark of a sloppy model, is that the basin encapsulating the minimum has significantly differing degrees of curvature along the principal axes. Therefore, the majority of interatomic potentials fitted in *potfit* will fall into the category of sloppy models. Specifically in sloppy models, in some directions the curvature varies incredibly slowly (i.e. they are sloppy directions), relative to in others. This means that without adjustment to traditional Markov chain Monte Carlo sampling

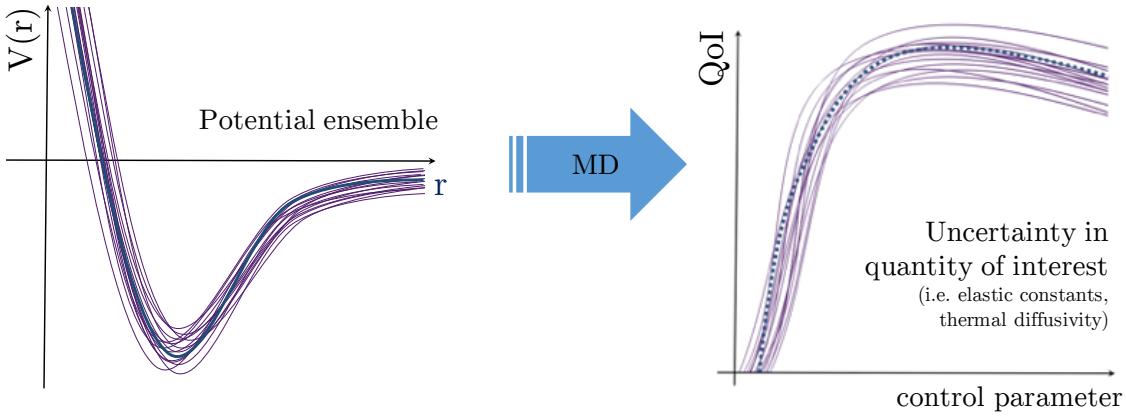

Figure 2.6: The left hand figure represents an ensemble of Lennard-Jones potentials sampled by performing MCMC within the cost minimum basin. By propagating each ensemble member through individual MD simulations (each calculating the quantity of interest) it is illustrated in the second figure how the uncertainties in the quantity may be propagated from the uncertainty in potential parameters.

methods, the slow exploration of sloppy directions result in inefficacious sampling of the space and painfully slow convergence to the underlying distribution. By using a MCMC step size which scales with the rate of curvature in the direction of movement, these issues can largely be circumvented. Investigation into the sampling of such sloppy models has been extensively undertaken by Brown and Sethna,[37] with a focus on their occurrence in systems biology. The approach has since been outlined for interatomic potentials[27] and forms the basis of the implementation of uncertainty quantification in *potfit*.

Having obtained samples from within the cost minimum basin, these can then be used to quantify the uncertainties within the fitted potential parameters. Having first assessed the convergence of the samples to the underlying landscape, through inspection of individual and $2d$-projections of the parameter distributions, the Markov chain is sub-sampled to acquire an ensemble of uncorrelated parameter sets, representing the spread of possible values in each parameter. This sub-sample contains alternative potentials, of reasonably low-cost, to the fitted potential for the given reference dataset. It is necessary to select decorrelated MCMC samples so as not to unrealistically influence the statistics. For a simple, two parameter Lennard-Jones potential, figure 2.6 illustrates how the potential ensemble may look. On the left side of the illustration, the lowest cost potential is shown in blue, with the alternate ensemble members in purple. Then, as in equation (2.31), the uncertainty in

the parameter ε, is quantified by the vertical spread (depth) of ensemble members at the minimum. With the uncertainty in σ, relating to the range of ensemble values resulting in the horizontal positioning of the potential minimum. As indicated by the arrow, the resulting uncertainties in quantities of interest are obtained through individual MD simulations, each using a different potential from the ensemble. The spread in QoI results then show the uncertainties which are propagated from the potential parameters.

2.5.1 The Hessian

In the cost landscape defined by the potential fitting procedure, there typically is significant covariance between potential parameters, hence the eigen-directions (principle axes of curvature) are used to define the basin curvature. The relative degrees of curvature are given by the eigenvalues of the Hessian matrix at the cost minimum

$$H_{ij}(\boldsymbol{\theta}^*) = \left.\frac{\delta^2 C(\boldsymbol{\theta})}{\delta\theta_i \delta\theta_j}\right|_{\boldsymbol{\theta}=\boldsymbol{\theta}^*} \tag{2.40}$$

where $\boldsymbol{\theta} = \{\theta_i\}_{i=1}^{\mathrm{N}}$ represents a set of interatomic potential parameters. The Hessian calculation is defined in terms of percentage change to each parameter, to overcome the issue of curvature across different length scales. Within the *potfit* implementation, numerical approximations of the partial derivatives, in equation (2.40), are calculated using a central finite difference method. Diagonal components of the Hessian require 2N cost evaluations, with $2\mathrm{N}(\mathrm{N}-1)$ cost calculations required for each diagonal element, where N is the number of fitted potential parameters. The numerical approximations of the Hessian elements are written

$$\begin{aligned} \frac{\delta^2 C}{\delta\theta_i^2} &= \frac{C(\theta_i + h) - 2C(\theta) + C(\theta_i - h)}{2h^2}, \\ \frac{\delta^2 C}{\delta\theta_i \delta\theta_j} &= \frac{[C(\theta_i + h, \theta_j + h) - C(\theta_i + h, \theta_j - h)] - [C(\theta_i - h, \theta_j + h) - C(\theta_i - h, \theta_j - h)]}{4h^2}. \end{aligned} \tag{2.41}$$

The eigenvectors of the Hessian provide a basis for the space, oriented along the principle axes of curvature. The degree of curvature along each principle axis is then indicated by the magnitude of the eigenvalues. Large eigenvalues indicate rapidly changing (steep)

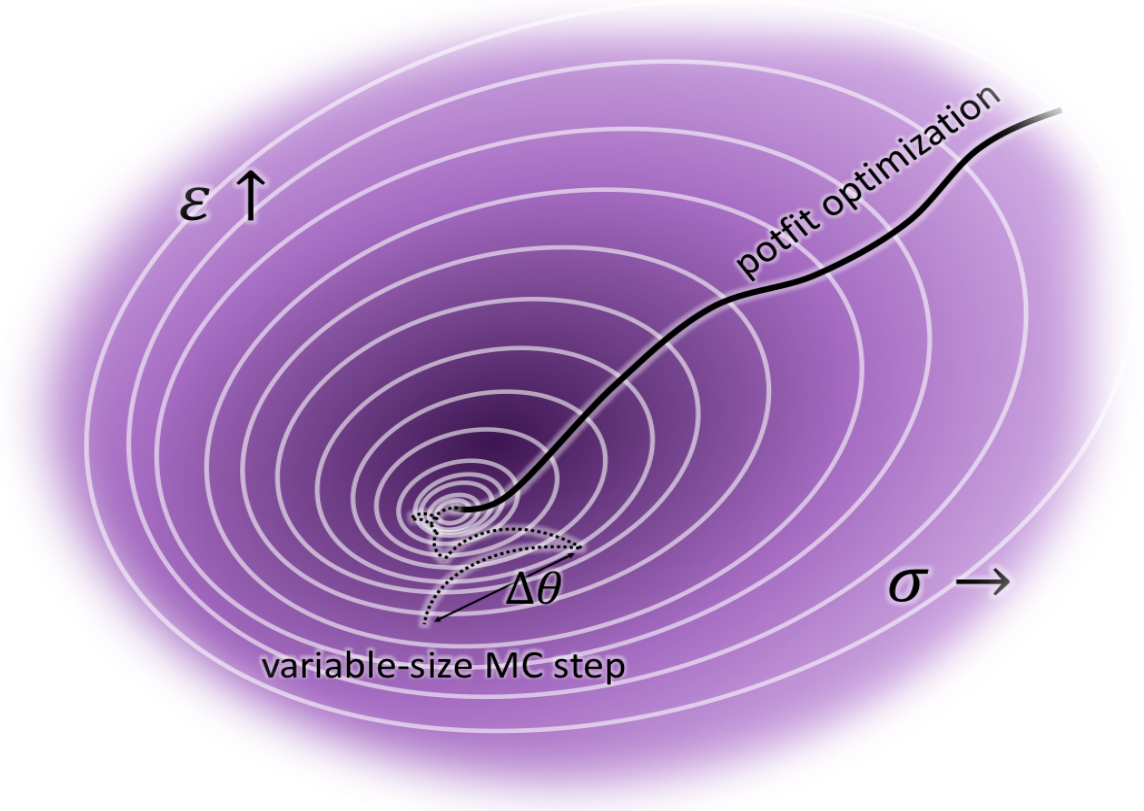

Figure 2.7: A visual illustration of the landscape exploration techniques within *potfit* within a hypothetical Lennard-Jones cost space. The optimisation algorithms within the fitting process are used to locate the minimum, followed by the exploration of the cost minimum basin using MCMC to generate a potential ensemble.

directions, and smaller eigenvalues indicate sloppy ones. Therefore, the eigenvectors can be used to translate between parameter space and the principle axes, with the eigenvalues providing the necessary step scaling information in each direction.

2.5.2 Markov Chain Steps

Candidate steps in the MCMC are generated using random displacements, taking into account information about the curvature, obtained from the eigenvalues of the minimum Hessian. In this way, larger steps are taken in sloppy directions (i. e. those associated with small eigenvalues), with smaller steps proposed in stiffer directions. The steps are taken in cost space, starting from the best fit parameter set, $\boldsymbol{\theta}^*$, by proposing a simultaneous perturbation to each parameter of the form

$$\Delta\theta_i = \sum_{j=1}^{\mathrm{N}} \sqrt{\frac{\mathrm{R}}{\max(\lambda_j, 1)}} V_{ij} r_j \tag{2.42}$$

where $\Delta\theta_i$ is the proposed change to each potential parameter i, R is a system-dependent scaling factor, and V_{ij} is the matrix of normalized eigenvectors of the Hessian. The parameter λ_j, is the j-th Hessian eigenvalue and r_j a normally distributed random number within the range [-1,1].

Figure 2.7 illustrates the general idea behind the *potfit* ensemble method for a hypothetical 2d Lennard-Jones cost landscape. The minimum of the landscape is found using existing *potfit* optimization algorithms, namely simulated annealing followed by Powell's method. After which the minimum basin is sampled using the MCMC steps given in equation (2.42), in order to build an ensemble of alternative candidate potentials for the uncertainty quantification.

As positive eigenvalues of the Hessian confirm that the best fit parameters sit at the minimum of the cost basin, it is unnecessary for the Markov chain to have a "burn-in" period, i.e. a period of initial samples which are later discarded. A "burn-in" is required to allow the Markov chain to move towards the centre of the underlying distribution. As the MCMC begins sampling from the minimum cost basin immediately by design, this requirement is already fulfilled. This can be further verified by observing convergence to the same parameter distributions of Markov chains with and without a "burn-in".

2.5.3 Sampling Temperature

The acceptance criteria for a MCMC step is set by a temperature, T, where the cost minimum is analogous to sampling at a temperature of $T = 0$. The sampling temperature is by default chosen to be the "natural" temperature $T_0 = 2C(\boldsymbol{\theta}^*)/\mathrm{N_p}$ as each mode in a harmonic model contributes an energy of $T/2$.[27] The acceptance probability of each the Monte Carlo move is calculated as

$$P_{\mathrm{acc}}(\boldsymbol{\theta}_{i+1}) = \begin{cases} 1 & \text{if} \quad C(\boldsymbol{\theta}_{i+1}) < C(\boldsymbol{\theta}_i), \\ \mathrm{e}^{-\frac{1}{\mathrm{T}}(\mathrm{C}(\boldsymbol{\theta}_{\mathrm{i+1}}) - \mathrm{C}(\boldsymbol{\theta}_{\mathrm{i}}))} & \text{otherwise.} \end{cases} \tag{2.43}$$

This ensures downhill moves are always accepted, and that MC moves to higher cost potentials are accepted with a probability decreasing exponentially with the increase in cost between potential parameter sets. As in simulated annealing, should it be necessary to further constrict samples to the minimum basin (e.g. the basin is very shallow), the sampling temperature can be reduced. Reducing this temperature restricts the number of accepted MCMC steps, by reducing the probability of acceptance in equation (2.43). The sampling parameter $\alpha \in (0, 1]$, is used to constrict the accepted samples as

$$T = \alpha T_0 = \frac{2C(\boldsymbol{\theta}^*)\alpha}{\mathrm{N_p}}. \tag{2.44}$$

2.6 Calculation of Quantities of Interest

Within chapter 4 the uncertainties in three potentials, fitted to nickel reference data, are quantified. The effect of this error is then propagated through MD simulations, to quantify the resulting uncertainties in predictions of the elastic constants and thermal expansion coefficient at 300 K.

2.6.1 The Elastic Constants

The elastic constants relate the strain generated within a material to the stress being applied. In the case of cubic crystals such as nickel, the symmetries within the structure simplify the number of independent elastic constants from thirty-six to only three: C_{11}, C_{12}, C_{44}. For stresses, σ_{ij}, and strain components ϵ_{ij} the elastic constant tensor is written as

$$\begin{pmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \sigma_{xy} \\ \sigma_{xz} \\ \sigma_{yz} \end{pmatrix} = \begin{pmatrix} C_{11} & C_{12} & C_{12} & 0 & 0 & 0 \\ C_{12} & C_{11} & C_{12} & 0 & 0 & 0 \\ C_{12} & C_{12} & C_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & C_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & C_{44} \end{pmatrix} \begin{pmatrix} \epsilon_{xx} \\ \epsilon_{yy} \\ \epsilon_{zz} \\ \epsilon_{xy} \\ \epsilon_{xz} \\ \epsilon_{yz} \end{pmatrix}. \tag{2.45}$$

Therefore to calculate each of the elastic constants, the strain induced as a result of small deformations to the crystalline material is measured, with $C_{ij} = \partial\sigma_i/\partial\epsilon_j$, as discussed in section 4.5.2.

2.6.2 The Thermal Expansion Coefficient

When subjected to a temperature change materials will deform. The expansion or contraction in one direction (i.e. length change) is written as a rate of change in length for a small temperature difference

$$\alpha = \frac{1}{L_0}\frac{dL}{dT}, \tag{2.46}$$

where L_0 is the initial length of the material. As the linear expansion coefficient of a material is not necessarily constant across all temperature ranges, the coefficient is usually reported as the value at a given temperature. In this work the linear thermal expansion coefficient is calculated for nickel at 300 K, and so five evaluations of the length change are conducted for temperatures equidistantly spaced within the range of $260 - 340$ K, with the final value of the coefficient obtained from the line of best fit.

3

Implementation of the Potential Ensemble Method

"Everything simple is false. Everything which is complex is unusable."

– Paul Valèry

The following chapter outlines the design considerations which underpin the uncertauinty quantification implementation, whilst detailing the technical aspects of the programming. The ensemble method implementation is included in *potfit* at compile time through the `--enable-uq` compile option. To include additional output details the `--debug` flag can also be selected, with the resulting additional information documented at the relevant sections throughout the chapter. The uncertainty quantification workflow is detailed, with the distinction highlighted between the functionality contained within *potfit*, and those aspects of the workflow which lie outwith. The interactions between the ensemble method implementation and the existing *potfit* code are outlined, as well as the input, and output, parameters and files. Finally, the chapter contains a series of investigations, which were per-

formed in the process of developing the implementation in order to better understand the nature of cost landscapes encountered in the potential fitting process. As a result of these investigations, a series of additional input parameters and functionality were introduced to the code, which trigger the programming required in order to reproduce the results of these analyses.

3.1 Design Principles

The ensemble method uncertainty quantification has been written in C, in line with the prior *potfit* development. Due to the generalised, flexible design of the implementation, the ensemble method, like the existing *potfit* code, is not restricted to a single potential model or material. The code has been designed so that it may sit alongside the existing code with minimal communication points between the ensemble work and the potential fitting, in order to minimise potentially costly data transfer. In fact, other than utilising the pre-existing data structures and information contained within them, the only historical function utilised by the ensemble implementation is the calculation of forces routine (`calc_forces()`), which evaluates the cost value at a specified location in the landscape. An existing design consideration within *potfit*, was to separate as much as possible, the optimisation routines from the force matching calculations. This is respected in the ensemble method implementation, as it maintains its independence from the details of the `calc_forces()` routine.

The predominant design consideration of the implementation, was to maintain the flexibility already established within *potfit*. It was important that any developments to the *potfit* software worked in a similar fashion to existing capabilities, both programatically and from a user standpoint. This influenced the design of input and output parameters, and files, for the ensemble implementation. Section 3.3 details the input parameters required to run the ensemble generation, as well as introducing the optional tuning parameters which can be used to improve the sampling. An in-depth discussion of each optional parameter, and their influence on the ensemble generation, is covered in subsequent sections. The ensemble output file is introduced in section 3.4, alongside an overview of the data analysis workflow used to quantify the uncertainties in quantities of interest.

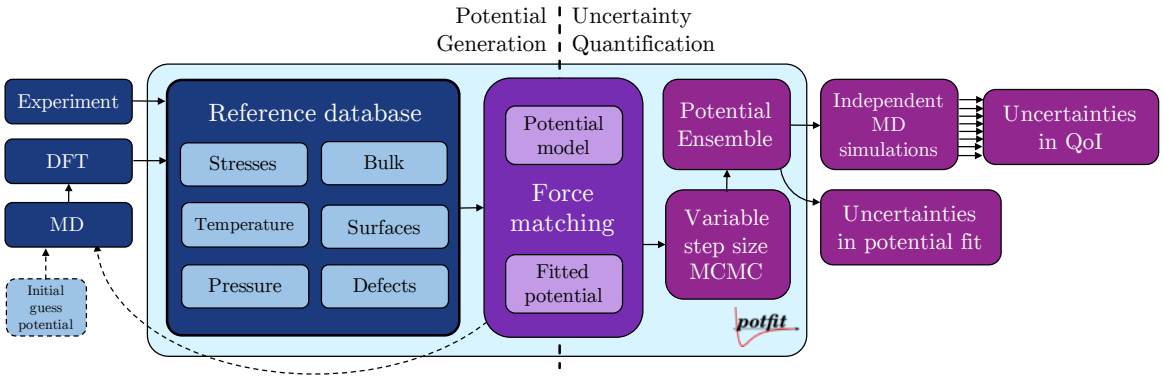

Figure 3.1: An overview of the uncertainty quantification workflow. The large light blue rectangle indicates the workflow section which is implemented within *potfit*.

Figure 3.1 gives a pictorial overview of the uncertainty quantification workflow. The section of the workflow implemented within *potfit* is indicated by the large light-blue background. The left side of the figure describes the best practice method of potential generation through force matching. This is an overview of *potfit* functionality prior to the ensemble implementation, as detailed in figure 2.3. Simulation snapshots under different conditions are generated from MD and DFT simulations to create the reference database. The right side of the figure describes the uncertainty quantification implementation, which focuses on two areas; the sampling of the potential parameter space, and the generation of uncertainties from ensembles of potentials. This builds upon the techniques outlined in the sloppy model method,[27] by investigating the representation of different potential models in the ensemble, and verifying the resultant uncertainties using this knowledge. It is important to highlight here, that the MCMC cost space sampling is performed within the *potfit* implementation, benefiting from the direct access to the existing function which calculates the parameter set cost. The subsequent data analysis however, which presents the uncertainties in potential parameters graphically, is created using a combination of user-specific bash scripting and python code outwith *potfit*, as these programming languages are better equipped to conduct the required analysis. The generation of subsequent MD simulations, and presentation of uncertainties, in the QoI workflow, is again largely automated using a series of bash scripts, which execute and extract results from the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS)[111] molecular dynamics software, and the Atomistic Simulation Environment (ASE)[112,113] python module.

3.2 Technical Details

The logic for the implementation is represented graphically by the flowchat in Figure 3.2, which highlights the essential steps in the process. The two main routines, which contain the majority of the more complex aspects of the development, are the self-descriptive `calc_hessian()` and `generate_mc_sample()` functions. Due to the additional implementation capabilities which are contained within `calc_hessian()`, this function is discussed in detail, with its own flowchart, in section 3.7. By default, the Hessian is calculated using a central finite difference algorithm,[90] perturbing each model parameter by small amount, by either the default perturbation, or a user specified amount.

Following the initial ensemble generation setup and Hessian calculation, the optional input parameter, `use_svd`, determines the method used in calculating the eigenvalues and eigenvectors. A discussion of both options (eigendecomposition and singular value decomposition) follows in section 3.6, and details the motivations behind the inclusion of the latter. The functions `dsyevx_` and `dgesvd_` are from the LAPACK (Linear Algebra Package) library, and respectively perform the eigenvalue decomposition and singular value decomposition of the Hessian. In the eigenvalue decomposition routine, upper and lower guess bounds for the eigenvalues are required. Should all of the eigenvalues fail to be found within the initial range, each bound is increased repeatedly by a factor, a maximum of 10 times until they are obtained, otherwise the user is produced with a warning, and singular value decomposition is used instead. Should one of the Hessian eigenvalues be negative, this implies that, on the scale of the Hessian curvature calculation, the best fit potential is not at the cost minimum, but a saddle-point. In this case, the user is given a warning yet the code continues, using the modulus of the negative eigenvalue instead. The decision was made to continue, despite the knowledge that the best-fit potential does not define a cost minimum, so that the MCMC sampling method could still be used to investigate the cost landscape,s and potentially assist in finding a lower cost parameter set, regardless of the now sub-efficient MC step size information. Similarly, if more than one eigenvalue is negative this could imply that the fit is unstable, and therefore there may be a greater benefit in improving the reference data than in proceeding.

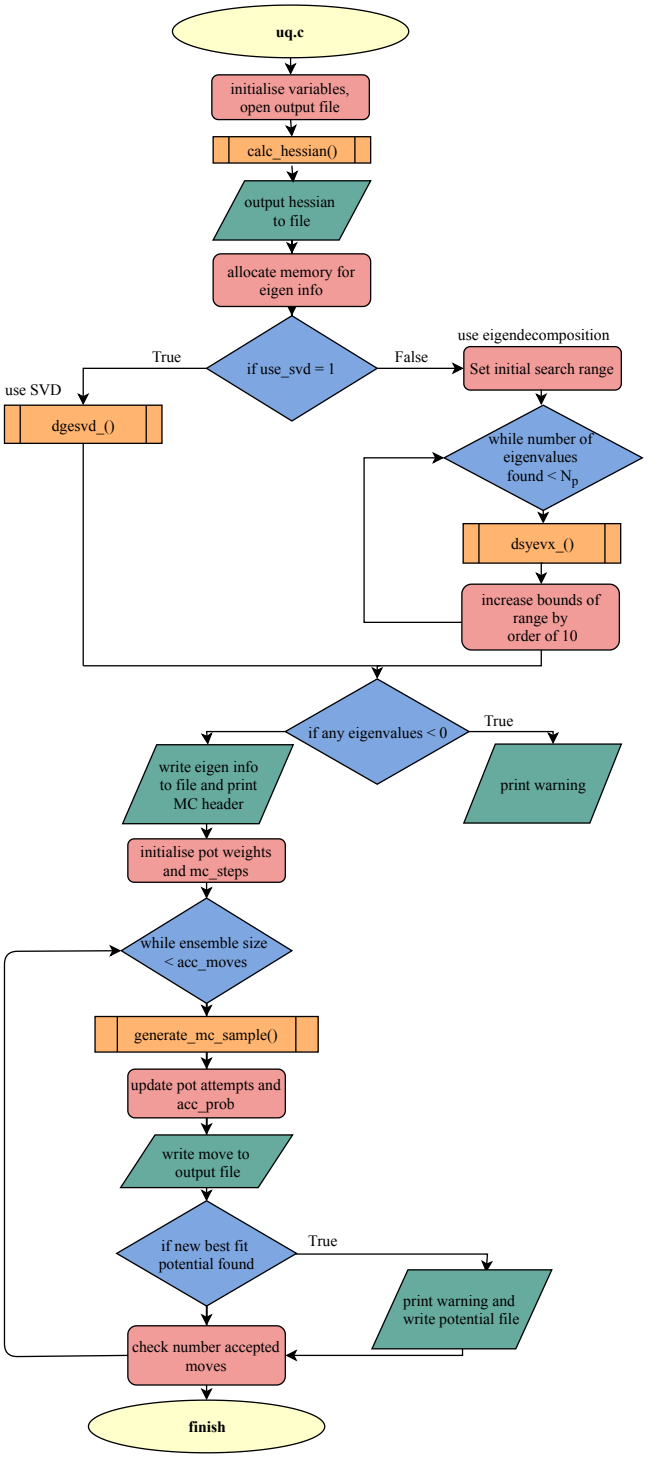

Figure 3.2: A pictorial description of the ensemble method implementation.

Finally, the Markov chain Monte Carlo simulation can begin. Trial moves are made until a step is accepted, at which point a weighting, equal to the number of trials required before a successful move, is assigned to the previous step. An updated acceptance rate, based on the running count of total moves made, is calculated, and the Markov chain continues until the desired number of potentials are selected. If at any point in the process a lower cost potential is found, a warning is displayed, and a *potfit* potential file is written for the new parameter set, along with all force matching error files. When compiled with the `--debug` option, the implementation outputs all trial moves, regardless of whether they are accepted or rejected. A flag in the output file is used to distinguish between which moves were successful and which were not.

As discussed in the following sections, the implementation includes a variety of optional input parameters, which provide information predominantly for diagnostic purposes, when one is beginning to understand the nature of a models cost landscape. Each of these options are not necessary in order to run the uncertainty quantification, but can be used to help understand the cause of any unexpected results. Therefore, it is recommended that any user of the ensemble method investigate their model, and build an understanding of the nuances of their simulation by using the additional tools, as opposed to simply using the compulsory parameters. By doing so, it is more likely that a user is able to correctly identify issues, or improvements relating to the types of reference data, analytic model choice, and minimisation algorithms in the potential fitting process, prior to embarking upon the uncertainty quantification.

3.3 Input Parameters

Table 3.1 lists the various parameters introduced in the ensemble method implementation. Only two additional *potfit* parameters are required by default to generate the ensemble: the first being the number of MCMC moves to be accepted (i.e. total number of ensemble members), and the second concerning the tuning of the R value in the MCMC step (equation (2.42)). The `ensemblefile` parameter specifies the filename of the implementation output, which is given the suffix ".uq". Should no filename be specified, the contents of the

Table 3.1: The input parameters for the ensemble method implementation. Required parameters are indicated by '⋆'.

	Parameter	Type	Default	Description
	`ensemblefile`	string	output_prefix.uq	Ensemble output filename
⋆	`acc_rescaling`	double	-	R value in equation (2.42)
⋆	`acc_moves`	int	-	Number of ensemble members
	`uq_temp`	double	1.0	Value of α in equation (2.44)
	`use_svd`	bool	0	Use SVD to find eigenvalues
	`hess_pert`	double	0.00001	θ^* perturbation in Hessian calculation
	`eig_max`	double	1.0	max(λ_j,eig_max) in equation (2.42)
	`write_ensemble`	bool	0	Write *potfit* potential file for each member

,

pre-exiting optional parameter, `output_prefix`, are used. Should `output_prefix` also be unspecified, the program will error.

The `uq_temp` parameter can be used to change the cost value which determines the probability of an uphill move being accepted. This ultimately relates to the how high above the minimum the cost basin is sampled, with the reference to "temperature" resulting from the similar description used in simulated annealing. Section 3.8 discusses this option, and includes a demonstration of the effect of `uq_temp` choice on the MCMC sampling.

The parameter `hess_pert` provides the perturbation to each parameter ($\Delta\theta_j$ = `hess_pert` $\times\ \theta_j$). If the user inputs a negative value for `hess_pert`, a bracketing algorithm is instead used to calculate the curvature on the same scale as the cost temperature. This is discussed in more detail in section 3.7.

The `eig_max` parameter can be used to alter the maximum contribution of the eigenvalues in the calculation of the step size. A small eigenvalue indicates a sloppy direction, of low curvature, along the related eigenvector. Therefore for an eigenvalue less than `eig_max` (by default equal to 1), the contribution to the parameter perturbation in equation (2.42) is bounded, in order to restrict the sloppy eigenvalue from contributing too large a component to the step size.

Table 3.2: The output parameters for the MCMC potential ensemble.

Name	Description
`param_`i	i^{th} potential parameter value
`cost`	Parameter set cost
`weight`	Number of trial attempts before the next accepted step
`accepted`	Accepted flag
`attempts`	Total number of trial steps
`acc_prob`	Percentage of steps accepted

3.4 Output Files

The implementation produces an output file, with name as specified by the `ensemblefile` parameter, which contains all necessary information for the parameter uncertainty quantification, as well as the eigenvalue and eigenvector data, which can be used in investigations of the cost landscape. The output file begins with a print out of the cost minimum Hessian, followed by the eigenvalues and their corresponding eigenvectors. The MCMC potential information is then output, alongside the running totals and percentages of accepted steps to trial moves. Table 3.2 lists the columns of data included in the MCMC section of the output file. The parameter header information lists the parameters in the order they are listed in the potential input file. The `accepted` column is only really necessary when compiling with the `--debug` option, as all accepted moves are indicated by '1's. The first potential output in this file is the best fit potential, along with its `weight`, indicating the number of trial moves attempted before the first step in the Markov chain is accepted.

The `write_ensemble` input parameter flag will output a *potfit* potential file for each accepted MC step if triggered. For large MCMC runs, this can quickly create an excessive number of files and so should be used with caution. This option can be useful when testing any additional scripting, for example, whilst automating the initialisation of simultaneous MD simulations used in the quantification of QoI uncertainties.

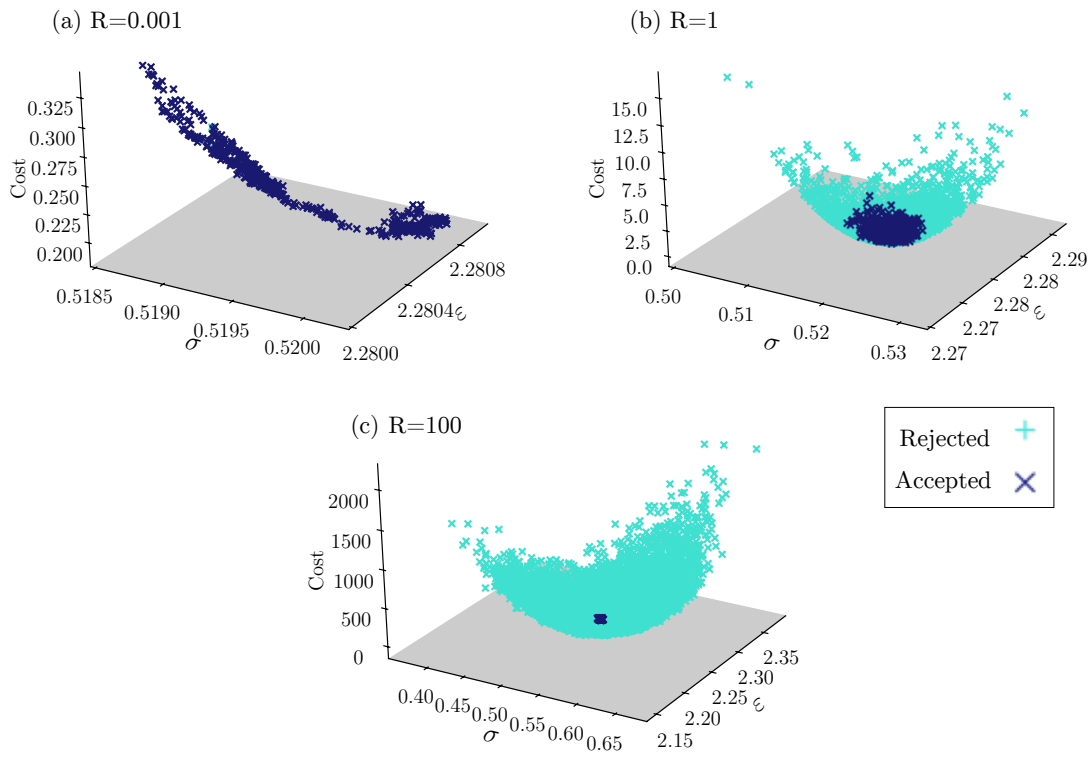

Figure 3.3: The difference in accepted and rejected steps for three magnitudes of step size, as controlled by R, in a Lennard-Jones ensemble generation. Note the difference in scale of the cost values sampled.

3.5 R Value Tuning

In figure 3.3, a Lennard-Jones potential ($\varepsilon = 0.520$, $\sigma = 2.809$, $h = 0.5$) has been fitted to six reference configurations. Each configuration is a snapshot from various molecular dynamics simulations, each using a Lennard-Jones potential ($\varepsilon = 0.519$, $\sigma = 2.808$). The reference data consists of two nickel surfaces, and four bulk nickel configurations, each simulated at a temperature between 300−2000 K. As the potential is fitted to MD forces which were generated using a LJ potential, the best fit potential is able to almost exactly reproduce them. The fitted *potfit* potential uses a smooth cutoff to 0 at the tail, unlike the LJ model used in the MD simulations. Therefore, due to the small difference in contributions from pair separations towards the tail of the fitted potential, the minimal cost is not exactly null ($C_0 = 0.189$). By changing the R value in equation (2.42), the effect of the R value contributions to the step size can be seen in figure 3.3, changing the MCMC acceptance

rate. For too small an R value, tiny steps in parameter directions are taken, resulting in a 100 % acceptance rate, whilst a fraction of the cost minimum basin is explored. As is clear in the scale of parameter sampling within figure 3.3(a) too small an R value also results in a slow traversing of the cost landscape, therefore resulting in longer decorrelation times, and ultimately requiring a significantly larger number of accepted moves until the Markov chain converges to the underlying distribution.

Conversely, as illustrated in figure 3.3(c), for unreasonably large R values, a tiny proportion of proposed steps are accepted. For too large a step size, the temperature (i.e. height) of the cost basin explored is so high that the majority of attempted steps have tiny acceptance probabilities. This also results in a significantly increased runtime, as many more trial steps are taken before a potential is eventually accepted.

When choosing the ideal R value for ensemble generation, it is advised that a variety of runs are conducted across a range of R values. In this work, 50 000 MCMC steps were accepted, for a variety of R values, until a step acceptance of close to 23 % was achieved for each model. As outlined in section 2.4.4, this acceptance rate has been documented as most efficient in balancing convergence with runtime when sampling multidimensional spaces.[110] When "tuning" the R value, the number of accepted steps for each test need only be large enough so that the percentage of accepted steps calculation has converged.

In summary, the R value which alters the Markov chain step size is chosen to best balance the sampling convergence, efficiency and step decorrelation time, against the total MCMC runtime. This is important, as the MCMC potentials are then sub-sampled in order to create an ensemble of uncorrelated potentials, representative of the underlying distribution, with which the uncertainties are quantified.

3.6 Singular Value Decomposition vs Eigenvalue Decomposition

Eigenvalue decomposition (EVD) of the Hessian, is used to find a linearly independent basis of vectors, which span the principle axes of curvature in the cost minimum basin. The eigenvectors and eigenvalues of the cost landscape are obtained by factorising the Hessian

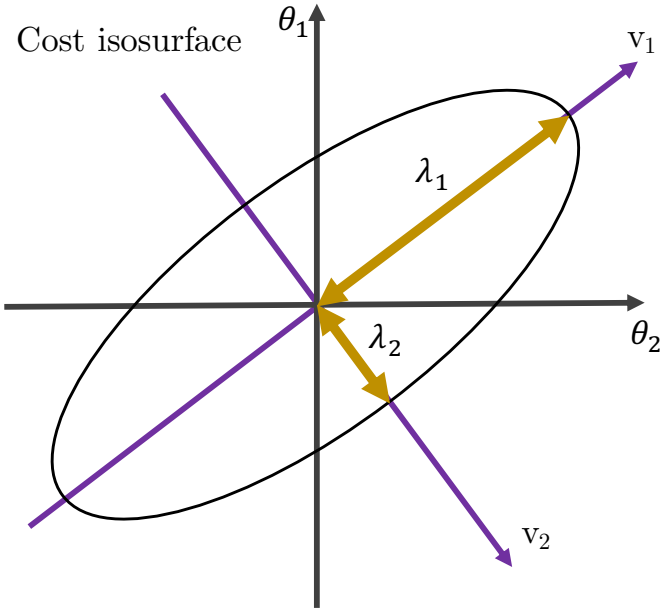

Figure 3.4: The relationship between the potential parameters and the eigenvector basis of a cost minimum in 2d. The ellipse indicates a constant cost contour in the landscape, with the principle axes of curvature indicated by the eigenvectors, v_1 and v_2. The degree of ellipse curvature is indicated by λ_1 and λ_2. The potential parameters, θ_1 and θ_2 indicate the original coordinate system of the landscape.

into a product of three matrices

$$H = V\Lambda V^{-1}, \tag{3.1}$$

where V is the matrix of eigenvectors, stored vertically within columns, and Λ a diagonal matrix of eigenvalues. Within the cost landscape, the coordinates of the space are given by the set of potential parameters, $\boldsymbol{\theta}$. The eigenvectors define another set of axes which are aligned along the principles axes of curvature within the basin. Figure 3.4 illustrates the relationship between the eigenvectors, eigenvalues, and parameters, for a 2d cost landscape along a constant cost contour. The matrix V, can be thought of as rotating the axes to orient them along the principle axes of the cost contour, with Λ a matrix encoding the relative scales of the curvature.

The generalisation of eigenvalue decomposition is known as singular value decomposition (SVD), and can be used in cases where the eigenvectors do not form a complete basis. SVD is included in the current implementation, to account for any unforeseen issues with, or future changes to, the Hessian calculation. In the current implementation of the

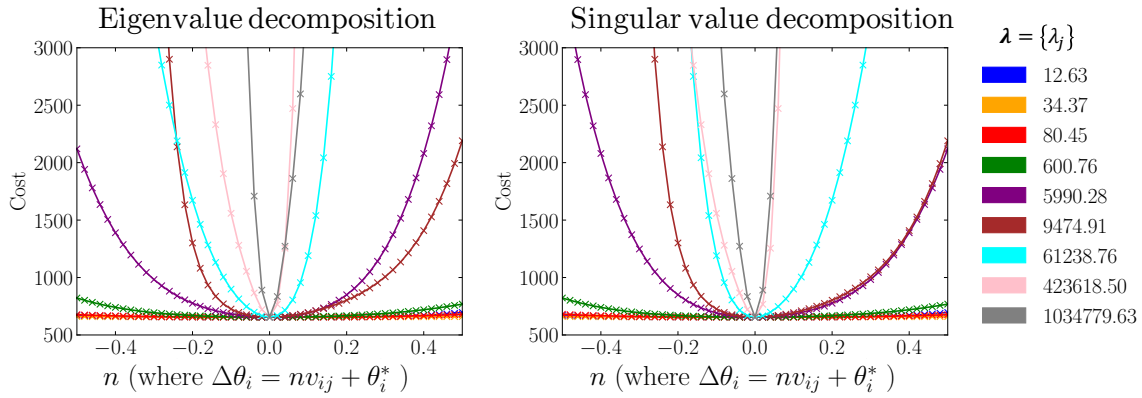

Figure 3.5: The cost values for potential parameter sets along the calculated eigenvectors for a 9-parameter fitted EAM potential. The lines are coloured according to the eigenvalues relating to each eigenvector, and are listed in increasing order. The results for eigenvalue decomposition are shown in (a), and the results of SVD in (b). Here θ_i are represents each potential parameter, with θ_i^* being the best fit (cost minimum) value for each.

minimum basin curvature, the assumption is that the curvature in both directions along a given principle axis is the same, i.e. the minimum basin is symmetric. This assumption ensures that the Hessian is always a (real) symmetric matrix (i.e. a Hermitian matrix), which is therefore almost always diagonalizable, i.e. possesses a complete basis of linearly independent eigenvectors. Should the Hessian be non-diagonalizable, i.e. defective, then SVD can be used to obtain an orthonormal basis for the cost landscape. The only case in which the Hessian may not be diagonalizable, is if there is degeneracy in the eigenvectors. This could, for example, be the result of an unnecessary potential parameter, which does not affect the cost value in the reference data fitting. Similarly, should the calculated Hessian perhaps not reside at the landscape minimum, e.g. is a saddle point, then the use of SVD will ensure that a basis is found, and the always positive singular vectors can be used to continue the MCMC sampling. Although obviously a far-from-ideal scenario, sampling of such a cost landscape region may be desirable in order to better understand the space.

The `use_svd` input flag uses SVD to calculate the Hessian eigenvalues and eigenvectors, instead of the standard eigendecomposition algorithm. This method is also employed if the eigendecomposition algorithm fails to find all eigenvalues within 10 attempts given a continually increasing range. In the case of positive-semidefinite, normal Hessians, the

output of the two algorithms are equal, meaning that all eigenvalues are positive. For symmetric matrices such as the Hessian, the singular values are equivalent to the square root of the eigenvalues, and the left singular vectors equal to the eigenvectors in all but sign.

Figure 3.5 illustrates the difference in sign of the basis vector found by each of the algorithms for the same Hessian. The data is produced by evaluating the cost along each eigenbasis vector for a variety of perturbations, n. In the eigenvalue decomposition case, these vectors are the eigenvectors, and in the case of SVD, they are the left singular vectors. The ith value of the jth column in the matrix of eigenvectors explains the relative change in each parameter θ_i required to move a distance of λ_j. There are 9 distinct eigenvalues (corresponding to each coloured line) as the data is for a 9-parameter EAM model fitted to some reference data. Figure 3.5(a) illustrates the cost of each parameter set for points along the eigenvectors when calculated using eigenvalue decomposition, with the results for singular value decomposition shown in figure 3.5(b). Although at first there may appear to be a difference in the curvatures along basis directions between the two plots, on closer inspection it becomes clear that for a given eigenvalue, some of the eigenvectors are pointing in the opposite direction.

3.7 The Hessian Bracketing Algorithm

The Hessian bracketing algorithm included in the ensemble method implementation is a consequence of an investigation into the asymmetry of cost basins. An implied assumption of the Hessian finite difference calculation, is that minimum cost basins are symmetric on the scale of the perturbations. In order to provide the user with the tools required to investigate and understand the specific nuances of the minimal cost basin defined by their model + reference data, the bracketing algorithm in figure 3.6 was developed to aid in any investigations of its shape in potential parameter directions. We speculate that the bracketing algorithm will yield the greatest insight in simulations where the harmonic approximation to the cost basin breaks down, in which case an alternative approach to the Hessian calculation could potentially be insightful. An anharmonic basin may fit into three theoretical categories: in situations where the cost basin is asymmetric along the principle

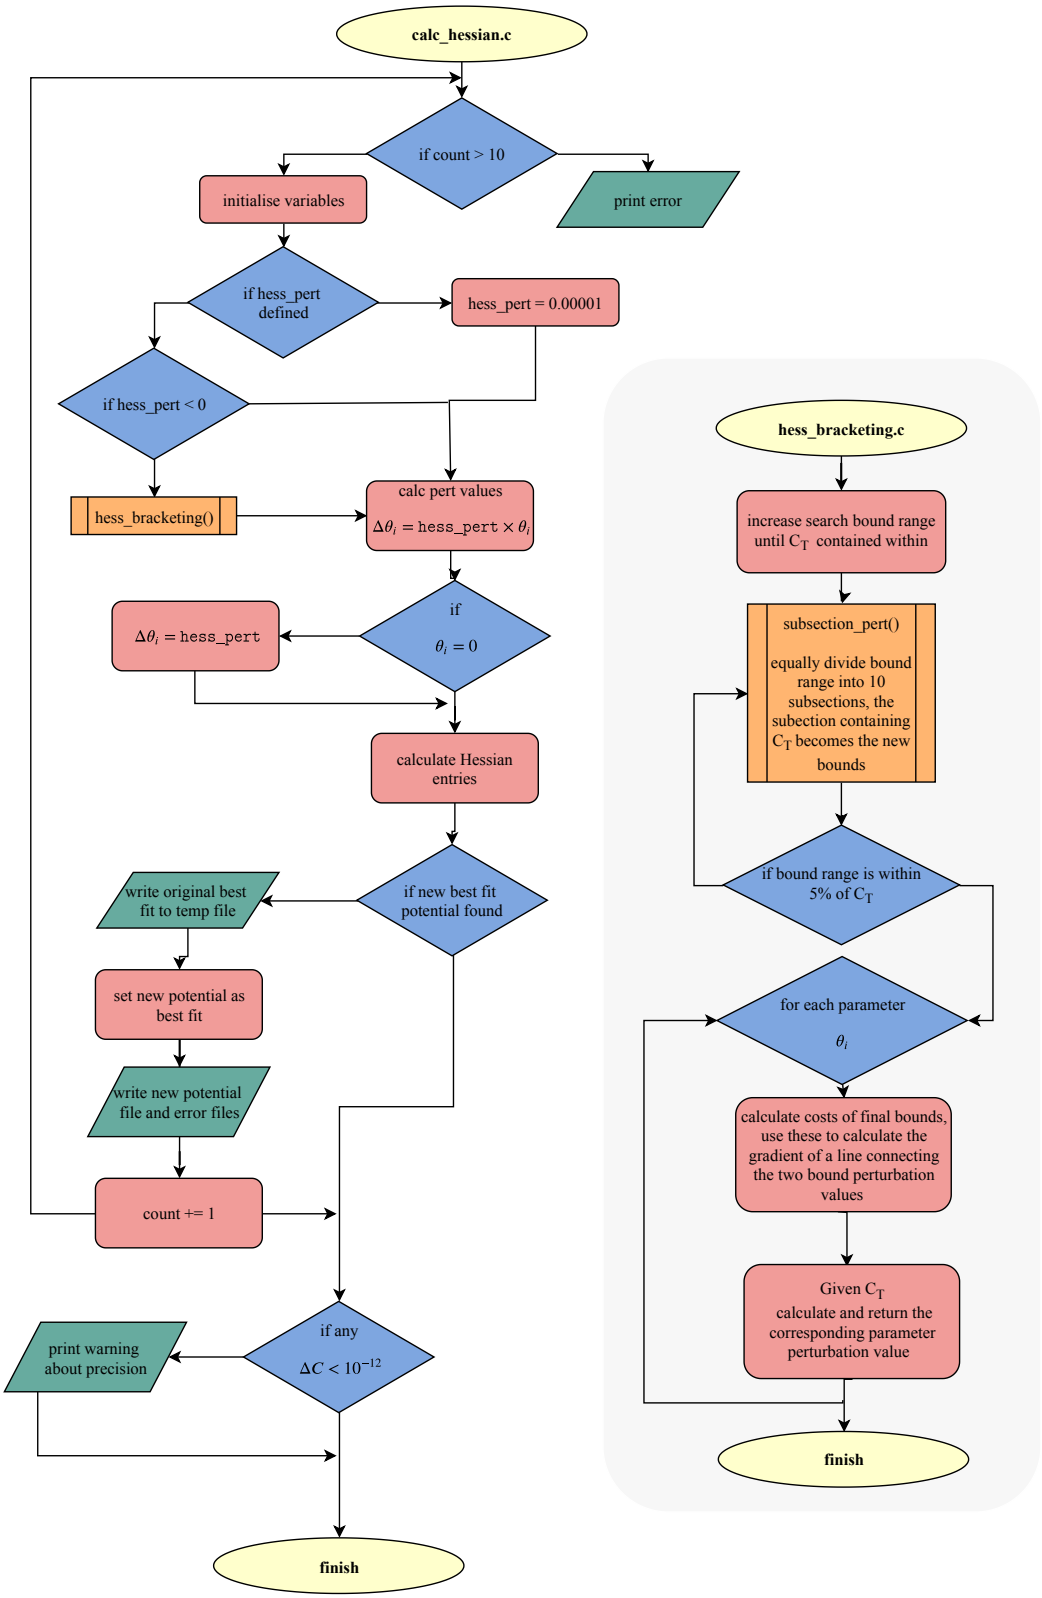

Figure 3.6: The `calc_hessian()` routine, and `hess_bracketing()` function employed when `hess_pert`< 0.

axes of curvature, where cost isosurfaces are non-convex (e.g banana-shaped), and finally in cases where the minimum basic is shallow for the scale of the Hessian calculation, or there is a significantly deep cost basin with a "wrinkled" bottom. It is theorised that these cases are most likely to contain noticeable asymmetries in the parameter directions of the cost basin, or result in negative Hessian eigenvalues due to saddle-points on the scale of the perturbations.

Figure 3.6 shows the logic within the `calc_hessian()` function which contains the `hess_bracketing()` routine. As standard, each parameter is multiplied by the `hess_pert` value, and added to the best fit parameter set to obtain curvatures in each direction. If the user specifies a negative `hess_pert` value, the bracketing algorithm is used to find parameter-specific perturbation values, which will calculate the curvature on the scale of the desired cost temperature. Each perturbation is found by calculating the corresponding cost values, for individual best fit parameters perturbed by a range of amounts. The perturbation chosen is the value which results in the closest cost to the desired isosurface. The desired cost value is given by the temperature relation in equation (2.44), which can be calculated as

$$C_T = C(\boldsymbol{\theta}^*) + T = C(\boldsymbol{\theta}^*) + \frac{2C(\boldsymbol{\theta}^*)\alpha}{\mathrm{N}}. \tag{3.2}$$

The `hess_bracketing` algorithm is then employed to calculate both the positive, and negative, perturbations in each parameter direction which yield C_T. If a minimum basin is asymmetric along a parameter direction, the smallest perturbation value is used. This means that in an asymmetric case, the Hessian calculated is a symmetric approximation embedded within the true minimum. This is necessary to ensure the Hessian is symmetric, and therefore that the left singular values are equivalent to the eigenvectors for both `use_svd` options.

To find each positive and negative parameter perturbation, an initial guess perturbation is made. If this is less than C_T, it is stored as the lower bound. The perturbation value is then continually increased by a factor of 10 until the result is greater than C_T. At this point, the perturbation value is stored as the upper bound, with the lower bound having been replaced previously by each failed perturbation value. If the initial guess was

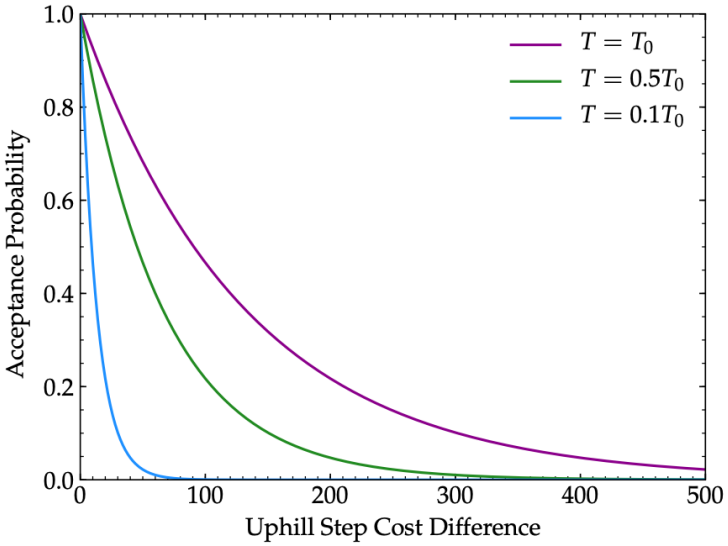

Figure 3.7: An example of the uphill step acceptance probability for a fitted EAM potential. Here the natural temperature, T_0, is defined as $2C_0/\mathrm{N}$, where the best fit cost, $C_0 = 655.2$, and number of parameters, $\mathrm{N} = 10$.

less than C_T, similar logic applies to finding the lower bound. Once the upper and lower bounds are initially set, the range is subdivided into 10 intervals and the bounds replaced by those of the section containing C_T. This process is repeated until the range is within 5 % of C_T. Finally, a line is fitted between the parameter bounds, and using the fitted gradient, the perturbation value calculated to yield C_T is found.

Understandably the bracketing algorithm should only be triggered when its limitations are understood, as its uninformed use could quickly lead to the brute force continuation through the workflow, for simulations with existing issues which should have been addressed. This is why the `hess_pert`< 0 input value is required, as a negative value is unlikely to be entered by a misinformed user. The algorithm should really be used to compare the per-parameter perturbation values, in order to understand how the difference in eigenvector curvature may affect the parameters in parameter coordinates. The algorithm can also be used to inform the user of the asymmetry in per-parameter curvature, highlighting potentially asymmetric cost basins which could then be inspected through 2d-projections of the MCMC sampling.

3.8 Temperature Reduction

From equations (2.43) and (2.44), uphill MCMC moves are accepted with probability

$$P_{\text{acc}} = e^{-\Delta C/T} = e^{-\Delta C N/2C_0\alpha}. \tag{3.3}$$

Figure 3.7 shows the probability of accepting an uphill step for a demonstrative, fitted, 10-parameter EAM potential, at given cost difference for three different temperatures. The purple line shows the acceptance probability for the "natural" temperature, T_0. This is the default acceptance rate for this system. The lines in green and blue show the reduction in acceptance rates for equivalently sized uphill moves. They correspond to optional `uq_temp` input parameter values of $\alpha = 0.5$, and $\alpha = 0.1$, respectively. For an uphill step corresponding to a cost increase of 100, the acceptance rates are 47 % ($\alpha = 1$), 22 % ($\alpha = 0.5$), and 0.05 % ($\alpha = 0.1$).

Figure 3.8 illustrates the resulting distribution of 500 cost values for MCMC subsamples of a 50 000 chain, for the above EAM system and three temperatures. As the temperature decreases (i.e. α is reduced) the range of cost values sampled above the minimum decreases. The reduction of MCMC sampling temperature can be invoked in cases where it is believed that the cost minimum basin is shallow, in order to restrict the sampling to within the basin of interest, as outside of this region the Hessian curvature calculation does not hold. This may arise if a potential is being fitted to a rarer material phenomenon, such as an unstable crystalline configuration, or a highly specific state, such as the movement of an interstitial in a non-periodic morphology. However, it is also imperative that the user ensures that a shallow minimum is not a result of a sparse or lacking reference data set.

Due to the interdependent influence of a variety of the optional input parameters on the ensemble generation, it would be perilous to discuss a reduction in sampling temperature without also mentioning the effect of potential fit quality on the acceptance probability. This crucial modulating factor in the acceptance rate, is the minimum (best fit) cost value. For an improved fit of potential to the reference data (i.e. a lower cost value), the acceptance rate of similarly sized steps decreases. This means that with increased confidence in the potential fit, there is a reduction in the likelihood that larger steps away from the minimum

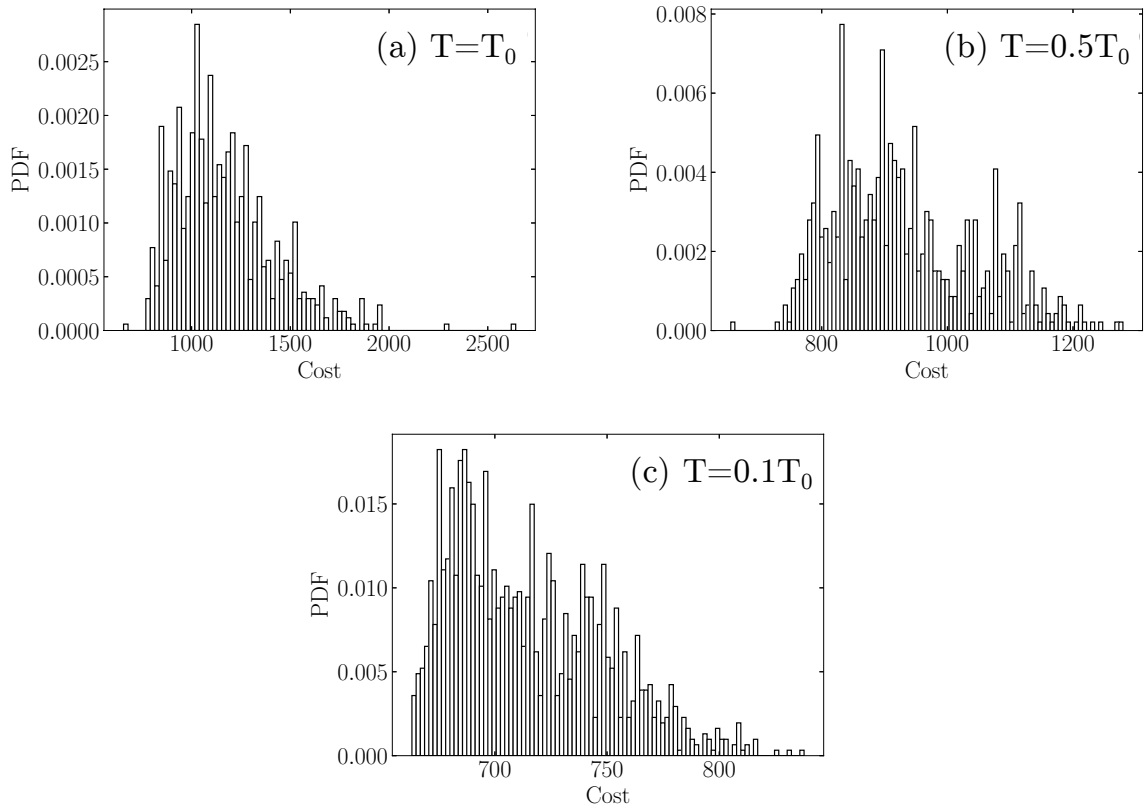

Figure 3.8: The cost distribution for 500 sub-samples of the EAM potential for $\alpha = 1, 0.5$ and 0.1. The cost range reduces with decreasing α as the size of accepted Markov chain steps is reduced.

are accepted. The resulting spread in the ensemble parameters of a good potential fit is therefore likely to be less than that of a poor fit.

3.9 Additional Scripting and Analysis

Having obtained the ensemble of potentials, the series of subsequent steps in the workflow rely predominantly on a set of directory-structure-specific bash and python scripts to automate the analysis. The downsampling of the ensemble should be performed after analysing the autocorrelation functions for each parameter, and selecting the samples at sufficiently decorrelated intervals. As in figure 3.2, the finalised potential ensemble can then be analysed to obtain the parameter uncertainties. This is done using a python script to create the box-plot figures shown in chapter 4.

To obtain uncertainties in quantities of interest, each potential ensemble member informs a MD simulation calculating the quantity. This process is automated using bash scripts to read the sub-sampled potential parameter sets, automatically create the required folder structures and input files for the MD simulations from template files, and initialise the running of the simulations, whilst piping the relevant outputs to a master file. For intensive calculations of thermodynamic quantities, such as the thermal expansion coefficient, scripting was produced to implement this workflow whilst additionally creating and submitting the individual MD simulations to run on high performance computing facilities, instead of executing the simulations locally.

Although the uncertainty propagation required multiple scripts for data processing, in this work, this part of the workflow was faster than the MCMC portion of the analysis. This is due to the quantities of interest simulated being relatively computationally inexpensive to compute, in contrast to the long autocorrelation times in the MCMC sampling, in which long runs were required to ensure decorrelated potential ensemble members. This distribution of computational time and cost may be different for other systems, as it is defined by the level of complexity of the quantities of interest simulations, and by the reference data + model set which has been fitted.

3.10 Summary

Following the principle design consideration of continuity with existing software, the ensemble implementation rigorously follows the existing *potfit* standards, whilst also providing additional debugging capabilities to the user through a variety of optional parameters. Within the discussions of the additional, optional capabilities, it has likely become apparent to the reader that many of the parameters affecting the sampling are interdependent on one another. Therefore, it should be clear that a successful uncertainty quantification using the ensemble method requires a thoughtful and vigilant analysis of data collected at each step in the process.

Through the investigations presented in this chapter, and suspicions of the existence of complex cost-space scenarios not yet encountered, an appreciation for the complexities

of the cost landscapes defined within potential fitting has been developed. In the following chapter the end-to-end uncertainty quantification process is demonstrated for three models of increasing complexity, including most interestingly, the quantification of uncertainties in MD quantities which, as of yet, have not been at the focus of discussions.

4

Uncertainty quantification of Three Potential Models for Nickel

4.1 Motivation

The application of the *potfit* uncertainty quantification implementation compares three potential models fitted to bulk nickel. The following work evolved from the initial testing phase of the implementation, and grew in complexity whilst attempting to understand and intuit the uncertainty quantification process. The chapter is structured as a sequential account of each stage in the uncertainty quantification process. Each section details the methodology underpinning the generation of the results, followed by a discussion and evaluation of each step in the process.

The chapter outline begins with section 4.2, in which the *potfit* reference database for the potential fitting is built. The reference database consists of DFT forces, energies and stresses; the quality of these calculations placing limits on the accuracy with which each

fitted potential can describe atomistic information. Section 4.3 details the analytic functions and fitted parameters for each potential model, along with details of the *potfit* fitting process and individual model results. In this section (4.3.2) the performance of the fitted models is evaluated for information included in the reference database, and the influence and limitations of the analytic forms are examined. Furthermore in 4.3.3, the cost space defined by the reference data + potential model set is discussed. With the cost landscape defined, the parameter uncertainties are probed through the creation of the MCMC potential ensembles in section 4.4. The process of selecting a subset of potentials to propagate in order to obtain the uncertainties in quantities of interest, is introduced in section 4.4.1, with the parameter projections of the selected potentials detailed in section 4.4.2. Penultimately, in section 4.5, the uncertainties in the model predictions of three quantities of interest are investigated and discussed. Section 4.5.1 focuses on the spread in predictions of the equilibrium lattice constant, followed by the methodology and results for the elastic constants in section 4.5.2, and concluding with predictions of the linear thermal expansion coefficient at 300 K in section 4.5.3. The chapter closes with a summary of the findings in section 4.6.

Three models of increasing complexity are selected to test the uncertainty quantification implementation, and to aid in the understanding of any pitfalls commonly encountered when sampling multidimensional cost landscapes. Testing was first performed with a two-parameter Lennard-Jones potential, as the cost landscape defined by this model is simple to conceptualise. Nickel has been modelled extensively in the past using pair potentials,[76,114–116] and so crystalline nickel was chosen as the material for which the potentials would model in the initial system testing.

Building upon the two parameter case, a Morse potential, which contains one further parameter, was introduced into the process. The two pair potentials were predominantly employed to test the MCMC sampling, and to build an initial understanding of the uncertainty quantification process. From the offset, the LJ and Morse models were not expected to provide a high quality description of nickel, due to their simplistic analytic forms. Therefore, to recreate a realistic potential fitting process within the testing, a popular model for nickel,[117–121] the 10−parameter embedded-atom method (EAM) potential, is explored.

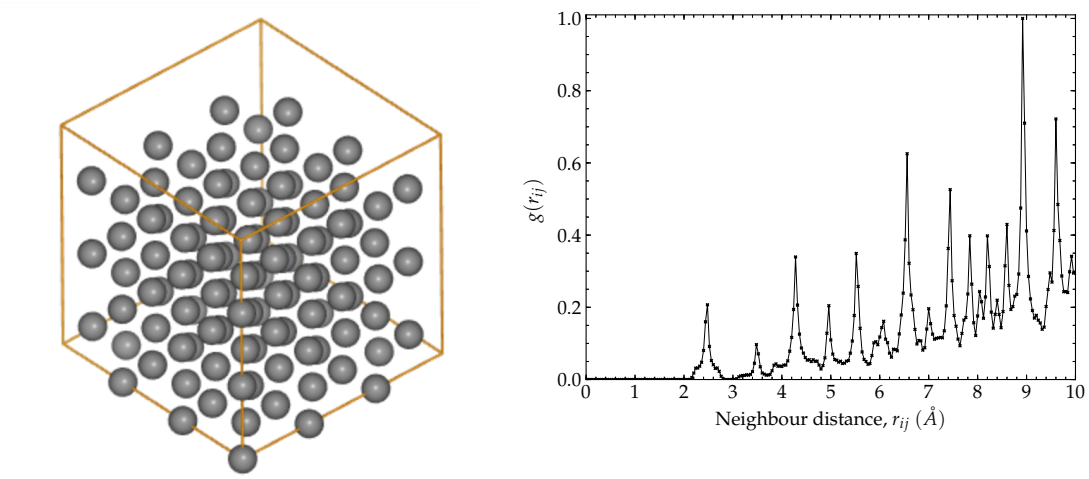

Figure 4.1: (a) A 108−atom $3 \times 3 \times 3$ nickel face-centred cubic supercell. (b) The radial distribution function (rdf) for the nickel reference configurations which are detailed in table 4.1.

4.2 The Reference Data

Nickel is a transition metal with 10 valence electrons ($3d^9 4s^1$), favouring a face-centred cubic (*fcc*) crystalline structure, with lattice parameter 3.499 Å.[122] The reference data to which the models are fit, is built to reflect these attributes, focusing on bulk *fcc*-nickel at a variety of temperatures, stresses and strains. The reference data contains atoms displaced from equilibrium positions at temperature, with stressed and strained configurations, to ensure that the fitted potentials have been exposed to a range of configuration space neighbourhoods.

In building the reference data the aim is to obtain a selection of realistic energies and forces for the systems the fitted potential will eventually be used to model. To create reference configurations at temperature and under strain for example, MD is often used to perturb the atoms from their crystalline positions. First principles MD could potentially be used of course, however it has not been explored in this work, and furthermore may not be suitable for reference data with large numbers of atoms due to the significant computational cost. In order to perform standard MD, some interatomic potential has to be chosen to represent the atomic interactions, in the absence of the bespoke potential being fitted. Whilst this is obviously not ideal, in order to create realistic reference configurations the

potential chosen should ideally be of a similar analytic form to that being fitted, as it is more likely to sample from a similar configuration space to the final potential. The reference data creation should follow a common-sense approach, where strange configurations are added to the reference dataset only if checked for errors and deemed necessary for the potential fitting. As it is not ideal to generate the reference database with the this interim potential, DFT singlepoint calculations should be performed to improve the accuracy of the interatomic forces, energies and stresses. In this way the fitted potential is effectively being trained on first-principles data, but with the perturbation of atoms (i.e. sampling of configuration space) having initially been provided by an MD simulation.

4.2.1 Reference Data Creation

When fitting potentials by force matching, as outlined by Wang and Van Voorhis,[123] a range of configuration space neighbourhoods should be represented. The benefits being twofold: firstly, by providing the fitting with a larger selection of forces, the resulting potentials are more robust to erroneous data, and secondly, the transferability of the fitted potentials in describing under-represented areas of configuration space is improved. This means that the reference data should contain MD, and even *ab initio* MD, information which will relate to areas of configuration space outside of the ideal 0 K atomic positions. Therefore, the reference data is based upon a selection of snapshots from MD simulations at temperature and under strain. However, the final force and energy information which makes it into the the reference dataset, is first improved by performing DFT singlepoint calculations on each MD snapshot.

After a series of iterations, 23 bulk nickel configurations were selected to act as the final reference database for the uncertainty quantification, consisting of 2484 atoms (7452 forces, 23 energies and 138 stresses). The 108-atom supercell ($3 \times 3 \times 3$ unit cells) illustrated in figure 4.1a, was geometry optimised using the DFT to form the first configuration. Using a published Lennard-Jones potential for nickel,[76] the remaining configurations were created with velocity-Verlet MD. The simulations consisted of combinations of temperature, isotropic and anisotropic stress applied to the cell as outlined in table 4.1. To ensure the reference data contained realistic forces and energies, singlepoint DFT calculations were

Table 4.1: Final reference dataset to which the Lennard-Jones, Morse and EAM models were fitted.

	ID	Snapshot (ps)	Temperature (K)	Isotropic (%)	Anisotropic x-y (%)	Energy (eV/atom)
temperature	1	geom opt	0	-	-	−5.241
	2	50	100	-	-	−5.214
	3	50	200	-	-	−5.210
	4	100	300	-	-	−5.207
	5	100	1000	-	-	−5.178
	6	90	2000	-	-	−5.145
isotropic deformation	7	32.5	300	90	-	−4.216
	8	22.5	300	92	-	−4.603
	9	12.5	300	94	-	−4.870
	10	82.5	300	96	-	−5.099
	11	62.5	300	98	-	−5.187
	12	22.5	300	100	-	−5.225
	13	52.5	300	102	-	−5.205
	14	72.5	300	104	-	−5.140
	15	4.25	300	106	-	−5.081
	16	3.25	300	108	-	−4.955
	17	92.5	300	110	-	−4.823
$xy-$deformation	18	8.25	300	-	92	−5.183
	19	42.5	300	-	94	−5.192
	20	20	300	-	96	−5.199
	21	16	300	-	98	−5.205
	22	9.25	300	-	102	−5.205
	23	18	300	-	104	−5.200

performed on all MD snapshots prior to their addition to the reference dataset.

The decision to use a LJ potential to generate the snapshots was made at the beginning of the implementation testing process, when initial uncertainty quantifications investigations for a simple LJ potential were being undertaken. The MD snapshots were generated by the LJ potential using parameters for nickel as reported in the literature[76] ($\varepsilon = 0.519$, $\sigma = 2.2808$, cutoff $= 10\,\text{Å}$). This was initially to minimise any additional error between the model and the reference set, as the LJ parameters should be easily recoverable by the cost minimisation algorithms, especially due to the simple parabolic nature of the potential. At this early stage of the testing, no DFT singlepoint calculations were performed on the reference configurations, and the so energy and forces being fitted to, were exactly

those resulting from the MD simulations. As expected in this initial test, the only discrepancy between the model and the reference data was due to the geometry optimised DFT snapshot, being the only reference configuration which did not rely on the LJ model for its creation. In order to define a cost landscape with a null cost minimum, a relaxed snapshot of a bulk *fcc* nickel lattice would be more suitable.

Subsequent tests were conducted for the LJ, Morse and EAM models, and required 8 versions of the reference dataset until the configurations in table 4.1 were decided upon. Factors considered in the selection of the final reference data were, for example: ensuring the snapshot had equilibrated to the desired temperature, minimal duplication of information in the reference data, and ensuring a variety of configurations. In later tests the energies, forces and stresses from the MD snapshots were replaced by those obtained from singlepoint DFT calculations of the configurations, to improve the accuracy and precision of the reference data to which the models were fit. To ensure a correct translation between the reference configuration DFT energy and the cohesive energies fitted by *potfit*, the DFT snapshot energies were subtracted from that of a single (DFT) nickel atom in a vacuum as

$$E_{\mathrm{coh}} = \frac{(E_{\mathrm{atom}} - E_{\mathrm{bulk}})}{\mathrm{N}_{\mathrm{atoms}}}. \tag{4.1}$$

DFT simulations were performed in CASTEP,[124] using the PBE functional, with a 400 eV planewave cut-off, and a Monkhurst-Pack k-point grid of 0.1 $\mathrm{Å}^{-1}$ spacing. This converges energies and forces to within a tolerance of 5×10^{-5} eV/atom, and 0.05 eV/Å, respectively. MD simulations were performed using the LAMMPS[111] simulator. To simulate a nickel crystal at temperature, atoms were initialised with velocity components drawn from a Maxwell-Boltzmann distribution, corresponding to twice the desired simulation temperature. The simulations were then run using a time-step of 0.001 ps in the microcanonical (NVE) ensemble, for a minimum of 4250 time-steps, to allow the kinetic and potential energy atom contributions to equilibrate. Snapshots of the atomic positions, energies and interatomic forces were taken at the times indicated in table 4.1.

4.2.2 The Radial Distribution Function

Figure 4.1b displays the radial, pair distribution function (also known as the pair correlation function), for the reference data in table 4.1. The function, $g(r_{ij})$, represents the likelihood of a neighbouring particle existing at a radial distance r_{ij}. The average density of particles is then given as $g(r_{ij})\rho = g(r_{ij})N/V$. Figure 4.1b is normalised with respect to crystalline nickel, where the largest peak at a neighbour distance of approximately 8.9 Å would contain 72 atoms of the crystal. Therefore for a perfect nickel crystal, at the lattice parameter separation of ∼3.5 Å, one would expect 6 atoms, giving a probability of $g(3.5\,\text{Å}) = 6/72 = 0.083$. However, the corresponding ∼3.5 Å peak observed in figure 4.1b, has a larger probability of 0.096, due to the contributions from the non-equilibrium reference configurations. Although $g(r_{ij})$ describes only the one aspect of the reference database (i.e. the atomic positions), it begins to illustrate the amount of information being provided to the potential model being fit. Were only configuration 1 in table 4.1 provided, the radial distribution function, in figure 4.1b, would only contain peaks at the ideal lattice positions. Understandably, this provides only 15 discrete atomic positions (with zero forces and stresses), for the potential fitting process. Therefore, the introduction of the temperature-dependent and stressed configurations (or the introduction of different morphologies, i.e. surfaces, interstitials, dislocations, etc.), is essential, in order to increase the number of datapoints in the fitting process. Such configurations, result in the observed broadening and overlapping of peaks, as information from perturbed atoms is included.

4.2.3 Reference Data Considerations

As will become clear throughout the discussion of the uncertainties accounted for by this method, it is difficult to isolate and quantify the effects of particular reference data choices on the reported parameter and QoI uncertainties. As discussed in section 2.3.7, the uncertainties quantified by this method, are those resulting from the choice of model parameters, not explicitly from the choice of model itself (although the total cost does give an indication of the model performance). Upon reflection, the time-evolved reference data could have been generated using published EAM potential parameters, in lieu of the LJ model used, as the

Table 4.2: The potential models and associated parameters.

Model	Analytic Function	Parameters
Lennard-Jones	$V_{LJ}(r) = 4\varepsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right]$	σ, ε
Morse	$V_M(r) = D_e\left[\left(1 - e^{-a(r-r_e)}\right)^2 - 1\right]$	D_e, a, r_e
EAM	$E_i = \frac{1}{2}\sum_{i\neq j}^{N} V_M(r_{ij}) + F(n_i)$ with $n_i = \sum_{j\neq i}^{N}\rho(r_{ij})$	D_e, a, r_e
	$\rho(r) = r^{-\beta}\left[1 + a_1\cos(\alpha r + \varphi)\right]$	a_1, α, φ, β
	$F(n) = F_0\left[1 - \gamma\ln n\right]n^{\gamma} + F_1 n$	F_0, γ, F_1
Smooth Cutoff	$V_{SC}(r) = \Psi\left(\frac{r-r_c}{h}\right)V(r)$ where $\Psi(x) = \frac{x^4}{1+x^4}$	$r_c = 10$ $h = 0.75$

LJ potential has been reported to favour *hcp* over *fcc* structures.[125,126] It is not obvious how this would affect the results, although due to the subsequent DFT calculations of the MD configurations, it was concluded that this difference would not be significant, as only the atomic positions would change. There would be cause for concern were the reference data simulating complex structures and morphologies, as unrealistic transitions could potentially occur. However, in this instance only crystalline bulk quantities are assessed, with no questionable restructuring observed. When working with an uncertainty-fought infrastructure, with the understanding that the effective potentials being fitted are already time-consuming approximations to the systems being modelled, it seems to be important to balance the time demands of creating flawless reference data with the perceived gains in predictive accuracy and precision.

It is also of note that the reference dataset was not particularly tailored to predict the elastic constants, or thermal expansion coefficient, in sections 4.5.2 and 4.5.3. The decision to quantify the resulting uncertainties in the predictions of these quantities, was made after the final reference dataset was selected. When fitting a potential for production use, the process typically involves multiple iterations of the reference dataset, and potential fitting, to tailor the regions of cost space explored for its intended use. The potentials generated are solely an illustration of the newly implemented potential ensemble method, and are by no means suggested as new production-grade potentials to be used outside of this work in

the modelling of nickel.

4.3 Potential Fitting by Force Matching

In addition to the simple LJ and Morse pair potentials, the 10-parameter EAM potential is investigated. The analytic functions for the potential models are detailed in table 4.2, with a $r_c = 10\,\text{Å}$ cut-off used for each potential model, and the tails smoothed to converge to zero over 0.75 Å. The hypothesis is, that due to the large parameter space (and encoded physics), the EAM model is expected to best reproduce the desired quantities of interest (QoI), with the least error. Furthermore, there are a variety of high quality EAM potentials for nickel in the literature.[3,119,121,127] As detailed in table 4.2, a Morse potential is chosen for the pair part of the EAM potential. For the embedding function, $F(n)$, the universal form proposed by Banerjea and Smith[118] is used, and the transfer function, $\rho(r)$, is of the oscillatory form necessary for cubic metals, as reported by Chantasiriwan and Milstein.[128]

The cut-off parameters impact the computational cost of the potential, which makes them indicators of the model complexity. Thus, it was decided to exclude r_c and h, in table 4.2, from the optimisation. The 10 Å cut-off distance was chosen to be slightly less than the supercell length in the reference snapshots ($3 \times 3.5\,\text{Å} = 10.5\,\text{Å}$), so that interactions from periodic repetitions are avoided. The smoothing of the tails over 0.75 Å was chosen by inspection of a variety of different values.

The weighting of the forces, energies and stresses in the *potfit* fitting process were chosen so that each contributed an equal amount to the cost function. With each of the 7452 force components possessing a weight of 1, the resulting 23 energies were consequently given weights of $7452/23 = 324$, and similarly a weighting of $7452/138 = 54$ for stress components.

Table 4.3: The fitted model parameter values.

Lennard-Jones		Morse		EAM pair		EAM transfer		EAM embedding	
σ	2.1120	D_e	0.2771	D_e	0.1734	a_1	−1.1918	F_0	−3.9433
ε	0.6830	a	0.8601	a	2.1640	α	2.9075	γ	3.4657
	—	r_e	3.5793	r_e	2.4988	φ	0.7785	F_1	−0.0008
	—		—		—	β	3.5218		—

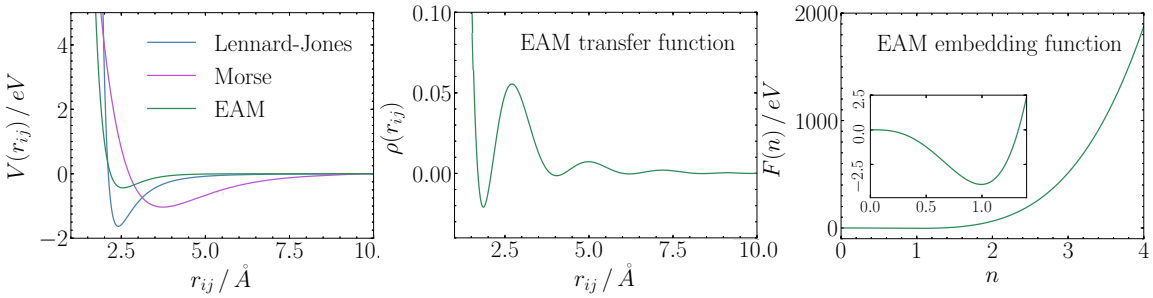

Figure 4.2: (a) The fitted pair potentials, where the EAM pair term uses the Morse functional form. (b) The fitted transfer function of the EAM potential. (c) The fitted embedding function of the EAM model.

For the fitting procedure, simulated annealing of the model parameters was first performed, followed by Powell's least squares, local gradient descent method. In later tests of the ensemble method, it was discovered that the initial Hessian calculation of the cost landscape occasionally yielded better fit parameter sets for models with slowly varying components. Therefore, the final parameter values reported in table 4.3 were iteratively found by performing a combination of Hessian calculations (input parameter `hess_pert`, all parameter perturbation), and simulated annealing (subsequent single parameter perturbations), with increasingly smaller values until no new minimal cost parameter set was found.

4.3.1 The Fitted Potential Models

The best fit potential models are shown in figure 4.2. The pair potentials are illustrated in the first subfigure, with the transfer and embedding functions of the EAM potential respectively, displayed in the remaining subfigures. Despite both using the Morse form, the EAM pair term in figure 4.2a differs significantly from the best fit Morse potential, in both minimum position and depth. The pair term of the EAM potential has a shallower minimum than that of the Morse potential, due to the additional energy contributions accounted for by the embedding function. Between the Morse and LJ potentials, the difference in both minimum depth and gradient are due, in part, to the additional parameter, a, in the Morse potential, controlling the gradient of the minimum well. The fitted Morse potential illustrates an inherent need for trepidation, when choosing to employ a potential without a true understanding of its fitting procedure. Were the Morse parameters to be used in another

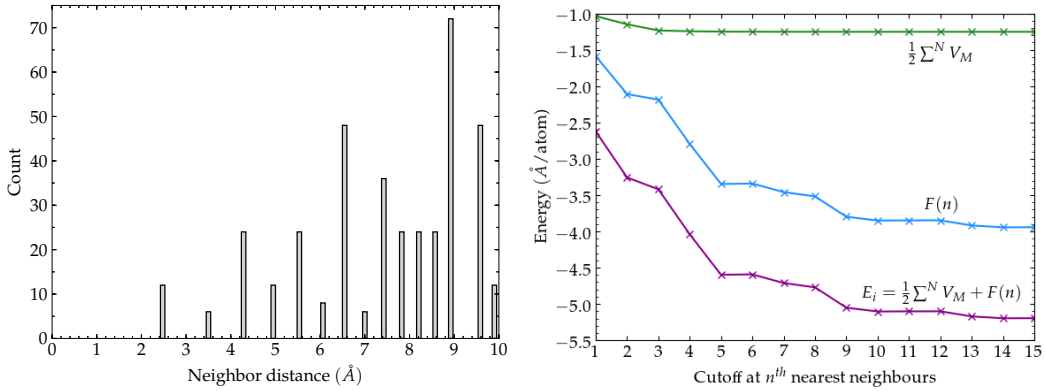

Figure 4.3: (a) The number of nearest neighbours for an atom in an *fcc* lattice with lattice constant of 3.5 Å. (b) The contribution of best fit EAM pair and embedding terms to total atomistic energy.

work, with a reduced cutoff value (i.e. of 3.5 Å, therefore including only the first and second nearest neighbours in crystalline nickel), with a ∼3.5 Å lattice constant as reported by experiments,[122] the potential predicts energies of 1.75 eV/atom, which is clearly incorrect and highly unrealistic. This highlights the importance of initiatives, such as the OpenKIM[39,129] and NIST[130] potential repositories, which provide full details of published model parameters available for use by the community.

Concerning the EAM embedding contribution to the interatomic energy, it is interesting to understand the magnitude of this contribution via a few a quick calculations. For a perfect $3 \times 3 \times 3$ supercell of crystalline nickel with a 3.5 Å lattice parameter, the number of nearest neighbours at a given interatomic distance is shown in figure 4.3a, up to the 10 Å cutoff to which the potential is fitted evaluated. In this crystalline configuration, each atom has 380 neighbours. The 12 nearest neighbours, commonly referred to as the bulk coordination number, are found at a distance of $3.5\,\text{Å}/\sqrt{2} = 2.47\,\text{Å}$. Figure 4.3b illustrates the breakdown in contributions to the total per-atom energy for a crystalline 3.5 Å system. The potential parameters used in the calculation are those of the best fit EAM potential, as listed in table 4.3, up to and including each nearest neighbour shell. It is clear for the best fit EAM potential, that the majority of the energy contribution is attributable to the embedding term. This explains the discrepancy in minimum depth found between the EAM pair term and Morse potential in figure 4.2a.

Table 4.4: The fitted model total error sums and their respective breakdowns.

Model	Total	Force (meV/Å)		Energy (meV)		Stress (MPa)	
	error sum	rms error	error	rms error	% error	rms error	% error
Lennard-Jones	6012.459	450.618	25.167 %	655.897	53.320 %	2.601	21.512 %
Morse	785.013	157.765	23.628 %	135.618	17.459 %	1.555	58.913 %
EAM	655.217	114.08	14.802 %	114.23	14.842 %	1.528	68.126 %

The ability of each fitted model to reproduce the reference data, is observed to improve as the number of model parameters increases. The fitting errors are reported in table 4.4, with the total error sum decreasing, with increasing model complexity. Pictorially, the total error sum defines the minimum offset in "height" of the cost landscape from the reference data parameter space. Therefore, this quantity is an indicator of comparative model error in the systems. The rms force errors for each fit were 0.45 eV/Å (LJ), 0.16 eV/Å (Morse), and 0.11 eV/Å (EAM). All of which are significantly greater than the DFT convergence tolerance of 0.05 eV/Å. This means that when quantifying the uncertainties in QoI, the significant sources of error in the system are predominantly due to the model parameters, not reference data error; the model error being that which this implementation is designed to elucidate. Similar trends were seen in the reproduction of reference energies and stresses, detailed in table 4.4.

4.3.2 The Best Fit Potential Performance

Before the uncertainty quantification process begins, the performance of the three fitted potentials is investigated. The results in figures 4.4a and 4.4b show the DFT energies for respective isotropic stretching, and $xy-$deformations, respectively of the corresponding 300 K reference configurations listed in table 4.1. The *potfit* calculations were performed for each configuration and model, without optimisation of potential parameters, to obtain the predictions for each fitted potential. Figure 4.4a illustrates a limitation of the LJ potential in predicting per-atom energies, for a variety of isotropically stretched crystalline nickel supercells, in line with the DFT results, due to the insufficiency of the analytic form. The Morse and EAM potentials perform surprisingly similarly in figure 4.4a, with the disagreements in predictions for a 6 % decrease (0.94 %) in equilibrium cell length. The difference

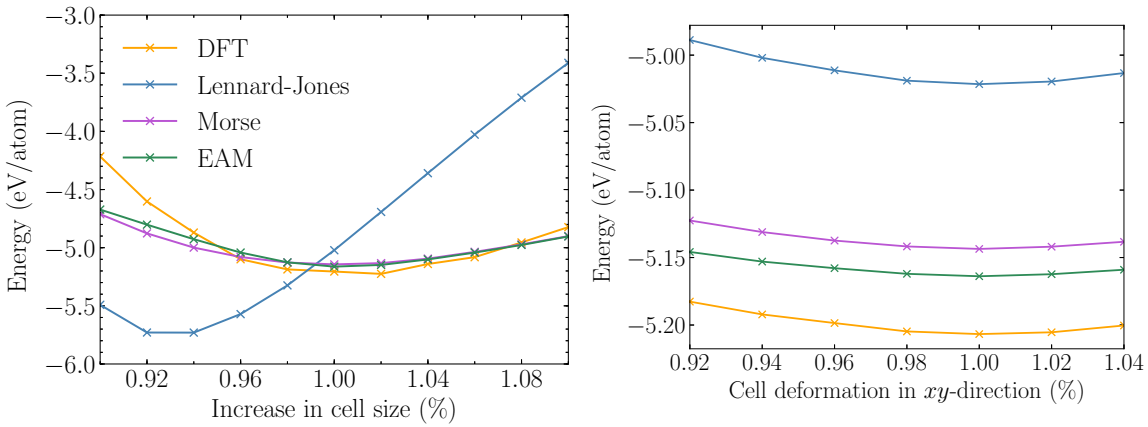

Figure 4.4: (a) The energies predicted by models for isotropically stretched (108-atom $3 \times 3 \times 3$) *fcc* cells. (b) The predicted energies of *fcc* cells subjected to an anisotropic strain in the xy-direction.

in "squashed" cell predictions between the DFT and Morse/EAM results may, in part, be due to the lack of fitting data for pair distances of $\sim 0.2-0.4$ Å less than the equilibrium atom distances in figure 4.1b. Therefore, an improvement to the reference database would be to supplement further with such configurations, should the potential parameter sets be designed to simulate related properties outside of this work.

In figure 4.4b, the effect of perturbations to atomic positions on a smaller scale is investigated. As only atoms along the $xy-$direction experience a significant change in positions, the range of predicted per-atom energies is significantly smaller than in figure 4.4a. On this scale, the difference in Morse and EAM predictions is clearer. The discrepancy in EAM and DFT predictions could be due to a systematic error in the EAM energy predictions, although this is unlikely, due to the crossing of the DFT and EAM results at 96 % of the relaxed lattice size, in figure 4.4a. A more likely explanation for the difference, is due to the comparatively huge number of forces versus energies in the reference dataset (324:1). Resulting in an EAM potential with a preference for reproducing forces very well, at the slight expense of total energies. In figure 4.4b, the maximum difference in the EAM energy predictions is in line with the reported rms error, of 0.114 eV, in table 4.4.

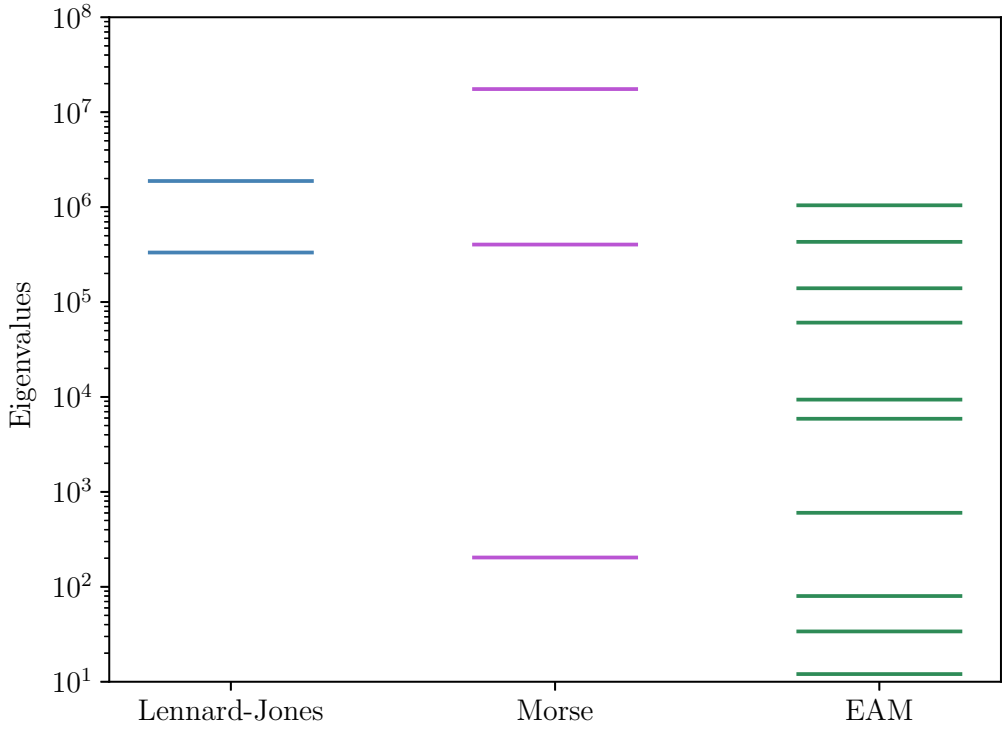

Figure 4.5: The spread in the eigenvalues of each Hessian for the potential models fitted to the reference set.

4.3.3 The Cost Landscape

Having investigated the performance, and make-up, of the best fit potentials, the next step in the uncertainty quantification process is to begin the ensemble generation. The eigenvectors of the Hessian denote the principles axes of curvature in the minimum cost basin, with the eigenvalues illustrating their relative lengths on the same cost contour as $\lambda^{-1/2}$. Therefore, directions with larger eigenvalues have a shorter widths and therefore a steeper curvature. Such directions are referred to as "stiff" eigendirections. Conversely, the smaller eigenvalues denote "sloppier" directions, whose predominant eigenvalue parameter contributions indicate the less sensitive potential parameters.

Figure 4.5 shows the spread in the eigenvalues obtained for the three fitted analytic forms. The eigenvalues are obtained from the calculation of the Hessian, as detailed in

equation (2.40). The degree of curvature, as denoted by the magnitudes of the eigenvalues, spans up to six orders of magnitude across models. Such vastly different curvature between eigendirections, illustrates the necessity for a sampling procedure which accounts for this variation, in order to sample the underlying distribution efficiently.

It can be seen in figure 4.5, that relative to the Morse cost landscape, the degree of curvature along the principle axes in the LJ cost space, are relatively similar. This implies that changes to either of the two LJ parameters contributes to a similar difference in cost. The primary parameter contribution to the "sloppiest" Morse eigendirection is r_e, defining the depth of potential minimum shown in figure 4.2a. It may not necessarily be intuitive that r_e is the "sloppiest" Morse parameter, when thinking about the change to the area bounded by the potential and the $x-$axis (D_e shifting up/down, a changing the curvature around the potential minimum and r_e shifting the potential minimum left/right). However, when considering fitting the potential to the reference data, it isn't too difficult to see that changes to r_e are likely to have least effect on the cost value, as the majority of the data, in figure 4.1b, relates to neighbour distances $> 5\,\text{Å}$.

The small eigenvalues in the EAM spread, indicate sloppy axes of curvature which result from insensitive model parameters. Through inspection of the relevant eigenvectors, it was found that the F_0 parameter in the embedding function, contributed most to the smallest eigenvalue. The second least sensitive parameter is γ, due to its predominant contribution to the second "sloppiest" eigendirection. This isn't particularly surprising when considering the relatively flat embedding function, in figure 4.2c for $n < 2$.

4.4 The Ensemble Generation

Table 4.5 displays the parameter values used to generate each potential ensemble. The values R and α refer to the step size tuning parameter, and to the sampling temperature parameter, from equations (2.42), and (2.44), respectively. The R value for each potential model was tuned over 50 000 accepted steps, so that 23 % of those proposed were accepted. This acceptance rate comes from guidance on the MCMC sampling of multidimensional landscapes.[110] One interesting benefit of the implementation, is that there is no "burn-in" period required

Table 4.5: The ensemble generation parameters for each model.

Parameter	Lennard-Jones	Morse	EAM
R (23 % steps accepted)	60000	80	0.9
α	1	1	0.05
Hessian perturbation (%)	0.00001	0.00001	0.00001
total steps	2×10^6	2×10^6	2×10^6
decorrelation time (steps)	50 000	50 000	50 000

for the Markov chain. This is due to the prior fitting of the potential parameters, meaning the sampling already begins from the "centre" of the underlying distribution.

In accordance with Frederiksen *et al.*,[27] sampling from the EAM potential landscape was performed at a reduced cost temperature of $0.05\,\mathrm{T}_0$, due to the Markov chain leaving the minimum basin for higher temperatures. Outside of the minimum basin, the Hessian calculated is no longer valid, resulting in inefficient sampling. Reducing the sampling temperature avoids these issues. However, the reduced sampling temperature for the EAM potential limits the insights from a direct comparison between model, but is not completely without merit; the scaled temperature results could be extrapolated to higher temperatures, assuming a roughly quadratic basin in cost function space. We believe that foremost, this behaviour is due to the selection of reference data the potential is fitted to. In the fitting of a potential for production use, the reference data is typically weighted, and complemented with configurations generated using iterative improvements of the fitted potential, which might alleviate the issue. In case that does not provide a remedy, a different sampling strategy would need to be employed, e.g. Riemann Monte Carlo methods,[131] or affine invariant samplers.[132] This might also allow incorporating prior information about the parameters into the ensemble generation process.

The curvature along cost space eigendirections for each model are illustrated in figure 4.6. The LJ curvatures are displayed first in blue, followed by the Morse results in purple, and finally the EAM eigendirections in green. C_1 refers to the smallest eigenvalue, hence the "sloppiest" eigendirection, which is reflected in the relatively small change in cost values. The magnitude of the eigenvalues increase with increasing C_i index. The cost change incurred in moving along eigendirections from the cost minimum is calculated by evaluating the cost for parameter sets along the relevant principle axes of curvature. For each

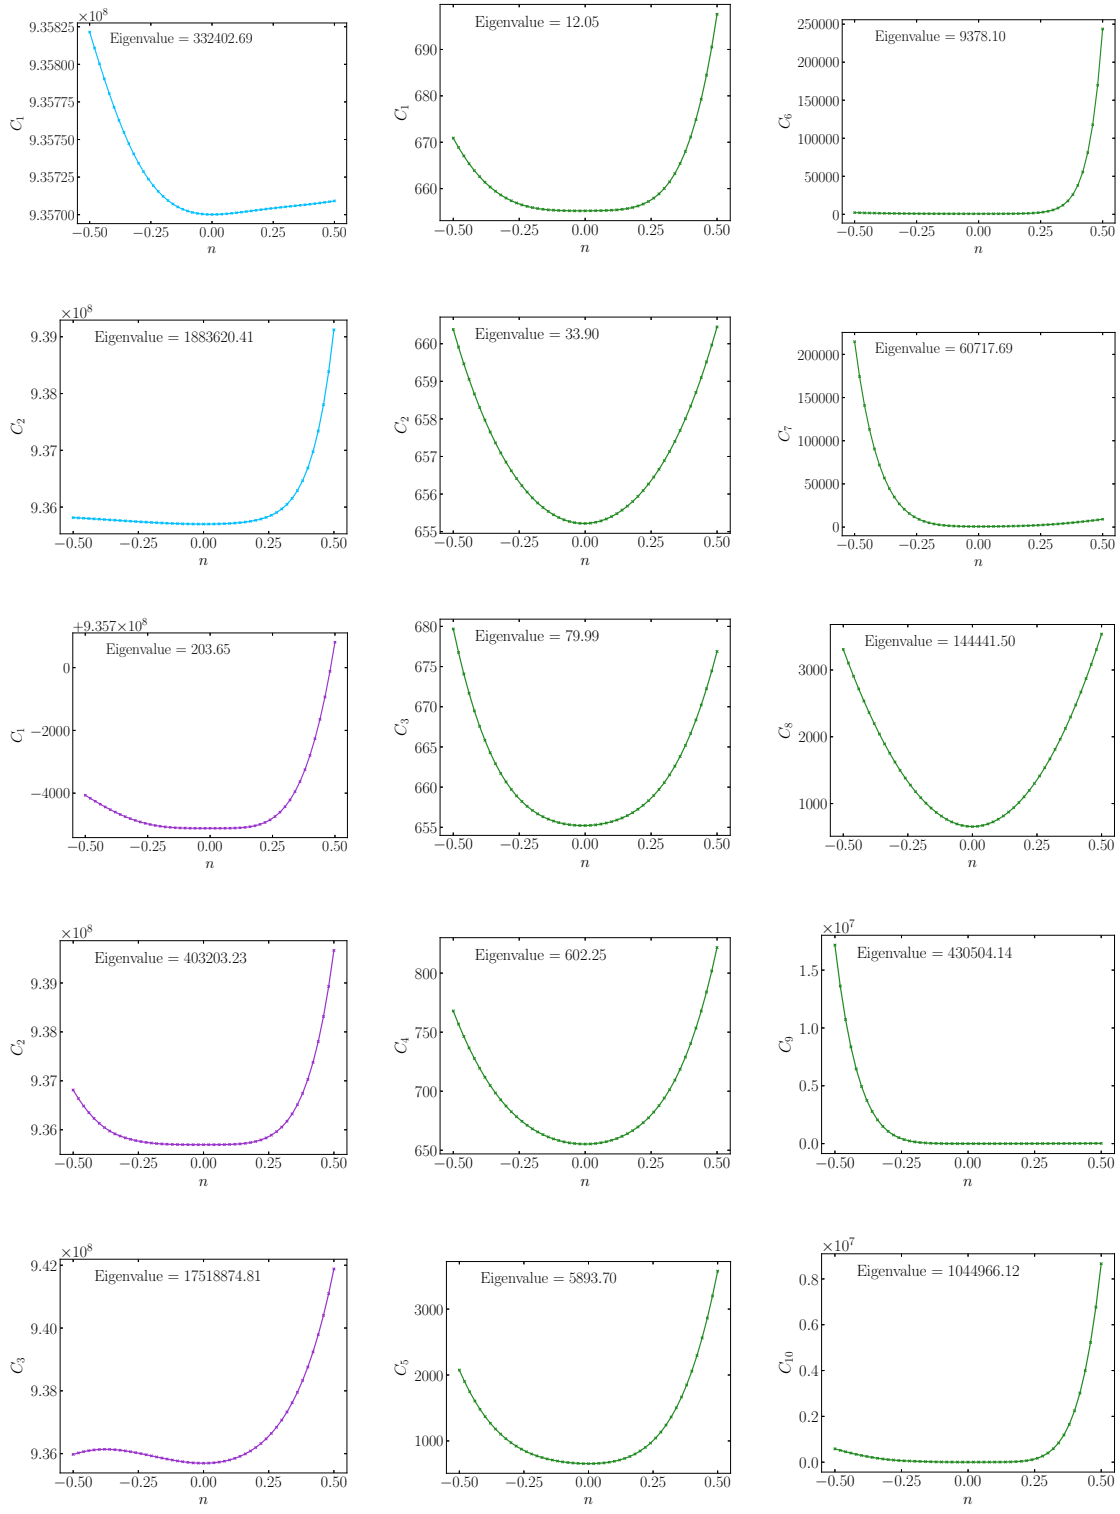

Figure 4.6: Parameter sets, $\boldsymbol{\theta} = \{\theta_i\}$, along eigendirections are calculated by perturbations, n, to the best fit parameter set θ_i^0 with $\theta_i = n\theta_i^0 v_{ij} + \theta_i^0$. Cost along eigenvectors for perturbation $C_j = C_j(\boldsymbol{\theta})$. The Lennard-Jones eigendirections are shown in blue, with Morse in purple and EAM in green.

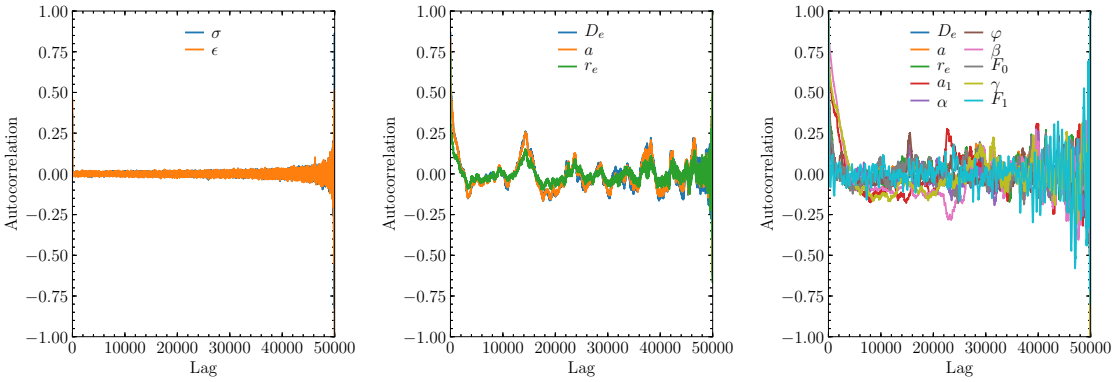

Figure 4.7: The individual parameter autocorrelation for 50000 runs, tuned to 23 % acceptance. The Lennard-Jones results are shown in (a), with the Morse potential results in (b), and finally the EAM in (c).

eigendirection, the sampling points are displayed in terms of a perturbation, n, to the best fit parameter set along each eigenvector.

The asymmetric cost along some eigendirections illustrates an anharmonic basin at large relative perturbations. The implementation is designed such that the harmonic approximation of the minimum basin is embedded into the true landscape minimum. For eigendirections with asymmetric curvature (or affected by a bifurcation point), this will result in the smoother direction curvature being over-approximated. Consequently, the sampling in the smoother direction will be less efficient (due to misleadingly small MCMC steps in that direction), and will therefore require a longer decorrelation time. Due to the understanding of this limitation, in section 4.4.1 what could be considered overly lengthy decorrelation times are used when selecting ensemble members.

4.4.1 Potential Parameter Autocorrelation

For each of the three fitted analytic potential forms, a 500-member, ensemble of potentials was obtained from MCMC samples, output from the *potfit* ensemble implementation. To ensure uncorrelated ensemble members, starting from the best fit potential, each sample was drawn after 50 000 accepted steps, with roughly 23 % of steps accepted for each analytic form, through tuning of the value of R in equation (2.42). The decorrelation time was assessed from the autocorrelation of each parameter in an initial MCMC run, with a con-

servative decorrelation time of 50 000 samples chosen. Sampling convergence was checked by ensuring reasonably smooth distributions in individual ensemble parameters, and in all $2d-$projections of 2×10^6 ensemble parameters.

The Markov chain decorrelation time, in table 4.5, was assessed through inspection of the correlograms shown in figure 4.7, for each fitted parameter set. As would be expected, the LJ parameters are quickest to decorrelate, being the space of lowest dimensionality, and containing eigendirections of relatively similar curvature. For large lags, approaching the total sample size of 50 000 MCMC steps, the observed increase in parameter autocorrelation is due to the decreasing number of values contained in the overlap, as a result of the finite sample size.

The Morse potential appears to have a slight sinusoidal quality to it, although the autocorrelation values never exceed 0.25. A possible explanation for this could be that the Markov chain begins to investigate an area of cost space with increasing rate of curvature in parameters, meaning that the steps accepted are those with smaller perturbations, resulting in more highly correlated parameter values.

As mentioned in section 4.4, longer than what may be considered necessary decorrelation times of 50 000 are used. As the evaluation of the MCMC steps is computationally cheap, there is little additional computational cost in doing so. Furthermore, having a larger number of accepted MCMC steps between samples is more likely to ensure decorrelation in areas of the cost landscape where there is a change in curvature, as touched upon above. The occurrence of such events is further discouraged through close monitoring of the Hessian calculation, accepted parameter sets and sampling temperature.

4.4.2 The Parameter Covariance

Figures 4.8 and 4.9 display the 500 parameter ensembles, and associated cost, in the LJ and Morse ensembles. The corresponding 10 EAM parameter histograms, and 45 $2d-$projections, have been excluded from this section for readability, and instead are included in appendix A. The histograms show projections of each ensemble for individual parameters. The heat-maps display the 2d-parameter projections, with the colour indicating the cost associated to with parameter set. In figure 4.8 of the LJ ensemble, the heat-map

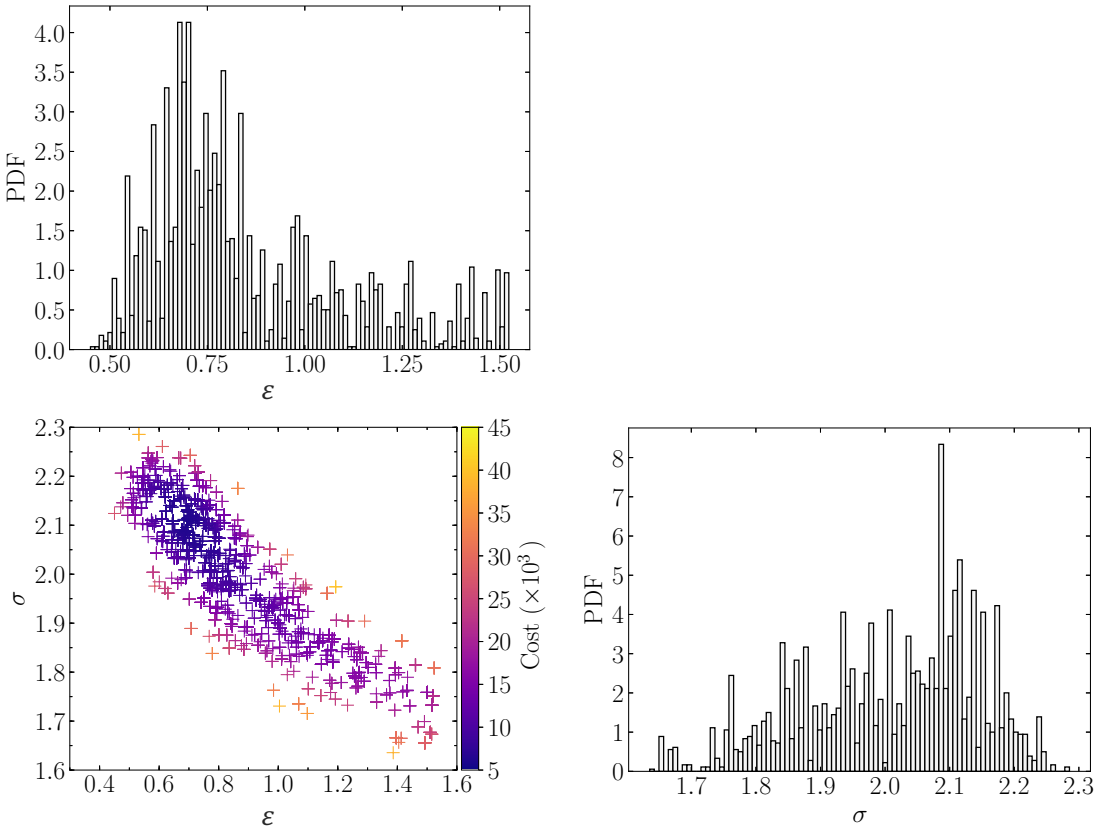

Figure 4.8: Lennard Jones parameter spread and covariance from 500 potential ensemble.

illustrates the simplistic basin structure expected from a distribution with a skewed asymmetric minimum basin. The individual parameter distributions, show approximate Gaussian distributions. The distributions are for the 500 ensemble members, which have been down-sampled from the 2×10^6 MCMC steps. From a visual inspection, the full MCMC parameter distributions are Gaussian, and so the deviation from the expected Gaussian distribution seen in figure 4.8, is due to the limited sub-sampled ensemble size.

Figure 4.9 displays the same parameter projections for the Morse potential. However, due to the additional parameter, a true visualisation of the space would be a 3d-heat map of the parameters and cost. In the projections of r_e vs D_e, and r_e vs a, it is clear that the basin has a non-ellipsoidal quality to the cost contours. The tails due to the large r_e values, are observed as significant contributions to the tail of the r_e histogram, and similarly in the corresponding plot for small a. There may be a connection between the sinusoidal autocorrelation in figure 4.7b and the curving cost basin. Another hypothesis could be

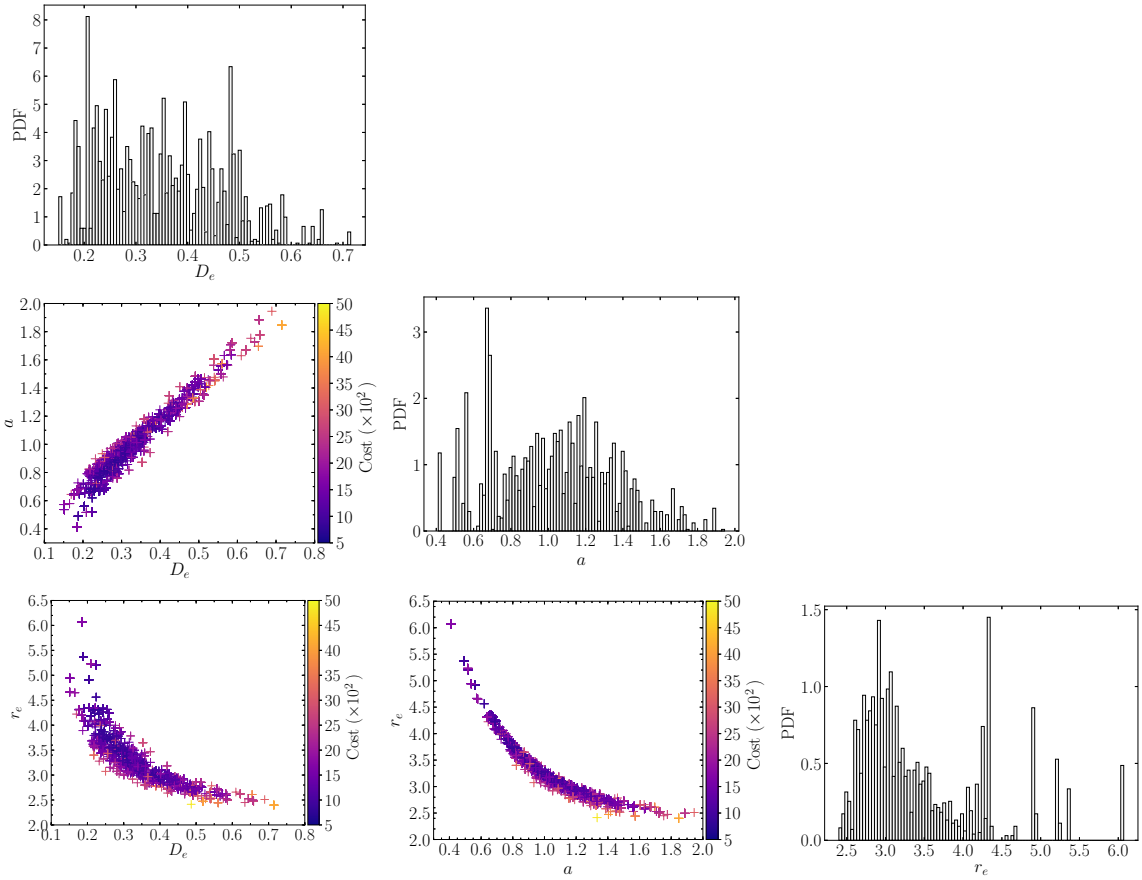

Figure 4.9: Morse parameter spread and covariance from 500 potential ensemble.

that the oscillating autocorrelation is a result of the Markov chain being confined to a small range of acceptable parameter sets, whilst exploring in the direction of the non-elliptical cost basin tails. To test this hypothesis an extensive examination would need to be conducted of the full dataset of MCMC steps taken, scanning for motifs in the pattern of exploration, corresponding to the autocorrelation in figure 4.7.

4.5 Demonstration of Uncertainty Quantification for Three Analytic Models

Finally the parameter uncertainties, introduced through the ensembles, are propagated to the quantities of interest. The performances of the potentials in reproducing the equilibrium lattice constant, the elastic constants C_{11}, C_{12} and C_{44}, and the thermal expansion coefficient

Table 4.6: Comparison of results for QoI with their associated uncertainties (IQR).

QoI	DFT/ exp.	Lennard-Jones median	Lennard-Jones IQR	Morse median	Morse IQR	EAM median	EAM IQR
a (Å)	3.51	3.14	[2.94, 3.26]	3.45	[3.33, 3.56]	3.52	[3.50, 3.54]
C_{11} (GPa)[133]	253	1503	[1236, 2355]	224	[171, 315]	222	[180, 259]
C_{12} (GPa)[133]	152	860	[704, 1347]	158	[123, 202]	180	[130, 213]
C_{44} (GPa)[133]	124	860	[704, 1347]	158	[123, 202]	90	[81, 98]
α (10^{-6} K^{-1})[134]	14.4	4.8	[3.9, 5.7]	8.2	[7.4, 9.1]	15.1	[12.8, 17.6]

at 300 K are investigated. Each analytic potential, from each of the three ensembles, was propagated through MD simulations to obtain the uncertainties for each potential model.

The resulting uncertainties, displayed in figures 4.10 and 4.11, are quantified using box-whisker diagrams, which illustrate the uncertainty in each quantity by the inter-quartile range (IQR). The box denotes the IQR, with whiskers extending 1.5×IQR beyond each quartile. The values obtained from the best fit (i.e. minimum cost) potential are shown in dashed purple. The notches indicate the confidence interval in the ensemble median (in green), and the ensemble means are indicated as dotted red. The dotted black lines, traversing the entirety of each figure, indicate the experimental values, detailed in Table 4.6. In table 4.6, the equilibrium lattice constant is denoted by a, and the DFT value for a is from the geometry optimised cell included in the reference data. The elastic constants are written in Voigt notation as C_{11}, C_{12}, C_{44}, and the linear thermal expansion coefficient at 300 K is denoted by α. Due to the presence of skewness and outliers, the uncertainties are reported using the resistant measures of sample median and IQR. Reporting uncertainties via the sample mean, and standard deviation, will bias the reported quantities towards misleading ensemble outliers, which tend to result from unlikely higher cost potential ensemble members.

4.5.1 The Equilibrium Lattice Constant

The lattice constants were found by minimising a *fcc* nickel lattice, with a starting lattice parameter guess of 3.5 Å, for each potential ensemble member. Table 4.6 reports the uncertainties for each fitted analytic potential.

Within table 4.6, the LJ ensemble is unable to capture the correct lattice constant within the IQR. This is largely due to higher temperature configurations in the reference

data, to which the potential was initially fit. The Morse ensemble does manage to capture the correct value, yet the larger spread in uncertainty, for a simple 0 K quantity, again implies the potential is limited by the higher temperature configurations in the reference data. It is important to note that poor potential performance due to reference data selection is not always the culprit; if a model is known to have limitations in its ability to reproduce certain physical quantities due to its simplistic design, then any issues which arise may be down to an insufficient model choice. This is illustrated by the best-fit (minimal cost) LJ potential, where the first warning is in the large rms force error (0.45 eV/Å). A second concern, is then the incorrect lattice constant prediction of 3.26 Å. This results in an ensemble of candidate potentials for an already insufficient fit, clearly indicating an insufficient model choice. Similarly, despite a more promising rms force error (0.16 eV/Å), the Morse potential performs poorly in the prediction of lattice parameter, which also alludes to an insufficient model choice.

The EAM potential accounts for a non-linear dependency on the local environment through the embedding term. It appears that this is essential for a realistic prediction of the lattice constant given the set of reference data used. The EAM ensemble, albeit sampled with a reduced temperature, not only captures the DFT value within its error but also has a considerably constricted spread of the ensemble predictions compared to the pair potentials.

4.5.2 The Elastic Constants

The elastic constants are investigated, in order to compare the restoring forces of the potentials, to small perturbations of atomic positions. In this work, the 0 K elastic constants are calculated by measuring the response of the strain tensor in (2.45), to small deformations in the crystalline nickel, i.e. $C_{ij} = \delta\sigma_i/\delta\epsilon_j$. To conduct finite-temperature elastic constant measurements, Zhou *et al.* propose using long time-averaged runs of each crystal deformation in order to reduce the thermal noise contribution to the total error.[135] As this work is focused on the error due to potential parameters, the 0 K measurements were used to ensure any measured error could be attributed wholly to the potential ensemble.

Pair potentials are known to be unable to resolve the differences in C_{12} and C_{44}, due to an insufficient number of parameters needed to describe the off-diagonal tensor compo-

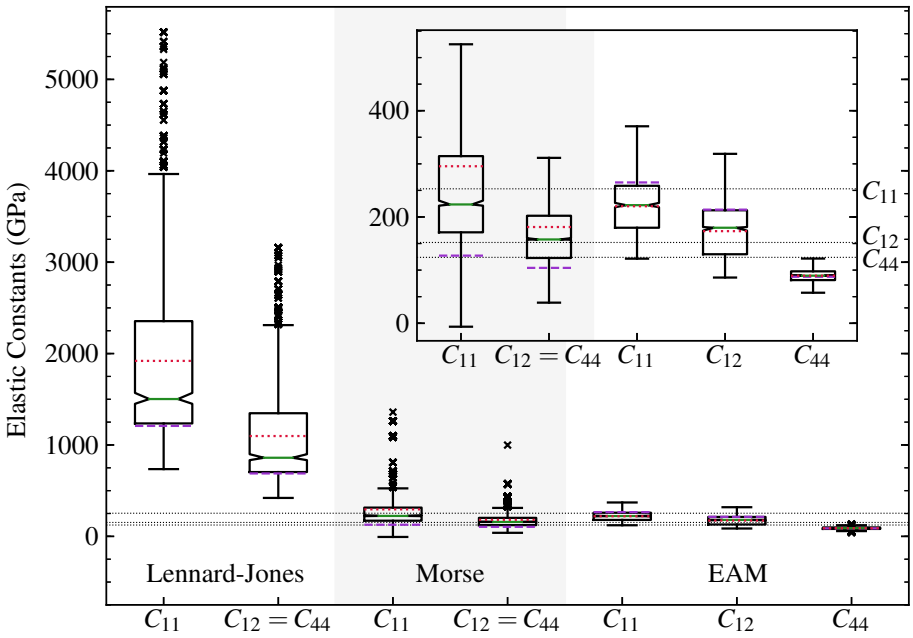

Figure 4.10: Elastic constants for each potential, with the Morse and EAM values shown without outliers in the inset for clarity. The box denotes the IQR, with whiskers extending to 1.5×IQR. The crosses outside of the IQR indicate the outlier values. The ensemble median is shown in solid green, with the ensemble mean in dotted red. The best fit potential prediction is shown in dashed purple.

nents. This is illustrated in figure 4.10, with only the EAM potential having distinct values for these elastic tensor components. It is noticeable that the best fit potential does not necessarily lie near the centre of the prediction interval. This highlights that there exist alternative potentials of competing suitability, which may shift predictions in a particular direction, away from the initially obtained best fit value.

The performance of the Lennard-Jones potential in the prediction of C_{11}, and C_{12}/C_{44} is poor, as expected. There is a very large spread in the predicted values for both quantities; failing to reproduce the experimental values is to be expected of an ill-fit two parameter potential. Due to the outliers in the fit, a stark difference between the mean and median values is observed. When dealing with a model which clearly fails to correctly reproduce the elastic constants, this tells us little more than there is a disagreement in predictions from ensemble members. Further to this, the best fit predictions are also vastly incorrect.

Together, these observations clearly demonstrate the known limitations of such a simple potential, illustrating results in line with an insufficient choice of model.

The Morse potential is able to capture the expected experimental values for both C_{11} and C_{44}, although still unable to resolve the C_{12}/C_{44} difference, due to the pair potential nature. The significant spread in the ensemble mean and median predictions, illustrate a disagreement in predictions from candidate potentials, which again alludes to insufficient model choice. This is another example of poor model selection, despite promising initial predictions of the diagonal elasticity components.

The EAM potential is able to capture the expected trend for the three constants, but it does not capture the expected value of C_{44} in the uncertainty. The initial failure of the fitted EAM potential to achieve the expected off-diagonal C_{44} component, could be rectified through improvements to the reference data. As discussed in section 4.3.2, when evaluating the performance of the best fit EAM potential against DFT predictions, for various cell deformations (figures 4.4a and 4.4b), the reference data contains less information for pair distances $\sim 0.2-0.4\,\text{Å}$ less than the ideal lattice separations. Again, this is illustrated in the radial distribution function shown in figure 4.1b. As the calculations of the elastic constants involve small compressive perturbations to the lattice, it is worth noting that the resulting EAM forces correspond to interatomic separations which have been fitted to a relatively limited subset of reference data. Irrespective of this observation, comparison of the best-fit tensor component predictions with the ensemble mean and median illustrate the performance of suitable alternative candidate potentials. Figure 4.10 focuses on demonstrating the output of the *potfit* uncertainty quantification. For a comparison of production-grade EAM potentials in predicting the elastic constants, the reader is referred to Rassoulinejad-Mousavi, Mao and Zhang.[120]

4.5.3 The Linear Thermal Expansion Coefficient at 300 K

Finally, the thermal expansion coefficient for nickel at 300 K was calculated. The linear expansion of the solid is investigated to compare the effects of energetic contributions to the system, provided by the potentials. The thermal expansion coefficient was calculated by evaluating the length change of crystalline nickel for five temperatures, symmetrically

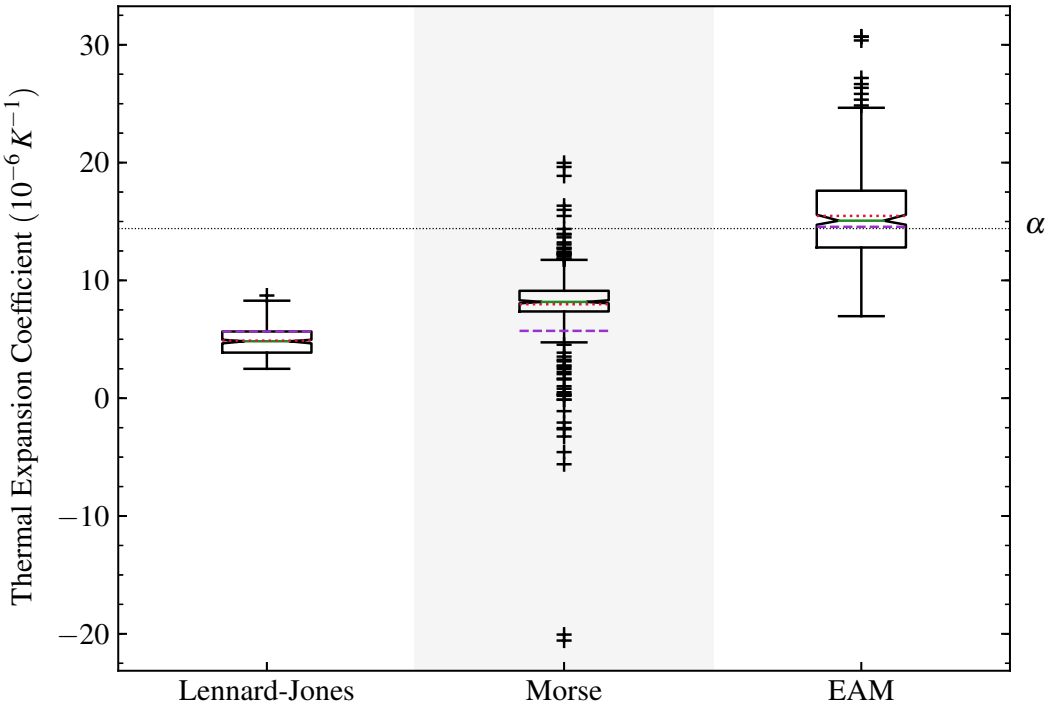

Figure 4.11: The uncertainties in thermal expansion coefficient at 300 K for each potential. The box denotes the IQR, with whiskers extending to 1.5×IQR. The crosses outside of the IQR indicate the outlier values. The ensemble median is shown in solid green, with the ensemble mean in dotted red. The best fit potential prediction is shown in dashed purple.

distributed around 300 K at 20 K intervals. A curve is then fitted to the results using regression, to find the thermal expansion coefficient at 300 K.

On first inspection, the results in figure 4.11 may misleadingly imply that the pair potentials outperform the EAM, due to the small spread in the ensemble members. However, upon closer inspection, the predicted pair values in fact fail to capture the correct value within the IQR, and even within the tails of the distribution. The significant outliers in the Morse predictions, again imply a disagreement in the predictions of ensemble members. The EAM potential however, does manage to bound the correct value within the uncertainty, although a large spread in the uncertainty is again likely a result of the choice of reference data to which it has been fit.

The uncertainties for the thermal expansion coefficient demonstrate the importance of looking at the predictions of a selection of relevant QoI, when evaluating the suitability of a potential model. In table 4.6, the predicted lattice constant, and uncertainty bounds, for the Lennard-Jones potential fail to capture the simple equilibrium quantity, despite a large uncertainty. This unsuccessful prediction, combined with the high relative uncertainty in the thermal expansion coefficient, illustrate an example of poor model selection. This leads to a caveat in the application of the ensemble framework: its application should not be used as a means to bypass an informed fitting of potentials. In the case of the Lennard-Jones potential, the initial best fit performed poorly in predicting the correct forces of the reference set. Furthermore, the incorrect prediction of the lattice parameter by the best fit potential, and the resulting ensemble, imply that the model is insufficient. Failure to assess the model suitability at each stage of the fitting process can result in misleading uncertainties, as we have attempted to illustrate in this case.

4.6 Summary

The initial testing of the implementation is undertaken for three models of bulk nickel, highlighting some important aspects of the potential fitting and uncertainty quantification process. Foremost, is the choice of material, and system being modelled, follwed by the effect of reference data quantity and quality on the defined cost landscape, closing with the potential precision of subsequent predictions. The choice of model defines both the dimensionality and complexity of the cost landscape, as well as affecting the accuracy of the resulting predictions. The eigenvalues and eigenvectors of the minimum cost Hessian provide vital information regarding the shape of the landscape from which MCMC samples will be drawn, as well as information on the relative contributions from model parameters. Tuning of the MCMC algorithm to accept 23 % of steps requires a significant number of MCMC calculations, in order to initially to distil the associated R parameter. Having assessed the parameters necessary for the MCMC cost landscape sampling, it is then imperative that the parameter autocorrelation and parameter projections are examined, to extract the required decorrelation time and ensure sample convergence. Finally, with the decorrelated ensemble,

the propagation of parameter uncertainties to QoI can begin.

As is clear, the workflow is iterative in nature, and at times requires a few steps before a seemingly simple question can be understood. However, despite being a relatively involved procedure, the uncertainty quantification process requires few decisions from the user in order to produce a significant quantity of information about the model + reference data system. In order to achieve successful potential fitting and uncertainty quantification results, it appears to be of greatest importance that the results from each stage of the workflow are astutely analysed, so that a deeper understanding of the system is obtained. As each reference data set and model combination defines an entirely new cost landscape, only through further use of the uncertainty quantification implementation will its limitations and subtle nuances be truly understood.

5

Outlook and Future Work

"Begin at the beginning," the King said, gravely, "and go on till you come to an end; then stop."

– Lewis Carroll, *Alice in Wonderland*

The open source *potfit* potential fitting workflow has been enhanced to generate an ensemble of potentials which encapsulate the uncertainty of the fitted parameters. This allows for propagation of the parameter uncertainty to quantities of interest, simulated using molecular dynamics. The ensemble is generated by sampling the cost landscape, using a Markov chain Monte Carlo sampling technique, developed for sloppy models. Additionally, functionality to investigate the cost landscapes encountered within the potential fitting process has been made available through optional parameters, in order to assist the user in diagnosing any issues which may arise. A demonstration of the uncertainty quantification workflow has been presented, based on the fitting of three potential models of increasing complexity to nickel reference data. Through this work, the propagation of parameter uncertainties to quantities of interest is detailed and illustrated.

The ensemble method can be used to build an understanding of the impact of parameter uncertainty on the precision of quantities of interest. As illustrated in the results for nickel, and reported by Pernot and Cailliez,[25,26] users must be aware that this method provides a lower bound of the error bars; some quantities might not be described well by a potential model. This is why it is mandatory that users of the ensemble method implementation diligently evaluate the suitability of the model throughout the fitting (as reported by rms errors in reproducing reference data), and the uncertainty quantification process.

This is a further puzzle piece towards reproducible and transparent MD simulations; an effective potential should not exist on its own, but rather together with its implementation (as e.g. provided by the OpenKIM framework[6]), its reference data, and its uncertainties (*potfit* + UQ). This integration is also a step toward predictive simulations, i.e. with error bounds determined *a priori*. Whilst progress has been made to introduce uncertainty quantification into the potential fitting process, there are many avenues for future work and further improvement. A crucial next step for the uncertainty quantification community, is the creation of best practise guidelines in the reporting of uncertainties, including, foremost, a standardised data structure for their reporting. As a result of this work,[42] there has been interest from the team behind the OpenKIM repository, in creating a UQ data structure in their API. In the following, suggestions for future development of the potential fitting UQ workflow are introduced and preliminary approaches to the research avenues proposed.

5.1 Technical Aspects of the Implementation

The first area for advancement of the work, pertains specifically to the technical aspects of the implementation. Unlike the pre-existing potential capabilities with *potfit*, the uncertainty quantification workflow involves a larger degree of data analysis. Therefore, whilst there remain benefits with regards to memory handling, and calculation time, in writing the ensemble method work in C, there are benefits in exposing certain *potfit* routines (i.e. `calc_forces()`) to python, in order to more seamlessly utilise some of its powerful data analysis modules. Additionally, allowing *potfit* to be called or interrogated from with a python script would allow for further integration with atomistic simulation packages (e.g.

the ASE module, DFT and MD software), as well as the ability to utilise other generalised uncertainty quantification modules.

Another, fairly easily implementable technical improvement, would be the parallelisation of the MCMC algorithm, to enable multiple Markov chains to sample the cost landscape at simultaneously. Parallelisation of any other aspect of the ensemble method is not possible, due to the serial nature of decisions at each step in the internal *potfit* process.The ability to run multiple MCMC chains in parallel, with the option to seed each with a different R value, i.e. to sample multiple R values with a single simulation. This would make the tuning of the step size much faster and less involved, as multiple samples could be generated from one simulation. Another consideration of this possible future work, would pertain the handling of the significantly greater amount of data produced in a single simulation. Working on the assumption that each Markov chain would output to its own file, one way of controlling the size of output files, would be to control the frequency of accepted steps written to output (i.e. only output the decorrelated sub-sample). This would significantly decrease the size of files generated in simulations with long autocorrelation times.

Finally, within the implementation, the most beneficial improvement would be the inclusion of alternative sampling algorithms, in order to ensure the sampling of each potential parameter distribution converges efficiently. In scenarios where the cost landscape is not well described by the Hessian (e.g. a banana-shaped minimum, or an asymmetric basin) it may be that a traditional Markov chain, without curvature information, would perform better. Therefore, for complicated cost landscapes, the addition of alternative sampling methods could prove to be a beneficial addition to the implementation. Some examples of alternative sampling strategies are Riemann Monte Carlo methods,[131] or affine invariant samplers[132]. These methods may also allow for the incorporation of prior information about the parameters into the ensemble generation process, potentially allowing for a fully Bayesian treatment.

Alternatively, the mathematical approach underpinning umbrella sampling techniques,[136–138] could provide a potentially interesting avenue with which to understand the cost landscape minimum, especially if modifications were made to restrict, as opposed to discourage, sampling to within the cost minimum basin. This could possibly be achieved

through the inclusion of a multidimensional Gaussian-like weighting term within the step acceptance calculation. If coupled with the simplification of the potential energy landscape (i.e. dimentionality reduction using sketch-map) proposed by Ceriotti *et al.*[139] a more complete understanding of the reference data-to-cost landscape mapping could be understood. As it is not immediately obvious how a restrictive Gaussian term may affect the statistics, as such a suggestion would require an investigation of possible mathematical formulations, and is therefore incredibly speculative as a possible avenue for future work at present.

5.2 Reference Data and Inherited Error

The second class of improvements and further work, relate to the uncertainties in the system inherited from the *ab initio* simulation snapshots which constitute the reference data. The error introduced through the reference data can be categorised into two types. The first, is error introduced by the suitability of the reference snapshots chosen (i.e. do they provide a good representation of the target configuration space?), and second is due to the parameters used and convergence tolerances of the reference data simulations.

It is not know if the morphologies and systems included in reference data are most suitable for the target quantities to be simulated by the fitted potential. This would be an example of an interatomic potential with poor transferability, and so a further avenue to explore would be quantification of the suitability of a given reference set for the desired quantities to be simulated. The weighted least squares approach from Zhang and Trinkle[31] is one example of how this may be achieved, by creating an inverse problem from the fitting process. This method would not only allow for the understanding of reference data suitability, but could also be used in conjunction with the ensemble method to explore alternative potentials to the lowest cost fit which may be more suitable in simulating the desired quantities.

The inherited error from the first principles simulations is not yet accounted for within the *potfit* uncertainty quantification workflow. As observed in chapter 4, the errors in forces, energies and stresses of reference data from DFT calculations should be significantly less those from the potential fitting. Therefore, the inherited systematic reference data

errors could be included through a simple first order addition.

Whilst there are a growing variety of tools available to quantify the *ab initio* simulation errors,[21–23] it is not immediately clear how they could be used to inform the potential fitting information in the current workflow. A possible approach could be to take inspiration from the mathematical construction of GAP,[8,10] and work towards a UQ workflow which deals in Gaussian processes. If the reference data quantification approach from Zhang and Trinkle[31] was also implemented, this would define a fitting landscape, similar to the cost landscape, but for the reference data. Sampling from this landscape to measure reference data error as in the methods above,[21–23] could then potentially allow for the fitting of Gaussian distributions, which could be propagated through the process. Alternatively, taking inspiration from neural network-fitted potentials,[12–16] the replacement of the end-to-end *potfit* UQ workflow by a neural network with latent variables or hidden layers could warrant further consideration.

5.3 Uncertainty Quantification Methodology

The final area of future work deals specifically with the choice of uncertainty quantification method. As mentioned in section 5.1, there are an array of alternative approaches to uncertainty quantification implemented within python modules. Providing the user with a variety of uncertainty quantification techniques when using *potfit* would allow for the comparison and evaluation of the results of each method. This could additionally produce insight into the limitations of various methods, and highlight the potential coupling of techniques which may be more powerful than a single method with known limitations.

The most desirable method to be implemented would be a fully Bayesian approach, using Gaussian process regression (GPR), to infer the uncertainties in quantities of interest without the need for the full set of MD simulations of each potential in an ensemble. As also mentioned in section 5.2, other machine learning approaches such as neural networks could be used to quantify the uncertainties. It is of the opinion of the author that the most favourable next steps with regards to the UQ of classical effective potentials would be to expose necessary routines within the *pofit* code in python so that alternative pre-

existing uncertainty quantification and data analysis modules could be exploited in addition to ensemble method presented within this work.

The caveat which should preface all of these future avenues, is that the choice of UQ method is always a question of cost. The future of this area of research relies upon the adoption of uncertainty quantification, and reporting of the results, by other researchers in the materials modelling community. In order to best facilitate the rapid uptake of UQ by our peers, it would be reasonable to expect that the cost of undertaking the UQ process should not exceed that of their simulations of interest. As has been discussed within this work, the ensemble method does require care to be taken, however our approach does not require domain knowledge from the user upfront, and so alternative approaches would be much more involved. In the work presented, each ensemble member is a potential in its own right, and therefore each maintains the simulation speedups that effective potentials were designed to enable. Additionally, the uncertainty quantification we have developed can be switched on or off, which may not be possible with highly complex implementation with an involved workflow from the beginning. Overall, the process of uncertainty quantification for effective potentials is complicated, and contains endless avenues for additional investigations within a single simulation. The implementation presented in this work is an effort to make uncertainty quantification of classical effective potentials more accessible, whilst remaining relatively uninvolved, yet informative.

A

Parameter Projections of the EAM Ensemble

The EAM ensemble results for nickel are illustrated in the following figures. Table A.1 details the EAM model and corresponding parameters for reference.

Table A.1: The EAM analytic functions and associated parameters.

Model	Analytic Function	Parameters
EAM	$E_i = \frac{1}{2}\sum_{i\neq j}^{N} V_M(r_{ij}) + F(n_i) \quad \text{with} \quad n_i = \sum_{j\neq i}^{N} \rho(r_{ij})$	D_e, a, r_e
	$\rho(r) = r^{-\beta}\left[1 + a_1 \cos(\alpha r + \varphi)\right]$	a_1, α, φ, β
	$F(n) = F_0\left[1 - \gamma \ln n\right] n^{\gamma} + F_1 n$	F_0, γ, F_1
Smooth Cutoff	$V_{SC}(r) = \Psi\left(\frac{r-r_c}{h}\right) V(r) \qquad \text{where} \quad \Psi(x) = \frac{x^4}{1+x^4}$	$r_c = 10$ $h = 0.75$

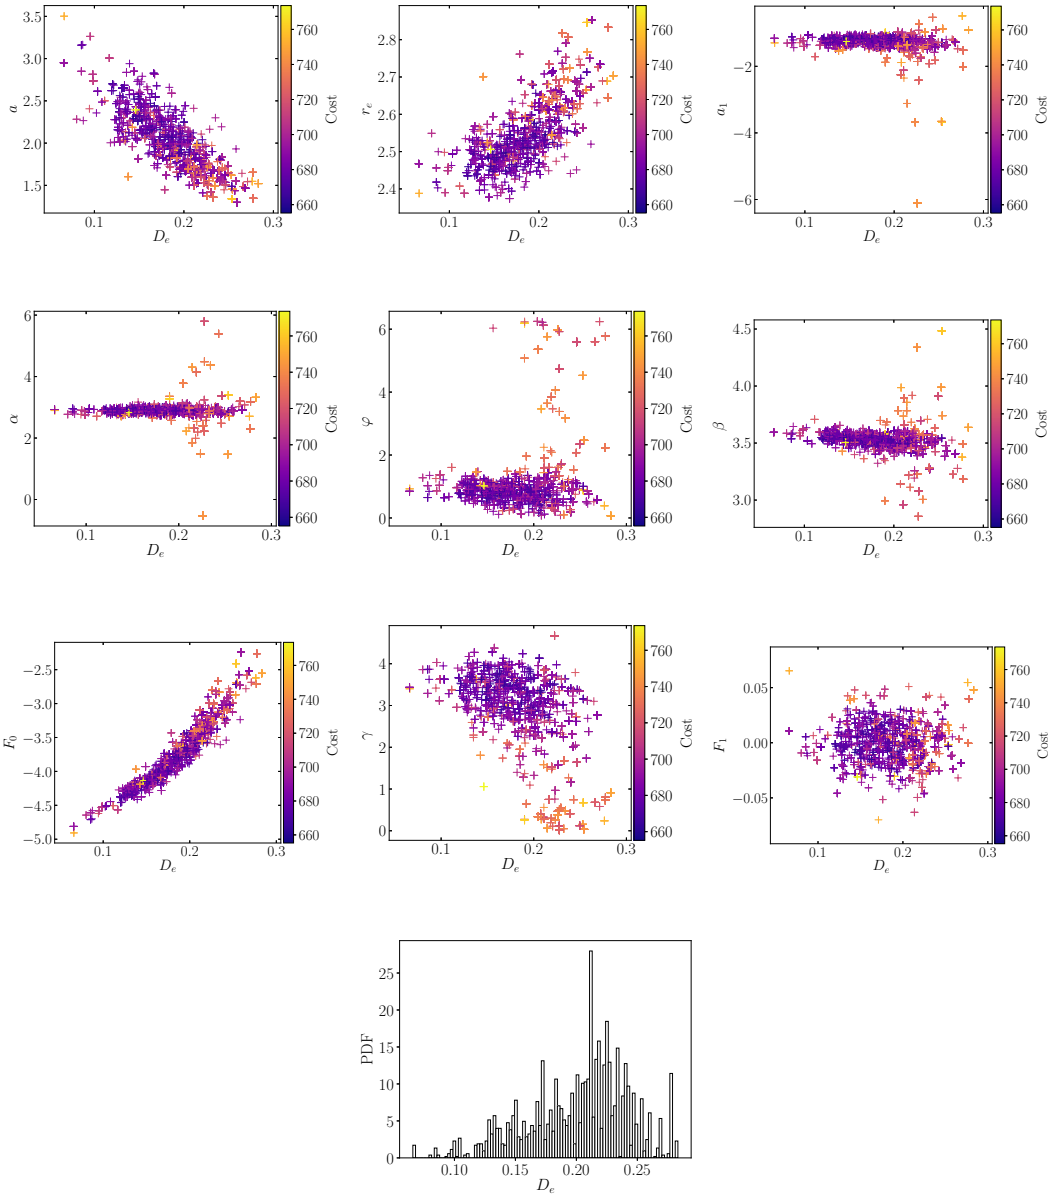

Figure A.1: EAM projections of D_e parameter spread and covariance of 500 potential ensemble.

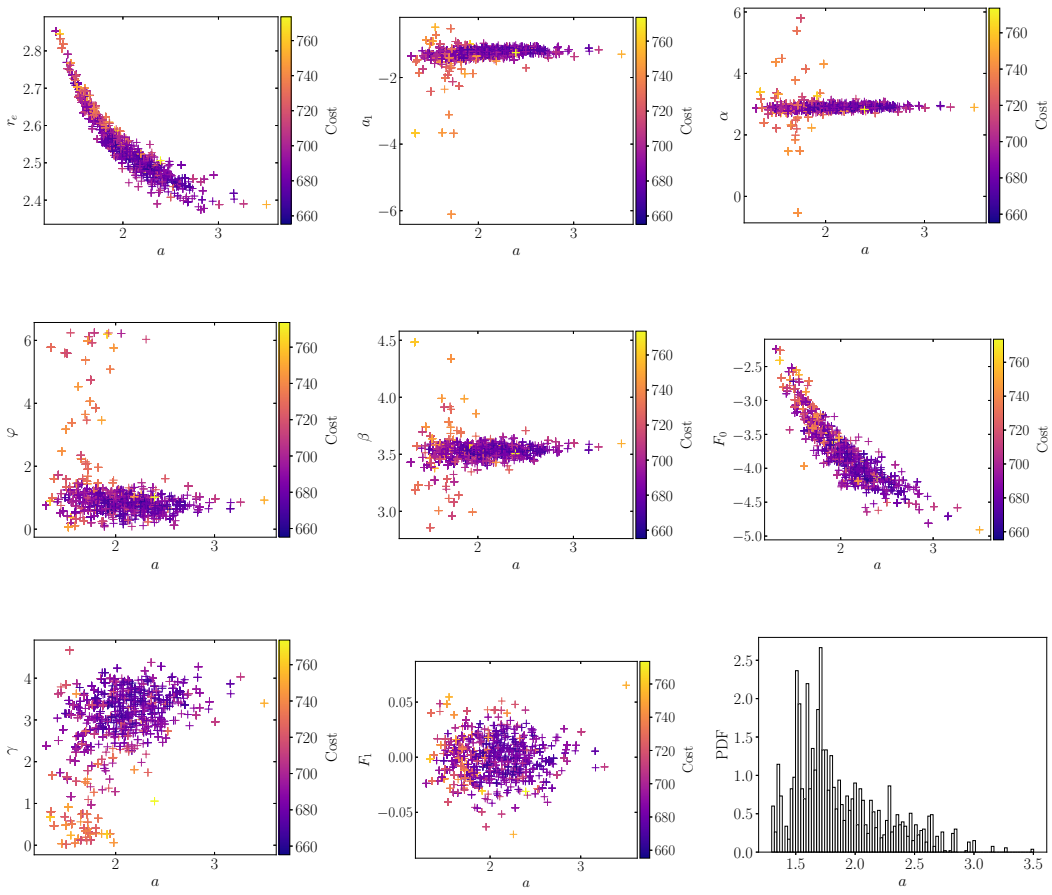

Figure A.2: EAM projections of a parameter spread and covariance of 500 potential ensemble.

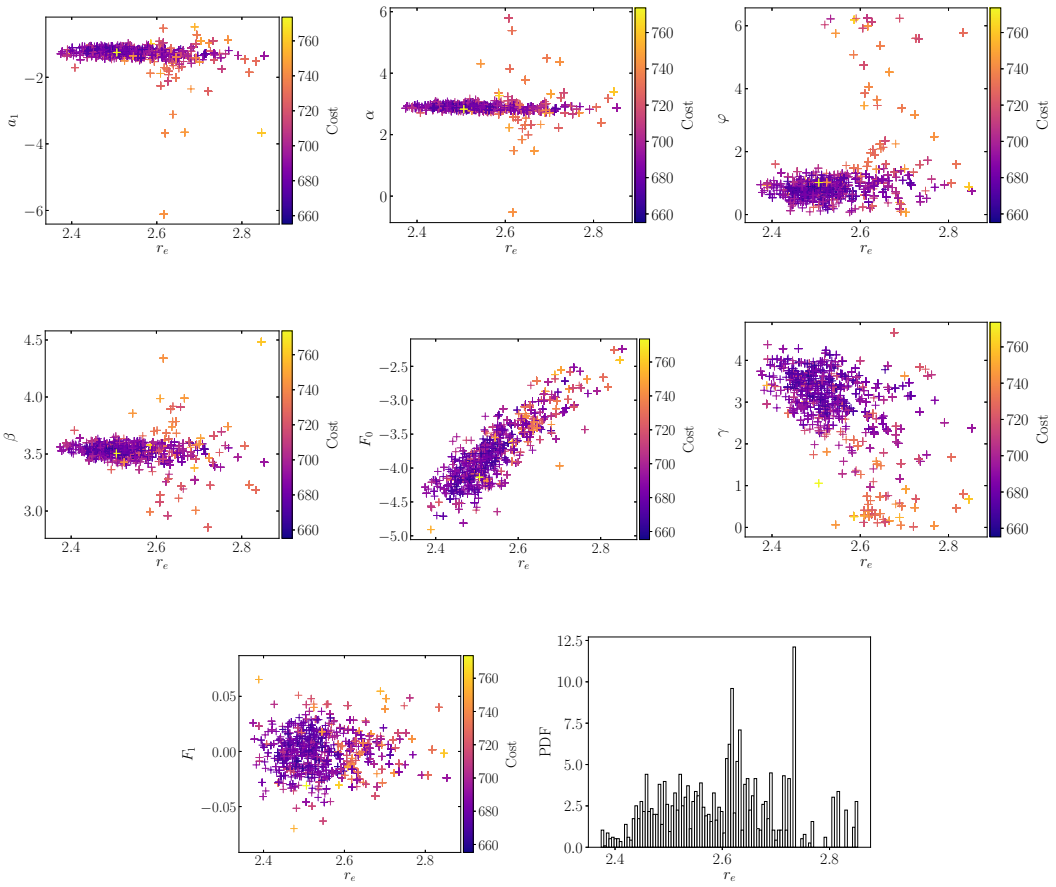

Figure A.3: EAM projections of r_e parameter spread and covariance of 500 potential ensemble.

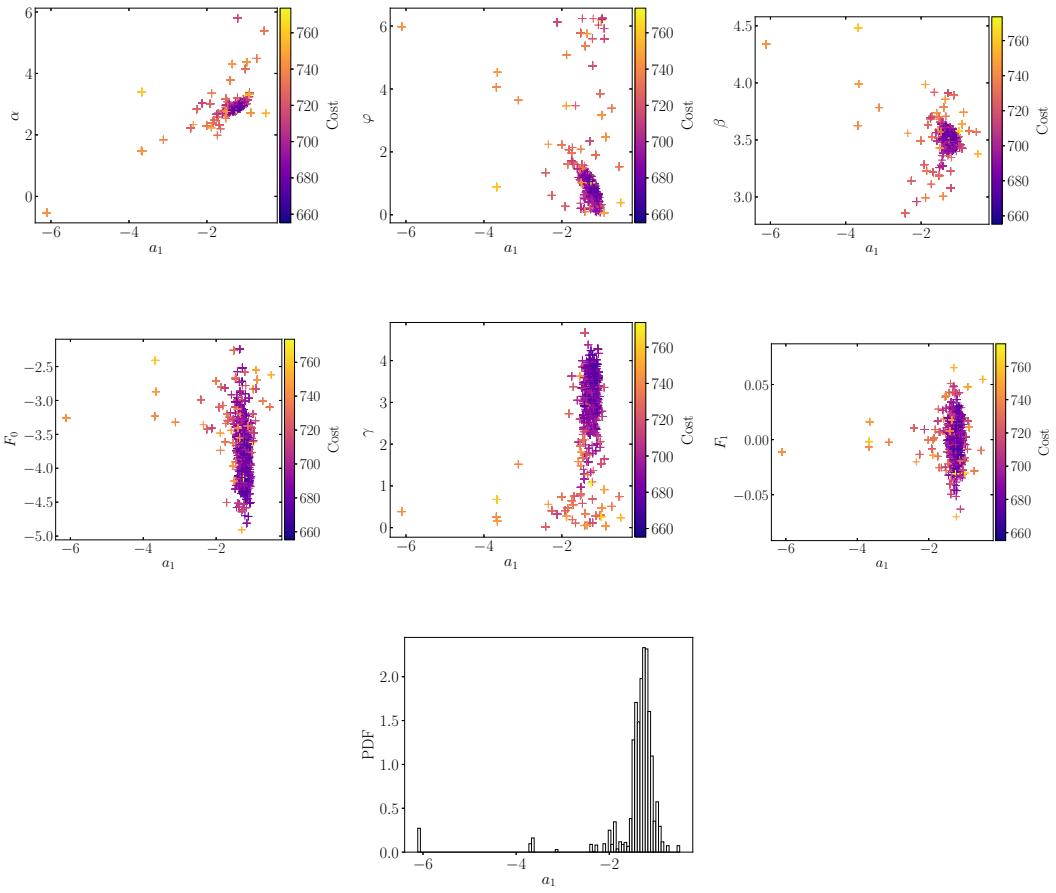

Figure A.4: EAM projections of a_1 parameter spread and covariance of 500 potential ensemble.

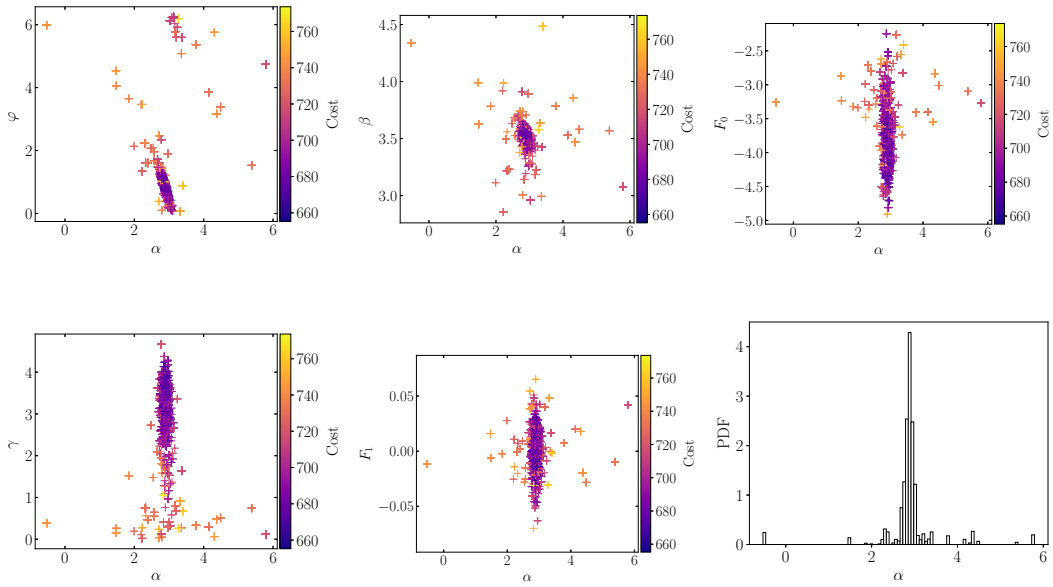

Figure A.5: EAM projections of α parameter spread and covariance of 500 potential ensemble.

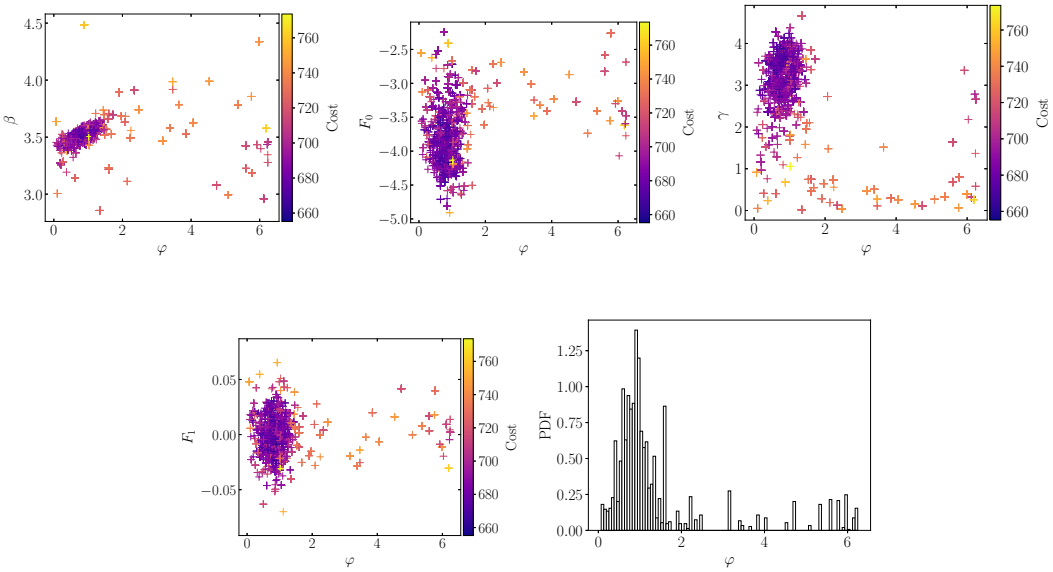

Figure A.6: EAM projections of φ parameter spread and covariance of 500 potential ensemble.

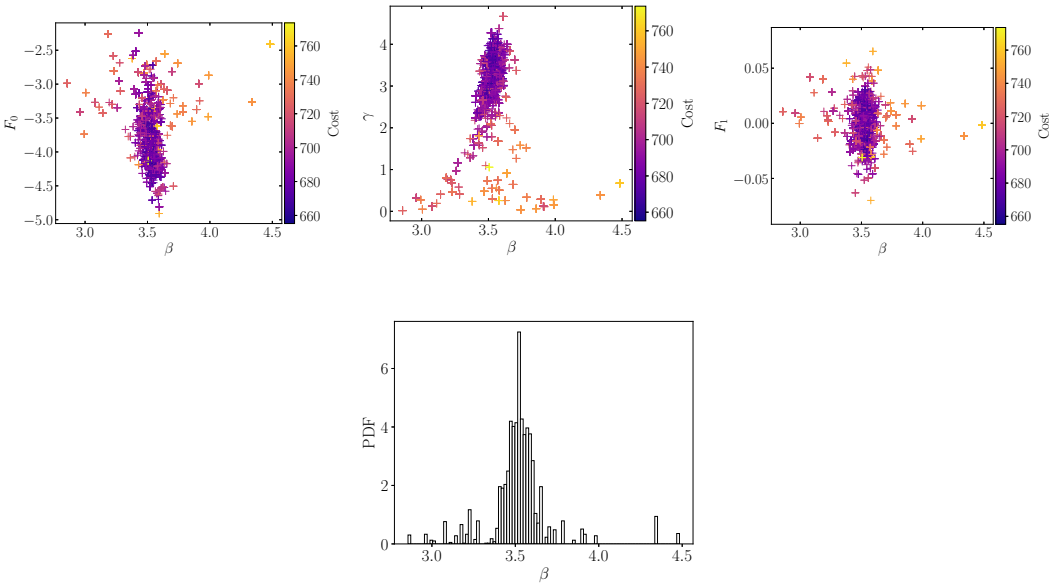

Figure A.7: EAM projections of β parameter spread and covariance of 500 potential ensemble.

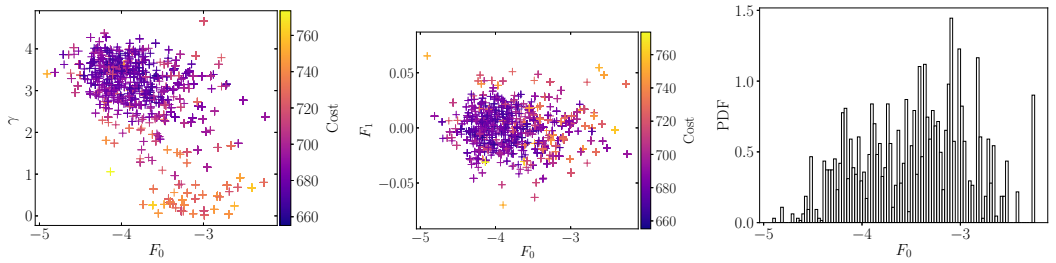

Figure A.8: EAM projections of F_0 parameter spread and covariance of 500 potential ensemble.

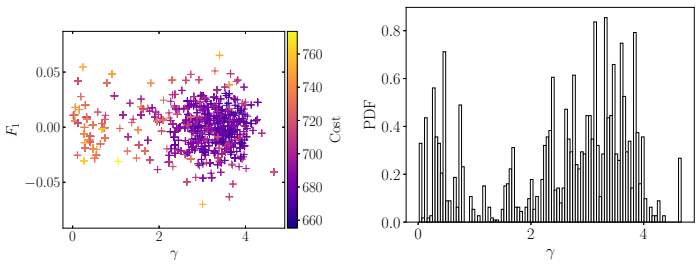

Figure A.9: EAM projections of γ parameter spread and covariance of 500 potential ensemble.

Bibliography

[1] Ercolessi F and Adams J B 1994 Interatomic Potentials from First-Principles Calculations: The Force-Matching Method *Europhysics Letters (EPL)* **26** 583–8

[2] Masia M, Guàrdia E and Nicolini P 2014 The force matching approach to multiscale simulations: Merits, shortcomings, and future perspectives *International Journal of Quantum Chemistry* **114** 1036–40

[3] Brommer P and Gähler F 2006 Effective potentials for quasicrystals from ab-initio data *Philosophical Magazine* **86** 753–8

[4] Brommer P and Gähler F 2007 Potfit: effective potentials from ab initio data *Modelling and Simulation in Materials Science and Engineering* **15** 295–304

[5] Brommer P, Kiselev A, Schopf D, Beck P, Roth J and Trebin H R 2015 Classical interaction potentials for diverse materials from ab initio data: a review of potfit *Modelling and Simulation in Materials Science and Engineering* **23** 074002

[6] Wen M, Li J, Brommer P, Elliott R S, Sethna J P and Tadmor E B 2017 A KIM-compliant potfit for fitting sloppy interatomic potentials: application to the EDIP model for silicon *Modelling and Simulation in Materials Science and Engineering* **25** 014001

[7] Chan H, Narayanan B, Cherukara M, Sen F, Sasikumar K, Gray S, Chan M and Sankaranarayanan S 2019 Machine Learning Classical Interatomic Potentials for Molecular Dynamics from First-Principles Training Data *The Journal of Physical Chemistry C* **123** 6941–57

[8] Bartók A P, Payne M C, Kondor R and Csányi G 2010 Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons *Physical Review Letters* **104** 136403

[9] Deringer V L and Csányi G 2017 Machine learning based interatomic potential for amorphous carbon *Physical Review B* **95** 094203

[10] Bartók A P, Kermode J, Bernstein N and Csányi G 2018 Machine Learning a General-Purpose Interatomic Potential for Silicon *Physical Review X* **8** 041048

[11] Byggmästar J, Hamedani A, Nordlund K and Djurabekova F 2019 Machine-learning interatomic potential for radiation damage and defects in tungsten *Physical Review B* **100** 144105

[12] Behler J and Parrinello M 2007 Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces *Physical Review Letters* **98** 146401

[13] Behler J 2010 Neural network potential-energy surfaces for atomistic simulations in M Springborg, editor, *Chemical Modelling: Applications and Theory Volume 7* The Royal Society of Chemistry 1–41

[14] Artrith N and Urban A 2016 An implementation of artificial neural-network potentials for atomistic materials simulations: Performance for TiO2 *Computational Materials Science* **114** 135–50

[15] Kobayashi R, Giofré D, Junge T, Ceriotti M and Curtin W A 2017 Neural network potential for Al-Mg-Si alloys *Physical Review Materials* **1** 053604

[16] Pun G P P, Batra R, Ramprasad R and Mishin Y 2019 Physically informed artificial neural networks for atomistic modeling of materials *Nature Communications* **10** 2339

[17] Lin H and Truhlar D G 2007 QM/MM: what have we learned, where are we, and where do we go from here? *Theoretical Chemistry Accounts* **117** 185–99

[18] Moras G, Choudhury R, Kermode J R, Csányi G, Payne M C and De Vita A 2010 The "Learn on the Fly" Molecular Dynamics Scheme in T Dumitrica, editor, *Trends in*

Computational Nanomechanics. Challenges and Advances in Computational Chemistry and Physics volume 9 Springer Netherlands, Dordrecht 1–23

[19] Li Z, Kermode J R and De Vita A 2015 Molecular Dynamics with On-the-Fly Machine Learning of Quantum-Mechanical Forces *Physical Review Letters* **114** 096405

[20] Schöberl M, Zabaras N and Koutsourelakis P S 2017 Predictive coarse-graining *Journal of Computational Physics* **333** 49–77

[21] Lundgaard K T, Wellendorff J, Voss J, Jacobsen K W and Bligaard T 2016 mBEEF-vdW: Robust fitting of error estimation density functionals *Physical Review B* **93** 235162

[22] Aldegunde M, Kermode J R and Zabaras N 2016 Development of an exchangeâĂŞcorrelation functional with uncertainty quantification capabilities for density functional theory *Journal of Computational Physics* **311** 173–95

[23] Tegner B E and Ackland G J 2012 Pseudopotential errors in titanium *Computational Materials Science* **52** 2–6

[24] Lejaeghere K, Bihlmayer G, Bjorkman T, Blaha P, Blugel S, Blum V, Caliste D, Castelli I E, Clark S J, Dal Corso A, de Gironcoli S, Deutsch T, Dewhurst J K, Di Marco I, Draxl C, Du ak M, Eriksson O, Flores-Livas J A, Garrity K F, Genovese L, Giannozzi P, Giantomassi M, Goedecker S, Gonze X, Granas O, Gross E K U, Gulans A, Gygi F, Hamann D R, Hasnip P J, Holzwarth N A W, Iu an D, Jochym D B, Jollet F, Jones D, Kresse G, Koepernik K, Kucukbenli E, Kvashnin Y O, Locht I L M, Lubeck S, Marsman M, Marzari N, Nitzsche U, Nordstrom L, Ozaki T, Paulatto L, Pickard C J, Poelmans W, Probert M I J, Refson K, Richter M, Rignanese G M, Saha S, Scheffler M, Schlipf M, Schwarz K, Sharma S, Tavazza F, Thunstrom P, Tkatchenko A, Torrent M, Vanderbilt D, van Setten M J, Van Speybroeck V, Wills J M, Yates J R, Zhang G X and Cottenier S 2016 Reproducibility in density functional theory calculations of solids *Science* **351** aad3000

[25] Pernot P 2017 The parameter uncertainty inflation fallacy *The Journal of Chemical Physics* **147** 104102

[26] Pernot P and Cailliez F 2017 A critical review of statistical calibration/prediction models handling data inconsistency and model inadequacy *AIChE Journal* **63** 4642–65

[27] Frederiksen S L, Jacobsen K W, Brown K S and Sethna J P 2004 Bayesian Ensemble Approach to Error Estimation of Interatomic Potentials *Physical Review Letters* **93** 165501

[28] Rizzi F, Najm H N, Debusschere B J, Sargsyan K, Salloum M, Adalsteinsson H and Knio O M 2012 Uncertainty Quantification in MD Simulations. Part II: Bayesian Inference of Force-Field Parameters *Multiscale Modeling & Simulation* **10** 1460–92

[29] Mishra A, Hong S, Rajak P, Sheng C, Nomura K i, Kalia R K, Nakano A and Vashishta P 2018 Multiobjective genetic training and uncertainty quantification of reactive force fields *npj Computational Materials* **4** 42

[30] Musil F, Willatt M J, Langovoy M A and Ceriotti M 2019 Fast and Accurate Uncertainty Estimation in Chemical Machine Learning *Journal of Chemical Theory and Computation* **15** 906–15

[31] Zhang P and Trinkle D R 2015 Database optimization for empirical interatomic potential models *Modelling and Simulation in Materials Science and Engineering* **23** 065011

[32] Angelikopoulos P, Papadimitriou C and Koumoutsakos P 2012 Bayesian uncertainty quantification and propagation in molecular dynamics simulations: a high performance computing framework. *J. Chem. Phys.* **137** 144103

[33] Rizzi F, Jones R E, Debusschere B J and Knio O M 2013 Uncertainty quantification in MD simulations of concentration driven ionic flow through a silica nanopore. I. Sensitivity to physical parameters of the pore *The Journal of Chemical Physics* **138** 194104

[34] Vohra M, Nobakht A Y, Shin S and Mahadevan S 2018 Uncertainty quantification

in non-equilibrium molecular dynamics simulations of thermal transport *International Journal of Heat and Mass Transfer* **127** 297–307

[35] Messerly R A, Knotts T A and Wilding W V 2017 Uncertainty quantification and propagation of errors of the Lennard-Jones 12-6 parameters for n -alkanes *The Journal of Chemical Physics* **146** 194110

[36] Messerly R A, Shirts M R and Kazakov A F 2018 Uncertainty quantification confirms unreliable extrapolation toward high pressures for united-atom Mie λ -6 force field *The Journal of Chemical Physics* **149** 114109

[37] Brown K S and Sethna J P 2003 Statistical mechanical approaches to models with many poorly known parameters *Physical Review E* **68** 021904

[38] Machta B B, Chachra R, Transtrum M K and Sethna J P 2013 Parameter Space Compression Underlies Emergent Theories and Predictive Models *Science* **342** 604–7

[39] Tadmor E B, Elliott R S, Sethna J P, Miller R E and Becker C A 2011 (KIM), Knowledgebase of Interatomic Potentials

[40] Tadmor E B, Elliott R S, Sethna J P, Miller R E and Becker C A 2011 The potential of atomistic simulations and the knowledgebase of interatomic models *Jom* **63** 17

[41] Tadmor E B, Elliott R S, Phillpot S R and Sinnott S B 2013 NSF cyberinfrastructures: A new paradigm for advancing materials simulation *Current Opinion in Solid State and Materials Science* **17** 298–304

[42] Longbottom S and Brommer P 2019 Uncertainty quantification for classical effective potentials: an extension to potfit *Modelling and Simulation in Materials Science and Engineering* **27** 044001

[43] Fock V 1930 Näherungsmethode zur Lösung des quantenmechanischen Mehrkörperproblems *Zeitschrift für Physik* **61** 126–48

[44] Hohenberg P and Kohn W 1964 Inhomogeneous Electron Gas *Physical Review* **136** B864–71

[45] Levy M 1979 Universal variational functionals of electron densities, first-order density matrices, and natural spin-orbitals and solution of the v-representability problem *Proceedings of the National Academy of Sciences* **76** 6062–5

[46] Kohn W and Sham L J 1965 Self-Consistent Equations Including Exchange and Correlation Effects *Physical Review* **140** A1133–8

[47] Perdew J P, Burke K and Ernzerhof M 1996 Generalized Gradient Approximation Made Simple *Physical Review Letters* **77** 3865–8

[48] Elliott P, Lee D, Cangi A and Burke K 2008 Semiclassical Origins of Density Functionals *Physical Review Letters* **100** 256406

[49] Furche F and Perdew J P 2006 The performance of semilocal and hybrid density functionals in 3d transition-metal chemistry *The Journal of Chemical Physics* **124** 044103

[50] Monkhorst H J and Pack J D 1976 Special points for Brillouin-zone integrations *Physical Review B* **13** 5188–92

[51] Pickett W E 1989 Pseudopotential methods in condensed matter applications *Computer Physics Reports* **9** 115–97

[52] Prandini G, Marrazzo A, Castelli I E, Mounet N and Marzari N 2018 A Standard Solid State Pseudopotentials (SSSP) library optimized for precision and efficiency (Version 1.1, data download)

[53] Kleinman L and Bylander D M 1982 Efficacious Form for Model Pseudopotentials *Physical Review Letters* **48** 1425–8

[54] Vanderbilt D 1990 Soft self-consistent pseudopotentials in a generalized eigenvalue formalism *Physical Review B*

[55] Hamann D R 2013 Optimized norm-conserving Vanderbilt pseudopotentials *Physical Review B* **88** 085117

[56] Blöchl P E 1994 Projector augmented-wave method *Physical Review B*

[57] Kresse G and Joubert D 1999 From ultrasoft pseudopotentials to the projector augmented-wave method *Physical Review B* **59** 1758–75

[58] Martin R M 2004 Pseudopotentials in *Electronic Structure: Basic Theory and Practical Methods* Cambridge University Press, Cambridge 204–32

[59] Hellmann H 1933 Zur Rolle der kinetischen Elektronenenergie für die zwischenatomaren Kräfte *Zeitschrift für Physik*

[60] Feynman R P 1939 Forces in molecules *Phys. Rev.* **56** 340–3

[61] Car R and Parrinello M 1985 Unified approach for molecular dynamics and density-functional theory *Physical Review Letters* **55** 2471–4

[62] Frenkel D and Smit B 2002 *Understanding Molecular Simulation* Elsevier

[63] Berendsen H J, Postma J P, Van Gunsteren W F, Dinola A and Haak J R 1984 Molecular dynamics with coupling to an external bath *The Journal of Chemical Physics*

[64] Braun E, Moosavi S M and Smit B 2018 Anomalous Effects of Velocity Rescaling Algorithms: The Flying Ice Cube Effect Revisited *Journal of Chemical Theory and Computation* **14** 5262–72

[65] Andersen H C 1980 Molecular dynamics simulations at constant pressure and/or temperature *The Journal of Chemical Physics* **72** 2384–93

[66] Snook I 2007 *The Langevin and Generalised Langevin Approach to the Dynamics of Atomic, Polymeric and Colloidal Systems* Elsevier

[67] Nosé S 1984 A unified formulation of the constant temperature molecular dynamics methods *The Journal of Chemical Physics* **81** 511–9

[68] Hoover W G 1985 Canonical dynamics: Equilibrium phase-space distributions *Physical Review A*

[69] Kim M, Kim E, Lee S, Kim J S and Lee S 2019 New Method for Constant- NPT Molecular Dynamics *The Journal of Physical Chemistry A* **123** 1689–99

[70] Finnis M 2003 *Interatomic Forces in Condensed Matter* Oxford University Press

[71] Behler J 2017 First Principles Neural Network Potentials for Reactive Simulations of Large Molecular and Condensed Systems *Angewandte Chemie International Edition* **56** 12828–40

[72] Kamath A, Vargas-Hernández R A, Krems R V, Carrington T and Manzhos S 2018 Neural networks vs Gaussian process regression for representing potential energy surfaces: A comparative study of fit quality and vibrational spectrum accuracy *The Journal of Chemical Physics* **148** 241702

[73] Bartók A P and Csányi G 2015 Gaussian approximation potentials: A brief tutorial introduction *International Journal of Quantum Chemistry* **115** 1051–7

[74] Morse P M 1929 Diatomic Molecules According to the Wave Mechanics. II. Vibrational Levels *Physical Review* **34** 57–64

[75] Lennard-Jones J E 1924 On the Determination of Molecular Fields *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **106** 463–77

[76] Filippova V P, Kunavin S A and Pugachev M S 2015 Calculation of the parameters of the Lennard-Jones potential for pairs of identical atoms based on the properties of solid substances *Inorganic Materials: Applied Research* **6** 1–4

[77] Daw M S, Foiles S M and Baskes M I 1993 The embedded-atom method: a review of theory and applications *Materials Science Reports* **9** 251–310

[78] van Dyke J 1978 SCF-Xα-MS calculations for H in metal clusters *Journal of Nuclear Materials* **69-70** 583–5

[79] Stott M J and Zaremba E 1980 Quasiatoms: An approach to atoms in nonuniform electronic systems *Physical Review B* **22** 1564–83

[80] Nørskov J K and Lang N D 1980 Effective-medium theory of chemical binding: Application to chemisorption *Physical Review B* **21** 2131–6

[81] Nørskov J K 1982 Covalent effects in the effective-medium theory of chemical binding: Hydrogen heats of solution in the 3d metals *Physical Review B* **26** 2875–85

[82] Nørskov J K 1983 Erratum: Covalent effects in the effective-medium theory of chemical binding: Hydrogen heats of solution in the 3d metals *Physical Review B* **28** 1138

[83] Daw M S and Baskes M I 1983 Semiempirical, quantum mechanical calculation of hydrogen embrittlement in metals *Physical Review Letters* **50** 1285–8

[84] Floudas C A and Gounaris C E 2009 A review of recent advances in global optimization *Journal of Global Optimization* **45** 3–38

[85] Hartke B 2011 Global optimization *Wiley Interdisciplinary Reviews: Computational Molecular Science* **1** 879–87

[86] Kirkpatrick S, Gelatt C D and Vecchi M P 1983 Optimization by Simulated Annealing *Science* **220** 671–80

[87] Storn R and Price K 1997 Differential evolutionâĂŞa simple and efficient heuristic for global optimization over continuous spaces *Journal of global optimization* 341–59

[88] Metropolis N, Rosenbluth A W, Rosenbluth M N, Teller A H and Teller E 1953 Equation of State Calculations by Fast Computing Machines *The Journal of Chemical Physics* **21** 1087–92

[89] Powell M J D 1965 A Method for Minimizing a Sum of Squares of Non-Linear Functions Without Calculating Derivatives *The Computer Journal* **7** 303–7

[90] Press W H, Teukolsky S A, Vetterling W T and Flannery B P 1992 *Numerical Recipes in C: The Art of Scientific Computing* Academic Press, Cambridge 2nd edition

[91] Heyden A, Peters B, Bell A T and Keil F J 2005 Comprehensive DFT study of nitrous oxide decomposition over Fe-ZSM-5 *Journal of Physical Chemistry B*

[92] Henkelman G and Jónsson H 1999 A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives *The Journal of Chemical Physics* **111** 7010–22

[93] Béland L K, Brommer P, El-Mellouhi F, Joly J F and Mousseau N 2011 Kinetic activation-relaxation technique *Physical Review E* **84** 046704

[94] N'Tsouaglo G K, Béland L K, Joly J F, Brommer P, Mousseau N and Pochet P 2015 Probing Potential Energy Surface Exploration Strategies for Complex Systems *Journal of Chemical Theory and Computation* **11** 1970–7

[95] Pickard C J and Needs R J 2011 Ab initio random structure searching *Journal of Physics: Condensed Matter* **23** 053201

[96] Drummond N D, Monserrat B, Lloyd-Williams J H, Ríos P L, Pickard C J and Needs R J 2015 Quantum Monte Carlo study of the phase diagram of solid molecular hydrogen at extreme pressures *Nature Communications* **6** 7794

[97] Robert C and Casella G 2011 A Short History of Markov Chain Monte Carlo: Subjective Recollections from Incomplete Data *Statistical Science* **26** 102–15

[98] Bishop C M 2007 Pattern Recognition and Machine Learning *Journal of Electronic Imaging* **16** 049901

[99] Kennedy M C and O'Hagan A 2001 Bayesian calibration of computer models *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **63** 425–64

[100] Owhadi H, Scovel C, Sullivan T J, McKerns M and Ortiz M 2013 Optimal Uncertainty Quantification *SIAM Review* **55** 271–345

[101] Mahoney M W and Jorgensen W L 2000 A five-site model for liquid water and the reproduction of the density anomaly by rigid, nonpolarizable potential functions *The Journal of Chemical Physics* **112** 8910–22

[102] Olson B S and Shehu A 2013 Rapid sampling of local minima in protein energy surface and effective reduction through a multi-objective filter *Proteome Science* **11** S12

[103] Hartmann C and Schütte C 2012 Efficient rare event simulation by optimal nonequilibrium forcing *Journal of Statistical Mechanics: Theory and Experiment* **2012** P11004

[104] Bolhuis P G, Chandler D, Dellago C and Geissler P L 2002 Transition Path Sampling: Throwing Ropes Over Rough Mountain Passes, in the Dark *Annual Review of Physical Chemistry* **53** 291–318

[105] Peters B 2015 Common Features of Extraordinary Rate Theories *The Journal of Physical Chemistry B* **119** 6349–56

[106] Farid B and Godby R W 1991 Cohesive energies of crystals *Physical Review B* **43** 14248–50

[107] Zen A, Roch L M, Cox S J, Hu X L, Sorella S, Alfè D and Michaelides A 2016 Toward Accurate Adsorption Energetics on Clay Surfaces *The Journal of Physical Chemistry C* **120** 26402–13

[108] Peters B 2006 Using the histogram test to quantify reaction coordinate error *The Journal of Chemical Physics* **125** 241101

[109] Hastings W K 1970 Monte carlo sampling methods using Markov chains and their applications *Biometrika*

[110] Roberts G O, Gelman A and Gilks W R 1997 Weak convergence and optimal scaling of random walk Metropolis algorithms *The Annals of Applied Probability* **7** 110–20

[111] Plimpton S 1995 Fast Parallel Algorithms for Short-Range Molecular Dynamics *Journal of Computational Physics* **117** 1–19

[112] Bahn S and Jacobsen K 2002 An object-oriented scripting interface to a legacy electronic structure code *Computing in Science & Engineering* **4** 56–66

[113] Hjorth Larsen A, Jørgen Mortensen J, Blomqvist J, Castelli I E, Christensen R, Dułak M, Friis J, Groves M N, Hammer B, Hargus C, Hermes E D, Jennings P C, Bjerre Jensen P, Kermode J, Kitchin J R, Leonhard Kolsbjerg E, Kubal J, Kaasbjerg K, Lysgaard S, Bergmann Maronsson J, Maxson T, Olsen T, Pastewka L, Peterson A, Rostgaard C, Schiøtz J, Schütt O, Strange M, Thygesen K S, Vegge T, Vilhelmsen L, Walter M, Zeng Z and Jacobsen K W 2017 The atomic simulation environment -

a Python library for working with atoms *Journal of Physics: Condensed Matter* **29** 273002

[114] Hildebrand F E and Abeyaratne R 2008 An atomistic investigation of the kinetics of detwinning *Journal of the Mechanics and Physics of Solids* **56** 1296–319

[115] Milstein F 1973 Applicability of exponentially attractive and repulsive interactomic potential functions in the description of cubic crystals *Journal of Applied Physics* **44** 3825–32

[116] Doye J P K and Wales D J 1997 Structural consequences of the range of the interatomic potential A menagerie of clusters *Journal of the Chemical Society, Faraday Transactions* **93** 4233–43

[117] Daw M S and Baskes M I 1984 Embedded-atom method: Derivation and application to impurities, surfaces, and other defects in metals *Phys. Rev. B* **29**

[118] Banerjea A and Smith J R 1988 Origins of the universal binding-energy relation *Physical Review B* **37** 6632–45

[119] Ackland G J, Tichy G, Vitek V and Finnis M W 1987 Simple N -body potentials for the noble metals and nickel *Philosophical Magazine A* **56** 735–56

[120] Rassoulinejad-Mousavi S M, Mao Y and Zhang Y 2016 Evaluation of copper, aluminum, and nickel interatomic potentials on predicting the elastic properties *Journal of Applied Physics* **119** 244304

[121] Mishin Y, Farkas D, Mehl M J and Papaconstantopoulos D A 1999 Interatomic potentials for monoatomic metals from experimental data and ab initio calculations *Physical Review B* **59** 3393–407

[122] Davey W P 1925 Precision Measurements of the Lattice Constants of Twelve Common Metals *Physical Review* **25** 753–61

[123] Wang L P and Van Voorhis T 2010 Communication: Hybrid ensembles for improved force matching *The Journal of Chemical Physics* **133** 231101

[124] Clark S J, Segall M D, Pickard C J, Hasnip P J, Probert M I, Refson K and Payne M C 2005 First principles methods using CASTEP *Zeitschrift fur Kristallographie*

[125] van de Waal B W 1991 Can the Lennard-Jones solid be expected to be fcc? *Physical Review Letters* **67** 3263–6

[126] Kihara T and Koba S 1952 Crystal Structures and Intermolecular Forces of Rare Gases *Journal of the Physical Society of Japan* **7** 348–54

[127] Angelo J E, Moody N R and Baskes M I 1995 Trapping of hydrogen to lattice defects in nickel *Modelling and Simulation in Materials Science and Engineering* **3** 289–307

[128] Chantasiriwan S and Milstein F 1996 Higher-order elasticity of cubic metals in the embedded-atom method *Physical Review B* **53** 14080–8

[129] Tadmor E B, Elliott R S, Phillpot S R and Sinnott S B 2013 NSF cyberinfrastructures: A new paradigm for advancing materials simulation *Current Opinion in Solid State and Materials Science* **17** 298–304

[130] Choudhary K, Congo F Y P, Liang T, Becker C, Hennig R G and Tavazza F 2017 Evaluation and comparison of classical interatomic potentials through a user-friendly interactive web-interface *Scientific Data* **4** 160125

[131] Girolami M and Calderhead B 2011 Riemann manifold Langevin and Hamiltonian Monte Carlo methods *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73** 123–214

[132] Goodman J and Weare J 2010 Ensemble samplers with affine invariance *Communications in Applied Mathematics and Computational Science* **5** 65–80

[133] Neighbours J R, Bratten F W and Smith C S 1952 The Elastic Constants of Nickel *Journal of Applied Physics* **23** 389–93

[134] Hidnert P 1957 Thermal expansion of some nickel alloys *Journal of Research of the National Bureau of Standards* **58** 89

[135] Zhou X W, Heo T W, Wood B C, Stavila V, Kang S and Allendorf M D 2017 Finite-temperature behavior of PdHx elastic constants computed by direct molecular dynamics *MRS Advances* **2** 3341–6

[136] Torrie G and Valleau J 1977 Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling *Journal of Computational Physics* **23** 187–99

[137] Kästner J and Thiel W 2006 Analysis of the statistical error in umbrella sampling simulations by umbrella integration *The Journal of Chemical Physics* **124** 234106

[138] Dickson A, Warmflash A and Dinner A R 2009 Nonequilibrium umbrella sampling in spaces of many order parameters *The Journal of Chemical Physics* **130** 074104

[139] Ceriotti M, Tribello G A and Parrinello M 2011 Simplifying the representation of complex free-energy landscapes using sketch-map *Proceedings of the National Academy of Sciences* **108** 13023–8