Computational Methods for Inferring the Structures of Amorphous Materials and Understanding Ionic Diffusivities

by

Xander A. Gouws

A thesis

presented to the University of Waterloo

in fulfillment of the

thesis requirement for the degree of

Master of Science

in

Chemistry

Waterloo, Ontario, Canada, 2025

© Xander A. Gouws 2025

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The simulation of solid-state electrolytes (SSEs) has allows researchers to directly observe the migration paths of lithium ions, and so has played a pivotal role in elucidating the mechanisms of ionic conduction. Performing these simulations requires three steps: Structure determination, simulation, and analysis. Here, we have developed and tested new computational methods to address key challenges in two of these steps. First, we develop a gradient-based optimization method for determining the structures of amorphous materials from total scattering data. Unlike traditional reverse Monte Carlo approaches that rely on random atomic movements and suffer from slow convergence, our gradient-based method moves atoms to directly minimize the χ^2 goodness-of-fit and potential energy. Our approach was tested on amorphous silicon and a nickel–niobium metallic glass. Convergence was achieved in on the order of 5000 steps, which is approximately one hundred times faster than existing hybrid Monte Carlo methods. Then, we introduce a method for detecting ion hopping events in SSEs without prior knowledge of site locations. This may be useful when simulating i) new materials, for which the positions of all lithium occupancy sites may be unknown, ii) structural changes (e.g. doping) that introduce local strains that shift site positions, or iii) amorphous materials, where lithium sites may be unknown prior to simulation. Testing our method on Li_6PS_5Cl and its BH_4-doped variant, we recover the cage-forming nature of lithium sites in argyrodite structures, and find that the correlation factor for hops between cages is greater than one, indicating a forward-bias for intercage hops.

Acknowledgements

First and foremost, I would like to express my sincere appreciation for my supervisor, Professor Conrard Giresse Tetsassi Feugmo, for proposing such interesting directions for study, for providing me guidance where needed, and for affording me the independence to explore my ideas and intuitions. I would also like to thank the members of my supervisory committee, Professors Pierre-Nicholas Roy and Linda Nazar, for lending me their time, expertise, and encouragement. I would also like to thank Professor Sjoerd Roorda for providing us with X-ray scattering data for amorphous silicon and assisting us in its processing.

I am also indebted to my many friends and colleagues for the countless conversations that helped shape this thesis. I owe special thanks to Nithin Aaron, Jacob Chaussé, Megan Dawson, Mohid Farooqi, Agnes Katai, Toby McConville, Azharuddin Mohammed, Anupom Roy, Arjun Varma, and Amelie Zahoruk.

And finally, I would like to thank my partner, Erika Kember, and my family. This journey would not have been possible without their unwavering support and belief in me, as well as the many sacrifices they have made to support my education.

Contents

List of Figures **x**

List of Abbreviations **xii**

1 Introduction **1**
1.1 Solid-state electrolytes . . . 1
1.2 The computational study of solid-state electrolytes . . . 3
1.2.1 Structure determination . . . 3
1.2.2 Simulation . . . 5
1.2.3 Analysis . . . 8
1.3 Research objectives . . . 10
1.3.1 Gradient-based structure determination . . . 10
1.3.2 Automated analysis of hopping in solid-state electrolytes . . . 11

2 Gradient-Based Structure Determination **14**
2.1 Differentiable evaluation of pair distribution functions . . . 14
2.2 Automatic differentiation . . . 17
2.3 Evaluation of distances . . . 18
2.4 Calculation details . . . 20
2.4.1 Amorphous silicon . . . 21
2.4.2 Nickel–niobium metallic glass . . . 22
2.5 Results and discussion . . . 24
2.5.1 Benchmarking . . . 24
2.5.2 Amorphous silicon . . . 25
2.5.3 Nickel–niobium metallic glass . . . 34

3 Automated Analysis of Hopping in Solid-State Electrolytes **42**
3.1 General approach . . . 42
3.2 Calculation details . . . 45
3.3 Results and discussion . . . 47
3.3.1 Analysis by mean squared displacement . . . 47
3.3.2 Analysis by proposed methods . . . 49

4 Conclusions **61**
4.1 Gradient-based structure determination 62
4.2 Automated analysis of hopping in solid-state electrolytes 65

Appendices **69**

A Appendix for Gradient-Based Structure Determination **69**
A.1 Calculating total scattering spectra 69
A.1.1 Neutron scattering 70
A.1.2 X-ray scattering 72
A.2 Using the `gbsd` package 73

B Appendix for the Analysis of Hopping in Solid-State Electrolytes **76**
B.1 Diffusion as a macroscopic phenomenon 76
B.1.1 The diffusion equation 76
B.1.2 Diffusivity 77
B.1.3 Solving the diffusion equation 78
B.2 Diffusion as a microscopic phenomenon 80
B.2.1 The theory of random walks 80
B.2.2 Relating random walks to diffusivity 82

Bibliography **83**

List of Figures

1.1 Unit cell of the Li_6PS_5X (X = Cl, Br, I) cubic polymorph. 12
1.2 Diagram showing the types of hops between lithium sites in Li_6PS_5X (X = Cl, Br, I). 13

2.1 Diagram showing a simplified computation graph for our method. 18
2.2 Diagram displaying the minimum image convention. 19
2.3 Times taken for a single central processing unit (CPU) and a single graphics processing unit (GPU) to compute the pair distribution function for different numbers of randomly placed atoms. 24
2.4 Pair distribution functions for melt-quench amorphous silicon, calculated with kernel density estimation. 25
2.5 Pair distribution functions for amorphous silicon and their uncertainties. . . 27
2.6 Per-atom potential energy and χ^2-loss over the course of optimizing $V + \lambda\chi^2$ for amorphous silicon, starting from a strained crystal structure. 29
2.7 Pair distribution functions calculated for our optimized amorphous silicon structures, starting from a strained crystal structure. 30
2.8 Distributions of (a) bond length, (b) bond angle, (c) dihedral angle, and (d) coordination number seen for our optimized amorphous silicon structures, starting from a strained crystal structure. 31
2.9 Per-atom potential energy and χ^2-loss over the course of optimizing $V + \lambda\chi^2$ for amorphous silicon, starting from a melt-quench structure. 33
2.10 Pair distribution functions calculated for our optimized amorphous silicon structures, starting from a melt-quench structure. 34
2.11 Distributions of (a) bond length, (b) bond angle, (c) dihedral angle, and (d) coordination number seen for our optimized amorphous silicon structures, starting from a melt-quench structure. 35
2.12 Average partial pair distribution functions and X-ray total pair distribution function, and the standard error in the X-ray total pair distribution function, calculated for one hundred $Ni_{62}Nb_{38}$ structures obtained by molecular dynamics (MD). 36
2.13 Per-atom potential energy and χ^2-loss over the course of optimizing $V + \lambda\chi^2$ for amorphous $Ni_{62}Nb_{38}$, starting from a random structure. 37

2.14 X-ray total pair distribution functions calculated for our optimized amorphous $Ni_{62}Nb_{38}$ structures, starting from a random structure. 38
2.15 (a)–(c) Partial pair distribution functions and (d) average coordination numbers, calculated for our optimized amorphous $Ni_{62}Nb_{38}$ structures, starting from a random structure. 39
2.16 Bond angle distributions calculated for our optimized amorphous $Ni_{62}Nb_{38}$ structures, starting from a random structure. 41

3.1 Mean-squared displacement (MSD) over time for lithium atoms in (a) Li_6PS_5Cl and (b) $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ at 400 K, 500 K and 600 K. 48
3.2 Diffusivity of lithium atoms in Li_6PS_5Cl and $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ as functions of reciprocal temperature. 48
3.3 Ridge plots showing the effect of the threshold parameter on the observed distributions of (a) hop lengths, (b) hop rates, (c) cosine angles between consecutive hops, and (d) the radius of gyration of sites for Li_6PS_5Cl at 600 K. . 51
3.4 Ridge plots showing the effect of the threshold parameter on the observed distributions of (a) hop lengths, (b) hop rates, (c) cosine angles between consecutive hops, and (d) the radius of gyration of sites for Li_6PS_5Cl at 600 K. . 52
3.5 Cross-sections of the trajectory of a single lithium atom in Li_6PS_5Cl at 600 K. 53
3.6 (a)–(c) Coordinates of a single lithium atom in Li_6PS_5Cl at 600 K and (d) its distance to the site centroid shown over time. 54
3.7 Mean and standard error for the (a) hop length, (b) hop rate, (c) correlation factor, and (d) diffusivity for lithium atoms in Li_6PS_5Cl and $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ at different temperatures, as extracted using a threshold distance of 5.5 Å. . 55
3.8 Ridge plots showing the effect of the threshold parameter on the observed distributions of (a) hop lengths, (b) hop rates, (c) cosine angles between consecutive hops, and (d) the radius of gyration of sites for Li_6PS_5Cl at 600 K. . 57
3.9 Cross-sections of the trajectory of a single lithium atom in $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ at 400 K. 58
3.10 (a)–(c) Coordinates of a single lithium atom in $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ at 400 K and (d) its distance to the site centroid shown over time. 59
3.11 Mean and standard error for the (a) hop length, (b) hop rate, (c) correlation factor, and (d) diffusivity for lithium atoms in Li_6PS_5Cl and $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ at different temperatures, as extracted using a threshold distance of 3.25 Å. . 60

List of Abbreviations

AIMD *Ab initio* molecular dynamics 7

ASE Atomic Simulation Environment 11

ASSLMB All-solid-state lithium-metal battery 1

CPU Central processing unit 20

FIRE Fast inertial relaxation engine 22

FWHM Full width at half maximum 16

GPU Graphics processing unit 17

KDE Kernel density estimation 15

LAMMPS Large-scale Atomic/Molecular Massively Parallel Simulator 21

MD Molecular dynamics 6

MISE Mean integrated square error 16

MLP Machine-learned potential 7

MSD Mean-squared displacement 8

NEB Nudged elastic band 5

NMR Nuclear magnetic resonance 65

RMC Reverse Monte Carlo 4

SSE Solid-state electrolyte 1

XAS X-ray absorption spectroscopy 69

Chapter 1

Introduction

1.1 Solid-state electrolytes

The growing demand for renewable energy has led to an urgent need for low-cost, high-performance, and environmentally-friendly energy storage [1]. Li-ion batteries have become the energy storage device of choice for a variety of applications including consumer electronics, electric vehicles, and grid-scale energy storage [2]. However, current Li-ion batteries suffer from safety risks associated with their liquid electrolyte and relatively low energy density. All-solid-state lithium-metal batteries (ASSLMBs) can potentially mitigate these issues by replacing the liquid electrolyte of Li-ion batteries with a solid electrolyte [3]. But to enable the widespread commercialization of ASSLMBs, they must be made both effective and economical.

One critical component of an ASSLMB is its solid-state electrolyte (SSE). An ideal SSE must fulfill many requirements. Firstly, it must allow lithium ions to move freely within it, i.e., it must have a high ionic conductivity. SSEs must also fulfill a number of other requirements, including being poor conductors of electrons, stable under operating conditions, compatible with other materials in the battery, and economical to produce.

The chemical space for possible SSEs is very broad; multiple chemistries have been proposed in multiple states. Yet, the search for an optimal SSE is ongoing, as none of these proposed materials has perfect performance in every category. For instance, sulphur-based SSEs typically have higher ionic conductivities than oxygen-based SSEs, but are also less stable and less compatible with lithium metal [3]. SSEs in the amorphous phase, in comparison to their crystalline counterparts, have been gaining increasing interest for their lack of resistive grain boundaries and potential for high interfacial contact [4].

The need for SSEs to be fast ionic conductors (with ionic conductivities on the order of at least $10\,\mathrm{mS\,cm^{-1}}$) is arguably the most constraining requirement, and will be the primary focus of this work. Solid-state ionic conductors are generally composed of mobile ions that hop between stable polyhedral lattice sites created by the other atoms in the material. Though many elements can be used as mobile ions, Li is of special interest to the application of ASSLMBs for its small size and monovalent nature [5]. As such, we are primarily concerned with Li-ionic conductors and their Li-ion conductivities $\sigma_{\mathrm{Li^+}}$.

Ionic conduction is ultimately driven by the mobile ions moving to minimize their electrochemical potential. In the absence of an electric field, this results in each mobile ion undergoing diffusion. For a mobile ion to move between sites, it must overcome an energy barrier with activation energy E_A [5]. So, the diffusivity of the mobile ions depends on temperature by an Arrhenius relation

$$D = D_0 \exp\left(-\frac{E_\mathrm{A}}{k_\mathrm{B}T}\right), \tag{1.1}$$

where D is the diffusivity, T is temperature, k_B is Boltzmann's constant, and D_0 is the pre-exponential factor. Then, the Nernst–Einstein relation [5, 6],

$$\sigma_{\mathrm{Li^+}} = \frac{nq^2 D}{k_\mathrm{B}T}, \tag{1.2}$$

can be used to relate the diffusivity to the ionic conductivity, where n is the number density of mobile ions and q is their ionic charge.

1.2 The computational study of solid-state electrolytes

Computational simulations have allowed us to directly observe the migration paths of lithium ions from first-principles, and so have played a pivotal role in elucidating the mechanisms underlying ionic conduction. These simulation techniques have also been able to predict, to good agreement with experiment, trends in the ionic conductivities of different materials [7]. In the remainder of this chapter, I will discuss the limitations faced at each step required in performing computational simulations of SSEs and propose methods to overcome these limitations. The remainder of this thesis then focuses on testing our proposed approaches.

1.2.1 Structure determination

Before the simulation of a material can be performed, its structure must be known. This poses relatively little difficulty for crystalline SSEs, which are composed of near-perfectly periodic unit cells. Amorphous SSEs, in contrast, have no long-range structural order. At first glance, seeking to determine the structure of an amorphous material may seem contradictory. However, while non-crystalline materials do not possess the long-range order found in crystals, they do maintain short-range order, manifesting as preferred coordination numbers, bond lengths, bond angles, and dihedral angles. And so, although a non-crystalline material will not have one single definite structure, certain atomic arrangements will certainly be more likely than others.

Nevertheless, determining the structures of amorphous materials is made challenging in several ways by their intrinsic structural disorder [8]. Computationally simulating these structures requires large simulation boxes with many atoms, to more accurately model their

long-range disorder. A related challenge is in accessing realistic timescales through simulations; experiments on amorphous materials can be on the timescale of minutes whereas molecular dynamics simulations are typically shorter than microseconds. Another barrier is interpreting the results of experimental characterization; information about preferred structure is only partially encoded in the structure factor $S(q)$ and pair distribution function $G(r)$, which can be measured with X-ray or neutron diffraction. Owing to the lack of long-range order in amorphous materials, their scattering spectra consist of broad, overlapping peaks. Importantly, however, these spectra can also be calculated computationally for a proposed structure [9], giving researchers a means to compare proposed structures to experimental data. It should be noted, though, that these spectra do not capture all structural information; therefore, a structure that matches the experimental data is not necessarily realistic. [10, 11].

A popular method used to find the structures of amorphous materials is reverse Monte Carlo (RMC) [12, 13]. This method modifies a guess structure with the goal of matching its calculated diffraction spectra to experimental diffraction data. This is done in an iterative stochastic process, inspired by Metropolis–Hastings Monte Carlo. It is typical to begin with a random configuration of atoms. A goodness-of-fit,

$$\chi^2 = \sum_i \frac{[\hat{y}(x_i) - y_i]^2}{\sigma_i^2}, \tag{1.3}$$

is determined, where y_i is a point in the measured diffraction spectra, σ_i is its uncertainty, and $\hat{y}_i$ is the corresponding point on the diffraction spectra calculated from the computational structure. One atom is then moved randomly. If this change improves the quality of the fit, the move is accepted. If it does not improve the quality of the fit, the move may be accepted with probability $\exp(-\Delta\chi^2/2)$, where $\Delta\chi^2$ is the change in the goodness-of-fit. If the move is not accepted, it is reverted. Note that, by accepting some moves that worsen the goodness-of-fit, RMC is able to escape local minima. After repeating many such moves, the structure

will ideally converge upon a configuration that matches experiment. Unfortunately, RMC as it is described above has two major issues. Firstly, the structures it generates are commonly too strongly disordered and may have unrealistically high energies [14]. Secondly, it can take days or weeks of computation time for a structure to converge—if it does at all [15].

The problem of energetically unrealistic configurations is solved by hybrid methods, which account for the system's potential energy during the optimization process. This is sometimes done by adding energy contributions to the calculation of χ^2 [16], in which case the energy is often estimated using simple classical force-fields (though this is still subject to the limitations of classical force-fields discussed below). Another approach is that of experimentally-constrained molecular relaxation and force-enhanced atomic refinement methods [14, 17], which repeatedly perform a number of RMC steps and structure-relaxation or molecular dynamics steps until convergence is reached. These methods have shown promising ability to generate structures that have realistic potential energies while matching experimental diffraction spectra [18, 19]. Other examples of hybrid methods are those that begin with a starting configuration created by some other method (e.g. melt-quench molecular dynamics or bespoke modelling) and then use RMC to refine the proposed structure towards experimental data [20].

1.2.2 Simulation

One technique that has been used to quantify ionic conductivity is the nudged elastic band (NEB) method. These calculations directly obtain the migration path and activation energy E_A associated with a mobile ion hopping between two equilibrium sites. The complete diffusivity and ionic conductivity can be extracted from kinetic Monte Carlo simulations with rates based on the results of NEB calculations. A major limitation of the NEB method is that it requires a well-defined migration path and so is not suitable for materials with disordered Li-ion sublattices [7].

One resolution to these issues is to use a molecular dynamics (MD) approach. In MD simulations, the positions of atoms are treated using classical mechanics. As such, each atomic nucleus i has a defined position $\mathbf{r}_i(t)$ and velocity $\mathbf{v}_i(t)$ at each discrete time t. The acceleration of each nucleus can be determined using Newton's second law,

$$\mathbf{a}_i(t) = \mathbf{F}_i(\mathbf{r}_1, ..., \mathbf{r}_N)/m_i. \tag{1.4}$$

Then, recalling that position is the integral of velocity and velocity is the integral of acceleration, the state of the system can be propagated by a small time-step Δt using a numerical integration technique (e.g., velocity Verlet integration). Assuming that all forces are conservative, the force on each atom can be calculated by taking the negative gradient of the system's potential energy,

$$\mathbf{F}_i(\mathbf{r}_1, ..., \mathbf{r}_N) = -\nabla_i V(\mathbf{r}_1, ..., \mathbf{r}_N). \tag{1.5}$$

Propagating the system as described above will conserve the system's total number of particles, volume, and energy (though some errors may be introduced by numerical integration), and so is consistent with the microcanonical NVE ensemble. The kinetic temperature can be calculated by relating it to the average kinetic energy;

$$\left\langle \frac{1}{2} m v^2 \right\rangle = \frac{3}{2} k_{\mathrm{B}} T. \tag{1.6}$$

Furthermore, kinetic theory tells us that, at equilibrium, the velocities will follow a Normal distribution with a mean of zero and standard deviation dependent on the temperature. However, in the case of melt-quench simulations, we instead wish to control the temperature directly. This can be accomplished with a thermostat, which rescales the velocities of atoms over some characteristic timescale. A simple example is the Andersen thermostat [21–23],

which resamples the velocities of particles from the desired velocity distribution. This requires introducing a timescale parameter that determines how frequently (on average) the velocity of each particle is resampled. While the Andersen thermostat is conceptually simple, it is sometimes not favourable that the velocities are discontinuously resampled. Alternatives are the Langevin and Nosé–Hoover thermostats [22, 23], which modify the velocities continuously. Additionally, we sometimes wish to control the pressure of the system. This can be done with barostats, which function similarly to thermostats, but modify the volume of the simulation box to control the pressure.

An important consideration is the potential (or force-field) used. For MD simulations of SSEs, it is most common to use *ab initio* molecular dynamics (AIMD), wherein forces are calculated using density functional theory [24]. While this is the most accurate method practically available, AIMD is quite computationally expensive, limiting users to both smaller system sizes and shorter simulations. These restrictions are especially detrimental when studying amorphous SSEs, as large simulation boxes are required to accurately model their lack of periodic structure and longer timescales are needed to replicate experiments [8]. Classical force-fields are an alternative that fit simple, empirical models to AIMD or experimental data. They are typically much faster and typically scale well for large systems, but are also less accurate. Additionally, they must be uniquely parameterized for each system, which adds another level of complexity to using these force-fields. While classical force-fields have been greatly effective in other applications, they have seen limited success in the modelling of SSEs [25, 26]. Universal machine-learned potentials (MLPs) [27–30], on the other hand, can maintain high levels of accuracy, can easily be applied to different systems, and remain much faster than AIMD. As such, these show great promise for modelling SSEs [26].

1.2.3 Analysis

The direct relationship between ionic conductivity and diffusivity greatly simplifies the computational study of solid-state electrolytes. Rather than simulating the effect of an electric field (on both ions and the electron density), we can simply simulate the ionic conductor in the absence of an electric field and observe the diffusion of lithium ions. An elementary result from the statistics of random walks is that the diffusivity of atoms in a substance is related to their mean-squared displacement (MSD) by

$$D = \lim_{t\to\infty} \frac{\mathrm{MSD}(t)}{2\,d\,t}, \tag{1.7}$$

where d is the dimension and t is time. The uncertainty in the measured diffusivity is directly related to the number of observed diffusion events [31]. And so, simulations are typically run at very high temperatures (500 K to 2000 K) and extrapolated to room temperature using Eq. (1.1). Then, this can be used to calculate the ionic conductivity using the Nernst–Einstein relation, Eq. (1.2).

The trajectories of lithium atoms have also been used to identify individual hopping events, from which statistics such as the mean-squared hop length, hopping rate, and correlation factor can be determined [32]. These statistics also play an important role in revealing the mechanisms behind diffusion. A particle undergoing a random walk with steps $\{\Delta\mathbf{r}_i\}$ and residence times $\{\Delta t_i\}$ has diffusivity

$$D = \frac{1}{2d} f \Gamma a^2, \tag{1.8}$$

where d is the dimension,

$$f = \frac{1 + \mathbb{E}[\cos\theta_{i,i+1}]}{1 - \mathbb{E}[\cos\theta_{i,i+1}]} \tag{1.9}$$

is the correlation factor or Haven ratio [33, 34] and $\theta_{i,i+1}$ is the angle between $\Delta\mathbf{r}_i$ and $\Delta\mathbf{r}_{i+1}$,

$$\Gamma = 1/\,\mathbb{E}[\Delta t_i] \tag{1.10}$$

is hopping rate, and

$$a^2 = \mathbb{E}[\Delta\mathbf{r}_i^{\mathrm{T}}\Delta\mathbf{r}_i] \tag{1.11}$$

is the mean-squared hop length. While f and a^2 do not vary significantly with temperature,

$$\Gamma = \Gamma_0 \exp\left(-\frac{E_{\mathrm{A}}}{k_{\mathrm{B}}T}\right) \tag{1.12}$$

is the origin of Arrhenius relationship in the diffusivity (Eq. 1.1), where Γ_0 is commonly called the attempt frequency. As such, the base diffusivity can be written as

$$D_0 = \frac{1}{2d} f\Gamma_0 a^2. \tag{1.13}$$

So, understanding the factors that affect the correlation factor, attempt frequency, and mean-squared hop length can help us understand changes in the base diffusivity.

A key challenge in this form of analysis is that the locations of hopping sites must typically be known beforehand. One solution has been to perform energy minimizations on structures every few timesteps to identify stable sites [35–37]. However, this approach can be prohibitively computationally expensive, particularly for larger structures where the number of diffusing particles increases and the total hopping rate is increased and so more frequent minimizations are required.

1.3 Research objectives

Computational methods have played an important role in elucidating the mechanisms underlying ionic conduction in SSEs. However, several challenges are faced when simulating these materials—particularly in the determination of amorphous structures and analysis of diffusion data. The goal of this thesis is to develop new methods towards overcoming these challenges. In this section, we discuss the specific issues we aim to tackle, as well as the proposed methodology used to do so.

1.3.1 Gradient-based structure determination

A major limitation of RMC is its slow convergence, owing to its stochastic nature. In Chapter 2, we develop a new approach that addresses the convergence issues inherent to RMC optimization. This is done by adopting a gradient-based approach to optimizing atomic configurations. Like RMC, our method uses a χ^2 goodness-of-fit to experimental diffraction data to find plausible structures. However, rather than moving atoms randomly one at a time, our method modifies structures using the gradient with respect to atomic positions $\nabla\chi^2$. This allows for the concerted movement of all atoms in a way that directly improves the goodness-of-fit. A potential concern with this approach could be its susceptibility to becoming trapped in local optima; however, it is well established that carefully designed gradient-based optimization schemes are able to escape local wells [38, 39].

Differentiable approaches to amorphous structure determination have been proposed previously, but their adoption has been limited due to their reliance on analytically calculated derivatives, which restricts their applicability primarily to structure factors [40]. Our approach overcomes this limitation through the use of automatic differentiation [41]. For this to be possible, the pair distribution function must be made differentiable. And so, rather than using the traditional histogram-based approach, we use kernel density estimation to construct a smooth histogram, which allows us to calculate derivatives in terms of the pair

distribution function.

We demonstrate the potential of our method by applying it to two different non-crystalline systems: Amorphous silicon (a-Si) and a nickel–niobium metallic glass ($Ni_{62}Nb_{38}$). These materials have been rigorously studied previously using different methods [42–45], and each presents its own unique challenges, making them ideal systems on which to test our approach. Our method is available as a Python package that integrates with the Atomic Simulation Environment (ASE) [46]. This enables hybrid approaches, allowing users to combine our differentiable structure refinement with quantum chemistry packages, classical force fields, and universal machine-learned potentials.

1.3.2 Automated analysis of hopping in solid-state electrolytes

A limitation in the analysis of simulation data was that advanced analysis (beyond just the calculation of the diffusion coefficient) typically requires *a priori* knowledge of all lithium sites. In Chapter 3, we work towards developing a method to detect hopping events that does not rely on any prior information. This may be useful for three key reasons. Firstly, when simulating new materials, the positions of all potential lithium occupancy sites may be unknown. Secondly, this approach ensures robust detection of hopping events despite structural changes; for instance, material doping typically introduces local strains that can shift lithium site locations. Finally, this method could enable hopping event detection in amorphous materials, where lithium sites may be fundamentally unknowable ahead of time.

The detection method we develop relies solely on lithium atom trajectories from molecular dynamics. We detect hops by identifying when lithium atoms move away from a rolling average of their position. We then test this approach using the argyrodite materials Li_6PS_5Cl and $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$. These systems provide an appropriate testing ground for our method because diffusion in argyrodite structures is both well-characterized and scientifically relevant.

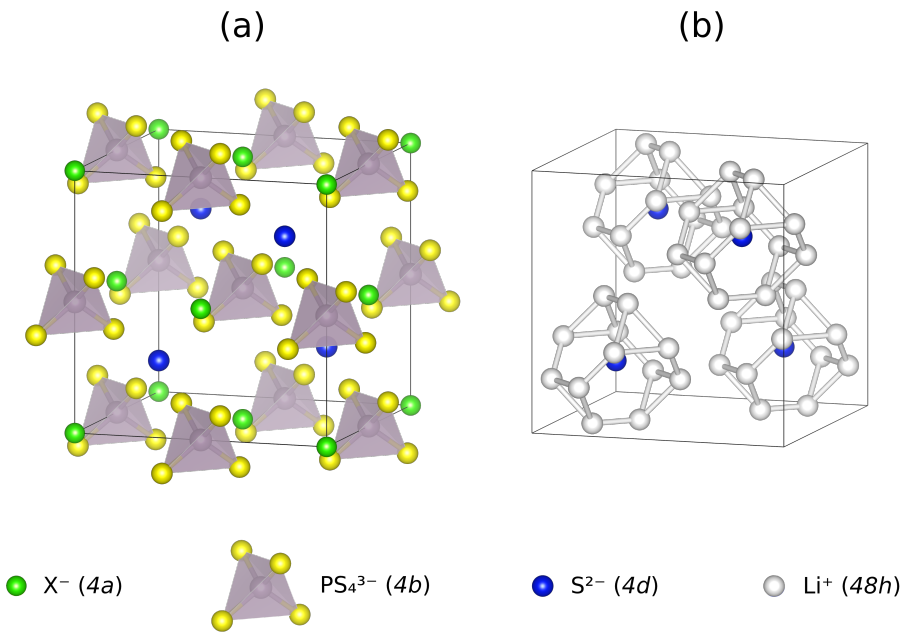


Figure 1.1: Unit cell of the Li_6PS_5X (X = Cl, Br, I) cubic polymorph. Shown separately are (a) the X^-, $PS_4{}^{3-}$ and S^{2-} anions, as well as (b) the Li^+ site cage network surrounding each free S^{2-} ion. These figures were prepared in VESTA [54] using the structure from Ref. [55].

At room temperature, Li_6PS_5X (X = Cl, Br, I) has a cubic crystal structure with the noncentrosymmetric space group $F\bar{4}3m$ [5, 47, 48], shown in Fig. 1.1. The X^- anions are distributed with an FCC lattice (Wyckoff position 4*a*), and the $PS_4{}^{3-}$ ions fill their octahedral voids (4*b*). Several free S^{2-} ions fill half of the tetrahedral voids (4*d*); however, disorder between the X^- and S^{2-} ions occurs commonly, and has been shown to contribute to decreased potential energy, increased entropy, and increased ion diffusivity [49–51]. The Li^+ ions are found in sites that form cages around each free S^{2-} ion (Fig. 1.1b). Li^+ ions are able to hop within cages (by intracage and doublet hops) and between cages (by intercage hops), shown in Fig. 1.2. Each cage contains twelve 48*h* sites (each approximately 50% occupied) and six 24*g* sites that mediate the doublet jump. Generally, a fully connected cage network is required for good lithium conductivity, with the overall jump rate being determined by its rate-limiting step [52, 53].

We test the doped structure $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ in addition to Li_6PS_5Cl because recent

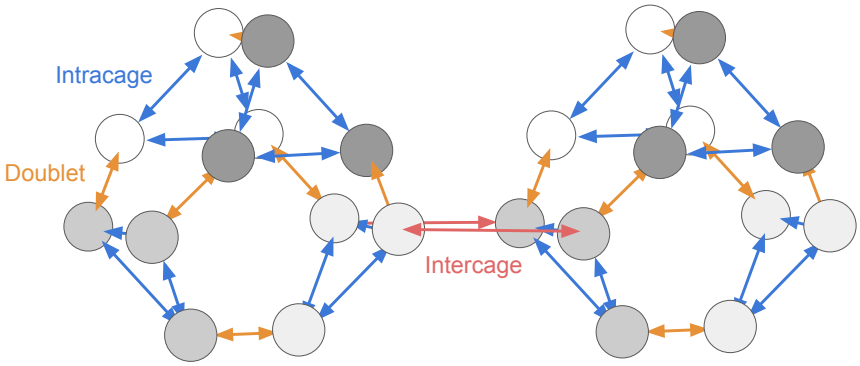


Figure 1.2: Diagram showing the types of hops between lithium sites in Li_6PS_5X (X = Cl, Br, I). These cages are composed of partially occupied Li^+ 48*h* sites, centred around S^{2-} 4*d* sites.

studies demonstrate improved ionic conductivity compared to the undoped material. It has been found that $Li_6PS_5Cl_{1-x}(BH_4)_x$ maintains an argyrodite structure for $x < 0.3$ and achieves optimal ionic conductivity for $x = 0.1$ [56, 57]. While this has been attributed to a reduced activation energy, their data may suggest that the increase in ionic conductivity more likely results from changes in base diffusivity. Examining individual hopping events could clarify the underlying mechanism—specifically, whether the enhancement stems from changes in attempt frequency, correlation factor, or hopping length.

Chapter 2

Gradient-Based Structure Determination

The goal of our approach is to use gradient-based optimization schemes rather than a Monte Carlo approach. There are two hurdles that must be overcome; firstly, the calculation of χ^2 must be performed in such a way that it is differentiable with respect to atomic coordinates, and secondly, its gradient should be calculated in a way that is both efficient and generalizable to different systems and forms of experimental data. In this section, I discuss the approaches taken in designing our method. In the following section, I discuss the procedures used to evaluate this method and its limitations.

2.1 Differentiable evaluation of pair distribution functions

It is clear that hybrid methods of structure determination—which incorporate both spectroscopic data and external chemical knowledge—are superior to those that only rely on experimental results. To this end, we create an effective potential

$$V_{\text{eff}}(\mathbf{R}) = V(\mathbf{R}) + \lambda\chi^2(\mathbf{R}), \tag{2.1}$$

where $\mathbf{R}$ is the collection of atomic coordinates, $V(\mathbf{R})$ is the potential energy of the system, as calculated by any method (so long as it is differentiable), $\chi^2(\mathbf{R})$ is the goodness-of-fit, and λ is a hyperparameter that controls the importance of χ^2 compared to V. Minimization is then performed using the effective force

$$\mathbf{F}_{\text{eff}}(\mathbf{R}) = -\nabla V_{\text{eff}}(\mathbf{R}) = \mathbf{F}(\mathbf{R}) - \lambda \nabla \chi^2(\mathbf{R}). \tag{2.2}$$

By including both V and χ^2 in our optimization target, we incentivize the optimizer to find structures that are both consistent with experimental data and energetically realistic.

The goodness-of-fit is calculated by

$$\chi^2(\mathbf{R}) = \sum_{i=1}^{N_{\text{data}}} \frac{[\hat{y}(x_i \mid \mathbf{R}) - y_i]^2}{\sigma_i^2}, \tag{2.3}$$

for some experimental spectrum with uncertainties $\{(x_i, y_i, \sigma_i)\}_{i=1}^{N_{\text{data}}}$. The estimate of the signal value $\hat{y}(x_i \mid \mathbf{R})$ is calculated at the corresponding x_i given the structure coordinates $\mathbf{R}$.

As is shown in Section A.1, the calculation of scattering spectra is done in terms of partial pair distribution functions $g_{\alpha\beta}(r)$, which requires a histogram $n_{\alpha\beta}(r)$ that counts the number of particles of species β at a distance between r and $r+\Delta r$ from a particle of type α (Eqs. A.3, A.4, A.16). Unfortunately, this results in a pair distribution function that is not differentiable with respect to the atomic positions; in general, an infinitesimal movement of an atom will result in either no change to the histogram or a discrete change, wherein the derivative is zero or undefined, respectively.

To alleviate this issue, we use kernel density estimation (KDE). Rather than having discrete

bins of width Δr, we evaluate

$$n_{\alpha\beta}(r) = \frac{1}{|\mathcal{D}_{\alpha\beta}|\sqrt{2\pi\sigma_{\mathrm{k}}^2}} \sum_{d \in \mathcal{D}_{\alpha\beta}} \exp\left[-\frac{(r-d)^2}{2\sigma_{\mathrm{k}}^2}\right], \tag{2.4}$$

where $\mathcal{D}_{\alpha\beta}$ is the set of observed distances and between atoms of species α and β, and $|\mathcal{D}_{\alpha\beta}|$ is its cardinality. The choice of a Gaussian kernel is arbitrary but natural, as Gaussian functions are simple to work with and are physically ubiquitous.

The kernel bandwidth σ_{k} is a free parameter, though several factors constrain its value. Firstly, it should not be much smaller than the Δr between experimental data points; otherwise, if two particles are at a distance in between two r values, it will have nearly no effect on any points of the computed computed spectrum $\hat{y}(x_i \mid \mathbf{R})$.

Another concern is that σ_{k} should be chosen so that our constructed signals fit the experimental curves well. If the parameter is too small, it will lead to unreasonably sharp peaks in the partial pair distribution functions. If it is too large, the partial pair distribution functions will be too diffuse. The optimal value will depend on the width of the distribution of the data as well as the number of samples. For normally distributed data, the mean integrated square error (MISE) is minimized for a kernel bandwidth of

$$\sigma_{\mathrm{k}} = \left(\frac{4\sigma^5}{3n}\right)^{1/5}, \tag{2.5}$$

where σ is the standard deviation of the distribution and n is the number of samples [58]. While our distances are not normally distributed, the first nearest-neighbour distances are approximately normal. To estimate the corresponding distribution width σ, we can use the full width at half maximum (FWHM) for the first nearest-neighbour peak from our experimental data,

$$\mathrm{FWHM} = 2\sqrt{2\ln 2}\,\sigma. \tag{2.6}$$

The number of samples in the first nearest-neighbour peak n is determined by the number of atoms in our computational structure N and their approximate coordination number z,

$$n = \frac{Nz}{2}, \tag{2.7}$$

as per the handshaking lemma. All-in-all, the MISE-optimal value for the kernel bandwidth is

$$\sigma_{\mathrm{k}} = \frac{1}{2\sqrt{2\ln 2}} \left(\frac{8}{3}\right)^{1/5} \frac{\mathrm{FWHM}}{(Nz)^{1/5}} \approx \frac{\mathrm{FWHM}}{2(Nz)^{1/5}}. \tag{2.8}$$

Of course, this only considers the first nearest-neighbour peak, and so some experimentation may be required.

2.2 Automatic differentiation

Now that $\chi^2(\mathbf{R})$ is differentiable, how should its gradient be calculated? Other approaches in the past have relied on analytical expressions, derived for use in specific cases [40]. However, such a method does not generalize easily. Here, we use automatic differentiation for its analytical correctness, computational simplicity, and generalizability [41].

In automatic differentiation, the operations applied to each variable are stored in a computation graph (Fig. 2.1). Then, gradients are calculated from this graph by backpropagation with the chain rule. This provides analytically correct values for gradients (up to floating point accuracy), is readily generalizable, and is computationally efficient for calculating gradients with respect to many variables. Additionally, this method is of critical importance for use in neural networks, and so has been robustly implemented for general use in libraries such as PyTorch, Tensorflow, and JAX. We have chosen to implement our code with PyTorch. These libraries also simplify the process of running code on graphics processing units (GPUs), which can make calculation times significantly shorter.

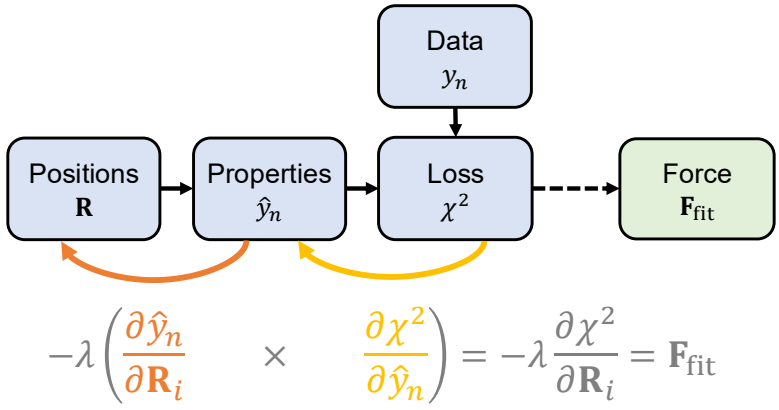


Figure 2.1: Diagram showing a simplified computation graph for our method. In reality, a new node is created after each mathematical operation (e.g., addition, multiplication, exponentiation, etc.). Note that the fitting force is calculated by taking the gradient of χ^2 by automatic differentiation using the chain rule, rather than by an explicitly defined analytical operation.

2.3 Evaluation of distances

Finally, I discuss how we have implemented the evaluation of the distances between particles. There are two critical considerations that were made; first, the use of the minimum image convention to handle periodic boundary conditions, and second, the implementation of this calculation to be highly parallelizable, i.e., rely heavily on tensor operations to take advantage of GPU architecture.

The minimum image convention is a protocol for determining the distances between two points under periodic boundary conditions, wherein the distance is chosen to be the smallest distance between all images of the points (Fig. 2.2). This is crucial because it essentially treats each particle as being in the centre of the simulation box. Our implementation of the minimum image convention must also work under any generalized supercell.

Let X_{mi} be the ith coordinate of the mth particle of species α, and let Y_{ni} be the ith coordinate of the nth particle of species β. The ordinary displacement vector between the

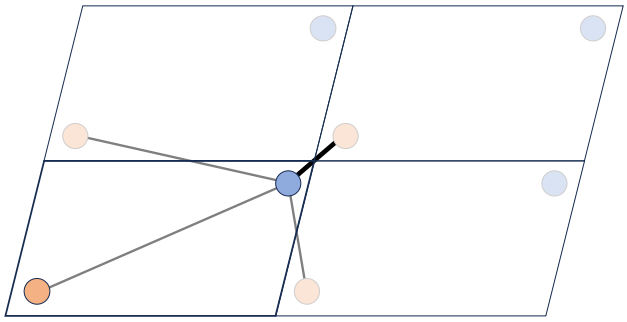


Figure 2.2: Diagram displaying the minimum image convention. A supercell and some of its periodic images are shown. Under the minimum image convention, the distance from the blue particle to the orange particle is given by the shortest distance among periodic images of the orange particle.

mth and nth particles has components

$$\Delta_{mni} = Y_{ni} - X_{mi}. \tag{2.9}$$

This forms a rank-3 tensor. However, our particle positions are under periodic boundary conditions within a supercell with lattice vectors $\{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3\}$. Let the cell matrix be

$$\mathbf{A} = \begin{bmatrix} \mathbf{a}_1^{\mathrm{T}} \\ \mathbf{a}_2^{\mathrm{T}} \\ \mathbf{a}_3^{\mathrm{T}} \end{bmatrix}, \tag{2.10}$$

so that A_{ij} is the jth component of the $\mathbf{a}_i$. Under the minimum image convention, the displacements between our particles is

$$\boldsymbol{\Delta}' = \boldsymbol{\Delta} - \mathrm{round}\left(\boldsymbol{\Delta} \cdot \mathbf{A}^{-1}\right) \cdot \mathbf{A}, \tag{2.11}$$

where we multiply rank-3 and rank-2 tensors such that their components are

$$[\mathbf{\Delta} \cdot \mathbf{B}]_{mni} = \sum_k \Delta_{mnk} B_{ki}. \tag{2.12}$$

The distance matrix under the minimum image convention is then

$$d_{mn} = \left[\sum_i (\Delta'_{mni})^2 \right]^{1/2}. \tag{2.13}$$

The set of observed distances $\mathcal{D}_{\alpha\beta}$ can then be constructed from the distance matrix d_{mn}. If α and β correspond to different species, then the set of observed distances is simply given by all components of the distance matrix. If α and β are the same species (i.e., $\alpha = \beta$), the distance matrix will be square, and the set of observed distances is given by the components above the diagonal; the components *on* the diagonal are self-distances, which are all zero, and the components *below* the diagonal are equivalent to those above by symmetry.

To validate the computational advantage of GPU acceleration, we benchmarked our pair distribution function calculation method on both central processing unit (CPU) and GPU architectures. Random atomic configurations of varying sizes (from 128 to 4096 atoms) were generated at constant density ($0.0496\,\text{Å}^{-3}$) with a minimum interatomic distance of 1.9 Å. For each system size, pair distribution function calculations were performed over 5 trials on both an Intel Silver 4216 Cascade Lake CPU (2.1 GHz) and an Nvidia V100 Volta GPU (32 GB memory). Different random structures were generated for each individual trial.

2.4 Calculation details

To test the efficacy of our method, we have applied it to two different systems — namely, amorphous silicon (a-Si) and a nickel–niobium metallic glass ($Ni_{62}Nb_{38}$). These materials have been rigorously studied previously using different methods [42–45], and each presents

its own unique challenges, making them ideal systems on which to test our approach.

2.4.1 Amorphous silicon

Experimental X-ray scattering data were provided by Laaziri *et al.* [59]. These data contained a scaled form of the pair distribution function $J(r) = 4\pi r^2 \rho\ g(r)$, and did not include uncertainties. The experimental $J(r)$ data were converted to $g(r)$ using the relation $g(r) = J(r)/(4\pi r^2 \rho)$ with number density $\rho = 0.0495\,\text{Å}^{-3}$. To ensure physical consistency, $g(r)$ values were set to zero for values of r definitively before the first peak ($r < 2.125\,\text{Å}$). Additionally, all negative values were flattened to zero. Since experimental uncertainties were not provided, they were estimated by i) smoothing the spectrum with a Bessel filter (cutoff frequency $0.125\,\text{Å}^{-1}$), ii) computing the absolute difference between original and smoothed spectra, and iii) applying another Bessel filter (cutoff frequency $0.4\,\text{Å}^{-1}$) to the logarithm of these absolute differences. The resulting uncertainty estimates capture the high-frequency noise characteristics of the experimental measurements and provide realistic error bars. To validate the qualitative behaviour of the uncertainties approximated with this method, they were compared to the distribution of $g(r)$ values found for melt-quench structures.

To determine how realistic our final structures were, it was important to generate comparison structures using melt-quench molecular dynamics. This simulation was performed in LAMMPS [60]. Starting from a $5 \times 5 \times 5$ crystalline silicon supercell with diamond lattice structure (lattice parameter 5.431 Å), the system was evolved using the modified Tersoff potential for silicon [61] with a timestep of 1 fs. The amorphization protocol consisted of four stages under NPT ensemble conditions at 1 bar of pressure: Initial heating from 300 K to 3500 K over 3.2 ps, equilibration at 3500 K for 0.5 ps, cooling back to 300 K over 3.2 ps, and final equilibration at 300 K for 1 ps. Temperature and pressure coupling used Nosé–Hoover thermostats and barostats with damping timescales of 0.1 ps and 1 ps, respectively. This procedure was performed 100 times with different random seeds to generate 100 independent

structures.

For our tests on amorphous silicon, we perform two optimizations from different starting structures. Our first test began from the structure of crystalline silicon, with an applied strain so that its density was that of amorphous silicon. Our second test was to evaluate the effectiveness of our method in refining a structural model by optimizing from a starting melt-quench structure [20]. The effective potential $V_{\text{eff}} = V + \lambda\chi^2$ was optimized using the FIRE optimizer [62]. The CHGNet universal machine-learned potential (MLP) [29] was used for the potential V, and was chosen for its generally high accuracy and speed. For both starting configurations, 2000 optimization steps were performed using a kernel bandwidth of $\sigma_{\text{k}} = 0.03\,\text{Å}$. For the optimization from the crystal structure, weightings of $\lambda = 1 \times 10^{-4}\,\text{eV},\ 4 \times 10^{-4}\,\text{eV}$ and $1 \times 10^{-3}\,\text{eV}$ were used, and for the optimizations from the melt-quench structure, weights of $\lambda = 1 \times 10^{-2}\,\text{eV},\ 3 \times 10^{-2}\,\text{eV}$ and $1 \times 10^{-1}\,\text{eV}$ were used.

To evaluate the accuracy of our final structures, their calculated pair distribution functions were compared to the experimental data, and their structural properties were compared to those of our melt-quench simulations. ASE was used to find the distribution of bond lengths, bond angles, dihedral angles, and coordination numbers in both melt-quench structures and our optimized structures [46].

2.4.2 Nickel–niobium metallic glass

As we were not able to access experimental spectra for the nickel–niobium metallic glass, we generated target spectra using molecular dynamics structures. The alloy was simulated in LAMMPS [60] using a $14 \times 14 \times 14$ BCC supercell with lattice parameter 3.30 Å, initially populated with nickel atoms. The composition was adjusted to $Ni_{62}Nb_{38}$ by randomly converting 38% of nickel atoms to niobium. Interatomic interactions were described using the embedded atom method potential of [44]. The system was equilibrated at 1447 K under NPT conditions at 1 bar of pressure using a Nosé–Hoover thermostat and barostat with

damping timescales of 0.1 ps and 1 ps, respectively. The simulation was run for 10 ps with a timestep of 1 fs. This procedure was performed 100 times with different random seeds to generate 100 independent structures. The total pair distribution function was calculated for each structure, and its mean value and standard error at each r were used for the target value and its uncertainty.

The starting configuration for our optimization test was generated using a random closest-approach algorithm. 1500 atoms (930 Ni and 570 Nb) were placed randomly within a cubic simulation cell while enforcing a minimum interatomic distance of 2.05 Å. The cell size was determined from the target number density of $0.0704\,\text{Å}^{-3}$. Atoms were placed sequentially, with each new position tested against all previously placed atoms using minimum image convention to handle periodic boundary conditions. This approach generated a starting structure that avoids unphysical atomic overlaps while maintaining the desired composition and density.

The effective potential $V_{\text{eff}} = V + \lambda\chi^2$ was optimized using the FIRE optimizer [62]. The CHGNet universal MLP [29] was used for the potential V. The χ^2 loss was calculated between our guess structure and the target pair distribution function. The simulated pair distribution function was calculated using approximate X-ray total scattering (Eq. A.16) with form-factor parameters from Ref. [63]. A kernel bandwidth of $\sigma_{\text{k}} = 0.05$ Å was used to calculate the pair distribution function. A total of 3,000 optimization steps were performed, and hyperparameter weights of $\lambda = 1 \times 10^{-4}\,\text{eV},\ 1 \times 10^{-2}\,\text{eV}$ and $1 \times 10^{-1}\,\text{eV}$ were used.

To evaluate the accuracy of our structure, we compared its total pair distribution function to that of the target. We also compared each of the partial pair distribution functions as well as the first nearest-neighbour coordination number, as given by the Eq. (A.5). Bond angles between all triplets of elements were also calculated using ASE [46], and compared between a melt-quench structure and our optimized structure.

2.5 Results and discussion

2.5.1 Benchmarking

I will begin by discussing the performance of our package on both the CPU and GPU. The times taken to calculate the pair distribution function for systems of different sizes on the CPU and GPU are shown in Fig. 2.3. Our results indicate that focusing on GPU parallelization has been both effective and worthwhile.

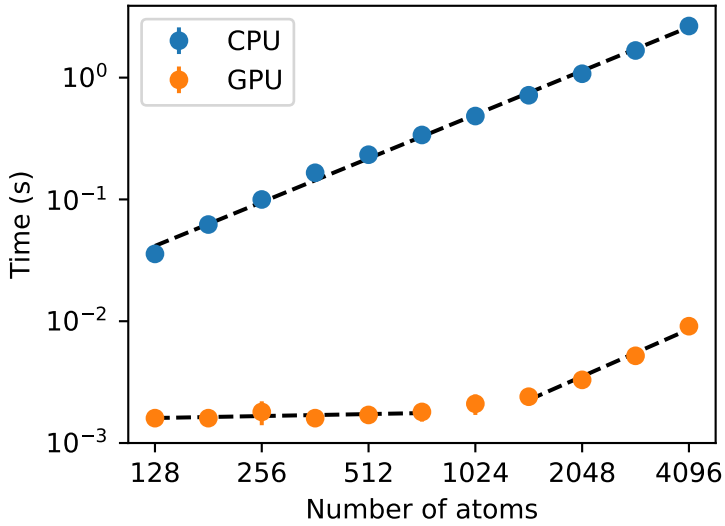


Figure 2.3: Times taken for a single central processing unit (CPU) and a single graphics processing unit (GPU) to compute the pair distribution function for different numbers of randomly placed atoms.

The plot shows that our calculations are indeed performed in parallel on a GPU. Computation time on the CPU scales as $\mathcal{O}(n^{\sim 1.2})$ regardless of system size. GPU computation time, in contrast, scales as $\mathcal{O}(1)$ for systems smaller than approximately 1000 atoms, as a result of calculations being performed in parallel. When the system is larger, there is increased overhead, and the GPU follows the same approximate scaling law as the CPU.

In addition to the advantageous scaling, computation on the GPU is significantly faster. The calculations took approximately 1 ms to 10 ms on the GPU and 50 ms to 3000 ms on the

CPU. As such, the GPU was roughly 50 to 300 times faster, depending on system size. Of course, the timings will depend on the specific hardware used; however, the improved scaling and significantly improved calculation time here suggest that the GPU will often be faster than the CPU for our package.

2.5.2 Amorphous silicon

We began by testing our method of calculating the partial pair distribution functions by testing it on a melt-quench structure for different values of kernel bandwidth σ_k (Fig. 2.4). Each curve exhibits many of the traits associated with the partial pair distribution function of amorphous silicon, including the sharp first nearest-neighbour peak at 2.35 Å, a second nearest-neighbour peak at 3.84 Å. Additionally, the first coordination number, as calculated by Eq. (A.5), is close to 4 for each curve. We also see that the limit of $g(r)$ as r approaches infinity is 1, as expected. This indicates that our method has been correctly implemented.

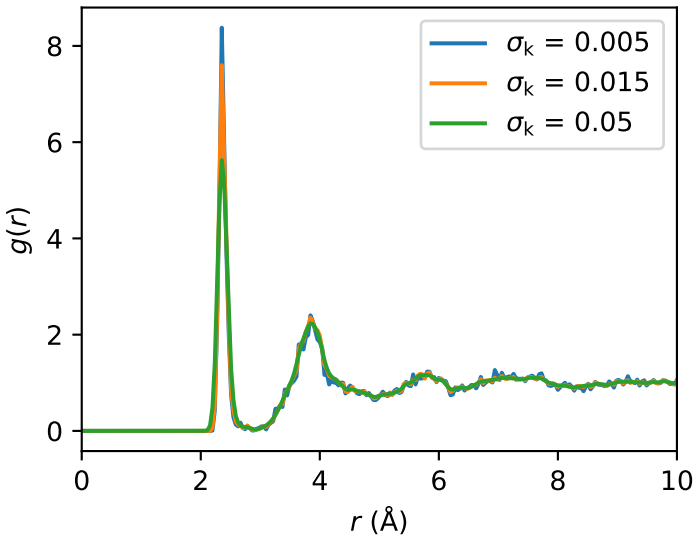


Figure 2.4: Pair distribution functions for melt-quench amorphous silicon, calculated with kernel density estimation. Different lines correspond to different kernel bandwidths σ_k, in units of Ångstroms.

Further, we wanted to investigate the effect of the kernel bandwidth parameter. As expected, larger kernel bandwidths result in smoother curves. While this is somewhat favourable for re-

ducing noise, it also leads to a much more diffuse first nearest-neighbour peak. The estimated mean integrated square error (MISE)-minimizing bandwidth is 0.015 Å; indeed, making the bandwidth smaller than this does not seem to drastically sharpen the first nearest-neighbour peak, but making the bandwidth larger does lead to noticeable broadening. However, the experimental data points we wish to compare against have a spacing of 0.04 Å, and using a bandwidth smaller than this spacing could results in instabilities while optimizing. We will use a kernel bandwidth of 0.03 Å as a compromise between the MISE-minimizer and the experimental bin width.

Uncertainties were not included in the experimental data. As such, they needed to be estimated. This was done by isolating high-frequency components of the signal and smoothing the result in logarithmic space. The result is shown in Fig. 2.5b. To validate our procedure, we ran 100 melt-quench simulations and calculated the mean pair distribution function and its standard error (Fig. 2.5a). This resulted in qualitatively similar curves with i) very small errors before the first peak, ii) local maxima around each peak in $g(r)$, and iii) a gradual fall off as r gets larger. Of course, we don't expect these curves to be identical. The objective was simply to increase the confidence in our procedure for estimating uncertainties.

Optimization from a strained crystal structure

Finally, we turn to our results from optimizations. We began by optimizing $V(\mathbf{R}) + \lambda\chi^2(\mathbf{R})$ starting from a strained crystal structure. The potential energy per atom and χ^2-loss are shown in Fig. 2.6. Both appear to converge within 1000 total steps. This is significantly faster than reverse Monte Carlo methods, which have taken 400,000 accepted RMC steps and 4,000 force evaluations using the same experimental data [42].

The χ^2-loss initially falls sharply and then stabilizes. The potential energy, on the other hand, initially *increases*, before eventually decreasing. The optimizer allows this because the resulting decrease in χ^2 is drastic enough that $V + \lambda\chi^2$ is still decreased; as such, larger λ

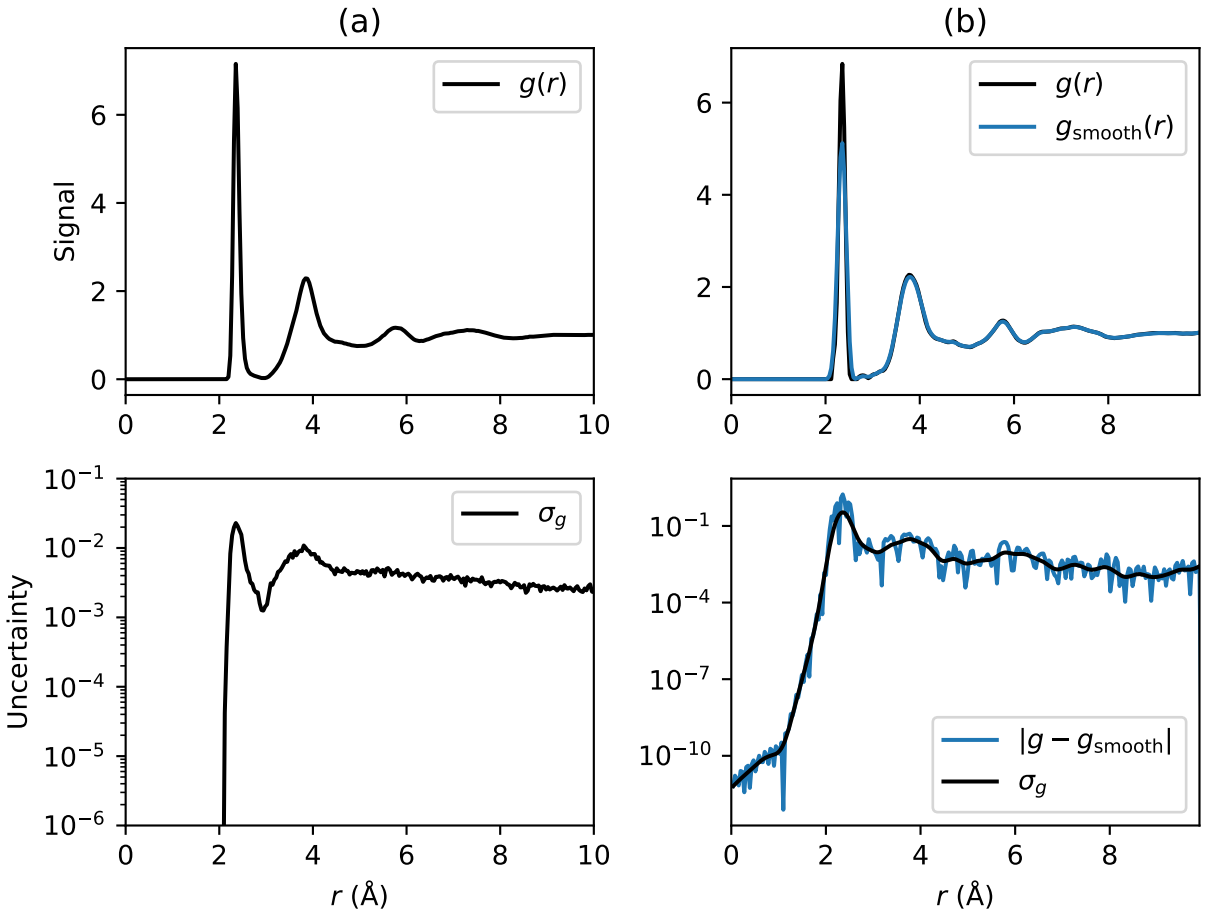


Figure 2.5: Pair distribution functions for amorphous silicon and their uncertainties. These were given by (a) melt-quench simulation and (b) experimental X-ray scattering.

values allow for larger peak energies before they turn around. It is possible that this behaves like an activation energy barrier into the amorphous state. The smallest of these peaks is at 0.2 eV/atom. This corresponds to a temperature of $2 \cdot 0.2\,\mathrm{eV}/3k_\mathrm{B} = 1550\,\mathrm{K}$, which is close to the melting point of silicon, 1414 K.

The final values of V and χ^2 converged to depend on the value of λ chosen. Larger values of λ result in smaller χ^2 values and larger energies. This is intuitive as λ controls the weighting of χ^2. For the hyperparameter value $\lambda = 4 \times 10^{-4}\,\mathrm{eV}$, both V and χ^2 find their final value within the range seen for melt-quench structures. While this does not necessarily indicate that our optimized structure is in an equivalent melt-quench state, it does suggest that our method is generally promising.

To give meaning to the χ^2-losses discussed above, we calculated pair distribution functions for our optimized structures and compared them to the experimental target (Fig. 2.7). In agreement with the results discussed above, the two larger values of λ lead to structures with closely agreeing pair distribution functions, while the smaller value of λ did not. Further, we see that the first nearest-neighbour peak is too sharp for all values of λ. This is possibly a consequence of the fact that we are partially minimizing the potential energy, which favours a preferred bonding length.

Next, we compared the structural properties of our optimized structures to those of a melt-quench structure (Fig. 2.8). This comparison indicates that our structure is distinct from the melt-quench structures in a few key ways.

While the observed bond lengths (Fig. 2.8a) are centred around similar values for both our optimized and melt-quench structures, those of our melt-quench structures follow a tighter distribution. This is in agreement with the increased sharpness of the first nearest-neighbour peak in Fig. 2.7. Again, this is likely because optimizing V will push the structure towards

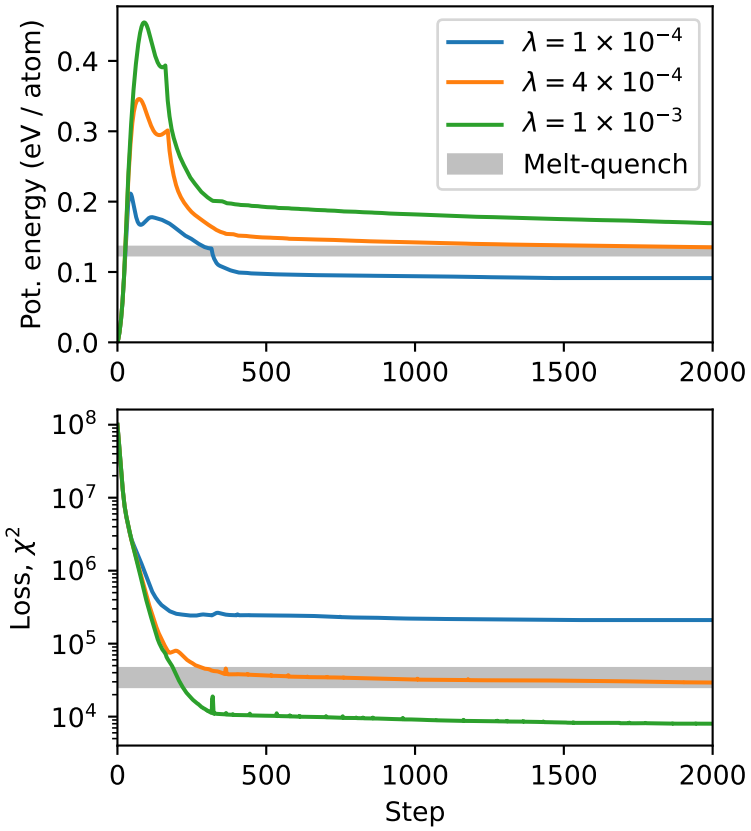


Figure 2.6: Per-atom potential energy and χ^2-loss over the course of optimizing $V + \lambda\chi^2$ for amorphous silicon, starting from a strained crystal structure. Different colours correspond to different values of λ in units of electronvolts. The grey shaded regions correspond to 95% confidence intervals on the potential energies and losses calculated for melt-quench structures.

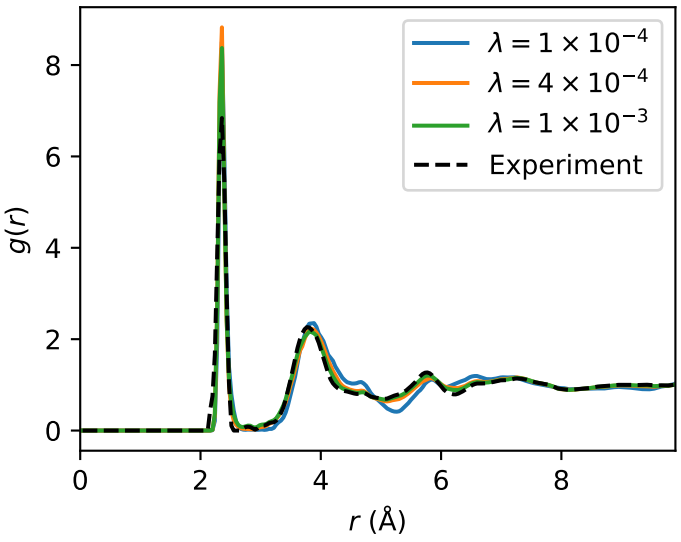


Figure 2.7: Pair distribution functions calculated for our optimized amorphous silicon structures, starting from a strained crystal structure. Different colours correspond to different values of λ in units of electronvolts. Also shown is the the experimental pair distribution function, which was used as an optimization target.

using its preferred bond length.

The bond angle distribution for our optimized structures is very close to that of the melt-quench structure (Fig. 2.8b). Note that for a crystalline structure, the bond angle distribution would be a sharp peak around 109.5°. The strong agreement seen here likely stems from the bond angle distribution being encoded by the first two peaks of the pair distribution function.[1] So, fitting the width of the first two peak results in the desired bond angle disorder.

In contrast, the dihedral angle distributions are much sharper than that of the melt-quench structure (Fig. 2.8c). The melt-quench dihedral density is approximately uniform across all angles with a slight preference towards 60°, 180° and 300°. This is due to steric hindrance, and is akin to the distribution of dihedral angles seen for ethane. For crystalline silicon,

[1]This can be seen by considering the locations of the first two peaks, $r_1 = 2.35\,\text{Å}$ and $r_2 = 3.84\,\text{Å}$. If two particles are at a distance of r_1 from a central atom and a distance r_2 from each other, the angle between them from the central atom is $\theta = 2\arcsin(r_2/2r_1) \approx 109.5°$.

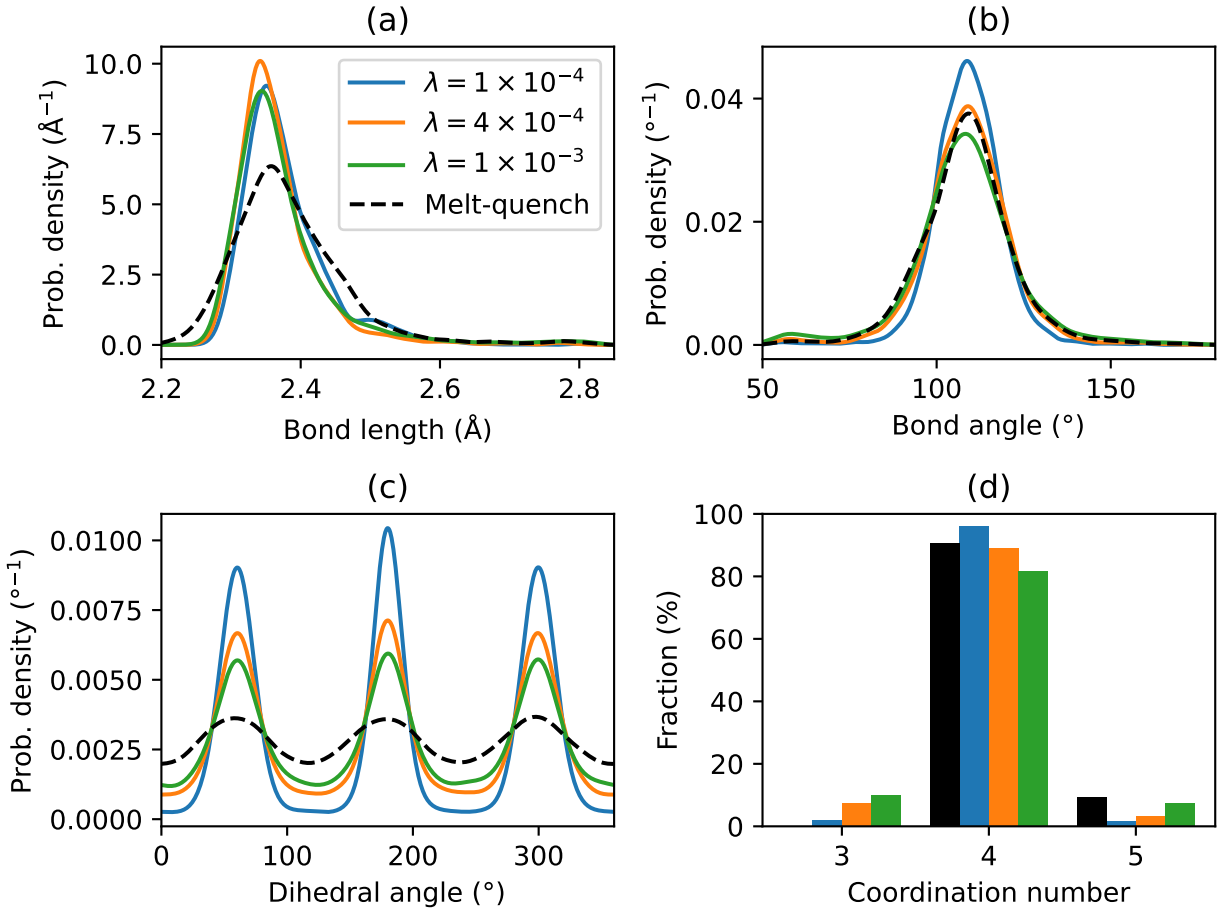


Figure 2.8: Distributions of (a) bond length, (b) bond angle, (c) dihedral angle, and (d) coordination number seen for our optimized amorphous silicon structures, starting from a strained crystal structure. Different colours correspond to different values of λ in units of electronvolts. Also shown are the distributions for a melt-quench structure.

the distribution consists of sharp peaks at the same values. For our optimized structure, the peaks are sharper than those of the melt-quench structure (though the sharpness does decrease for increased values of λ). This points to a more fundamental issue with our structure than the bond length distribution because bond lengths are easy to change, but dihedral angles are not. The issues with the dihedral angle distribution could likely only be alleviated by better exploration of the state space, i.e., because we are starting in a near crystal structure and performing local optimization, our final structure is still similar to that of crystal silicon.

Similarly, the distribution in the coordination number of atoms also points to an issue. For our melt-quench structure, approximately 90% of the atoms have a coordination number of 4, approximately 10% have a coordination number of 5, and none have a coordination number of 3. For our optimized structures, smaller values of λ give larger fractions of atoms with coordination number of 4, which is consistent with the fact that smaller λ values put more emphasis on the potential energy. For $\lambda = 4 \times 10^{-4}\,\mathrm{eV}$, the fraction of atoms with a coordination number of 4 appears to be correct. However, it appears that atoms have a slight preference for coordination number 3 compared to 5. It is not clear what causes this.

Optimization from a melt-quench structure

Next, we review the results of optimizing $V + \lambda\chi^2$, beginning this time from a melt-quench structure. This serves two purposes; firstly, it tests how our model would behave if it were to find a suitable amorphous state, and secondly, it tests the ability of our model to optimize melt-quench structures. Ideally, this should not deviate significantly in energy from the starting structure, but should improve the χ^2-loss.

The potential energy and losses for this optimization (Fig. 2.9) behave qualitatively similar to that starting from a strained crystal structure (Fig. 2.6); the losses drop dramatically throughout the optimization and the potential energies rise and then fall. However, the

quantitative scales of these changes are very different. The potential energy increased only by a maximum of 0.02 eV, whereas it increased by 0.4 eV for the crystalline structure. Furthermore, the loss decreases by only 1–2 orders of magnitude rather than 3–4. These changes can be attributed to the starting structure being closer to the amorphous state. Another interesting difference is that, when optimizing from the crystalline state, both the potential energy and loss equilibrated to values that were within the range seen for the melt-quench structures. Here, the energy remains within that range but the loss becomes much lower. This suggests that our model could be used to optimize a melt-quench structure towards a structure that fits experimental data.

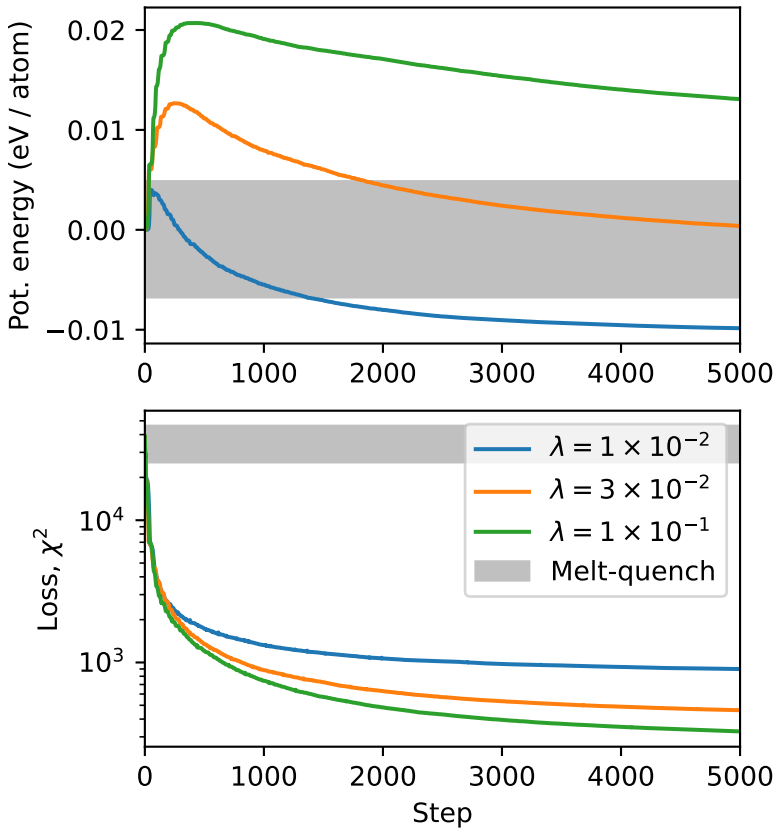


Figure 2.9: Per-atom potential energy and χ^2-loss over the course of optimizing $V + \lambda\chi^2$ for amorphous silicon, starting from a melt-quench structure. Different colours correspond to different values of λ in units of electronvolts. The grey shaded regions correspond to 95% confidence intervals on the potential energies and losses calculated for melt-quench structures.

All structures fit the experimental pair distribution function well (Fig. 2.10). However, the first nearest-neighbour peak is much closer for the larger values of λ, and much closer than was seen for structures optimized from the strained crystal (Fig. 2.7). The bond length distribution, bond angle distribution, and dihedral angle distribution remain close to those of the melt-quench structure (Fig. 2.11). However, the coordination number distribution has similar oddities as the structure optimized from a strained crystal, though to a lesser extent.

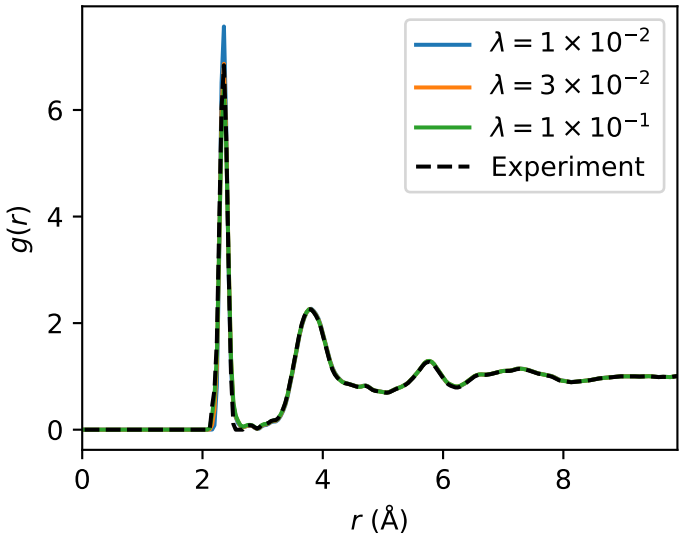


Figure 2.10: Pair distribution functions calculated for our optimized amorphous silicon structures, starting from a melt-quench structure. Different colours correspond to different values of λ in units of electronvolts. Also shown is the the experimental pair distribution function, which was used as an optimization target.

2.5.3 Nickel–niobium metallic glass

Next, we turn to our results for $Ni_{62}Nb_{38}$ metallic glass. As we did not have experimental data for this structure, we performed one hundred molecular dynamics simulations to generate our simulated target data (Fig. 2.12). We calculated the partial pair distribution functions and total X-ray pair distribution function for each simulated structure. The average total X-ray pair distribution function was used as a training target, and the standard error was used as the uncertainty.

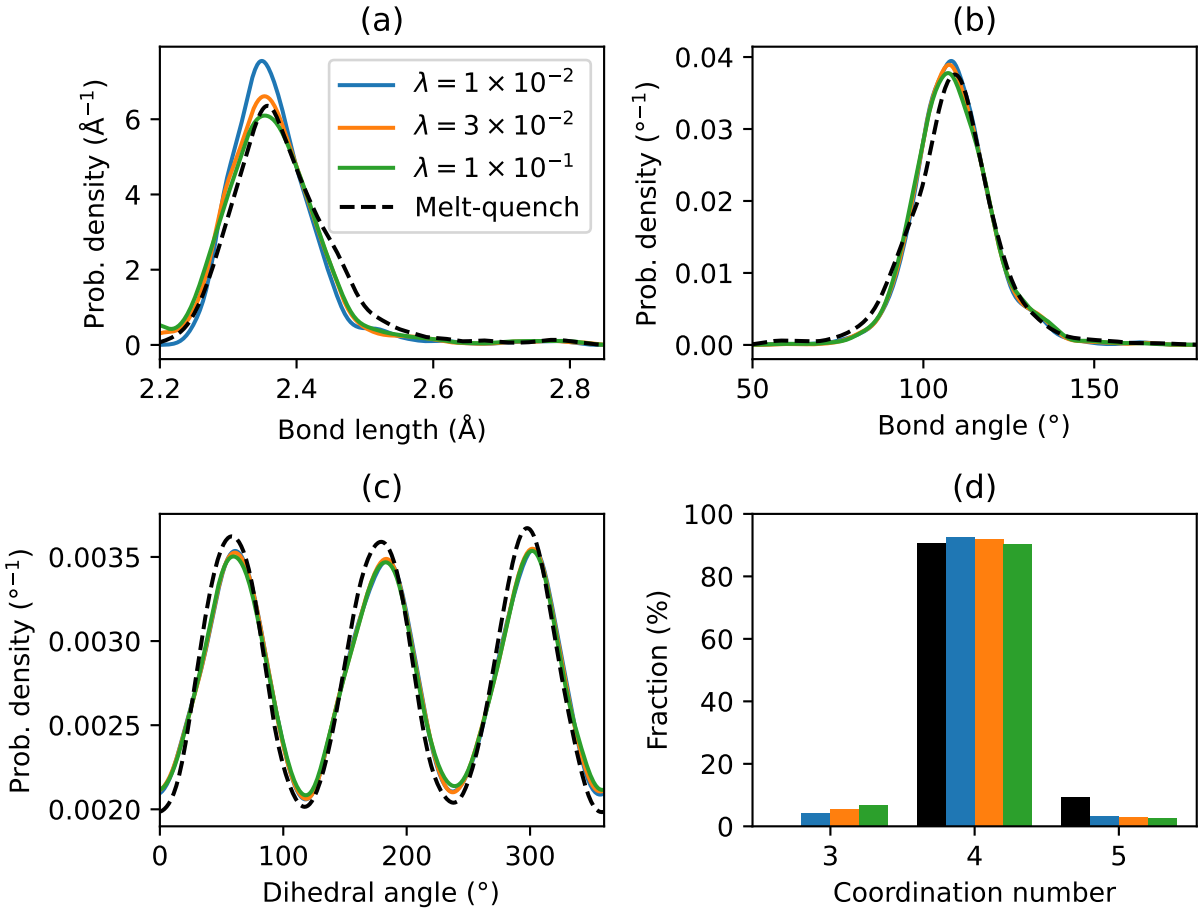


Figure 2.11: Distributions of (a) bond length, (b) bond angle, (c) dihedral angle, and (d) coordination number seen for our optimized amorphous silicon structures, starting from a melt-quench structure. Different colours correspond to different values of λ in units of electronvolts. Also shown are the distributions for an unoptimized melt-quench structure.

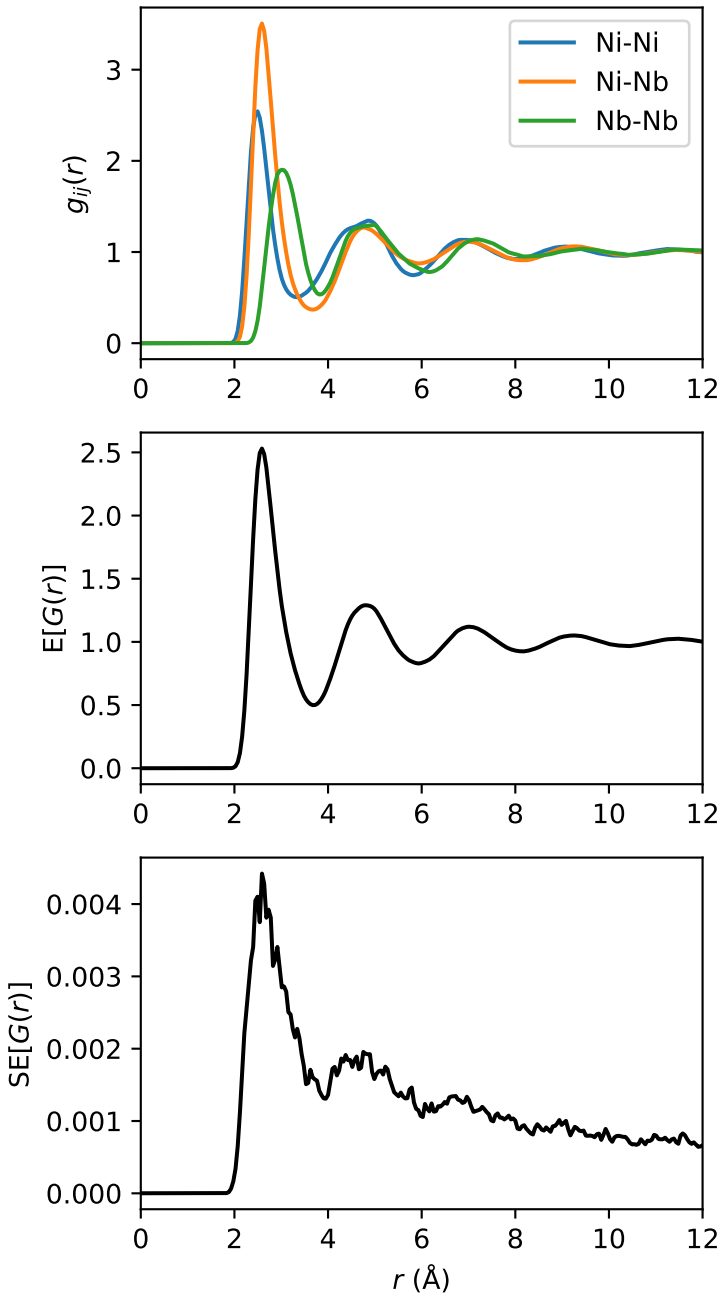


Figure 2.12: Average partial pair distribution functions and X-ray total pair distribution function, and the standard error in the X-ray total pair distribution function, calculated for one hundred $Ni_{62}Nb_{38}$ structures obtained by molecular dynamics (MD).

Optimization from a random starting structure

We optimized our structure from a random starting configuration. The potential energy and χ^2-loss are shown in Fig. 2.13. Each quantity had stabilized after about 2000 steps. The potential energy quickly decreases by more than 1 eV/atom. For the largest value of λ tested, the χ^2-loss decreases continually; however, for the smaller values, it decreases initially and then begins to increase again. The final potential energies are close to those seen for our molecular dynamics structures, but are lower for the smaller values of λ. Furthermore, the losses are smaller than those seen for the molecular dynamics structures for all but the smallest value of λ tested.

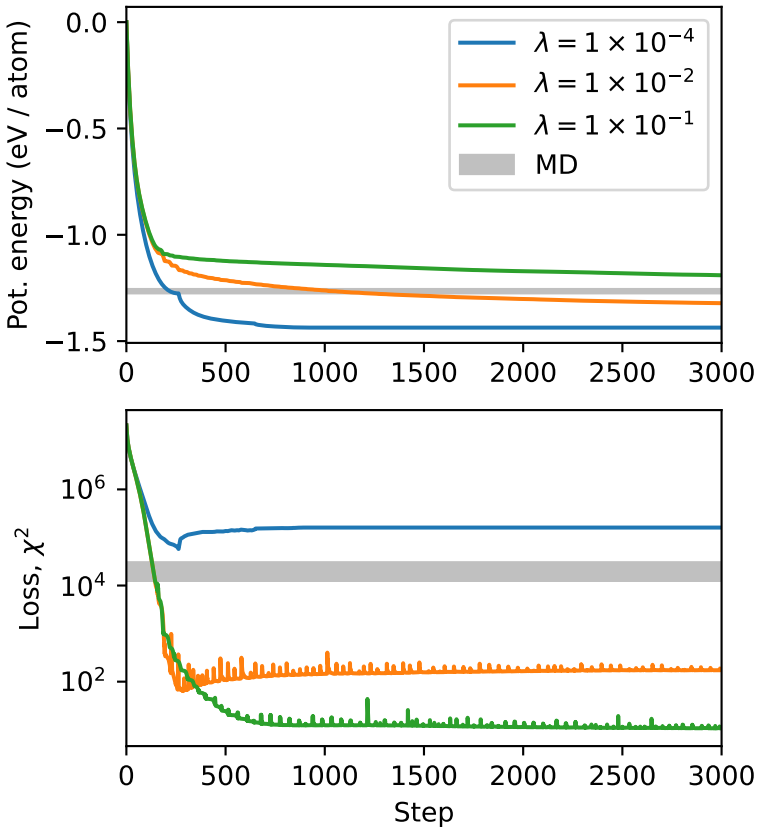


Figure 2.13: Per-atom potential energy and χ^2-loss over the course of optimizing $V + \lambda\chi^2$ for amorphous $Ni_{62}Nb_{38}$, starting from a random structure. Different colours correspond to different values of λ in units of electronvolts. The grey shaded regions correspond to 95% confidence intervals on the potential energies and losses calculated for amorphous $Ni_{62}Nb_{38}$ structures obtained through molecular dynamics (MD).

Comparing the total pair distribution functions of our optimized structures to the target function, we see remarkable fit for our larger values of λ (Fig. 2.14). The smallest value of λ does not fit the first peak very well, but does fit other parts of the curve. This suggests that the inclusion of the χ^2-loss does indeed play an important role in our optimization.

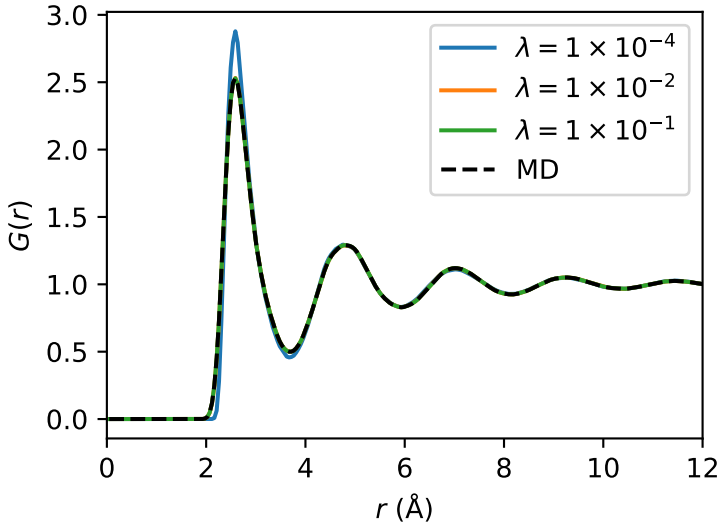


Figure 2.14: X-ray total pair distribution functions calculated for our optimized amorphous $Ni_{62}Nb_{38}$ structures, starting from a random structure. Different colours correspond to different values of λ in units of electronvolts. Also shown is the the average X-ray total pair distribution function for amorphous $Ni_{62}Nb_{38}$ structures obtained through molecular dynamics (MD), which was used as an optimization target.

Next, we compare the structural properties of our molecular dynamics simulations with our final structures. All partial pair distribution functions (Fig. 2.15) have peaks in the correct locations, but are best fit by the structure with the largest λ value. Note that the partial pair distribution functions were not used directly as training targets. The Nb-Nb pair distribution function has the worst fit, likely because it is the signal with the least data (due to the lower concentration of Nb than Ni in $Ni_{62}Nb_{38}$). However, we still find that the structure from the optimization with the largest λ value has the closest Nb-Nb coordination number to the molecular dynamics structures.

The bond angle distributions are shown in Fig. 2.16. Firstly, we note that the bond angle

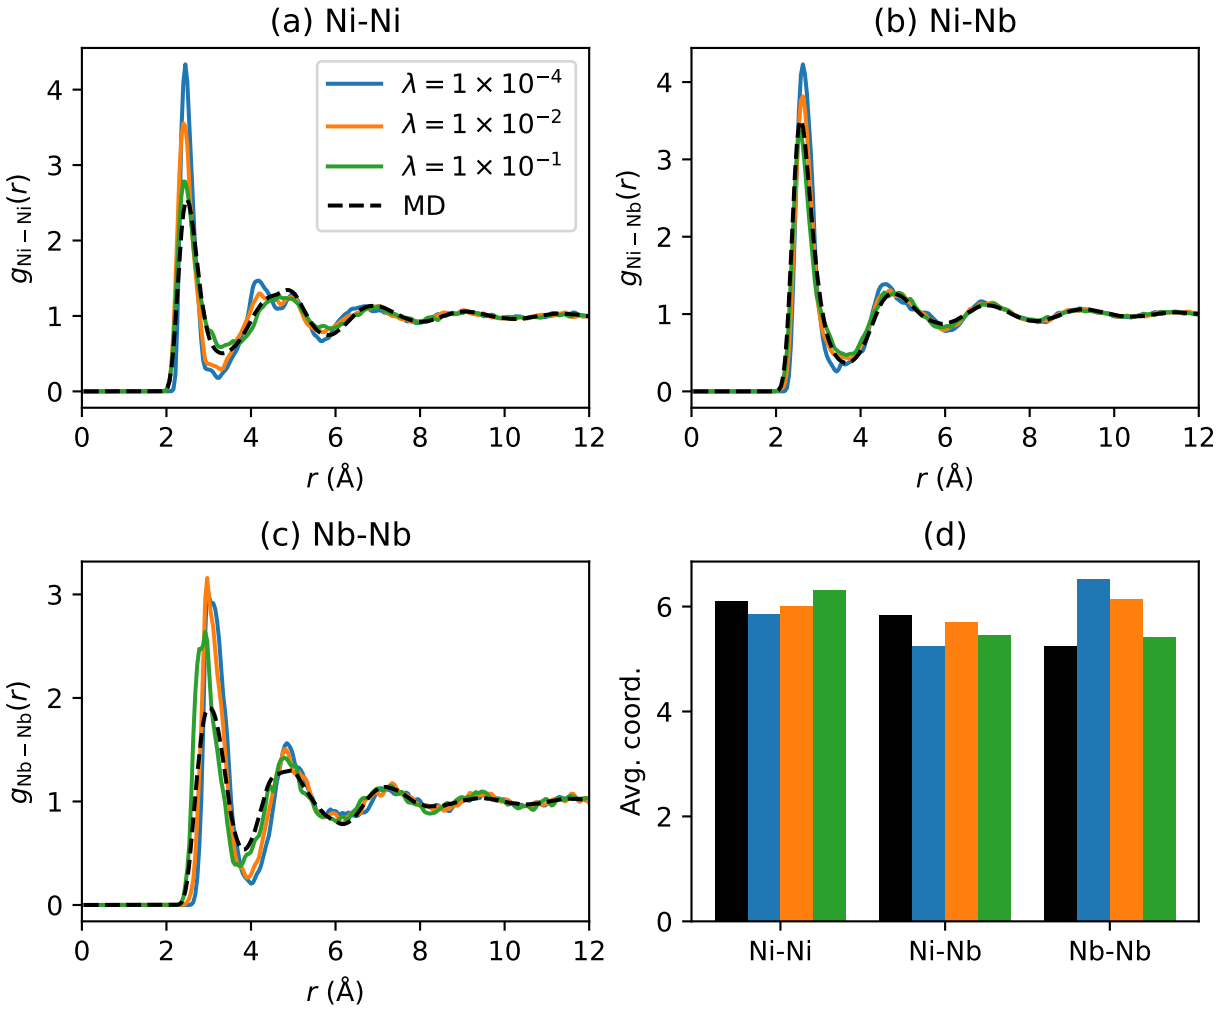


Figure 2.15: (a)–(c) Partial pair distribution functions and (d) average coordination numbers, calculated for our optimized amorphous $Ni_{62}Nb_{38}$ structures, starting from a random structure. Different colours correspond to different values of λ in units of electronvolts. Also shown are the average partial pair distribution functions and coordination numbers for one hundred amorphous $Ni_{62}Nb_{38}$ structures obtained through molecular dynamics (MD). Note that these partial pair distribution functions were *not* directly used as optimization targets.

distributions of our optimized structures generally match those of the molecular dynamics structure and those that have been seen in the literature [44]. While general trends are difficult to discern, the structure with the largest λ value often has the closest fit to the molecular dynamics structure. This indicates that the optimization with just the machine learned potential is not enough to faithfully reproduce the bond angle distribution, and that information from the total pair distribution function plays an important role.

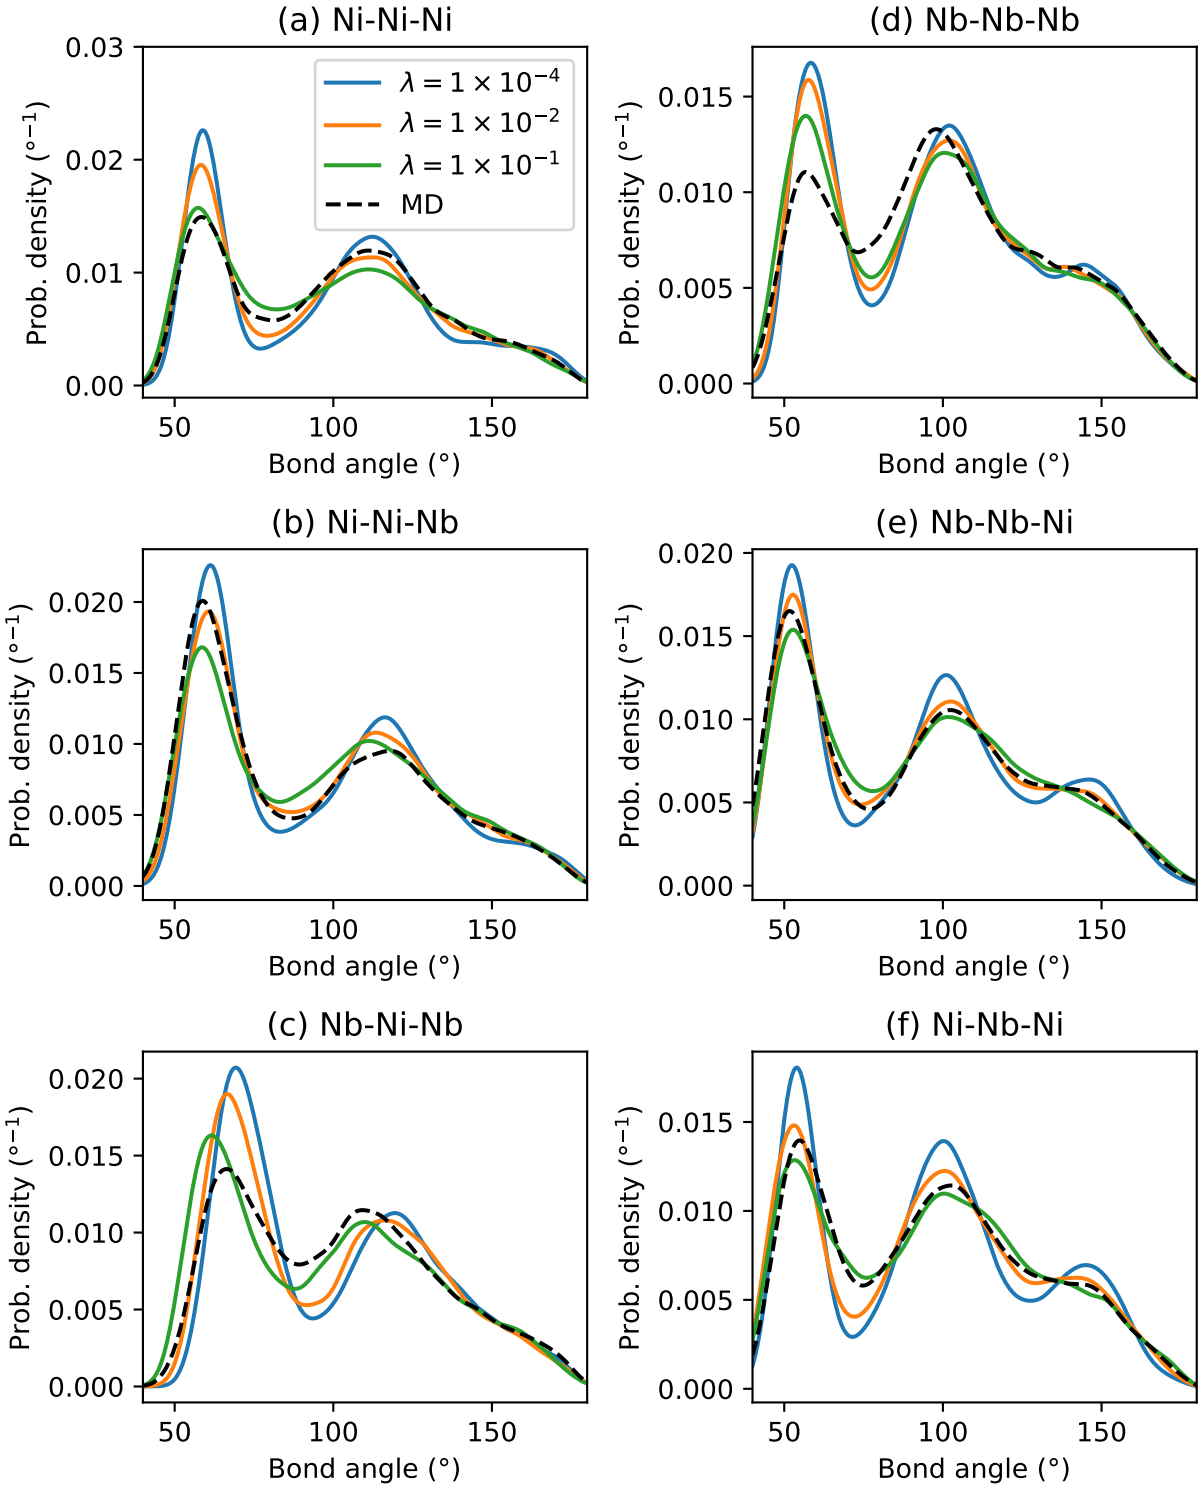


Figure 2.16: Bond angle distributions calculated for our optimized amorphous $Ni_{62}Nb_{38}$ structures, starting from a random structure. Different colours correspond to different values of λ in units of electronvolts. Also shown are the bond angle distributions for a amorphous $Ni_{62}Nb_{38}$ structure obtained through molecular dynamics (MD).

Chapter 3

Automated Analysis of Hopping in Solid-State Electrolytes

3.1 General approach

The problem of identifying different sites from trajectory data resembles an unsupervised machine learning task, where clusters can be identified by examining relative distances between points in space. However, conventional clustering methods require careful parametrization and have inherent limitations for trajectory analysis. For example, k-means clustering requires prior knowledge of the number of clusters [64]; there are methods to get around this, but they are not robust, and would themselves require careful parametrization. Density-based clustering algorithms (e.g., DBSCAN [65] and OPTICS [66]) require specifying a density threshold for cluster identification. But, for trajectory analysis, the ideal density parameter would depend not only on the material properties but also on simulation parameters such as the timestep. Furthermore, these approaches neglect the inherently temporal nature of our data, though this limitation can be partially addressed by including time in the state vectors and carefully choosing distance metrics. These global clustering approaches analyze the entire signal simultaneously to define clusters.

An alternative, but also flawed, approach would be to consider the position momentarily at each time instance. A low-pass filter, such as a Bessel filter with a carefully chosen cutoff frequency, can be applied to smooth noise around each moment in time. The derivative of the filtered signal would yield the velocity of the diffusive particle, ideally producing sharp peaks in speed when the particle transitions between sites. Peaks exceeding a predetermined threshold can then be identified as hop events. However, this method requires careful selection of the cutoff frequency—too low and noise remains unfiltered, too high and the diffusive motion itself becomes smoothed, reducing peak speeds. There is a clear trade-off between the choice of cutoff frequency and the speed threshold.

Evidently, neither of the above approaches is ideal. Our procedure combines elements of both global and local methods. We employ a cumulative average as a low-pass filter. We denote the averaged signal, termed the centroid position, as $\mathbf{r}_{\mathrm{c}}(t)$. A cumulative average naturally adapts to provide the best estimate of the local cluster centre based on all positions observed since the last hop event. We use the distance between the original signal and its cumulative average to identify hop events. Let $\tau_0 = 0$ and τ_i represent the time of the ith hop. The cumulative average of $\mathbf{r}(t)$ is defined as:

$$\mathbf{r}_{\mathrm{c}}(\tau_i < t \leq \tau_{i+1}) = \frac{1}{t - \tau_i} \int_{\tau_i}^{t} \mathbf{r}(u) \, \mathrm{d}u. \tag{3.1}$$

We define τ_{i+1} as the first time $t > \tau_i$ for which

$$\| \, \mathbf{r}(t) - \mathbf{r}_{\mathrm{c}}(t) \, \| > \lambda, \tag{3.2}$$

where λ is a threshold distance. A key advantage of this approach is that it requires only a single parameter, λ, which has a clear physical interpretation as the minimum displacement threshold for identifying a hop event. This threshold depends solely on the characteristic length scale of the hopping process of interest, making the method straightforward to

parameterize for different materials and systems.

The above equations are written in the continuous limit. Of course, molecular dynamics data are discrete, with positions $\mathbf{r}(t_n)$ at times $\{t_n\}$ where $t_n = n\Delta t$. The rolling centroid is calculated recursively

$$\mathbf{r}_{\mathrm{c}}(t_n) = \frac{(m-1)\mathbf{r}_{\mathrm{c}}(t_{n-1}) + \mathbf{r}(t_n)}{m}, \tag{3.3}$$

where m is the number of timesteps since the last hop.

Once the hopping times $\{\tau_i\}$ are identified, we can extract key quantities that characterize the hopping dynamics. For instance, time time spent in each site i is

$$\Delta\tau_i = \tau_{i+1} - \tau_i. \tag{3.4}$$

The hopping vector represents the displacement between consecutive sites;

$$\Delta\mathbf{r}_i = \mathbf{r}_{\mathrm{c}}(\tau_{i+1}) - \mathbf{r}_{\mathrm{c}}(\tau_i). \tag{3.5}$$

To characterize the correlation between successive hops, we can calculate the angle between consecutive hop vectors;

$$\cos\theta_{i,i+1} = \frac{\Delta\mathbf{r}_i \cdot \Delta\mathbf{r}_{i+1}}{\|\Delta\mathbf{r}_i\| \, \|\Delta\mathbf{r}_{i+1}\|}. \tag{3.6}$$

To assess the spatial extent and consistency of each identified cluster, we also calculate the radius of gyration for site i;

$$R_{\mathrm{g},i} = \sqrt{\frac{1}{M_i}\sum_{m=1}^{M_i} \| \, \mathbf{r}(\tau_{i-1} + m\Delta t) - \mathbf{r}_{\mathrm{c}}(\tau_i) \, \|^2} \tag{3.7}$$

where M_i is the number of timesteps during which the particle resides at site i, and Δt is the simulation timestep. This quantity measures the root-mean-squared distance of trajectory points from the site centre, providing insight into the compactness and consistency of the

identified sites.

From these fundamental quantities, we can compute several parameters of physical interest. The mean-squared hop length provides a measure of the characteristic length scale of the hopping process;

$$a^2 = \frac{1}{N_{\text{hop}}} \sum_{i=1}^{N_{\text{hop}}} \|\Delta \mathbf{r}_i\|^2, \tag{3.8}$$

where N_{hop} is the total number of observed hops. The average hopping rate characterizes the frequency of hop events;

$$\Gamma = \left(\frac{1}{N_{\text{hop}}} \sum_{i=1}^{N_{\text{hop}}} \Delta \tau_i \right)^{-1}. \tag{3.9}$$

The correlation factor quantifies the degree of correlation between successive hops and is defined as

$$f = \frac{1 + \frac{1}{N_{\text{hop}}-1} \sum_{i=1}^{N_{\text{hop}}-1} \cos\theta_{i,i+1}}{1 - \frac{1}{N_{\text{hop}}-1} \sum_{i=1}^{N_{\text{hop}}-1} \cos\theta_{i,i+1}}. \tag{3.10}$$

3.2 Calculation details

The initial Li_6PS_5Cl unit cell was obtained from the Materials Project database (mp-985592) [67]. A $2\times2\times5$ supercell was constructed to create the undoped reference structure. In this structure file, all lithium atoms are initially placed in the $24g$ sites, which are the intermediates of the doublet jump; during the equilibration phase, lithium atoms move into the $48h$ sites, as they are able to diffuse. The doped structure was generated by randomly substituting 10% of Cl atoms with BH_4 subunits using ASE [46]. Both structures were geometrically optimized using the FIRE optimizer [62] implemented in ASE with variable cell parameters, using the CHGNet machine-learned potential (MLP) [29].

Each structure was equilibrated through NPT ensemble simulations at temperatures of 400 K, 500 K and 600 K, again using CHGNet and ASE. The equilibration simulations were

conducted for 50 ps with a timestep of 1 fs. Initial velocities were sampled from the Maxwell–Boltzmann distribution corresponding to the target temperature. Temperature and pressure control were achieved using a Berendsen thermostat and inhomogeneous barostat [68], respectively, with an equilibrium pressure of 1 atm. The thermostat time constant was set to 100 fs and the barostat time constant to 1000 fs. The bulk modulus required for the barostat was determined by fitting the Birch–Murnaghan equation of state with pymatgen [69, 70]. Trajectory data were recorded every 10 timesteps during equilibration.

For our final trajectories, dynamics were simulated using NVE ensemble simulations, initialized from the final configurations of the corresponding NPT equilibration runs. Each simulation used a 1 fs timestep and ran for 100 ps, with trajectory data stored every 5 timesteps to capture the lithium ion dynamics with sufficient temporal resolution.

Diffusion coefficients were calculated from the MSD of lithium atoms in each simulation. Linear regression was performed on the MSD as a function time, excluding the first 20 ps. The resulting diffusion coefficients were plotted as a function of inverse temperature in an Arrhenius plot, from which the activation energy and pre-exponential factor were extracted through linear fitting. We then compared these resulting values to those from the motivating experimental work [57] and other computational simulations [49, 71]. Comparison to Ref. [57] required using the Nernst–Einstein equation (Eq. 1.2) to convert their conductivity values to diffusivities. As the experimental densities were not reported, we fit a polynomial model for both the undoped and doped argyrodite volumes as functions of temperature. These density models were fit using the equilibrium cell volumes at each temperature from our computational simulations, taking the form $V = A + BT^3$, where the parameters A and B were adjusted for each composition (undoped: $A = 20\,710.4\,\text{Å}^3$, $B = 0.019\,22\,\text{Å}^3/\text{K}^3$; doped: $A = 20\,781.3\,\text{Å}^3$, $B = 0.019\,62\,\text{Å}^3/\text{K}^3$). The number densities were then given by $n = 480/V$ (as there are 480 lithium atoms in our simulation supercells), resulting temperature-dependent densities were then used in the Nernst–Einstein relation to convert

the experimental conductivity data to diffusivities for direct comparison with our simulation results.

The hopping analysis method described in the previous section was applied to each lithium trajectory. To determine the effect of the threshold parameter, we systematically varied λ and examined the resulting distributions of hop distances $\|\Delta\mathbf{r}_i\|^2$, residence times $\Delta\tau_i$, angular correlations $\cos\theta_{i,i+1}$, and radii of gyration $R_{\mathrm{g},i}$ for each lithium atom. Ridge plots were used to visualize how these distributions evolved with changing λ values. Based on this analysis, λ values of interest were selected for detailed examination of individual lithium trajectories. Finally, the temperature dependence of the averaged hopping parameters was compared between the undoped and doped materials to assess the effect of BH_4^- substitution on lithium ion transport.

3.3 Results and discussion

3.3.1 Analysis by mean squared displacement

In Fig. 3.1, we plot the mean-squared displacement (MSD) of lithium atoms at different temperatures in both our undoped and doped material. These data are from our NVE MD simulations. We see a clear linear trend in all our data. Though the data for our doped materials have larger fluctuations, it still follows a linear trend. We have extracted the diffusivities from each line and plotted them in an Arrhenius plot as functions of the reciprocal temperature (Fig. 3.2). The Arrhenius relationship appears to be obeyed, as evidenced by the linear relationship between the logarithm of the diffusivity and reciprocal temperature.

From our Arrhenius plot, we extract the base diffusivity D_0 and activation energy E_A (Table

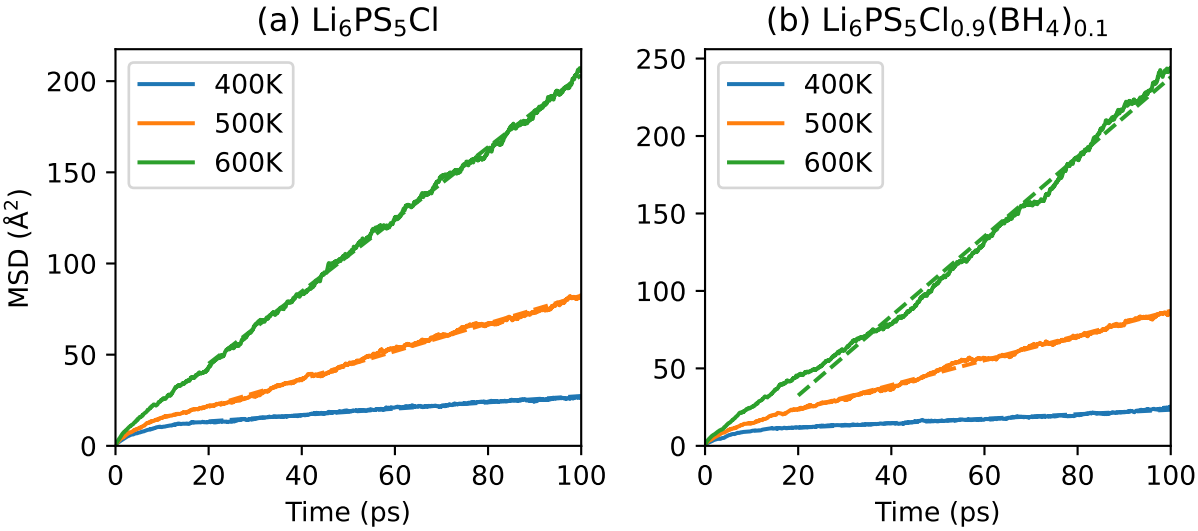


Figure 3.1: Mean-squared displacement (MSD) over time for lithium atoms in (a) Li_6PS_5Cl and (b) $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ at 400 K, 500 K and 600 K. Lines of best fit are calculated from data after 20 ps and are shown as dashed lines.

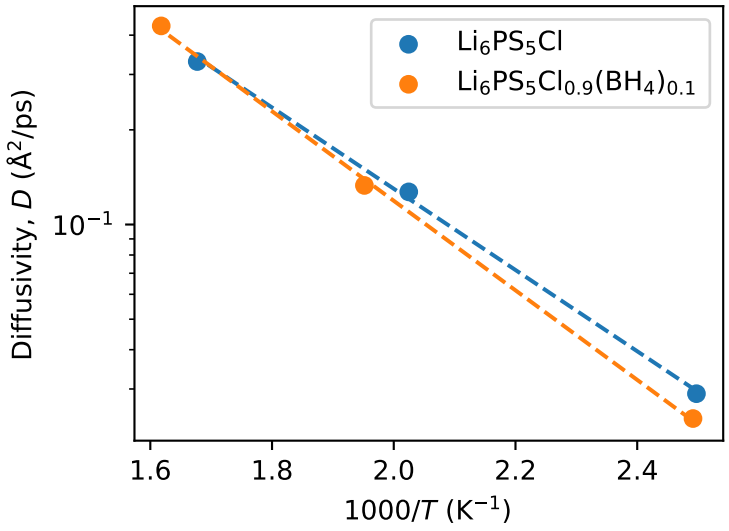


Figure 3.2: Diffusivity of lithium atoms in Li_6PS_5Cl and $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ as functions of reciprocal temperature. Diffusivity is shown on a logarithmic scale.

3.1). In comparison to other computational literature, our results seem reasonable. Li_6PS_5Cl has been previously studied by *ab initio* molecular dynamics, finding D_0 values of $5\,\text{Å}^2/\text{ps}$ [49] and $72\,\text{Å}^2/\text{ps}$ [71], and E_A values of 310 meV [49] and 520 meV [71]. Our use of a machine learned potential could produce a source of bias, but our results are always closer to at least one of these reported values than they are to each other.

We can also compare our results to experimental studies of both the undoped and doped material [57]. General trends hold between our and their results; namely, we both find that the base diffusivity of the doped structure is approximately 1.5 times greater than that of the undoped structure, and that the activation energies for both materials are relatively close. It is unsurprising that the values we find are not identical; each experiment is unique in the precise details of how structures are prepared, computational methods always require some simplification such as neglecting grain boundaries or assumptions regarding the potential. Furthermore, each measurement holds inherent uncertainty—calculations of D_0, in particular, are notoriously noisy.

	D_0 (Å^2/ps)	E_A (meV)
Li_6PS_5Cl	50 ± 11	256 ± 9
$Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$	84 ± 15	282 ± 8

Table 3.1: The base diffusivity D_0 and activation energy E_A for Li_6PS_5Cl and $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$, as found by molecular dynamics.

3.3.2 Analysis by proposed methods

Threshold selection

We now turn to the results from our proposed method of hopping analysis. In Fig. 3.3, we show the distributions of several quantities of interest, as functions of the threshold distance, for $2 \leq \lambda < 10$.

In Fig. 3.3a, we find that the hop length generally centres around one of two modes. We will soon show that this is because, when λ is small, it finds hops between specific sites, and

when λ is large it finds hops between cages. The switch occurs when λ is approximately 4.5 Å. We also note that, though the mode is relatively stable, the prevalence of larger hop lengths leads to an increasing mean.

The distribution of hopping rates (Fig. 3.3b) decreases as λ increases. This is because our method is less likely to identify a hop when the threshold is larger. We also notice that, although the mean of this distribution monotonically decreases with increasing λ, the rate of decrease appears to flatten in in the 5 Å to 7 Å range.

For the cosine angle between consecutive hops (Fig. 3.3c), the distributions themselves don't appear to be particularly helpful. However, we do see a sharp change in the mean of the distributions around $\lambda = 4$ Å. This suggests that the correlation factor $f > 1$ for hops between cages and $f < 1$ for hops between sites. Other methods of analysis have shown that intercage jumps have a significantly higher correlation factor than doublet and intercage jumps, but have not found it to be greater than one [51].

In the distribution of radii of gyrations (Fig. 3.3d), the distribution centres on larger radii of gyration as λ increases between 2 Å and 5 Å. For values of λ between 5 Å and 6.5 Å, the distribution mode is stable, but becomes more positively skewed, increasing its mean. Beyond $\lambda = 7$ Å, the distribution mode shift further right.

We have found two regions for our λ parameter and identified how they should between understood; when λ is small, our method identifies hops between lithium sites, and when λ is large, our method identifies hops between cages. We now our focus our analysis within these two regions.

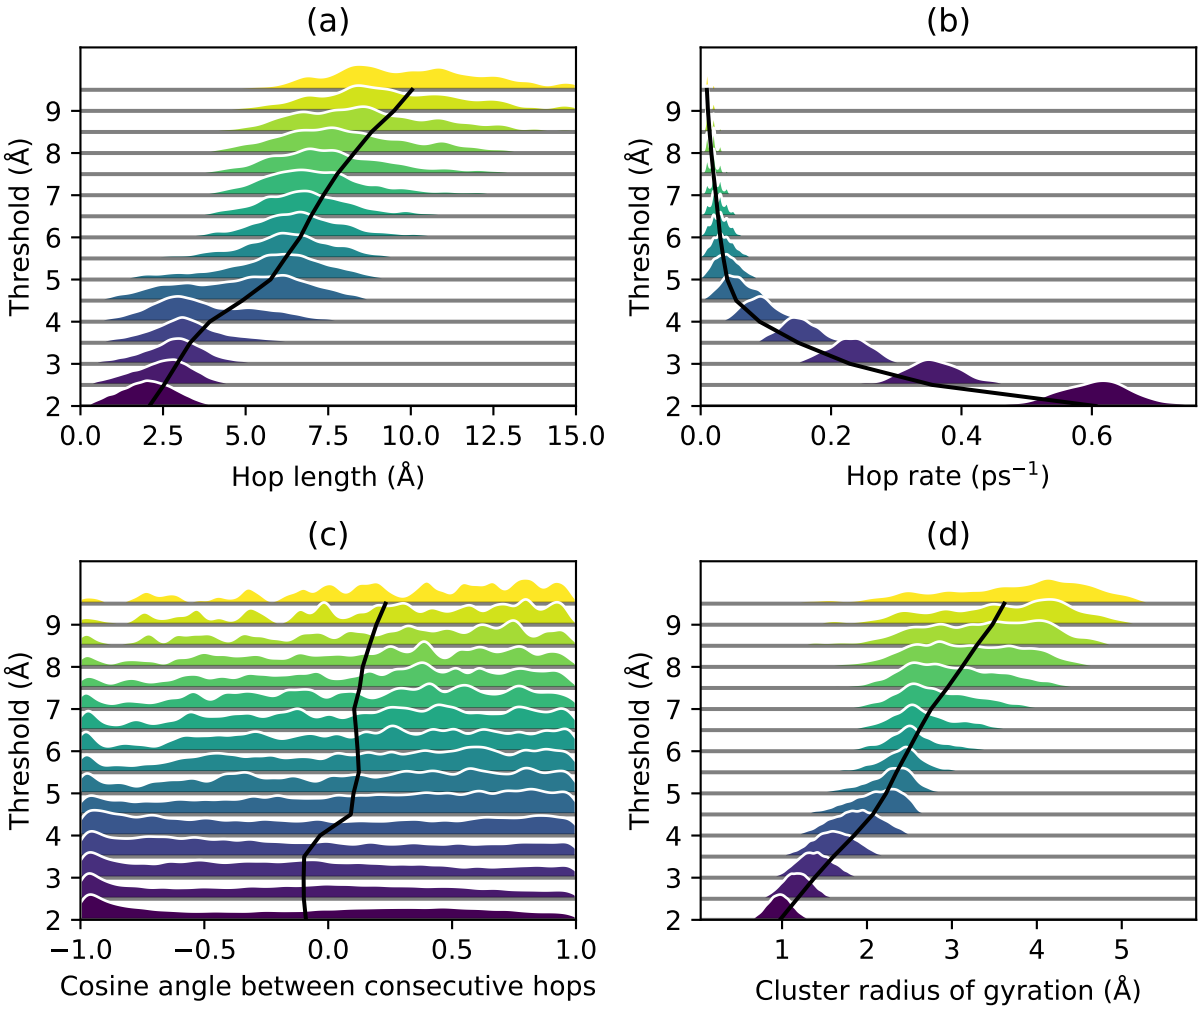


Figure 3.3: Ridge plots showing the effect of the threshold parameter on the observed distributions of (a) hop lengths, (b) hop rates, (c) cosine angles between consecutive hops, and (d) the radius of gyration of sites for Li_6PS_5Cl at 600 K. A solid black line is shown on top of each ridge plot indicating the mean value as a function of the threshold parameter. These plots show a wide range of threshold values, from 2 Å to 9.5 Å.

Hops between cages

We first wish to see whether our method is robust to changes in λ by showing a ridge plot of distributions of interest for a smaller range of λ values (Fig. 3.4). Evidently, the values do not change significantly within this region. We will take $\lambda = 5.5\,\text{Å}$ as our threshold in this subsection. This value was chosen because it is the value for which the mean hop length passes through the mode of its distribution.

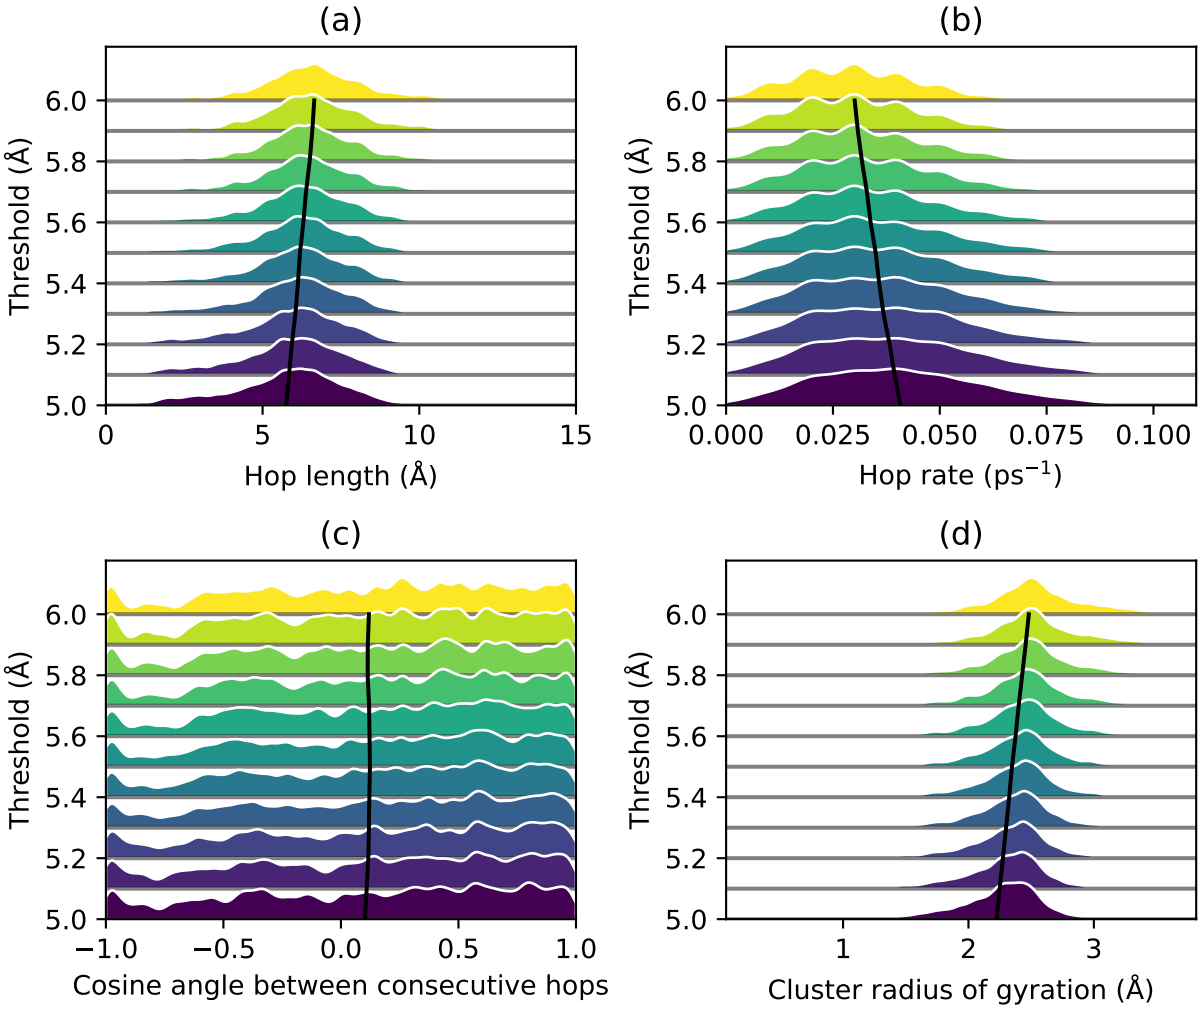


Figure 3.4: Ridge plots showing the effect of the threshold parameter on the observed distributions of (a) hop lengths, (b) hop rates, (c) cosine angles between consecutive hops, and (d) the radius of gyration of sites for Li_6PS_5Cl at 600 K. A solid black line is shown on top of each ridge plot indicating the mean value as a function of the threshold parameter. These plots show a narrow range of large threshold values, from 5 Å to 6 Å.

We show the trajectory of a single lithium atom in space (Fig. 3.5) and as a function of time (Fig. 3.6a–c). Evidently, we see that our method is separating out distinct cages and is working as expected. We also plot the distance from the particle to our current centroid (Fig. 3.6d). There is clearly some noise in this signal, and it gets close to the threshold at several points in time, which are nearly identified as hops. Of course, there is a fine balance between incorrectly detecting a hop and incorrectly ignoring one.

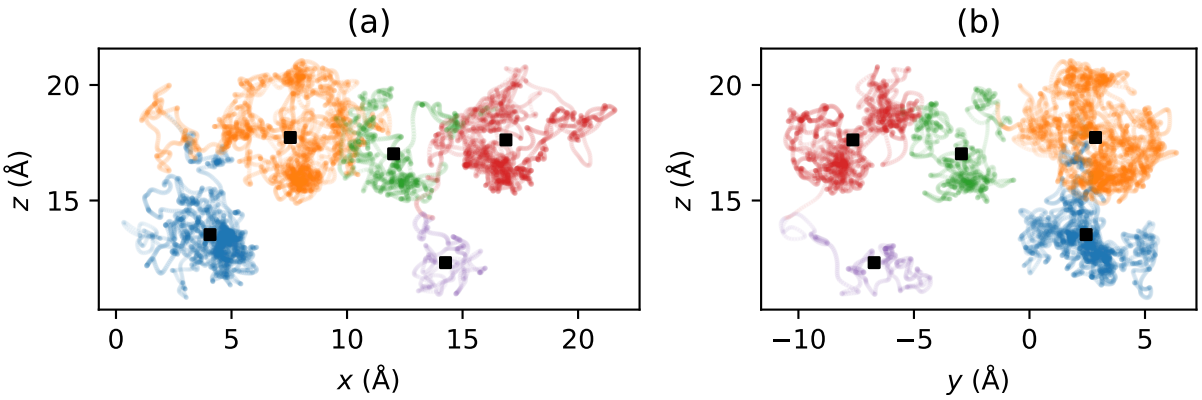


Figure 3.5: Cross-sections of the trajectory of a single lithium atom in Li_6PS_5Cl at 600 K. The colour indicates a change in the identified site. The centre of sites are shown as black squares. Hop detection was performed using our method at a threshold of 5.5 Å.

We then extract the mean parameters of interest and display them as functions of the reciprocal temperature (Fig. 3.7). It seems that the hop length and correlation factor are approximately invariant with temperature, though it is possible that there is a relationship that would only be clear with more data, i.e., statistical tests could be done to verify independence, but our tests would have very little statistical power with only three data points. We also see that the hopping rate and diffusivity follow Arrhenius relationships. Note that the diffusivity here has been calculated by Eq. (1.8), but retains similar values to those found by analysis of the mean squared displacement, thus providing some additional confidence in the results of our method. We also note that the hopping rate found here at 400 K is in close agreement with the intercage hopping rate found by Chen *et al.* [51], extrapolated to no S/Cl anion disorder; both are $3 \times 10^{-3}\,\mathrm{ps}^{-1}$.

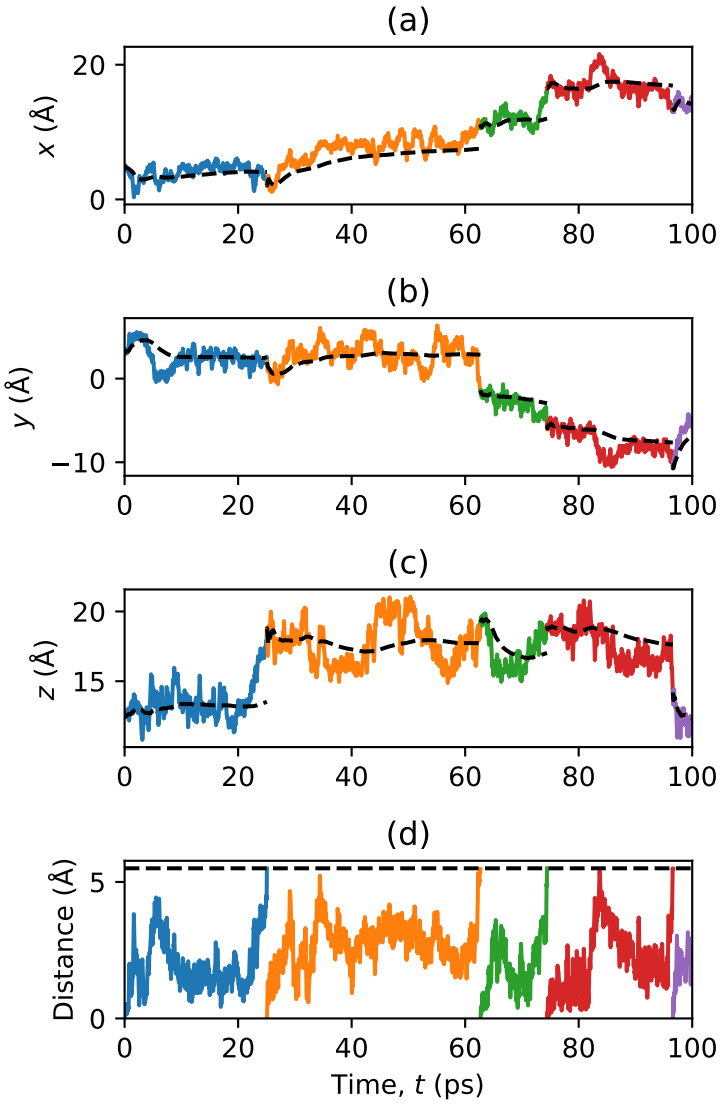


Figure 3.6: (a)–(c) Coordinates of a single lithium atom in Li_6PS_5Cl at 600 K and (d) its distance to the site centroid shown over time. The colour indicates a change in the site, as identified by our method. In (a)–(c), the rolling estimates of the site centroid are shown as dashed lines, and in (d) the black dashed line shows the threshold distance of 5.5 Å.

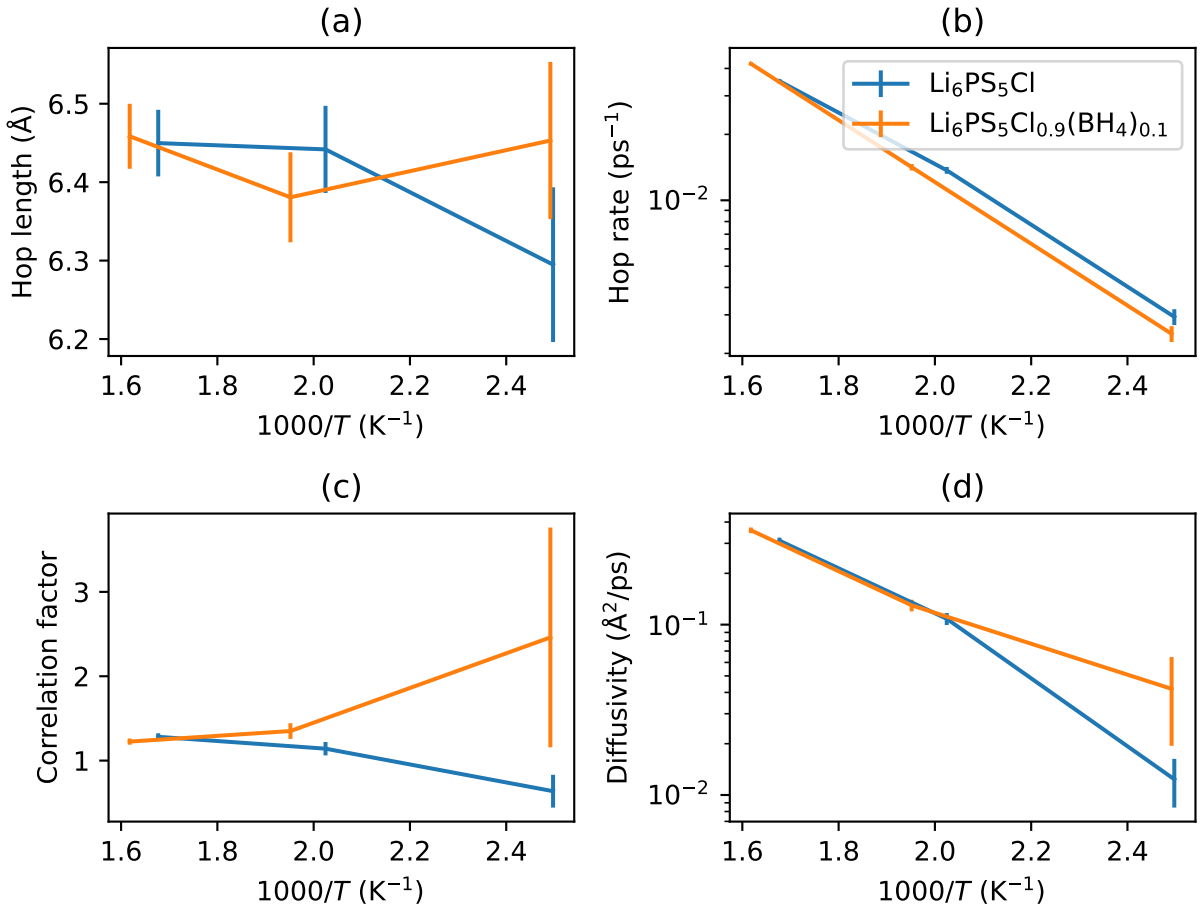


Figure 3.7: Mean and standard error for the (a) hop length, (b) hop rate, (c) correlation factor, and (d) diffusivity for lithium atoms in Li_6PS_5Cl and $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ at different temperatures, as extracted using a threshold distance of 5.5 Å.

Hops between sites

We now turn to the use of a smaller threshold λ, to identify hops between lithium sites rather than between cages. We show the distributions for properties of interest for λ values in the range of 2 Å and 4.5 Å (Fig. 3.8). The distribution for both hop lengths and cosine angles between consecutive hops are both relatively stable, but those for the hop rate and site radius of gyration change drastically in this range. The sensitivity to λ in this case is worrisome. For the sake of completeness, we will attempt to perform analysis in this range with $\lambda = 3.25$ Å.

The trajectory of a single lithium atom in space (Fig. 3.9) for this value of λ are not particularly helpful; it is difficult to separate the sites even by eye. When viewed as a function of time (Fig. 3.10), however, we can see that several moments that appear to correspond to hopping events are not identified as such. This suggests that a lower threshold would be better. However, we also see that the noise level is another consideration; the size of fluctuations in the distance are close to the values that might be wanted as thresholds.

When the mean parameters of interest are extracted and plotted as functions of the reciprocal temperature (Fig. 3.11), an unexpected temperature-dependence is seen for the mean hop length and correlation factor. It is highly unlikely that any 'real' effects are being observed in these plots.

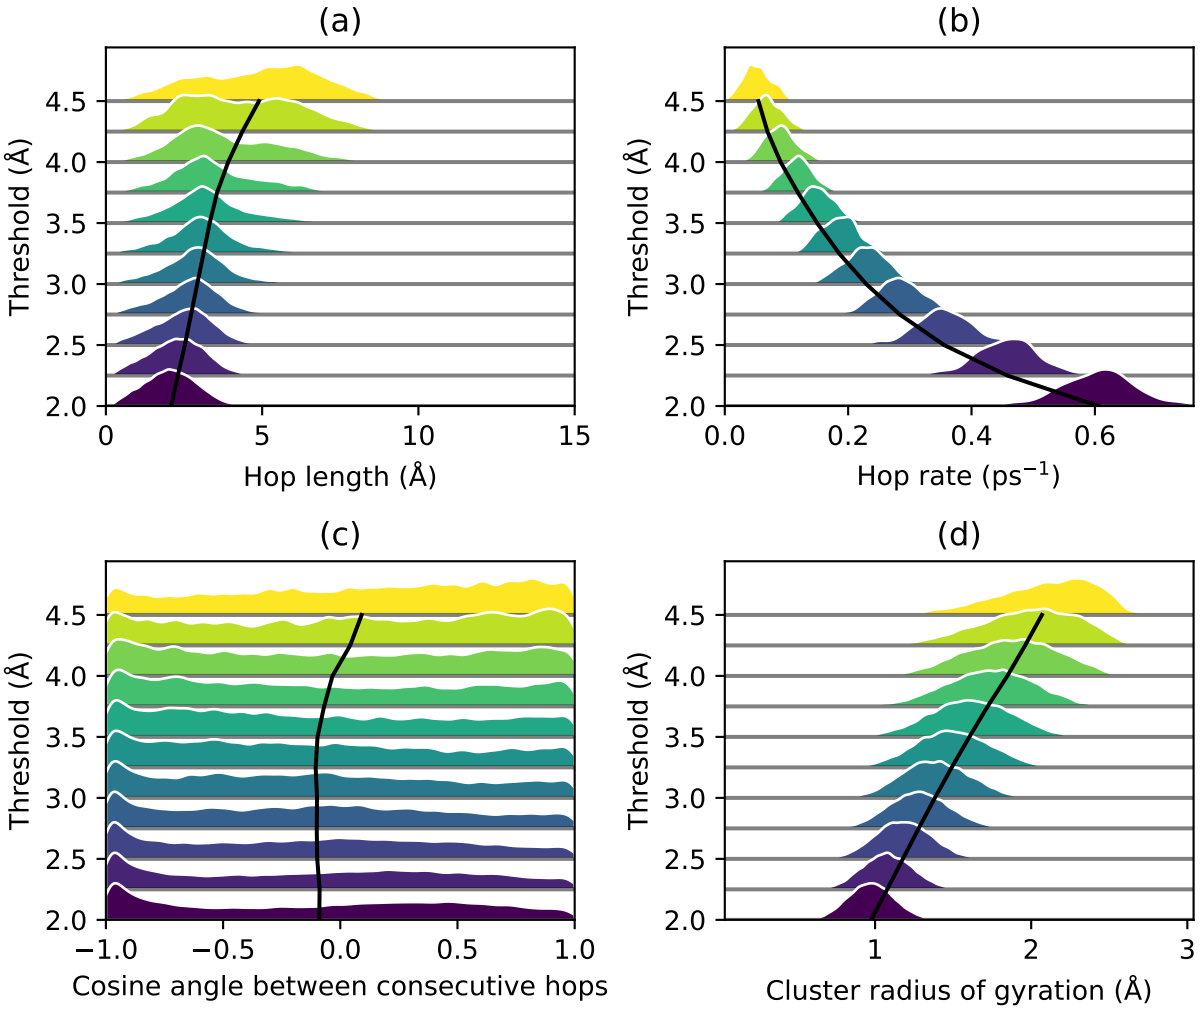


Figure 3.8: Ridge plots showing the effect of the threshold parameter on the observed distributions of (a) hop lengths, (b) hop rates, (c) cosine angles between consecutive hops, and (d) the radius of gyration of sites for Li_6PS_5Cl at 600 K. A solid black line is shown on top of each ridge plot indicating the mean value as a function of the threshold parameter. These plots show a narrow range of small threshold values, from 2 Å to 4.5 Å.

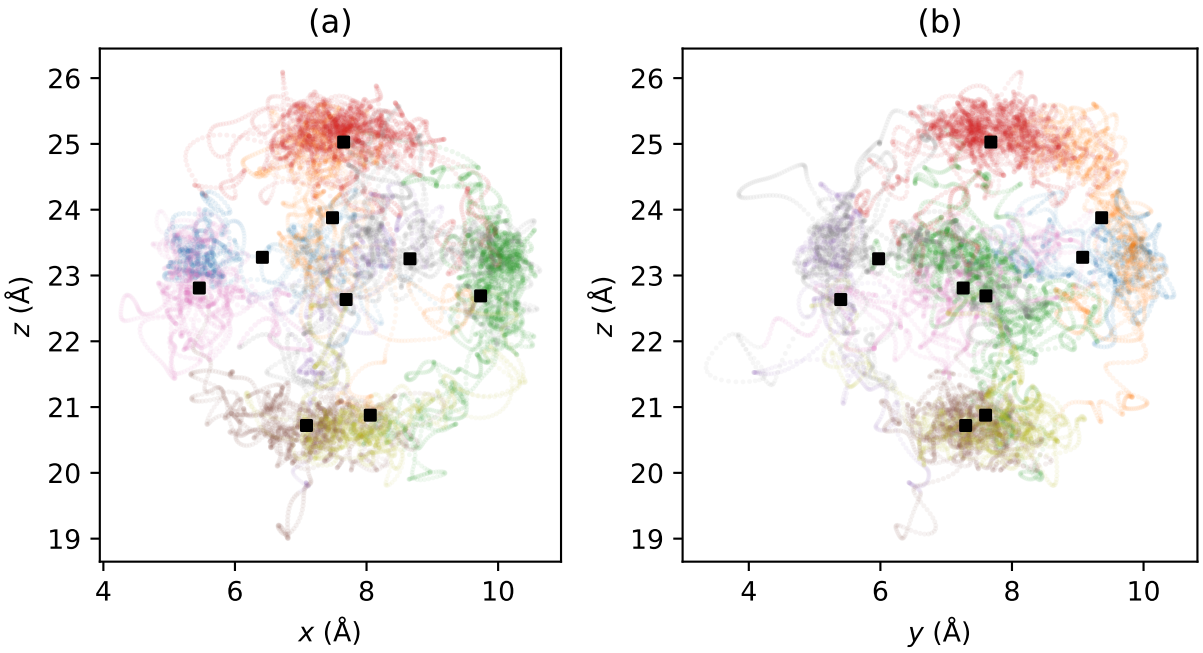


Figure 3.9: Cross-sections of the trajectory of a single lithium atom in $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ at 400 K. The colour indicates a change in the identified site. The centre of sites are shown as black squares. Hop detection was performed using our method at a threshold of 3.25 Å.

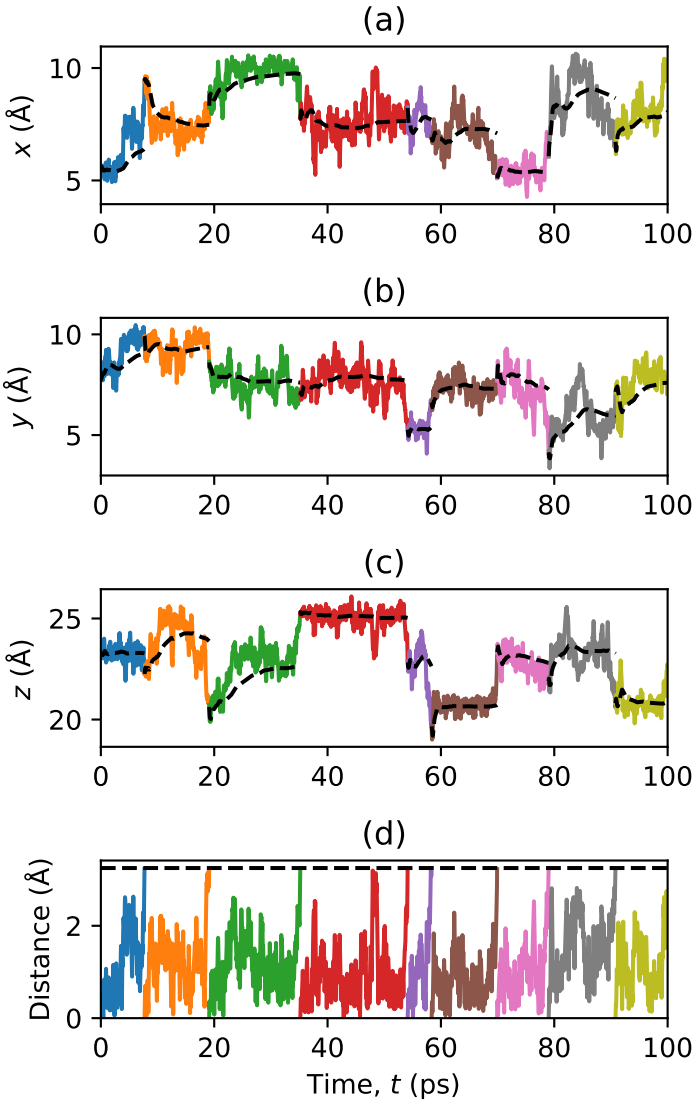


Figure 3.10: (a)–(c) Coordinates of a single lithium atom in $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ at 400 K and (d) its distance to the site centroid shown over time. The colour indicates a change in the site, as identified by our method. In (a)–(c), the rolling estimates of the site centroid are shown as dashed lines, and in (d) the black dashed line shows the threshold distance of 3.25 Å.

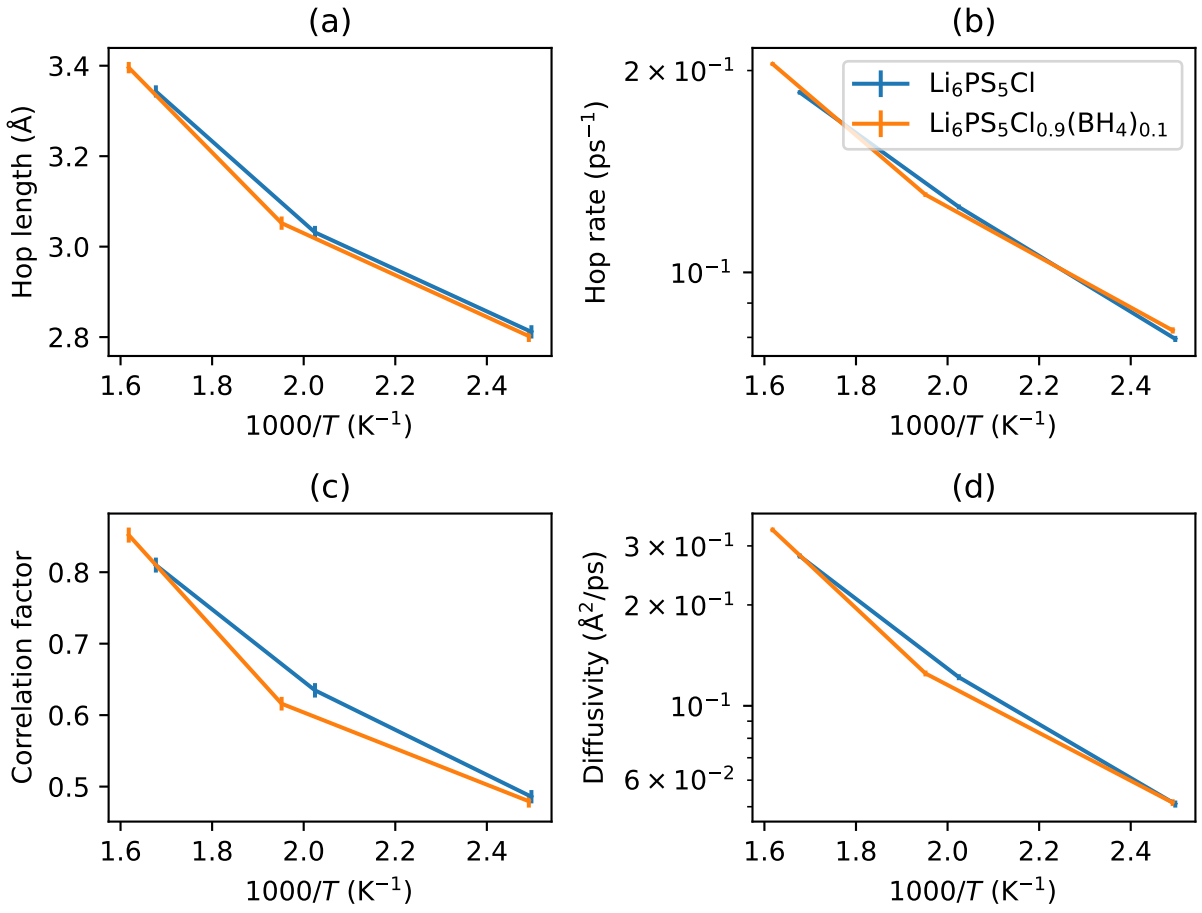


Figure 3.11: Mean and standard error for the (a) hop length, (b) hop rate, (c) correlation factor, and (d) diffusivity for lithium atoms in Li_6PS_5Cl and $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$ at different temperatures, as extracted using a threshold distance of 3.25 Å.

Chapter 4

Conclusions

The simulation of solid-state electrolyte (SSE) materials has allowed researchers to directly observe the migration paths of lithium ions, and so has played a pivotal role in elucidating the mechanisms behind ionic conduction [7]. Performing these simulations requires three steps: Structure determination, simulation, and analysis. Here, we have developed and tested new computational methods to address key challenges in the first and final steps.

Structure determination is especially challenging for amorphous materials—owing to their lack of long-range spatial structure. Reverse Monte Carlo (RMC), which randomly moves atoms in accordance with a Metropolis–Hastings scheme to lower the χ^2 loss, suffers from slow (or even a lack of) convergence and energetically unrealistic structures. We address these by using gradient-based optimization to march atoms directly in the direction that improves fit to experiment, rather than relying on random movement. Our method is available as a Python package, which can be installed from the Python Package Index [72]. The details of its usage are described in Appendix A.

The trajectories of lithium atoms have also been used to identify individual hopping events, from which statistics such as the mean squared hop length, hopping rate, and correlation factor have been determined [32]. We have worked towards developing a method to detect

hopping events without prior knowledge of lithium site locations. This was done by keeping track of a rolling average centroid for each lithium atom $\mathbf{r}_c(t)$, and identifying a hop as the first moment in time where the lithium atom is at a distance greater than λ from its centroid, where λ is some set parameter. We have tested our method on Li_6PS_5Cl and $Li_6PS_5Cl_{0.9}(BH_4)_{0.1}$. We chose these materials for our analysis because the doped structure has had an increased base diffusivity D_0, and our analysis method could be used to explain why, by isolating the factors that affect the base diffusivity.

4.1 Gradient-based structure determination

We tested our method on two systems: Amorphous silicon (a-Si) and a nickel–niobium metallic glass ($Ni_{62}Nb_{38}$). Our method was able to fit the χ^2-loss and potential well in all cases, achieving losses and energies on par or lower than those of molecular dynamics structures. Yet, this does not mean that our structures were necessarily realistic. A good example of this is our optimization from crystalline silicon; though its energies and losses were both in the ranges seen for melt-quench structures, our material still had a strongly crystalline structure. This simply points to how challenging the problem of amorphous structure determination is. Because the pair distribution function does not contain all information about a structure, it's never possible to know for certain whether a found structure is indicative of those used in experiment.

One promising result was that the number of steps required for convergence was fewer than 3,000 steps for all of our tests. In comparison, other hybrid methods have required on the order of 400,000 accepted Monte Carlo steps and 4,000 force calculations for amorphous silicon [42]. Evidently, the use of gradient-based optimization does indeed yield faster convergence times than reverse Monte Carlo. For this reason, it does seem that our method is sufficient for structure refinement [20], as evidenced by our optimization from a melt-quench structure for a-Si; the χ^2-loss was improved dramatically, energy was not changed drasti-

cally, and the optimization converged after only thousands of steps. We also found that our method was well-suited to GPU implementation, which enabled fast computation times per step.

Another general result was that the qualities of our final structure depended on the chosen λ parameter in a predictable and rational way. Larger values of λ resulted in smaller losses, sometimes at the cost of higher energies. This behaviour is expected because larger values of λ give greater prioritization of χ^2 in our optimization. Another key detail is that our results were not strongly sensitive to the specific value of λ; order-of-magnitude changes were needed to see significantly different qualities in the final structure. This makes it relatively simple to use this method in practice; λ could be changed by orders of magnitude until researchers find a structure that has a suitable balance between energy and loss, without the need for careful fine-tuning. We also note that algorithms like the penalty method or augmented Lagrangian method could be used to incrementally increase the importance of reducing χ^2 [73].

We have also shown that our method can generalize to multi-component systems, as it was able to reproduce the individual partial pair distribution functions for $Ni_{62}Nb_{38}$ when given only the total pair distribution function as a target. Interestingly, increased values of λ lead to improved estimation of not just the total pair distribution function, but also the partial pair distribution functions. As the total pair distribution function is a weighted sum of each pair pair distribution function, it is not possible to directly extract each term from the total pair distribution function, and so we would expect the potential to play an important role in determining the structure. However, it seems that the total pair distribution function was more important than the potential.

Another important takeaway is that the initial structure plays a very important role for achieving realistic final structures. This is clear in comparing the resulting structures for optimizing starting from crystalline silicon compared to starting from a melt-quench config-

uration; the former struggled to escape its crystalline structure, but the latter refined the existing amorphous structure well. For our tests on $Ni_{62}Nb_{38}$, a completely random starting structure was sufficient—likely because a random configuration is a *good guess* for a structure without covalent bonds.

We also observed that beginning from a stable structure would lead to an energy transition during the optimization. The *activation energy* of this transition depended on the configuration and the chosen λ value. When beginning with crystalline silicon, it was about 0.2 eV/atom, corresponding to a temperature of 1550 K, which is close to the melting temperature of silicon. It is still not fully clear what this transition corresponds to. It is unlikely to be the transition into an amorphous state because the dihedral angles were still strongly peaked for our optimization from c-Si, indicating the final structure was still strongly crystalline. We note that no transition was seen when optimizing from a random structure, because its initial energy was already quite high.

Our findings suggest that another approach may be better suited to finding amorphous structures using the gradient-based method. Rather than simultaneously optimizing both the potential V and the χ^2-loss, we could alternate between performing molecular dynamics simulation steps using the potential and performing targeted gradient-based optimization steps on χ^2. This would address two issues with our current approach. Firstly, the direct optimization of V was found to push our configuration towards structures conducive to a lower temperature, which is inconsistent with the target pair distribution function. Secondly, we found that the starting structure is of critical importance, and by alternating between MD simulations and χ^2 optimizations, we are essentially able to try many different starting configurations. Indeed, this approach is more inline with how hybrid RMC methods have operated [14]. We expect that using the gradient-based optimization steps in this fashion will maintain the fast convergence observed here but will increase the accuracy of the final structures. We also note that, in this approach, we would not have a λ parameter, but

would instead have to decide how many steps of molecular dynamics and optimization are performed per cycle.

Another important direction for future work would be to test our method on different materials. Though SSEs have been the motivation for this work, these methods apply generally to any amorphous materials. Further testing could be performed on a multicomponent system that also contains covalent bonds, such as amorphous silica (a-SiO_2). This system is ideal as it is both well understood and challenging to determine the structure of [74]. A more challenging system would be lithium thiophosphate glasses ($Li_2S-P_2S_5$). These materials have been well-characterized because of their application as solid-state electrolytes — this includes experimental studies by nuclear magnetic resonance (NMR) [75] and X-ray total scattering [76], as well as by classical molecular dynamics [25]. They are also more challenging to study because they are composed of multiple covalently bonded polyatomic ions. It would be interesting to see what degree of constraints would be needed to result in reasonable structures. The ultimate challenge for our method, of course, would be to determine the structure of a material that has been resistant to study with RMC. Additionally, further analysis could then be performed on our resulting structures, such as by inspecting their network topology [77] or self-sorted local atomic motifs [78].

4.2 Automated analysis of hopping in solid-state electrolytes

From analysing the mean-squared displacement, we find the base diffusivities are $50\,\text{Å}^2/\text{ps}$ and $84\,\text{Å}^2/\text{ps}$ for the undoped and doped structures, respectively, and that their activation energies are 256 meV and 282 meV. These are similar to the trends in the base diffusivity and activation energy to those of experiment [57]. For instance, we find that the base diffusivity D_0 is approximately 1.5 times larger for the doped structure. Both our results and theirs

show roughly similar values for these activation energies. Our results also seem reasonable in comparison to other computational literature. Li_6PS_5Cl has been previously studied by *ab initio* molecular dynamics, finding D_0 values of $5\,Å^2/ps$ [49] and $72\,Å^2/ps$ [71], and E_A values of 310 meV [49] and 520 meV [71].

We have also shown that our method was able to identify the cage-forming nature of hopping in Argyrodite structures. Further, our results suggest that the correlation factor f is less than 1 for hops between sites and greater than 1 for hops between cages. This indicates, somewhat surprisingly, that when a Li atom moves between cages, it is less likely to move back to its previous cage than it is to move into another one. This is in agreement with the literature; other methods of analysis have shown that intercage jumps have a significantly higher correlation factor than doublet and intercage jumps, but have not found it to be greater than one [51].

We also found that our method was only able to robustly identify hops between cages; the detection of hops between sites was quite sensitive to the value of λ chosen. We extracted the correlation factor f, hopping rate Γ, and mean-squared hop length $\langle a^2 \rangle$ for hops between cages. The changes in $\langle a^2 \rangle$ and f as functions of temperature and between material appeared to be insignificant. The hopping rate Γ followed an Arrhenius relationship. We also note that the hopping rate found here at 400 K for Li_6PS_5Cl is in close agreement with the intercage hopping rate found by Chen *et al.* [51], extrapolated to no S/Cl anion disorder; both are $3 \times 10^{-3}\,ps^{-1}$.

A major limitation of this work is that more data is needed to be able to make stronger conclusions. This would ideally be done by performing more independent simulations at similar temperatures or at different temperatures. Additionally, we have only performed our simulations for a single instance of the doped structure; the placements of BH_4^- subunits were chosen once initially, but more accurate results could be achieved by performing simulations

with many different placements of the BH_4^- subunits. Furthermore, the S/Cl anion disorder has not been taken into account here; this has been shown to increase the diffusivity and is more energetically and entropically stable. We also did not test the limitations of using CHGNet on these structures. This could be addressed by comparing forces calculated by CHGNet to those of *ab initio* molecular dynamics (AIMD) and, if necessary, fine tuning the model with that data.

Appendices

Appendix A

Gradient-Based Structure Determination

A.1 Calculating total scattering spectra

Many experimental methods can be used to study non-crystalline materials, including NMR, X-ray absorption spectroscopy (XAS), and neutron and X-ray total scattering. NMR and XAS are well suited for deriving information about the local environment around atoms of particular elements [79]; however, it seems that local detail is not sufficient by itself to deduce the full structures of amorphous materials [80]. Scattering methods, in contrast, reveal information at all length-scales. For this reason, scattering data has been especially useful for determining the structures of non-crystalline materials [9].

Scattering methods measure how incident particles (e.g., neutrons or X-rays) are diffracted after interacting with a sample. Although the theoretical and experimental details of scattering are outside the scope of this thesis, here we give a brief review of some important results and common notation used to describe the results of neutron and X-ray scattering experiments, based largely on a seminal review written by David A. Keen [9].

A.1.1 Neutron scattering

Neutrons are electrically neutral and so only scatter off the atomic nuclei in a system. The amount of scattering done by a nucleus of element α is quantified by the mean neutron scattering length $\bar{b}_\alpha$, where the average is taken over different isotopes and nuclear spin states. As it is not essential to our discussion, we will disregard the overline and refer simply to b_α. Note that the scattering cross-section is given by $4\pi b_\alpha^2$. These scattering lengths and cross-sections have been tabulated for many atoms and isotopes [81, 82].

In a neutron scattering experiment, one measures the total-scattering structure $F(q)$. This is often transformed into the total pair distribution function by

$$G(r) = \frac{1}{8\pi^3\rho_0} \int_0^\infty 4\pi q^2 F(q) \frac{\sin(qr)}{qr}\, dq, \tag{A.1}$$

where ρ_0 is the overall number density of the system. One can also calculate $F(q)$ from $G(r)$ by the inverse transform

$$F(q) = \rho_0 \int_0^\infty 4\pi r^2 G(r) \frac{\sin(qr)}{qr}\, dr. \tag{A.2}$$

It is helpful to think in terms of the pair distribution function $G(r)$ because it can be directly related to the positions of atoms in the system, and hence can be calculated computationally for a given structure. The total pair distribution function is given by

$$G(r) = \sum_{\alpha=1}^{n} \sum_{\beta=1}^{n} c_\alpha c_\beta b_\alpha b_\beta \left[g_{\alpha\beta}(r) - 1 \right], \tag{A.3}$$

where the sum is over the n elements present in the system, c_α is the composition fraction of element α, b_α is the mean neutron scattering length of nuclei of element α, and $g_{\alpha\beta}(r)$ is the partial pair distribution function between elements α and β.

These partial pair distribution functions are often calculated in terms of a histogram;

$$g_{\alpha\beta}(r) = \frac{n_{\alpha\beta}(r)}{4\pi r^2 \, \Delta r \, c_\beta \rho_0}, \tag{A.4}$$

where ρ_0 is the overall number density and $n_{\alpha\beta}(r)$ gives the number of particles of type β at a distance between r and $r + \Delta r$ from a particle of type α.

The partial pair distribution function can be understood as normalizing the distance histogram $n_{\alpha\beta}(r)$ by the distance histogram that would be seen for a perfectly homogeneous fluid of the same density. As such, $g_{\alpha\beta}(r)$ has peaks corresponding to different nearest-neighbour shells around atoms of the element α. The average coordination number of atoms of type β between distances r_1 and r_2 of an atom of type α is given by

$$c_\beta \rho_0 \int_{r_1}^{r_2} 4\pi r^2 g_{\alpha\beta}(r) \, dr. \tag{A.5}$$

It is clear that $g_{\alpha\beta}(0) = 0$ due to steric hindrance — otherwise, it would indicate that two particles are directly on top of each other. At the other extreme, as r approaches infinity, the system is indistinguishable from a homogeneous fluid (as amorphous systems do not have long-range order); hence, $g_{\alpha\beta}(r \to \infty) = 1$.

As shown above, the total pair distribution function $G(r)$ is an offset weighted sum of the partial pair distribution functions (Eq. A.3). By the limits given on each $g_{\alpha\beta}(r)$, it is clear that $G(0) = -\left(\sum_\alpha c_\alpha b_\alpha\right)^2$ and $G(r \to \infty) = 0$. It is sometimes desired for $G(r)$ to be given the same limiting behaviour as each $g_{\alpha\beta}(r)$, in which case it can be rescaled to

$$G'(r) = 1 + \frac{G(r)}{\left(\sum_{\alpha=1}^{n} c_\alpha b_\alpha\right)^2}. \tag{A.6}$$

Correspondingly, the normalized total-scattering structure factor is

$$S(q) = 1 + \frac{F(q)}{\left(\sum_{\alpha=1}^{n} c_\alpha b_\alpha\right)^2}. \tag{A.7}$$

A.1.2 X-ray scattering

X-ray scattering is conceptually very similar to neutron scattering; however, rather than scattering off atomic nuclei, X-rays scatter off electrons, which are much more diffuse. As such, the scattering lengths are no longer constant, but are dependent on the scattering vector $\mathbf{q}$. These atomic form factors are given by

$$f(\mathbf{q}) = \int_{\mathbb{R}^3} \rho(\mathbf{r})\ e^{i\mathbf{q}\cdot\mathbf{r}}\ d^3\mathbf{r}, \tag{A.8}$$

where $\rho(\mathbf{r})$ is the electron density. In general, the atomic form factor will depend on the environment of each atom, but can be approximated as the sum of four isotropic Gaussians with a constant shift;

$$f(q) = \sum_{i=1}^{4} a_i \exp\left[-b_i \left(\frac{q}{4\pi}\right)^2\right] + c. \tag{A.9}$$

The parameters a_i, b_i, and c have been tabulated for many atoms at different oxidation states [63, 83].

The X-ray total-scattering structure factor is

$$F^{\mathrm{X}}(q) = \rho_0 \sum_{\alpha=1}^{n}\sum_{\beta=1}^{n} c_\alpha c_\beta f_{\alpha\beta}(q) \int_0^\infty 4\pi r^2 \left[g_{\alpha\beta}(r) - 1\right] \frac{\sin(qr)}{qr}\ dr, \tag{A.10}$$

where

$$f_{\alpha\beta}(q) = \frac{f_\alpha(q) f_\beta(q)}{\left[\sum_{\gamma=1}^{n} c_\gamma f_\gamma(q)\right]^2}, \tag{A.11}$$

and $f_\alpha(q)$ is the atomic form factor for atoms of element α. The X-ray total pair distribution

function is

$$G^{\mathrm{X}}(r) = \frac{1}{8\pi^3\rho_0}\int_0^\infty 4\pi q^2 F^{\mathrm{X}}(q)\frac{\sin(qr)}{qr}\,dq. \tag{A.12}$$

Notice that, because of the q-dependence of the form factors, $G^{\mathrm{X}}(r)$ must be calculated as a double integral, which compounds numerical errors. To avoid this, it can be advantageous to make the approximation that all form factors are different scalings of the same fundamental curve, i.e.,

$$f_\alpha(q) \approx K_\alpha f_{\mathrm{e}}(q), \tag{A.13}$$

where $K_\alpha = f_\alpha(0)$ is the effective number of electrons around atoms of species α, and

$$f_{\mathrm{e}}(q) = \frac{\sum_{\alpha=1}^n c_\alpha f_\alpha(q)}{\sum_{\alpha=1}^n c_\alpha K_\alpha} \tag{A.14}$$

is the average scattering factor per electron. Under this approximation,

$$f_{\alpha\beta}(q) \approx \frac{K_\alpha K_\beta}{\left(\sum_{\gamma=1}^n c_\gamma K_\gamma\right)^2} \tag{A.15}$$

becomes independent of q, and we have that

$$G^{\mathrm{X}}(r) \approx \sum_{\alpha=1}^n\sum_{\beta=1}^n c_\alpha c_\beta \frac{K_\alpha K_\beta}{\left(\sum_{\gamma=1}^n c_\gamma K_\gamma\right)^2}\left[g_{\alpha\beta}(r) - 1\right]. \tag{A.16}$$

Now, $G^{\mathrm{X}}(r)$ can be calculated directly from the partial pair distribution functions, without performing an integral — though the validity of this approximation will be system-dependent.

A.2 Using the `gbsd` package

Our method has been made available on the Python Package Index [72].

```
pip install gbsd
```

Here we describe its basic usage. Begin by importing useful packages and functions.

```
import torch
from ase.optimize import FIRE

from gbsd.calculator import GBSDCalculator
from gbsd.losses import chi_squared
from gbsd.spectra import get_approximate_xray_total_rdf
from gbsd.structure import generate_random_closest_approach
```

We can create an initial structure (or could read one from a file using ASE).

```
element_counts = {'Si': 1_000}
number_density = 0.0496 # per cubic Angstrom
closest_approach = 2.05 # Angstrom

atoms = generate_random_closest_approach(
    element_counts,
    number_density,
    closest_approach,
)
```

We must load our experimental data to perform optimizations against, and define our loss function (or goodness-of-fit).

```
R_BINS = torch.load('./data/r_bins.pt', weights_only=True)
G_EXP = torch.load('./data/g_exp.pt', weights_only=True)
G_UNC = torch.load('./data/g_unc.pt', weights_only=True)

def get_loss(
    positions: torch.Tensor,
    chemical_symbols: list[str],
    cell: torch.Tensor,
) -> torch.Tensor:
    g_atoms = 1 + get_approximate_xray_total_rdf(
        chemical_symbols,
        positions,
        cell,
        R_BINS,
        kernel_width = 0.03, # Angstrom
    )
```

```
    return chi_squared(g_atoms, G_EXP, G_UNC)
```

We then assign the calculator for our atoms object and perform optimization.

```
atoms.calc = GBSDCalculator(get_loss)

optimizer = FIRE(atoms)
optimizer.run(steps=2_500)
```

This is a very minimal example, focusing only on the basic aspects of our package. This can be extended using mixed calculators, callbacks, dynamics, and more [46]. It is also trivial to perform calculations on the GPU by specifying the device when loading data and creating the `GBSDCalculator` object.

Appendix B

Analysis of Hopping in Solid-State Electrolytes

B.1 Diffusion as a macroscopic phenomenon

B.1.1 The diffusion equation

At a macroscopic scale, diffusion manifests as the evolution of some concentration field over time. Let $c(\mathbf{r}, t)$ denote the concentration of diffusing particles. Fick's first law states that the the diffusive flux $\mathbf{j}$ is proportional to the concentration gradient by

$$\mathbf{j} = -\mathbf{D}\nabla c, \tag{B.1}$$

where $\mathbf{D}$ is the diffusivity tensor, which can be represented with a $d \times d$ matrix [33, 84]. Here, we assume that diffusivity is homogeneous, but may be anisotropic, i.e., diffusion may be faster in some directions than others, but does not vary as a function of position.

For the total concentration to be conserved, it must obey the continuity equation [33, 84],

$$\partial_t c = -\nabla \cdot \mathbf{j}. \tag{B.2}$$

This equation simply relates the diffusive flux to the change in concentration, by its definition. Combining Eqs. (B.1) and (B.2) yields the diffusion equation [33, 84],

$$\partial_t c = \nabla \cdot (\mathbf{D} \nabla c). \tag{B.3}$$

In Appendix B, we show that the solution to the diffusion equation for a single point source is given by a multivariate Gaussian distribution centred at the origin with covariance matrix $2\mathbf{D}t$,

$$c(\mathbf{r}, t) = \frac{1}{(2\pi)^{d/2} \det(2\mathbf{D}t)^{1/2}} \exp\left[-\frac{1}{2} \mathbf{r}^{\mathrm{T}} \left(2\mathbf{D}t\right)^{-1} \mathbf{r} \right]. \tag{B.4}$$

We now discuss key properties of the diffusivity tensor that will greatly simplify calculations.

B.1.2 Diffusivity

For entropy to increase, the diffusivity tensor must have only positive eigenvalues. If this were not the case, there would exist a concentration gradient that causes the diffusive flux to align with the gradient itself, which would amplify the concentration gradient rather than diminish it. It can also be shown, via the Onsager reciprocal relations, that the diffusivity tensor must be symmetric [33, 84]. As a result of these facts, the diffusivity tensor must be positive-definite. This implies that there exists some choice of *principal coordinates* for which the diffusivity tensor is diagonal;

$$\mathbf{D} = \mathrm{diag}(D_1, D_2, ..., D_d). \tag{B.5}$$

Using this coordinate frame greatly simplifies the process of solving the diffusion equation.

The average diffusivity is defined as

$$D = \frac{1}{d} \mathrm{Tr}(\mathbf{D}). \tag{B.6}$$

This definition follows clearly from Eq. (B.5). As the trace is invariant to coordinate rotations, so is the average diffusivity. This notion of the average diffusivity can be made more rigorous by considering an integral over all rotation matrices,

$$\langle \mathbf{D} \rangle = \frac{\int_{\mathrm{SO}(d)} \mathbf{R}^{\mathrm{T}} \mathbf{D} \mathbf{R} \, \mathrm{d}\mu(\mathbf{R})}{\int_{\mathrm{SO}(d)} \mathrm{d}\mu(\mathbf{R})}, \tag{B.7}$$

where μ is the Haar measure on $\mathrm{SO}(d)$ [85]. Taking the trace of both sides yields $\mathrm{Tr}(\langle \mathbf{D} \rangle) = \mathrm{Tr}(\mathbf{D})$. This means that the average diffusivity is equal to the average diffusivity of the isotropically averaged diffusion tensor. Evidently, Eq. (B.6) is a natural choice for the average diffusivity.

B.1.3 Solving the diffusion equation

We can solve the diffusion equation by considering the Fourier transform of the concentration [33, 84]. In d dimensions, the non-unitary Fourier transform and inverse Fourier transform are

$$\tilde{c}(\mathbf{k}, t) = \int_{\mathbb{R}^d} c(\mathbf{r}, t) \; e^{-i\mathbf{k}\cdot\mathbf{r}} \, \mathrm{d}\mathbf{r}, \text{ and} \tag{B.8}$$

$$c(\mathbf{r}, t) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \tilde{c}(\mathbf{k}, t) \; e^{i\mathbf{k}\cdot\mathbf{r}} \, \mathrm{d}\mathbf{k}. \tag{B.9}$$

Substituting Eq. (B.9) into Eq. (B.3) gives

$$\frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \partial_t \tilde{c}(\mathbf{k}, t) \; e^{i\mathbf{k}\cdot\mathbf{r}} \, \mathrm{d}\mathbf{k} = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \tilde{c}(\mathbf{k}, t) \; \nabla \cdot \left(\mathbf{D} \nabla e^{i\mathbf{k}\cdot\mathbf{r}} \right) \, \mathrm{d}\mathbf{k} \tag{B.10}$$

The right-hand side can be simplified by noting that

$$\nabla e^{i\mathbf{k}\cdot\mathbf{r}} = i\mathbf{k} e^{i\mathbf{k}\cdot\mathbf{r}}, \tag{B.11}$$

and that

$$\nabla \cdot (i\mathbf{D}\mathbf{k}e^{i\mathbf{k}\cdot\mathbf{r}}) = -\mathbf{k}^{\mathrm{T}}\mathbf{D}\mathbf{k}e^{i\mathbf{k}\cdot\mathbf{r}}. \tag{B.12}$$

Additionally, for the above integrals to be equal for all functions $\tilde{c}(\mathbf{k}, t)$, it must be the case that their integrands are equal. Therefore,

$$\partial_t \tilde{c}(\mathbf{k}, t) = -\mathbf{k}^{\mathrm{T}}\mathbf{D}\mathbf{k}\ \tilde{c}(\mathbf{k}, t). \tag{B.13}$$

This is now effectively an ordinary differential equation, which can be solved directly as

$$\tilde{c}(\mathbf{k}, t) = \tilde{c}(\mathbf{k}, 0) \exp\left(-\mathbf{k}^{\mathrm{T}}\mathbf{D}\mathbf{k}\ t\right). \tag{B.14}$$

This is in terms of our initial condition $\tilde{c}(\mathbf{k}, 0)$. We are interested in the case of a single particle starting at the origin. In this case, the initial condition in real space is

$$c(\mathbf{r}, 0) = \delta(\mathbf{r}), \tag{B.15}$$

where $\delta(\mathbf{r})$ is the Dirac delta function. In reciprocal space, this condition becomes

$$\tilde{c}(\mathbf{k}, 0) = \int_{\mathbb{R}^d} \delta(\mathbf{r})\ e^{-i\mathbf{k}\cdot\mathbf{r}}\ \mathrm{d}\mathbf{r} = 1. \tag{B.16}$$

Therefore,

$$\tilde{c}(\mathbf{k}, t) = \exp\left(-\mathbf{k}^{\mathrm{T}}\mathbf{D}\mathbf{k}\ t\right). \tag{B.17}$$

Now, using Eq. (B.9), we find that

$$c(\mathbf{r}, t) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \exp\left(-\mathbf{k}^{\mathrm{T}}\mathbf{D}\mathbf{k}\ t\right) e^{i\mathbf{k}\cdot\mathbf{r}}\ \mathrm{d}\mathbf{k}. \tag{B.18}$$

This is the well-known inverse Fourier transform of a multivariate Gaussian distribution,

which gives

$$c(\mathbf{r},t) = \frac{1}{(2\pi)^{d/2}\det(2\mathbf{D}t)^{1/2}} \exp\left[-\frac{1}{2}\mathbf{r}^{\mathrm{T}}\left(2\mathbf{D}t\right)^{-1}\mathbf{r}\right], \tag{B.19}$$

a multivariate Gaussian distribution centred at the origin with covariance matrix $2\mathbf{D}t$.

B.2 Diffusion as a microscopic phenomenon

B.2.1 The theory of random walks

In the solid-state, diffusion occurs through hops between distinct sites, and so can be accurately modelled as a random walk [33]. Let $\mathbf{r}_n$ be the position of a particle after n hops, and let t_n be the time at which the nth hop occurred.

Choosing our coordinates such that our particle is initially at the origin, we can write

$$t_n = \sum_{i=1}^{n} \Delta t_i, \text{ and} \tag{B.20}$$

$$\mathbf{r}_n = \sum_{i=1}^{n} \Delta \mathbf{r}_i, \tag{B.21}$$

where Δt_i and $\Delta \mathbf{r}_i$ are the time taken and displacement moved for the ith hop. These are random variables, and supposing that they are independent and identically distributed (in the most ideal case), we can invoke the central limit theorem to say that both t_n and $\mathbf{r}_n$ are approximately normally distributed for large values of n.

The expected value for t_n, i.e., the expected amount of time for n hops to occur, is then simply

$$\mathbb{E}[t_n] = n\,\mathbb{E}[\Delta t_i] = n/\Gamma, \tag{B.22}$$

where

$$\Gamma = 1/\,\mathbb{E}[\Delta t_i] \tag{B.23}$$

is average hopping rate. The standard deviation in t_n is given by

$$\mathrm{Std}[t_n] = \sqrt{n}\,\mathrm{Std}[\Delta t_i]. \tag{B.24}$$

Note that the standard deviation scales as $\mathcal{O}\left(\sqrt{n}\right)$ whereas the mean scales as $\mathcal{O}(n)$. For large values of n, i.e., after many hops, the distribution for t_n becomes sharp around its mean. As such, we will simply write $t_n = n/\Gamma$.

The average position of a particle after n hops is

$$\mathbb{E}[\mathbf{r}_n] = n\,\mathbb{E}[\Delta\mathbf{r}_i]. \tag{B.25}$$

However, it must be the case that $\mathbb{E}[\Delta\mathbf{r}_i] = 0$; otherwise, the particle would drift in one specific direction over time. As such, we have that

$$\mathbb{E}[\mathbf{r}_n] = 0. \tag{B.26}$$

The covariance matrix for $\mathbf{r}_n$ is given by

$$\mathrm{Var}[\mathbf{r}_n] = n\,\mathrm{Var}[\Delta\mathbf{r}_i] = n\boldsymbol{\Sigma}, \tag{B.27}$$

where $\boldsymbol{\Sigma} = \mathrm{Var}[\Delta\mathbf{r}_i] = \mathbb{E}[\Delta\mathbf{r}_i\Delta\mathbf{r}_i^{\mathrm{T}}]$ is the covariance matrix for $\Delta\mathbf{r}_i$. Substituting n for Γt, we have

$$\mathrm{Var}[\mathbf{r}_n] = \Gamma\boldsymbol{\Sigma}t \tag{B.28}$$

B.2.2 Relating random walks to diffusivity

As $\mathbf{r}_n$ follows a multivariate normal distribution with mean 0 and covariance matrix $n\boldsymbol{\Sigma} = \Gamma\boldsymbol{\Sigma}t$ (for large n), it should be distributed with density

$$c(\mathbf{r}_n) = \frac{1}{(2\pi)^{d/2}\det(\Gamma t\boldsymbol{\Sigma})^{1/2}} \exp\left[-\frac{1}{2}\mathbf{r}_n^{\mathrm{T}}\left(\Gamma t\boldsymbol{\Sigma}\right)^{-1}\mathbf{r}_n\right]. \tag{B.29}$$

Comparing this with Eq. (B.4), we see that

$$\mathbf{D} = \frac{1}{2}\Gamma\boldsymbol{\Sigma}. \tag{B.30}$$

Further, we can find the average diffusivity with Eq. (B.6);

$$D = \frac{1}{2d}\Gamma a^2, \tag{B.31}$$

where

$$a^2 = \mathbb{E}[\Delta\mathbf{r}_i^{\mathrm{T}}\Delta\mathbf{r}_i] \tag{B.32}$$

is the mean-squared hop length.

In the above derivation, we have assumed that all hops are independent. Accounting for correlations, one can write

$$D = \frac{1}{2d}\Gamma f a^2, \tag{B.33}$$

where

$$f = \frac{1+\mathbb{E}[\cos\theta_{i,i+1}]}{1-\mathbb{E}[\cos\theta_{i,i+1}]}, \tag{B.34}$$

is the correlation factor or Haven ratio and $\theta_{i,i+1}$ is the angle between $\Delta\mathbf{r}_i$ and $\Delta\mathbf{r}_{i+1}$ [33, 34].

Bibliography

[1]D. Larcher and J.-M. Tarascon, "Towards greener and more sustainable batteries for electrical energy storage", Nature Chemistry **7**, 10.1038/nchem.2085 (2015).

[2]G. Zubi, R. Dufo-López, M. Carvalho, and G. Pasaoglu, "The lithium-ion battery: State of the art and future perspectives", Renewable and Sustainable Energy Reviews **89**, 10.1016/j.rser.2018.03.002 (2018).

[3]T. Schmaltz, F. Hartmann, T. Wicke, L. Weymann, C. Neef, and J. Janek, "A roadmap for solid-state batteries", Advanced Energy Materials **13**, 10.1002/aenm.202301886 (2023).

[4]Z. A. Grady, C. J. Wilkinson, C. A. Randall, and J. C. Mauro, "Emerging role of non-crystalline electrolytes in solid-state battery research", Frontiers in Energy Research **8**, 10.3389/fenrg.2020.00218 (2020).

[5]J. C. Bachman, S. Muy, A. Grimaud, H.-H. Chang, N. Pour, S. F. Lux, O. Paschos, F. Maglia, S. Lupart, P. Lamp, L. Giordano, and Y. Shao-Horn, "Inorganic solid-state electrolytes for lithium batteries: mechanisms and properties governing ion conduction", Chemical Reviews **116**, 10.1021/acs.chemrev.5b00563 (2016).

[6]A. J. Bard, L. R. Faulkner, and H. S. White, *Electrochemical Methods, 3rd Ed.* (John Wiley & Sons, 2022).

[7]A. M. Nolan, Y. Zhu, X. He, Q. Bai, and Y. Mo, "Computation-accelerated design of materials and interfaces for all-solid-state lithium-ion batteries", Joule **2**, 10.1016/j.joule.2018.08.017 (2018).

[8]A. Madanchi, E. Azek, K. Zongo, L. K. Béland, N. Mousseau, and L. Simine, "Is the future of materials amorphous? Challenges and opportunities in simulations of amorphous materials", ACS Physical Chemistry Au **5**, 10.1021/acsphyschemau.4c00063 (2025).

[9]D. A. Keen, "A comparison of various commonly used correlation functions for describing total scattering", Journal of Applied Crystallography **34**, 10.1107/S0021889800019993 (2001).

[10]O. Gereben and L. Pusztai, "Structure of amorphous semiconductors: Reverse Monte Carlo studies on a-C, a-Si, and a-Ge", Physical Review B **50**, 10.1103/PhysRevB.50.14136 (1994).

[11]P. M. Maffettone, W. J. K. Fletcher, T. C. Nicholas, V. L. Deringer, J. R. Allison, L. J. Smith, and A. L. Goodwin, "When can we trust structural models derived from pair distribution function measurements?", Faraday Discussions **255**, 10.1039/D4FD00106K (2025).

[12]R. L. McGreevy and L. Pusztai, "Reverse Monte Carlo Simulation: A New Technique for the Determination of Disordered Structures", Molecular Simulation **1**, 10.1080/08927028808080958 (1988).

[13]R. L. McGreevy, "Reverse Monte Carlo modelling", Journal of Physics: Condensed Matter **13**, 10.1088/0953-8984/13/46/201 (2001).

[14]A. Pandey, P. Biswas, and D. A. Drabold, "Inversion of diffraction data for amorphous materials", Scientific Reports **6**, 10.1038/srep33731 (2016).

[15]G. King, "Reverse Monte Carlo modeling of PDF data", Candian Light Source and Brockhouse Beamlines XRD School, 2022.

[16]M. M. J. Treacy and K. B. Borisenko, "The local structure of amorphous silicon", Science **335**, 10.1126/science.1214780 (2012).

[17]P. Biswas, D. N. Tafen, and D. A. Drabold, "Experimentally constrained molecular relaxation: the case of glassy $GeSe_2$", Physical Review B **71**, 10.1103/PhysRevB.71.054204 (2005).

[18]Q. Zhou, T. Du, L. Guo, M. M. Smedskjaer, and M. Bauchy, "New insights into the structure of sodium silicate glasses by force-enhanced atomic refinement", Journal of Non-Crystalline Solids **536**, 10.1016/j.jnoncrysol.2020.120006 (2020).

[19]Q. Zhou, Y. Shi, B. Deng, T. Du, L. Guo, M. M. Smedskjaer, and M. Bauchy, "Revealing the medium-range structure of glassy silica using force-enhanced atomic refinement", Journal of Non-Crystalline Solids **573**, 10.1016/j.jnoncrysol.2021.121138 (2021).

[20]D. A. Keen, "Refining disordered structural models using reverse monte carlo methods: Application to vitreous silica", Phase Transitions **61**, 10.1080/01411599708223732 (1997).

[21]H. C. Andersen, "Molecular dynamics simulations at constant pressure and/or temperature", The Journal of Chemical Physics **72**, 10.1063/1.439486 (1980).

[22]D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Academic Press, 2002).

[23]R. J. Sadus, *Molecular Simulation of Fluids: Theory, Algorithms, Object-Orientation, and Parallel Computing* (Elsevier, 2023).

[24]R. Iftimie, P. Minary, and M. E. Tuckerman, "*Ab initio* molecular dynamics: Concepts, recent developments, and future trends", Proceedings of the National Academy of Sciences **102**, 10.1073/pnas.0500193102 (2005).

[25]S. Ariga, T. Ohkubo, S. Urata, Y. Imamura, and T. Taniguchi, "A new universal force-field for the Li_2S–P_2S_5 system", Physical Chemistry Chemical Physics **24**, 10.1039/D1CP05393K (2022).

[26]Z. Zhu and Y. Zhu, "Developing classical interatomic potentials for solid electrolytes", Accounts of Materials Research **3**, 10.1021/accountsmr.2c00184 (2022).

[27]J. Behler, "Perspective: Machine learning potentials for atomistic simulations", The Journal of Chemical Physics **145**, 10.1063/1.4966192 (2016).

[28]C. Chen and S. P. Ong, "A universal graph deep learning interatomic potential for the periodic table", Nature Computational Science **2**, 10.1038/s43588-022-00349-3 (2022).

[29]B. Deng, P. Zhong, K. Jun, J. Riebesell, K. Han, C. J. Bartel, and G. Ceder, "CHGNet as a pretrained universal neural network potential for charge-informed atomistic modelling", Nature Machine Intelligence **5**, 10.1038/s42256-023-00716-3 (2023).

[30]Y. Park, J. Kim, S. Hwang, and S. Han, "Scalable parallel algorithm for graph neural network interatomic potentials in molecular dynamics simulations", Journal of Chemical Theory and Computation **20**, 10.1021/acs.jctc.4c00190 (2024).

[31]X. He, Y. Zhu, A. Epstein, and Y. Mo, "Statistical variances of diffusional properties from *ab initio* molecular dynamics simulations", npj Computational Materials **4**, 10.1038/s41524-018-0074-y (2018).

[32]N. J. J. de Klerk, E. van der Maas, and M. Wagemaker, "Analysis of diffusion in solid-state electrolytes through MD simulations, improvement of the Li-ion conductivity in β-Li_3PS_4 as an example", ACS Applied Energy Materials **1**, 10.1021/acsaem.8b00457 (2018).

[33]H. Mehrer, *Diffusion in Solids* (Springer Berlin, 2007).

[34]K. Compaan and Y. Haven, "Correlation factors for diffusion in solids", Transactions of the Faraday Society **52**, 10.1039/TF9565200786 (1956).

[35]A. Heuer, "Exploring the potential energy landscape of glass-forming systems: From inherent structures via metabasins to macroscopic transport", Journal of Physics: Condensed Matter **20**, 10.1088/0953-8984/20/37/373101 (2008).

[36]A. S. Keys, L. O. Hedges, J. P. Garrahan, S. C. Glotzer, and D. Chandler, "Excitations are localized and relaxation is hierarchical in glass-forming liquids", Physical Review X **1**, 10.1103/PhysRevX.1.021013 (2011).

[37]B. J. Morgan, "Mechanistic origin of superionic lithium diffusion in anion-disordered Li_6PS_5X argyrodites", Chemistry of Materials **33**, 10.1021/acs.chemmater.0c03738 (2021).

[38]S. Goedecker, "Minima hopping: An efficient search method for the global minimum of the potential energy surface of complex molecular systems", The Journal of Chemical Physics **120**, 10.1063/1.1724816 (2004).

[39]D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization", arXiv, 10.48550/arXiv.1412.6980 (2014).

[40]M. Ishida and T. Ohba, "Hybrid Reverse Molecular Dynamics Simulation as New Approach to Determination of Carbon Nanostructure of Carbon Blacks", Scientific Reports **10**, 10.1038/s41598-020-60372-0 (2020).

[41]A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, "Automatic Differentiation in Machine Learning: A Survey", Journal of Machine Learning Research **18**, 10.48550/arXiv.1502.05767 (2018).

[42]A. Pandey, P. Biswas, B. Bhattarai, and D. A. Drabold, "Realistic inversion of diffraction data for an amorphous solid: The case of amorphous silicon", Physical Review B **94**, 10.1103/PhysRevB.94.235208 (2016).

[43]V. L. Deringer, N. Bernstein, A. P. Bartók, M. J. Cliffe, R. N. Kerber, L. E. Marbella, C. P. Grey, S. R. Elliott, and G. Csányi, "Realistic atomistic structure of amorphous silicon from machine-learning-driven molecular dynamics", The Journal of Physical Chemistry Letters **9**, 10.1021/acs.jpclett.8b00902 (2018).

[44]Y. Zhang, R. Ashcraft, M. Mendelev, C. Z. Wang, and K. F. Kelton, "Experimental and molecular dynamics simulation study of structure of liquid and amorphous $Ni_{62}Nb_{38}$ alloy", The Journal of Chemical Physics **145**, 10.1063/1.4968212 (2016).

[45]L. J. Lewis, "Fifty years of amorphous silicon models: The end of the story?", Journal of Non-Crystalline Solids **580**, https://doi.org/10.1016/j.jnoncrysol.2021.121383 (2022).

[46]A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, and K. W. Jacobsen, "The atomic simulation environment—a python library for working with atoms", Journal of Physics: Condensed Matter **29**, 10.1088/1361-648X/aa680e (2017).

[47]R. Schlenker, A.-L. Hansen, A. Senyshyn, T. Zinkevich, M. Knapp, T. Hupfer, H. Ehrenberg, and S. Indris, "Structure and Diffusion Pathways in Li_6PS_5Cl Argyrodite from Neutron Diffraction, Pair-Distribution Function Analysis, and NMR", Chemistry of Materials **32**, 10.1021/acs.chemmater.0c02418 (2020).

[48]L. Zhou, N. Minafra, W. G. Zeier, and L. F. Nazar, "Innovative Approaches to Li-Argyrodite Solid Electrolytes for All-Solid-State Lithium Batteries", Accounts of Chemical Research **54**, 10.1021/acs.accounts.0c00874 (2021).

[49]T. Jeon, G. H. Cha, and S. C. Jung, "Understanding the anion disorder governing lithium distribution and diffusion in an argyrodite Li_6PS_5Cl solid electrolyte", Journal of Materials Chemistry A **12**, 10.1039/D3TA06069A (2024).

[50]M. Jang, K. Park, H.-G. Jung, K. Y. Chung, J. H. Shim, O. Kwon, and S. Yu, "Unraveling li-ion transport mechanisms in high-entropy anion-disordered argyrodites via machine-learned interatomic potentials", Journal of Materials Chemistry A **13**, 10.1039/D5TA02205C (2025).

[51]J. Chen, M. Fang, Q. Wu, S. Tang, J. Zheng, C. Wei, X. Cao, Y. Shi, N. Xu, and Y. He, "Insights into the Atomic Mechanism of Lithium-Ion Diffusion in Li_6PS_5Cl via a Machine Learning Potential", Chemistry of Materials **37**, 10.1021/acs.chemmater.4c01152 (2025).

[52]N. J. J. de Klerk, I. Rosłoń, and M. Wagemaker, "Diffusion mechanism of Li argyrodite solid electrolytes for Li-Ion batteries and prediction of optimized halogen doping: The effect of li vacancies, halogens, and halogen disorder", Chemistry of Materials **28**, 10.1021/acs.chemmater.6b03630 (2016).

[53]P. Wang, H. Liu, S. Patel, X. Feng, P.-H. Chien, Y. Wang, and Y.-Y. Hu, "Fast ion conduction and its origin in $Li_{6-x}PS_{5-x}Br_{1+x}$", Chemistry of Materials **32**, 10.1021/acs.chemmater.9b05331 (2020).

[54]K. Momma and F. Izumi, "*VESTA 3* for three-dimensional visualization of crystal, volumetric and morphology data", Journal of Applied Crystallography **44**, 10.1107/S0021889811038970 (2011).

[55]C. Liu, B. Chen, T. Zhang, J. Zhang, R. Wang, J. Zheng, Q. Mao, and X. Liu, "Electron Redistribution Enables Redox-Resistible Li_6PS_5Cl towards High-Performance All-Solid-State Lithium Batteries", Angewandte Chemie International Edition **62**, 10.1002/anie.202302655 (2023).

[56]H. Wang, L. Gao, Z. Lu, Y. Tang, D. Ye, G. Zhao, H. Zhao, and J. Zhang, "Borohydride substitution effects of Li_6PS_5Cl solid electrolyte", ACS Applied Energy Materials **4**, 10.1021/acsaem.1c02892 (2021).

[57]L. Liu, H. Wang, D. Ye, H. Zhao, J. Zhang, and Y. Tang, "Fluorine-like BH_4-doped Li_6PS_5Cl with improved ionic conductivity and electrochemical stability", ACS Applied Materials & Interfaces **16**, 10.1021/acsami.4c05403 (2024).

[58]B. W. Silverman, *Density Estimation for Statistics and Data Analysis* (CRC Press, 2018).

[59]K. Laaziri, S. Kycia, S. Roorda, M. Chicoine, J. L. Robertson, J. Wang, and S. C. Moss, "High resolution radial distribution function of pure amorphous silicon", Physical Review Letters **82**, 10.1103/PhysRevLett.82.3460 (1999).

[60]A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in 't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, and S. J. Plimpton, "LAMMPS — a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales", Computer Physics Communications **271**, 10.1016/j.cpc.2021.108171 (2022).

[61]T. Kumagai, S. Izumi, S. Hara, and S. Sakai, "Development of bond-order potentials that can reproduce the elastic constants and melting point of silicon for classical molecular dynamics simulation", Computational Materials Science **39**, 10.1016/j.commatsci.2006.07.013 (2007).

[62]E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, and P. Gumbsch, "Structural relaxation made simple", Physical Review Letters **97**, 10.1103/PhysRevLett.97.170201 (2006).

[63]Graz Center of Physics, *Atomic form factors*, https://lampz.tugraz.at/~hadley/ss1/crystaldiffraction/atomicformfactors/formfactors.php (visited on 07/05/2025).

[64]C. Chen, D. Chen, and F. Ciucci, "A molecular dynamics study of oxygen ion diffusion in A-site ordered perovskite $PrBaCo_2O_{5.5}$: Data mining the oxygen trajectories", Physical Chemistry Chemical Physics **17**, 10.1039/C4CP05847J (2015).

[65]M. Ester, H.-P. Kriegel, J. Sander, X. Xu, et al., "A density-based algorithm for discovering clusters in large spatial databases with noise", in KDD, Vol. 96, 34 (1996).

[66]M. Ankerst, M. M. Breunig, H.-P. Kriegel, and J. Sander, "OPTICS: Ordering points to identify the clustering structure", in Proceedings of the 1999 ACM SIGMOD International Conference on Management of Data (1999).

[67]A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson, "Commentary: The Materials Project: A materials genome approach to accelerating materials innovation", APL Materials **1**, 10.1063/1.4812323 (2013).

[68]H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, and J. R. Haak, "Molecular dynamics with coupling to an external bath", The Journal of Chemical Physics **81**, 10.1063/1.448118 (1984).

[69]F. Birch, "Finite elastic strain of cubic crystals", Physical Review **71**, 10.1103/PhysRev.71.809 (1947).

[70]S. P. Ong, W. D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V. L. Chevrier, K. A. Persson, and G. Ceder, "Python materials genomics (pymatgen): A robust, open-source python library for materials analysis", Computational Materials Science **68**, 10.1016/j.commatsci.2012.10.028 (2013).

[71]Z. Deng, Z. Zhu, I.-H. Chu, and S. P. Ong, "Data-Driven First-Principles Methods for the Study and Design of Alkali Superionic Conductors", Chemistry of Materials **29**, 10.1021/acs.chemmater.6b02648 (2017).

[72]*Python Package Index — PyPI*, https://pypi.org/ (visited on 07/14/2025).

[73]M. R. Hestenes, "Multiplier and gradient methods", Journal of Optimization Theory and Applications **4**, 10.1007/BF00927673 (1969).

[74]D. A. Keen, "Reverse Monte Carlo refinement of disordered silica phases", in *Local structure from diffraction*, edited by S. J. L. Billinge and M. F. Thorpe (Springer US, 2002), pp. 101–119.

[75]C. Dietrich, D. A. Weber, S. J. Sedlmaier, S. Indris, S. P. Culver, D. Walter, J. Janek, and W. G. Zeier, "Lithium ion conductivity in $Li_2S–P_2S_5$ glasses — building units and local structure evolution during the crystallization of superionic conductors Li_3PS_4, $Li_7P_3S_{11}$ and $Li_4P_2S_7$", Journal of Materials Chemistry A **5**, 10.1039/C7TA06067J (2017).

[76]K. Ohara, A. Mitsui, M. Mori, Y. Onodera, S. Shiotani, Y. Koyama, Y. Orikasa, M. Murakami, K. Shimoda, K. Mori, T. Fukunaga, H. Arai, Y. Uchimoto, and Z. Ogumi, "Structural and electronic features of binary $Li_2S–P_2S_5$ glasses", Scientific Reports **6**, 10.1038/srep21302 (2016).

[77]S. Kohara, M. Shiga, Y. Onodera, H. Masai, A. Hirata, M. Murakami, T. Morishita, K. Kimura, and K. Hayashi, "Relationship between diffraction peak, network topology, and amorphous-forming ability in silicon and silica", Scientific Reports **11**, 10.1038/s41598-021-00965-5 (2021).

[78]A. Rahemtulla, B. Tomberli, and S. Kycia, "Visualization of self-sorted local atomic motifs in disordered solids", MRS Advances **3**, 10.1557/adv.2018.434 (2018).

[79]A. F. Harper, S. P. Emge, P. C. M. M. Magusin, C. P. Grey, and A. J. Morris, "Modelling amorphous materials via a joint solid-state NMR and X-ray absorption spectroscopy and DFT approach: Application to alumina", Chemical Science **14**, 10.1039/D2SC04035B (2023).

[80]A. Di Cicco and F. Iesari, "Advances in modelling X-ray absorption spectroscopy data using reverse Monte Carlo", Physical Chemistry Chemical Physics **24**, 10.1039/D1CP05525A (2022).

[81]V. F. Sears, "Neutron scattering lengths and cross sections", Neutron News **3**, 10.1080/10448639208218770 (1992).

[82]NIST Center for Neutron Research, *Neutron scattering lengths and cross sections*, (2021) https://www.ncnr.nist.gov/resources/n-lengths/ (visited on 07/05/2025).

[83]P. J. Brown, A. G. Fox, E. N. Maslen, M. A. O'Keefe, and B. T. M. Willis, "Chapter 6.1. Intensity of diffracted intensities", in *International tables for crystallography*, Vol. C (Wiley & Sons, Limited, John, 2006) Chap. 6.1.

[84]J. Crank, *The Mathematics of Diffusion, 2nd Ed.* (Oxford University Press, 1975).

[85]P. Diaconis and M. Shahshahani, "On the eigenvalues of random matrices", Journal of Applied Probability **31**, 10.1017/S0021900200106989 (1994).