universität wien

MASTERARBEIT | MASTER'S THESIS

Titel | Title

Microscopic mechanism of the superionic transition of copper sulfide

verfasst von | submitted by

Besnik Deliu BSc

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien | Vienna, 2025

Studienkennzahl lt. Studienblatt | degree programme code as it appears on the student record sheet:	UA 066 876
Studienrichtung lt. Studienblatt | degree programme as it appears on the student record sheet:	Masterstudium Physics
Betreut von | Supervisor:	Univ.-Prof. Mag. Dr. Christoph Dellago

Inhaltsverzeichnis

Acknowledgements

I would like to express my special thanks to Univ.-Prof. Mag. Dr. Christoph Dellago for providing the opportunity to write a master's thesis within his research group. I would also like to thank Dott. Mag. Salvatore Romano, Pablo Montero de Hijes, PhD. and Alessandro Coretti, PhD. for helpful discussions, for their assistance in introducing me to the world of computational physics and for helping me with any questions. Additionally, I would like to extend my gratitude to Dr. Andreas Singraber for providing his research data.

Zusammenfassung

Die vorliegende Arbeit zielt darauf ab, den mikroskopischen Mechanismus des superionischen Phasenübergangs von Kupfersulfid Cu_2S zu verstehen, indem der kürzlich veröffentlichte Algorithmus zur Dimensionsreduktion, genannt "Pairwise Controlled Manifold Approximation Projection"(PaCMAP), verwendet wird. Kupfersulfid durchläuft bei etwa 103.5°C einen Phasenübergang von der γ-Phase (Low-Chalkosit, γCh) zur β-Phase (High-Chalkosit, βCh), wobei das Schwefel-Teilgitter seine hexagonale Kristallstruktur beibehält, während die Kupferatome an Mobilität gewinnen, wodurch die β-Phase eine Fest-Flüssig-Hybridphase darstellt. Wir stellen fest, dass der Phasenübergang von Cu_2S nicht über einen lokalisierte Nukleationsprozess erfolgt, sondern durch eine delokalisierte, teilweise synchronisierte Umwandlung der atomaren Kupferumgebungen [2].

Abstract

This thesis aims to understand the microscopic mechanism of the superionic transition of copper sulfide Cu_2S, using the recently published algorithm for dimension reduction called Pairwise Controlled Manifold Approximation Projection (PaCMAP). Copper Sulfide undergoes a phase transition from γ phase (low-chalcocite, γCh) to β phase (high-chalcocite, βCh) at approximately 103.5°C, where the sulfur sublattice keeps its hexagonal crystal structure while the copper atoms gain mobility making the β phase a solid-liquid hybrid phase. We find that the transition from low- to high-chalcocite in Cu_2S does not proceed via a localized nucleation front but instead through a delocalized, partially synchronized transformation of copper atomic environments [2].

1 Introduction

Since the beginning of modern science it has been crucial to completely understand phenomena, physical properties, mechanism etc. in order to make use of them in a way beneficial to our lives. Regarding copper sulfide Cu_2S, which is an intensly studied material (studied for almost two centuries), it has been shown that due to its unique properties, for example the small band gap of 1.2eV [5, 6] and many others, it is capable of being used in photovoltaics, sensors, batteries, biomedical imaging et cetera [5, 6]. However, without fully understanding the material, the use in those areas will be arduous. Dennler et al. showed [5] that in the case of Cu_2S the increasing ion migration (copper mobility) results in stability issues. These problems remain even in the LC phase, since copper atoms still migrate especially when applying electric fields or a temperature gradient. Therefore understanding the microscopic mechanism of the superionic transition of Cu_2S is crucial as already mentioned [5]. To study the microscopic mechanism, we need to introduce methods such as dimension reduction algorithms and properties like the diffusion coefficient. In the first chapter we will introduce some theoretical backround including the properties of Cu_2S and the operating principle of High Dimensional Neural Network Potentials (HDNNPs), in our case the second generation of HDNNP proposed by Behler and Parrinello [8]. In this thesis one relies on the Cu_2S HDNNP trained by Andreas Singraber in order to carry out Molecular Dynamics Simulations (MD-Simulations) [1]. Concerning the analysis of the MD-Simulation we use the recently introduced dimension reduction algorithm PaCMAP to distinguish between certain clusters. In our case, these clusters represent the atomic environment of copper and sulfur during different timesteps.

2 Theoretical Backround

2.1 Copper Sulfide Cu_2S

There exists a lot of copper sulfides $Cu_{2-x}S$, in this thesis we will treat the special stochiometric compound Cu_2S (chalcocite) which has three different phases. The γ-phase (low-chalcocite, γCh) exists for temperatures below 105°C, the β-phase (high-chalcocite, βCh) between 105°C and 425°C and the α-phase for temperatures above 425°C. They phase transition between low-chalcocite and high-chalcocite happens at approximately 103.5°C [2].

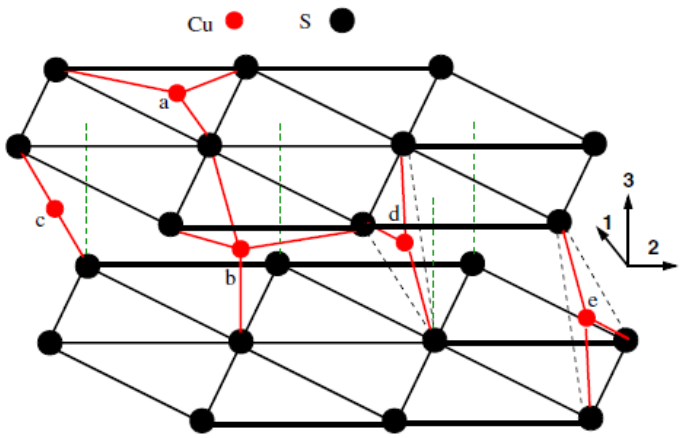

Figure 2.1: Schematic view of the Cu_2S structure. Adapted with permission from [2]

The best description for γCh comes from H.T.Evans [10] describing that compound as a monoclinic unit cell, space group $P2_1/c$, characterized by lattice paramters a = 15.246(4) Å, b = 11.884(2) Å, c = 13.494(3) Å, β = 116.35(1)°C containing 48 Cu_2S units. In contrast the βCh phase, space group $P6_3/mmc$, is described as a solid-liquid hybrid structure, where the sulfur atoms preserve a hexagonal crystal structure, while the copper atoms migrate between different sites exhibiting diffusion behaviour similar to that in a liquid [2]. Lin-Wang Wang found out that, at least for zero temperature and using an unrelaxed high-chalcocite model, that after relaxation the twofold and fourfold (c and b in Fig 2.1) coordinates switch to the threefold coordinates (d and e in Fig 2.1). This was also shown

for the low-chalcocite phase by displacing a copper atom from the threefold coordinates to two or fourfold coordinates and then relaxing them. The copper atom relaxes into the threefold coordinates [2]. This suggests that the threefold coordinates seem to be energetically prefered.

2.2 High-Dimensional Neural Network Potentials (HDNNPs)

The content of Section 2.2 is heavily based on the work of (Behler and Parrinello, 2007 [8]) and (Behler, 2021 [3]), which provided the foundational framework for the subsequent analysis. Those wishing to gain a deeper insight and further explore the methodology and findings are advised to consult the original paper.

The field of atomistic simulations has witnessed considerable activity over the past few decades, encompassing a diverse range of methodologies, including ab-initio approaches based on density functional theory (DFT) [11] and machine learning potentials. While ab-initio calculations provide authentic potential-energy surfaces (PES), the computational effort required is prohibitively expensive for larger systems. Furthermore, the simulation range of picoseconds renders certain research impractical. Consequently, the development of methods to handle larger systems while decreasing the overall computing time has become imperative. This was achieved when Behler and Parrinello introduced the generalized neural-network representation of high-dimensional potential-energy surfaces, a significantly more efficient method for creating DFT-based PESs [8, 3].

2.2.1 Neural Networks

The earliest application of machine learning (ML) techniques to model potential energy surfaces (PES) for atomistic simulations is represented by first-generation neural network potentials (NNPs). The application of machine learning to model potential energy surfaces began in the early 1990s. Notably, Sumpter and Noid [15] utilized neural networks in 1992 to analyze vibrational spectra derived from macromolecular force fields, marking a significant milestone in this field. A significant milestone was reached in 1995 with the creation of the first ML-based PES by Doren and colleagues [14]. This was achieved by utilising density functional theory (DFT) data to model H_2 adsorption on a silicon surface. These initial studies demonstrated the potential of neural networks in capturing the intricate relationships between atomic configurations and their corresponding energies [3].

The architectural design of the initial generation of NNPs seen in Figure 2.2 predominantly relied on single feed-forward neural networks (FFNNs), which processed a fixed

set of input coordinates to predict the global energy of a system. These FFNNs were designed with a minimal number of layers and neurons, rendering them well-suited for low-dimensional systems such as small molecules or rigid systems with fixed atomic arrangements. FFNNs are versatile tools capable of modeling complex relationships, including potential energy surfaces (PES), with minimal reliance on predefined physical assumptions. This adaptability allows them to accurately describe the energetic landscapes of diverse chemical systems. These neural network potentials (NNP) were employed in the modelling of small molecular systems, reactive molecular environments, and molecules interacting with fixed surfaces. For instance, they were used to model hydrogen adsorption on Si(100) and other low-dimensional molecular systems [14]. Such use cases demonstrated the potential of neural networks to provide efficient and accurate energy evaluations [3].

However, the initial generation of NNP models encountered considerable difficulties. A significant constraint was their requirement for fixed input dimensionality, which rendered them unsuitable for systems comprising a variable number of atoms. This limitation impeded their scalability and applicability to larger or more complex systems. Additionally, these early NNPs were deficient in robust descriptors that could guarantee invariance to translation, rotation, and permutation of atomic configurations, which is an indispensable prerequisite for precise PES modelling. Moreover, as the scale of the neural network expanded to accommodate more intricate systems, the computational expense associated with training and evaluating these models increased exponentially. The generation of sufficient training data from high-level electronic structure calculations constituted an additional challenge [3, 16, 17].

In order to address these limitations, researchers investigated a number of potential strategies. The use of multi-body expansions incorporating combinations of neural networks was introduced as a means of describing interactions in larger systems. In order to enhance the invariance properties, symmetry-adapted coordinates were developed. Furthermore, variable-size neural networks were proposed in order to handle systems with varying atom counts. Notwithstanding these developments, the initial generation of NNPs remained predominantly confined to low-dimensional systems, and their broader deployment was constrained by technical and computational limitations [3, 16, 17].

In summary, the initial generation of NNPs represented a significant advancement in the fields of computational chemistry and materials science. They demonstrated the feasibility of utilising machine learning to model potential energy surfaces with high accuracy and efficiency, thereby establishing a foundation for the more advanced and scalable methods that followed in subsequent generations. These initial developments established the foundation for the advent of high-dimensional neural network potentials, which addressed many of the limitations encountered by their predecessors [3].

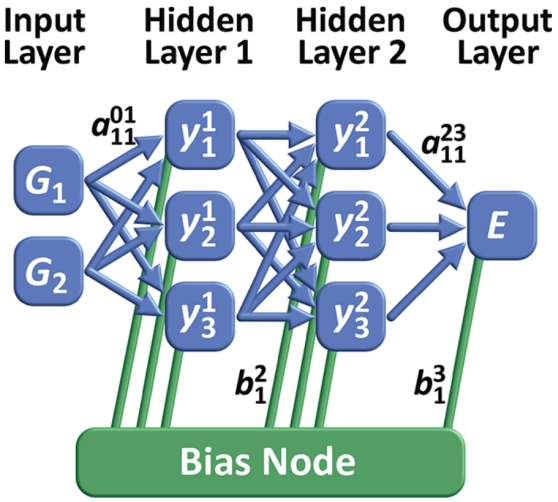

Figure 2.2: Structure of a feed-forward neural network. Adapted with permission from [3]

.

The structure depicted above comprises two input nodes, designated G_1 and G_2, respectively. The node situated at the output layer is associated with the potential energy, which is a function of the input nodes, the nodes located in the hidden layers and the corresponding weights. In the present case, two hidden layers are considered, although the number of hidden layers may vary. The nodes in these layers are not associated with any physical meaning. Each node in each layer is connected by specific weights, which are initially randomly selected. In this context, we define a connection between a neuron j in layer l (y_j^l) and a neuron i in layer k (y_i^k) by the weight a_{ij}^{kl}, where $l = k + 1$. The input layer is designated with the index 0, while the output layer is designated with the index $N+1$, where N represents the number of hidden layers. Furthermore, a bias node is present, connecting each node in each layer with the exception of the input layer. The bias node provides a bias weight b_j^i to the nodes in layer i. Both the primary weights a_{ij}^{kl} and the bias weights b_j^i are updated and optimised during the training of the FFNN. The bias node acts as an offset. Consequently, the final node, representing the potential energy as a function of the input coordinates [3], can be written as

$$E = f_1^3 \left(b_1^3 + \sum_{l=1}^{3} a_{l1}^{23} f_1^2 \left(b_l^2 + \sum_{k=1}^{3} a_{kl}^{12} f_k^1 \left(b_k^1 + \sum_{j=1}^{2} a_{jk}^{01} G_j \right) \right) \right), \tag{2.1}$$

where f is a nonlinear activation function in order to make the neural network a nonlinear model [3].

2.2.2 Second-Generation Neural Network Potentials

Second-generation neural network potentials (NNPs), which facilitate the modelling of high-dimensional systems with enhanced efficiency and accuracy, were first introduced by Behler and Parrinello in 2007 [8]. The locality approximation represents a major innovation, asserting that most atomic interactions can be effectively described by examining the immediate chemical surroundings of each atom. This approach permits the total potential energy to be expressed as a sum of contributions from these local environments, thereby significantly simplifying the computations [3]. The potential energy is expressed as a summation of atomic energy contributions, mathematically represented as [3]

$$E_{short} = \sum_{i=1}^{N_{elem}} \sum_{j=1}^{N_{atom}^{i}} E_{j}^{i}. \tag{2.2}$$

Here, N_{elem} represents the number of elements present within the system, whereas N_{atom}^{i} denotes the number of atoms belonging to element i. The term E_{short} is used to denote the short-range energy, which is a result of interactions occurring within a specific cutoff radius R_c [3].

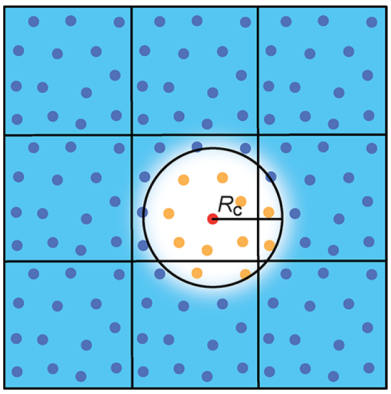

Figure 2.3: Cutoff radius representation. Adapted with permission from [3]

The illustration above 2.3 depicts the cutoff sphere, which represents the radius of the cutoff around a specific atom highlighted in red. The orange dots represent the local atomic environment.

In order to guarantee invariance with regard to translation, rotation and permutation of atoms, atom-centred symmetry functions (ACSFs) were introduced. These descriptors capture the local atomic environment and serve as inputs to atomic neural networks,

thereby ensuring a consistent and accurate energy representation [3].

Furthermore, the computational efficiency is enhanced by the utilisation of a cutoff radius, which limits the range of interactions considered for each atom, focusing solely on the nearest neighbours within this radius. This approach results in a reduction in complexity without any compromise to the accuracy required for many systems. To accommodate systems of disparate sizes and compositions, distinct atomic neural networks are trained for each element type. This modular architecture enables the model to scale effectively, rendering it applicable to systems with variable numbers of atoms. Furthermore, active learning techniques are employed to construct training datasets in an efficient manner, ensuring the inclusion of crucial configurations to achieve accurate potential energy surfaces. These developments address the limitations of first-generation NNPs, facilitating their application to large and complex systems while maintaining computational feasibility [3].

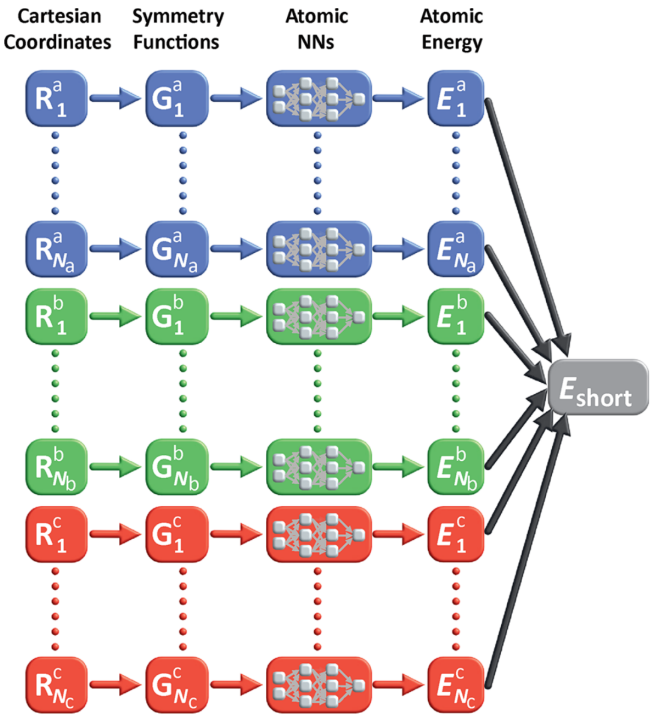

Figure 2.4: Structure of a second-generation high-dimensional neural network potential. Adapted with permission from [3]

.

Figure 2.4 illustrates the structure of a second-generation high-dimensional neural network potential, as developed by Behler and Parrinello in 2007 [8]. The various colours employed in the aforementioned illustration represent the chemical elements in ques-

tion. Thus, the elements a, b, c are represented by the colours blue, green and red, respectively, while the numbers N_a, N_b, N_c indicate the number of atoms present in each element. The structure operates as follows: the Cartesian coordinates R_i^ν of each atom are transformed into symmetry functions G_i^ν, which inherit information about the local atomic environment of atom i of element ν. The symmetry functions serve as the input for the neural network, which yields the atomic energy contribution, denoted as E_i^ν. The sum of all atomic energy contributions results in the short-range energy, denoted as E_{short}. It should be mentioned that only coordinates within the cutoff sphere R_c are considered for the symmetry functions [3].

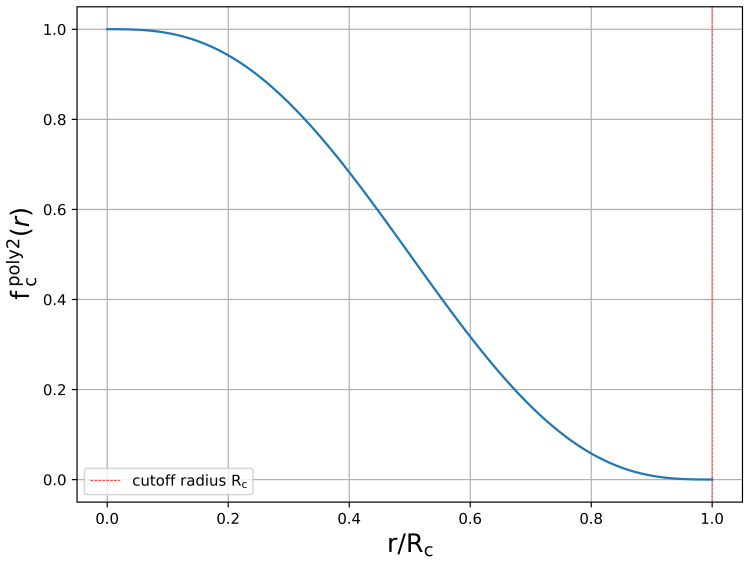

Figure 2.5: Cutoff function used in the atom-centered symmetry functions

In the figure above 2.3 we see the cutoff function which takes the form of a polynomial

$$f_c^{poly2}(r) = x^3(x(15 - 6x) - 10) + 1, \tag{2.3}$$

where $x = \frac{r}{R_c}$.

Cutoff functions may assume a variety of forms, including trigonometric functions. However, the polynomial presented above was employed by Singraber [1], which serves as the foundation for this work. Following the introduction of a cutoff function, the subsequent step would be to describe the position of neighbouring atoms within the cutoff sphere using radial and angular ACSFs [3].

For the radial atom-centered symmetry functions we use the sum of Gaussians and cutoff functions seen below [3] for all atoms within the cutoff sphere

$$G_{i,\nu}^{rad} = \sum_{j \neq i}^{N_{atom} \in R_c} e^{-\eta(R_{ij}-R_s)^2} f_c(R_{ij}). \tag{2.4}$$

The combination of Gaussians with the cutoff function enables the symmetry function to reach zero at the cutoff radius, as illustrated in the following diagram. The sum of the two ensures that a single function value is obtained, which is then required as the input for the neural network. In this instance, the variable η is responsible for the spatial spread of the symmetry function, while R_s represents the shift parameter, which shifts the centre of the Gaussian to a specific distance from the central atom [3].

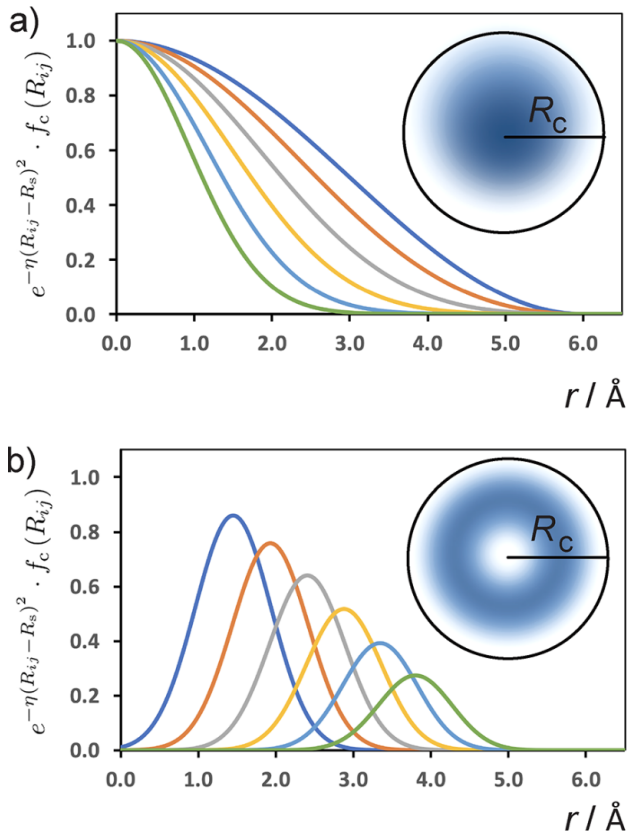

Figure 2.6: Radial symmetry functions with varying parameters and a cutoff value of 6Å. Adapted with permission from [3]

.

In a) we see the unshifted ($R_s = 0$) radial symmetry functions with $\eta \in \{0.0, 0.03, 0.08, 0.16, 0.3\}$Å^{-2} from green to dark blue, and in b) the shifted radial symmetry functions with $\eta = 0$ and $R_s \in \{1.5, 2.0, 2.5, 3.0, 3.5, 4.0\}$Å [3].

In order to describe the positions of the surrounding atoms within the cutoff sphere, it is also necessary to define angular symmetry functions [3], which take the form

$$G_{i,\nu}^{ang} = 2^{1-\zeta} \sum_{i,j,k} (1 + \lambda \cos\theta_{ijk})^{\zeta} \, e^{-\eta\left(R_{ij}^2 + R_{ik}^2 + R_{jk}^2\right)} f_c(R_{ij}) f_c(R_{ik}) f_c(R_{jk}). \tag{2.5}$$

θ_{ijk} is the angle between the central atom i and its neighbouring atoms j and k. λ can be assigned values of $+1$ or -1, is employed to position the maxima of the cosine terms at either $\theta_{ijk} = 0°$ or $\theta_{ijk} = 180°$. In the angular symmetry functions, the exponent ζ controls the angular resolution where a larger ζ yields a narrower peak centered at $\theta_{ijk} = 0°$ for $\lambda = +1$ and at $\theta_{ijk} = 180°$ for $\lambda = -1$. The prefactor $2^{1-\zeta}$ normalizes the term so that $2^{1-\zeta}(1 + \lambda \cos\theta_{ijk})^{\zeta}$ attains a constant maximum value, independent of ζ, hence ζ changes only the width/shape of the angular term, not its overall scale (see Figure 9, Behler 2021 [3]).

It may be beneficial to incorporate additional angle terms. This could be achieved by incorporating a less restrictive or a more expansive angular symmetry function. If the constraint $R_{jk} < R_c$ is omitted, it is possible to include additional angle terms, as illustrated in the figure below 2.7 [3].

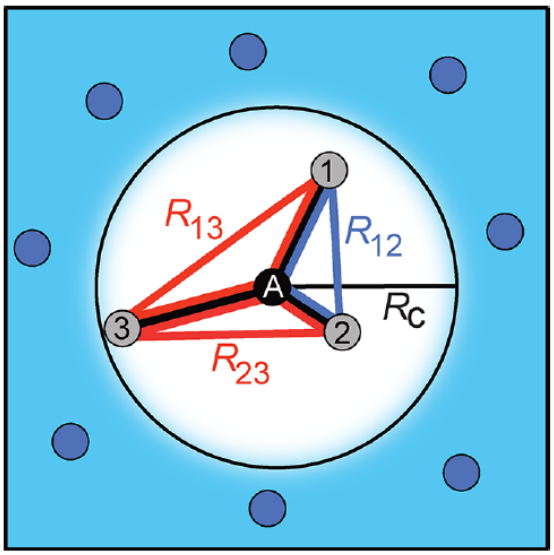

Figure 2.7: Angular terms in the less restrictive case. Adapted with permission from [3]

.

In the case of the common angular symmetry function the red connections would not be considered since $R_{13} > R_c$ and $R_{23} > R_c$. In contrast, in the less restrictive case, these would be allowed since the constraint $R_{jk} < R_c$ was omitted [3]. So in general we differ between two angular symmetry functions, the narrow angular symmetry function [1]

$$G_{i,\nu}^{ang.n} = 2^{1-\zeta} \sum_{\substack{j,k \neq i \\ j<k}} (1 + \lambda \cos\theta_{ijk})^{\zeta} \, e^{-\eta\left(R_{ij}^2 + R_{ik}^2 + R_{jk}^2\right)} f_c(R_{ij}) f_c(R_{ik}) f_c(R_{jk}) \tag{2.6}$$

and the wide angular symmetry function [1]

$$G_{i,\nu}^{ang.w} = 2^{1-\zeta} \sum_{\substack{j,k \neq i \\ j<k}} (1 + \lambda \cos\theta_{ijk})^{\zeta} \, e^{-\eta\left(R_{ij}^2 + R_{ik}^2\right)} f_c(R_{ij}) f_c(R_{ik}). \tag{2.7}$$

Once the energy has been calculated in accordance with the specifications set forth in equation 2.2, it can then be employed in the computation of the forces utilized in molecular dynamics simulations. Here, $F_{\mu\alpha}$ represents the force acting on atom μ in direction $\alpha = \{x, y, z\}$ [3],

$$\begin{aligned} F_{\mu\alpha}^{short} &= -\frac{\partial E_{short}}{\partial R_{\mu\alpha}} = -\sum_{i=1}^{N_{elem}} \sum_{j=1}^{N_{atom}^i} \frac{\partial E_j^i}{\partial R_{\mu\alpha}} \\ &= -\sum_{i=1}^{N_{elem}} \sum_{j=1}^{N_{atom}^i} \sum_{k=1}^{N_{sym}^i} \frac{\partial E_j^i}{\partial G_{jk}^i} \frac{\partial G_{jk}^i}{\partial R_{\mu\alpha}}. \end{aligned} \tag{2.8}$$

N_{sym}^i is the number of symmetry functions.

3 Methods

3.1 Software

We used the neural network potential package software n2p2 with the LAMMPS interface [12] to run MD simulations. The already trained Cu_2S HDNNP [1] was used to implement the properties of Cu_2S in LAMMPS. Multiple simulations were performed in the NVT and NpT ensembles. In order to obtain a similar result as Singaber [1] for the phase transition in terms of changes in energy, density, pressure, and mean square displacement, we used identical initial conditions. The unit cell of low chalocite consists of 144 atoms [9]. We enlarged the system by replicating the unit cell three times in all dimensions. The final structure consists of 3888 atoms. The initial thermostat temperature T_i, prior to heating, was set to $350K$, while the final value T_f reached $410K$, allowing capture of the superionic transition. For the NpT ensemble, a target pressure of $1.5GPa$ was used, while for the NVT Ensemble the experimental value of $5790kg/m^3$ was used [1].

3.2 PaCMAP

To visualize high-dimensional atomic environments datasets and study structural changes during the superionic phase transition of copper sulfide, we used PaCMAP (Pairwise Controlled Manifold Approximation Projection), a dimensionality reduction technique that, according to the authors, preserves both local and global data structures more effectively than traditional methods such as t-SNE or UMAP [7].

Graph construction PaCMAP constructs a graph in which edges represent the pairwise relationships between data points. These edges are categorised into three distinct types: neighbour pairs, mid-near pairs, and further pairs. Each of these categories is designed to capture a different aspect of the data's geometric structure.

In order to construct neighbour pairs, PaCMAP selects the nearest neighbours n_{NB} for each data point in the high-dimensional space. The selection is based on a scaled distance metric, adapted from TriMap, defined as [7]:

$$d^2_{ij,\text{select}} = \frac{\|x_i - x_j\|^2}{\sigma_{ij}}, \quad \text{where} \quad \sigma_{ij} = \sigma_i \cdot \sigma_j \tag{3.1}$$

In PaCMAP, the local scaling parameter σ_i is defined as the mean Euclidean distance from point i to its 4th through 6th nearest neighbors. To build mid-near pairs, the algorithm randomly draws a small set of candidate points (typically six) for each observation and then chooses the second-closest of those candidates as the partner. This procedure introduces connections that are neither strictly local nor fully global, thereby capturing intermediate structure in the dataset [7]. The total number of mid-near pairs is proportional to the number of neighbour pairs, determined via the hyperparameter $n_{MN} = MNratio \cdot n_{NB}$.

In a similar manner, further pairs are formed by means of randomly sampling $n_{FP} = FPratio \cdot n_{NB}$ distant points for each observation. These random sampling strategies reduce computational cost while maintaining representative global structure.

This combination of pair types, locally dense neighbour edges, moderately distant mid-near pairs, and widely separated further pairs enables PaCMAP to effectively preserve both local continuity and global separability in the resulting low-dimensional embedding.

Loss function The fundamental principle of PaCMAP's embedding quality can be attributed to its pairwise loss function, which is minimised to learn the low-dimensional representation $y \in \mathbb{R}^2$ for each high-dimensional point x. The loss function consists of contributions from the three pair types outlined below:

$$\mathcal{L} = \sum_{(i,j)\in\mathcal{P}_n} w_n \cdot \ell_{ij}^{(n)} + \sum_{(i,j)\in\mathcal{P}_m} w_m \cdot \ell_{ij}^{(m)} + \sum_{(i,j)\in\mathcal{P}_f} w_f \cdot \ell_{ij}^{(f)}, \tag{3.2}$$

where the sets of near, mid-near and further pairs are denoted by $\mathcal{P}_n$, $\mathcal{P}_m$ and $\mathcal{P}_f$, respectively. In the context of the training process, the weights w_n, w_m and w_f assigned to each pair type are subject to dynamic adjustment. Each distinct pairwise loss function has been determined to be

$$\ell_{ij}^{(n)} = \frac{\widetilde{d_{ij}}}{\widetilde{d_{ij}} + 10}, \quad \ell_{ij}^{(m)} = \frac{\widetilde{d_{ij}}}{\widetilde{d_{ij}} + 1000}, \quad \ell_{ij}^{(f)} = \frac{1}{\widetilde{d_{ij}} + 1}, \tag{3.3}$$

where $\widetilde{d_{ij}} = \|y_i - y_j\|^2 + 1$.

Optimization Strategy PaCMAP implements a three-stage optimization scheme with an adaptive weighting schedule in order to balance the preservation of global and local structure [7]:

- Phase 1: (iterations from $\tau_1 = 1$ to $\tau_2 - 1$) Here the primary focus is the establishment of a meaningful global layout, whilst simultaneously ensuring the retention of local detail. This objective is realised through the allocation of substantial weights

to mid-near pairs, thereby providing guidance to the preliminary structure of the embedding. Subsequently, the mid-near pair weights undergo a gradual reduction, enabling a natural transition towards local refinement.

- Phase 2: (iterations from τ_2 to $\tau_3 - 1$) The second phase aims to enhance the local structure of the embedding, while preserving the global structure acquired in Phase 1. Mid-near pairs are still included, but with significantly reduced weight.
- Phase 3: (iterations from τ_3 to $n_{iterations}$) Phase 3 involves the elimination of mid-near and further pairs, with a subsequent focus on near pairs. The culminating phase of this process entails the refinement of local neighbourhoods and the culmination of the embedding process.

Initialization In order to facilitate convergence and circumvent the randomness-induced artefacts observed in methodologies such as t-SNE, PaCMAP employs a PCA-based initialisation. This approach ensures a globally consistent starting point that reflects the major variance directions of the input data, while still allowing for non-linear deformation during optimization [7].

In the initial phase of this work, the focus was placed on analyzing the local environments of copper atoms, which exhibit significant mobility during the superionic phase transition. To this end, atom-centered symmetry functions (ACSFs) were computed based on a descriptor setup inspired by Singraber et al. [1], using a combination of radial and angular terms (both narrow and wide). The descriptor sets were carefully filtered to exclude physically redundant or uninformative symmetry functions. These include, for instance, radial functions centered on copper with a shift parameter $r_s = 1.5$Å, as well as certain angular triplets such as Cu–Cu–Cu combinations with excessive angular resolution (e.g., $\eta = 0.2222, \lambda = 1, \zeta = 6$) that provided no significant differentiation between phases.

The exclusion of these functions was implemented directly in the descriptor generation pipeline by matching the generated feature names with a predefined exclusion list. The descriptor matrix was then truncated accordingly before saving, ensuring both computational efficiency and semantic clarity.

3.3 Neural Network Classification of Atomic Environments

In order to provide a complementary analysis of the superionic phase transition in Cu_2S, a binary neural network is employed. The classification model was trained to differentiate between atomic environments in low-chalcocite (LC) and high-chalcocite (HC) contexts. The presence of chalcocite (HC) phases is indicated. The objective of this study was to identify the distinctive features of atoms in each phase. The subsequent stage of the process is the detection of intermediate or transitional states.

The training of neural networks was based on atom-centred symmetry functions (ACSFs) as input descriptors, which encode the local environment of each atom. The computation of these ACSFs was conducted for each atom in the selected timesteps of molecular dynamics trajectories.

For each atom, the ACSFs were calculated using the DScribe library. The selection of symmetry function parameters was conducted in accordance with the methodologies outlined in previous research by Singraber [1], incorporating radial, angular narrow, and angular wide terms. Functions of symmetry which are known to be physically irrelevant or redundant were excluded from consideration. The objective is twofold: firstly, to enhance the efficiency of the training process and secondly, to ensure the accuracy of the model.

The vector of symmetry functions for atom i is given by:

$$\mathrm{s}^{(i)} = \begin{pmatrix} s_1^{(i)} \\ s_2^{(i)} \\ \vdots \\ s_d^{(i)} \end{pmatrix} \in \mathbb{R}^d \tag{3.4}$$

Rolling average smoothing was applied to atomic trajectories prior to the computation of ACSFs, with the objective of mitigating thermal fluctuations. This preliminary processing stage did not compromise the efficacy of the classification process, however, it may have served to mitigate minor fluctuations in the data.

ACSFs were computed for the entire 3888-atom system across a range of timesteps. For the copper model, ACSFs were computed from timestep 104,000 to 1,000,000 for the

LC phase and from 5,004,000 to 6,100,000 for the HC phase. These specific timesteps for LC and HC were chosen because the trajectory file obtained through MD simulations was determined after 100,000 NVT equilibration steps. After applying a rolling average to the unwrapped coordinates with a window size of five frames, we started with timestep 104,000 instead of 100,000. Up to approximately timestep 2,500,000, the system is still in the low chalcocite phase. At this point, selecting the LC range is justified. The same applies to the high chalcocite phase. The training dataset was constructed using these labels: 0 (LC) and 1 (HC), which we denote as:

$$\ell^{(i)} \in \{0, 1\} \tag{3.5}$$

The data were standardized using a StandardScaler, and then used to train a simple FFNN implemented in PyTorch. The architecture included three layers: a 64-unit hidden layer, a 32-unit hidden layer, and a final sigmoid-activated output node. The network was trained using the binary cross-entropy loss function with the Adam optimizer.

The neural network predicts a probability that atom i belongs to the LC or HC phase based on its ACSFs:

$$\mathcal{Y}(\mathbf{s}^{(i)}, \boldsymbol{\theta}) \in [0, 1] \tag{3.6}$$

Here θ denotes the weights and biases of the neural network. During training, the model showed clear convergence, with the loss decreasing over 30 epochs. The binary cross-entropy loss function minimized during training is given by:

$$\mathcal{L}(\boldsymbol{\theta}) = -\sum_{i=1}^{N} \left[\ell^{(i)} \log \mathcal{Y}(\mathbf{s}^{(i)}, \boldsymbol{\theta}) + (1 - \ell^{(i)}) \log \left(1 - \mathcal{Y}(\mathbf{s}^{(i)}, \boldsymbol{\theta}) \right) \right] \tag{3.7}$$

The trained model was then applied to all ACSF files from timestep 1500000 to 5000000 in steps of 50000. For further analysis, the model's output can be correlated with local diffusion constants D_i of individual atoms, which are computed over a time lag τ. A plot of D_i versus predicted probability p_i may highlight if atoms in transitional states (p_i near 0.5) also show distinct mobility. The diffusion constant for each atom is computed using:

$$D_i = \frac{1}{6\tau} \left\langle |\vec{r}_i(t + \tau) - \vec{r}_i(t)|^2 \right\rangle \tag{3.8}$$

This will allow us to assign atoms to intermediate states in a more physically meaningful manner.

4 Results

4.1 MD Simulations

In Figure 4.1 we see four graphs with the x axis being the simulation time t and the temperature T. In the first plot we see the potential energy per atom displayed in both ensembles. Both energy trajectories show a significant increase in potential energy per atom, starting at approximately 376.5K, and slowly stabilizing at 400K, indicating a phase transition. The lower plots (density and pressure) also both show a phase transition-like behavior. Further in the last plot we computed the mean square displacement (MSD) of both atomy types in the NVT ensemble, where sulfur seems to show no mobility at all, while copper atoms show liquid-like diffusive behaviour. The critical temperature, T, at which the phase transition occurs, is 103.5 degrees Celsius, as illustrated by the green dotted line. Due to thermal fluctuations we used an average rolling with a window size of 5 frames for E_{atom} and a window size of 15 frames for density ρ and pressure p. The results are in accordance with Singrabers outcome [1], validating his Cu_2S HDDNP.

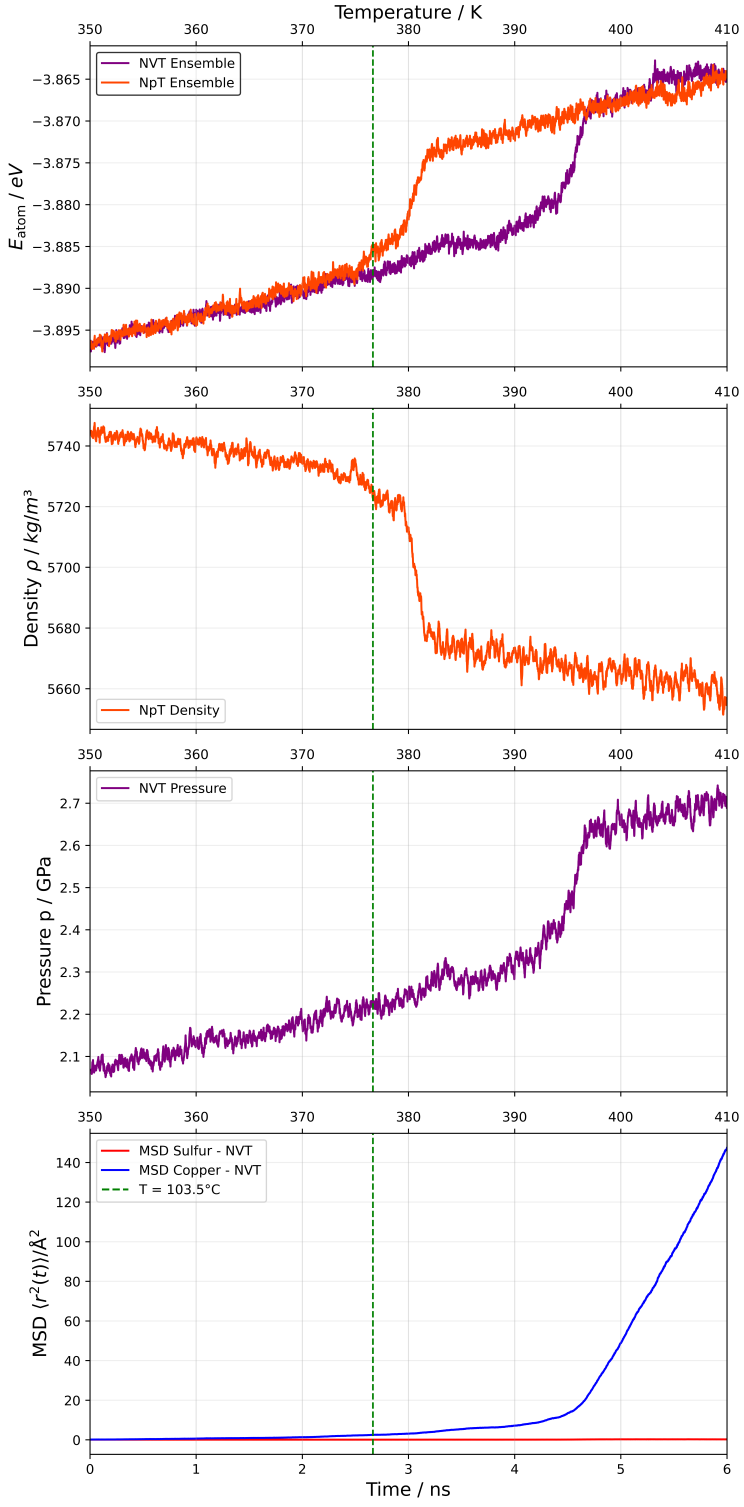

Figure 4.1: Thermodynamic properties of copper sulifde during the phase transition computed using the HDNNP.

4.2 PaCMAP results

In order to investigate the structural evolution during the superionic phase transition of Cu_2S, PaCMAP dimensionality reduction was applied to atom-centered symmetry functions (ACSFs) across a series of overlapping time intervals. A total of 12 PaCMAP projections were generated, with each one representing the low-dimensional embedding of atomic environments for a selected range of timesteps (see Fig. 4.2). The windows in question span almost the entire trajectory from 200000 to 5000000 timesteps, thus capturing the transition from the low-chalcocite to the high-chalcocite phase.

Each plot visualizes the spatial distribution of local atomic environments in a two-dimensional latent space, where distinct clusters correspond to structurally similar atomic neighborhoods. The colour scheme reflects the clustering results obtained via DBSCAN, while the progression of plots provides insight into how structural motifs emerge, evolve, and reorganise as the system approaches and surpasses the critical transition point near timestep 2.6 million, which corresponds to a temperature of approximately 376.5K.

The overlapping intervals were selected to guarantee temporal resolution and facilitate the tracking of cluster continuity or transformation across adjacent windows. This approach aims to facilitate a detailed analysis of nucleation phenomena, the coexistence of phases, and the potential emergence of intermediate or transient atomic configurations.

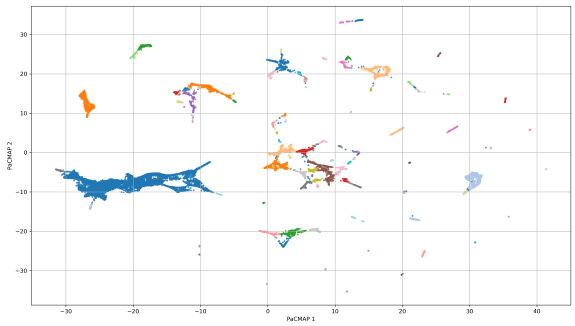

(a) Timestep 200k - 500k

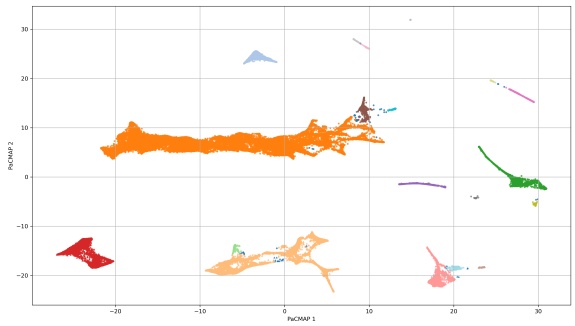

(b) Timestep 2.6m - 2.9m

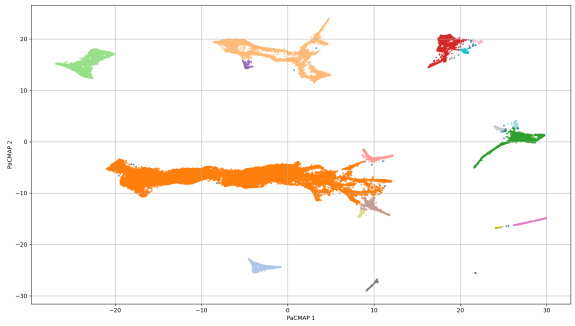

(c) Timestep 2.8m - 3.1m

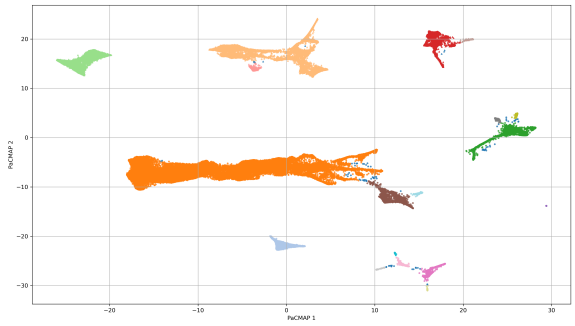

(d) Timestep 3m - 3.3m

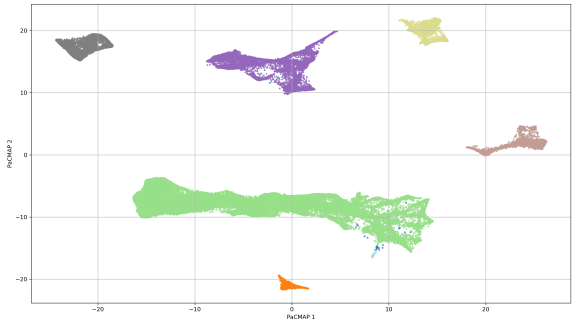

(e) Timestep 3.4m - 3.7m

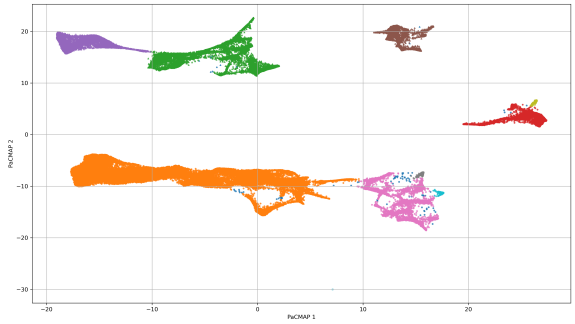

(f) Timestep 3.6m - 3.9m

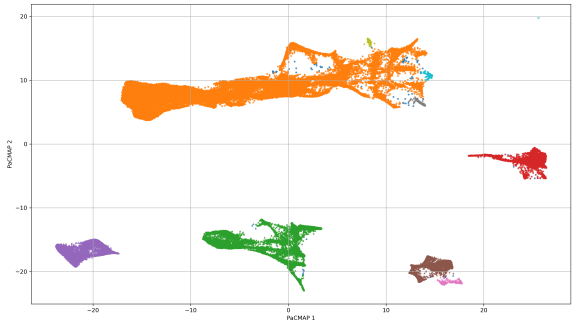

(g) Timestep 3.8m - 4.1m

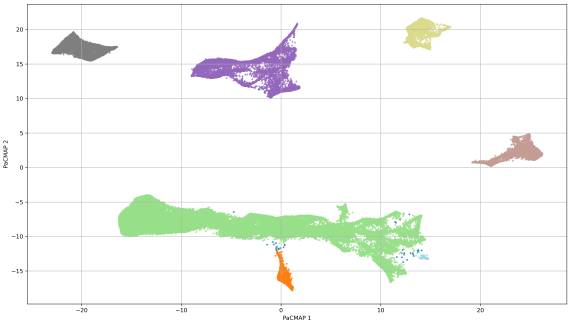

(h) Timestep 4m - 4.3m

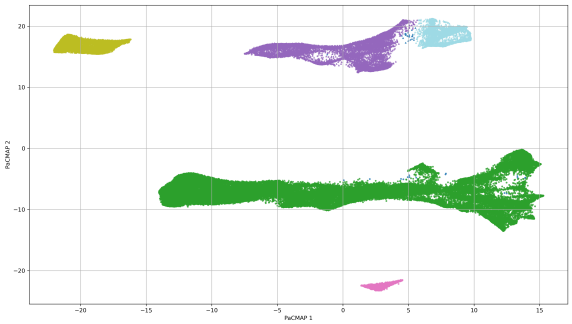

(i) Timestep 4.2m - 4.5m

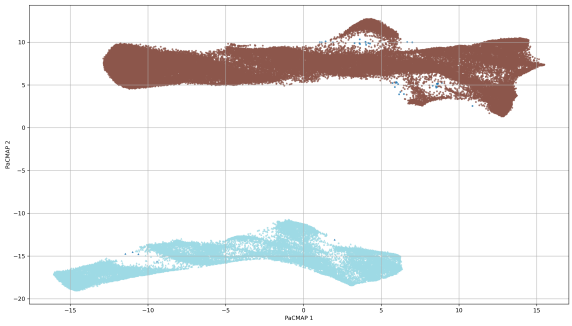

(j) Timestep 4.4m - 4.7m

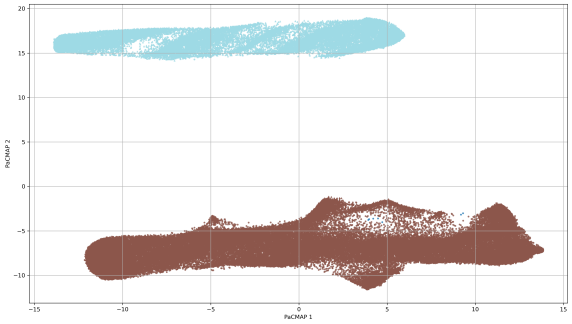

(k) Timestep 4.6m - 4.9m

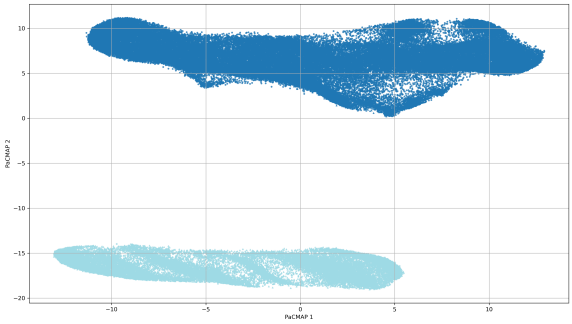

(l) Timestep 4.8m - 5m

Figure 4.2: PaCMAP-results for various timesteps

Figure 4.2a presents the PaCMAP projection of atomic environments in the low-chalcocite (-phase) regime, covering timesteps 200000–500000. Despite the system being in its thermodynamically stable low-temperature phase, the analysis reveals a surprisingly high number of distinct clusters (85). Most clusters consist exclusively of either copper or sulfur atoms, as identified by DBSCAN. The reasoning behind these vast clusters will be discussed in detail in the subsequent Chapter 5.

Figure 4.2 presents the PaCMAP projections of local atomic environments in three overlapping windows during the transition region: 4.2b 2.6m–2.9m, 4.2c 2.8m–3.1m, and 4.2d 3.0m–3.3m. These visualisations highlight the progressive structural changes in the system as it approaches the superionic state.

Across all three time intervals, the emergence of a dominant copper cluster (orange cluster) is observed, containing over 44000 atoms in each snapshot. This particular cluster is distinguished by its elongated, interconnected structure within the PaCMAP space. It is conceivable that this configuration may correspond to copper atoms that commence occupation of a more fluid-like, thermally accessible configuration space. The continuous shape of this cluster indicates that many Cu atoms share similar, but not identical local environments, consistent with the onset of disorder and increased mobility. It is evident that the cluster sizes remain consistent across all three configurations, irrespective of the atom type.

In the figures 4.2e - 4.2i we see projections of atomic environments: these are from five overlapping windows from 3.4 million to 4.5 million timesteps, with each covering a 300000-timestep interval. The clusters identified by DBSCAN (ranging from 5 to 11 per window) reflect distinct local atomic configurations of copper and sulfur atoms, captured via ACSF descriptors.

Across all five intervals, the copper atoms predominantly populate one large, elongated cluster, which consistently spans the centre of each PaCMAP projection. This dominant cluster, which is visually stable in shape and size, accounts for approximately 63,000–66,000 atoms, thus confirming its structural robustness. It is plausible that this corresponds to the high-chalcocite (H-chalc.) phase, which becomes more prevalent after the nucleation onset.

As illustrated in figures 4.2e - 4.2h, a light brown/red cluster containing approximately 8000 copper atoms can be observed, undergoing dissolution in timestep 4.2m-4.5m 4.2i. Concurrently, the primary elongated cluster experiences an augmentation of approximately 8000 - 9000 copper atoms. In accordance with predictions, sulfur atoms predominantly maintain their positions within the fixed lattice, exhibiting negligible cluster redistribution. The clustering patterns of these systems remain relatively static, with large groups of sulfur atoms (approximately 17000) tightly bound within one or two distinct clusters across all intervals. This observation serves to reinforce the structural importance of the sulfur

sublattice and lends further support to the earlier findings that only copper atoms actively participate in the phase transition dynamics.

The final three PaCMAP projections 4.2j - 4.2l, covering timesteps 4.4m to 5.0m, demonstrate a marked structural stabilisation of the system following the primary transition period. In contrast to the preceding intervals, which exhibited between seven and 11 clusters, these final three windows contain solely two stable clusters, thereby signifying that the structural rearrangement has been predominantly accomplished. This spatial separation is indicative of the emergence of two well-defined, thermodynamically stable local environments for Cu and S atoms. The copper cluster exhibits a compact and dense structure with minimal substructure or fragmentation, suggesting that the atoms are fully embedded in the mobile superionic phase. It is noteworthy that a small number of outliers (predominantly copper) persist, though their number has decreased from 45 to 1, thereby reinforcing the system's increasing homogeneity.

The sulfur atoms, previously exhibiting a partial blend with copper-rich environments, now form a visually distinct and internally cohesive cluster. The sulfur sublattice, which remains structurally fixed throughout the transition, now exhibits tighter grouping and minimal spread in descriptor space. This finding lends further credence to the hypothesis that sulfur atoms experience consistent local environments and are largely decoupled from the fluctuating copper subsystem.

4.3 Temporal Evolution of the High-Chalcocite Phase Fraction

In order to monitor the progress of the superionic phase transition quantitatively, we computed the relative fraction of copper atoms residing in the high-chalcocite phase as a function of simulation time. This was achieved by applying a supervised neural network classifier, which was trained using atom-centred symmetry functions (ACSFs), to distinguish between the local atomic environments that are characteristic of the low-temperature and high-temperature phases. For each timestep, the model outputs a probability p_i for each copper atom, indicating its likelihood of belonging to the low or high-chalcocite phase. Atoms with a $p_i > 0.9$ were categorised as belonging to the high-chalcocite phase, while those with a $p_i < 0.1$ were categorised as belonging to the low-chalcocite phase. Intermediate values were excluded to ensure high-confidence classifications.

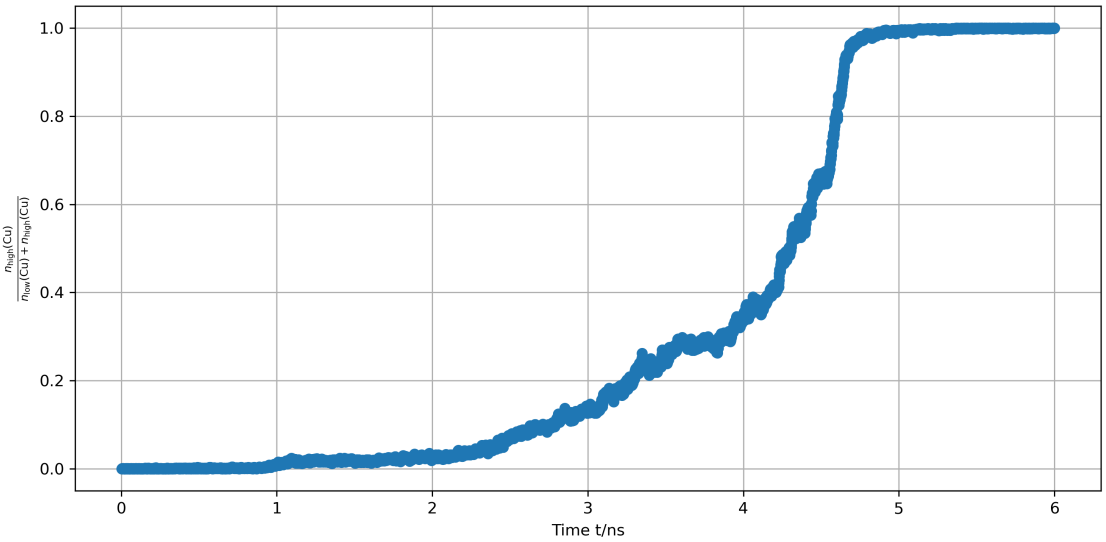

Figure 4.3: Fraction of High-Phase Cu Atoms Over Time

Figure 4.3 presents the time evolution of the normalized phase fraction, defined as

$$\frac{n_{\text{high}}(\text{Cu})}{n_{\text{low}}(\text{Cu}) + n_{\text{high}}(\text{Cu})}, \tag{4.1}$$

where $n_{\text{high}}(Cu)$ and $n_{\text{low}}(Cu)$ denote the number of copper atoms confidently classified as high- and low-phase respectively.

Initially, the system remains predominantly in the low-chalcocite phase, with a negligible fraction of high-phase atoms. However, commencing at approximately timestep 3 million, a gradual increase in the high-phase fraction is observed, indicating the onset of nucleation events. This growth accelerates significantly between 4.5 million and 5.2 million, where the transition exhibits a sigmoidal shape, suggesting a first-order transition characterised by the coexistence of two distinct phases, and a rapid growth of the high-temperature phase once nucleation has surpassed a critical threshold. Beyond timestep 5.2 million, the fraction saturates at unity, confirming the complete transformation into the high-chalcocite phase. The results obtained provide compelling evidence that the phase transition in copper sulfide is not continuous but proceeds via a nucleation and growth mechanism, consistent with a first-order transition.

4.4 Temporal Evolution of Copper Diffusion During the Phase Transition

In order to gain quantitative insight into the dynamical behaviour of copper atoms across the phase transition, the atom-wise diffusion constants, D_i, were computed over a fixed lag time window of $\tau = 500ps$ for multiple starting timesteps throughout the molecular dynamics trajectory. This analysis exclusively considers copper atoms, in accordance with their critical role in the superionic behaviour of the material. For each selected initial time, the displacement of each atom after a time interval τ was calculated from the unwrapped trajectory. This was then used to determine the individual diffusion constants, D_i, via the following equation:

$$D_i = \frac{|\vec{r}_i(t+\tau) - \vec{r}_i(t)|^2}{6\tau} \tag{4.2}$$

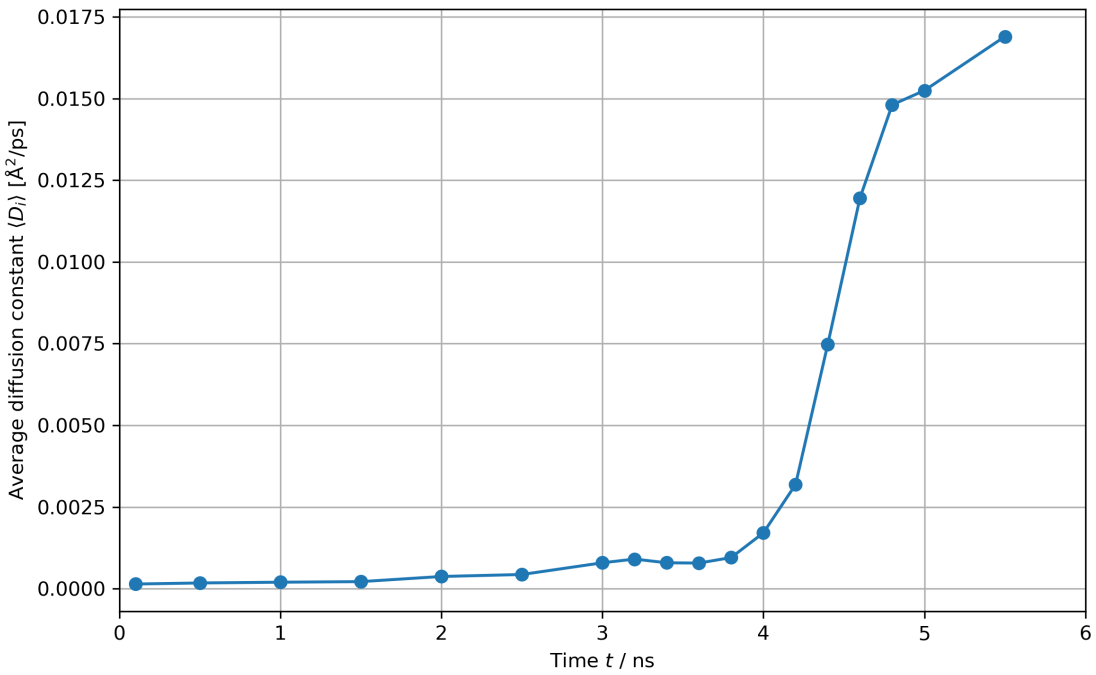

Figure 4.4: Average Diffusion constant D_i as a function of time t

The average diffusion constant $<D_i>$ was then obtained by averaging over all copper atoms at each time window. The temporal evolution of $<D_i>$ is illustrated in Figure 4.4. As demonstrated in the figure, the system persists in a low-diffusivity state up to approximately $t \approx 4.2ns$, with average diffusion constants below $0.002\text{Å}^2/ps$, which is indicative

of a structurally ordered phase where copper atoms are predominantly localized. This plateau is succeeded by a sharp and nonlinear increase in $<D_i>$, which peaks at values in excess of $0.015\text{Å}^2/ps$ at $t = 5.5ns$. The abrupt onset of high mobility indicates a sudden transition from the low-chalcocite phase to the high-chalcocite (superionic) phase.

4.5 Atomwise Diffusion vs. Neural Network Phase Probability

In order to continue the investigation of the microscopic dynamics of copper diffusion in relation to the phase transition, the atom-wise diffusion constant, D_i, was analysed as a function of the phase probability, p_i. This analysis was predicted by a neural network classifier that had been trained on ACSF descriptors. The NN output $p_i \in [0, 1]$ functions as a probabilistic indicator of an atom belonging to the high-temperature (superionic) phase. A value near 0 corresponds to the low chalcocite phase, while values close to 1 denote a high likelihood of being in the high chalcocite phase.

The diffusion constants, D_i, were computed over a time interval of $\tau = 500ps$, and the analysis was performed for selected timesteps before, during, and after the phase transition.

Low-Temperature Regime As demonstrated in Figure 4.5, at a time step of 200k (a point well in advance of the phase transition), copper atoms manifest exceedingly small diffusion constants, with $D_i < 0.01\text{Å}^2/ps$. This finding is consistent with a vibrationally dominated motion in the ordered low chalcocite phase. The corresponding NN probabilities, denoted by pi, are found to be nearly zero for the vast majority of atoms, thereby confirming the homogeneous nature of the low-temperature phase. Even at the timestep 1m, diffusion remains suppressed and probabilities continue to cluster around zero, indicating that the system remains in the low phase.

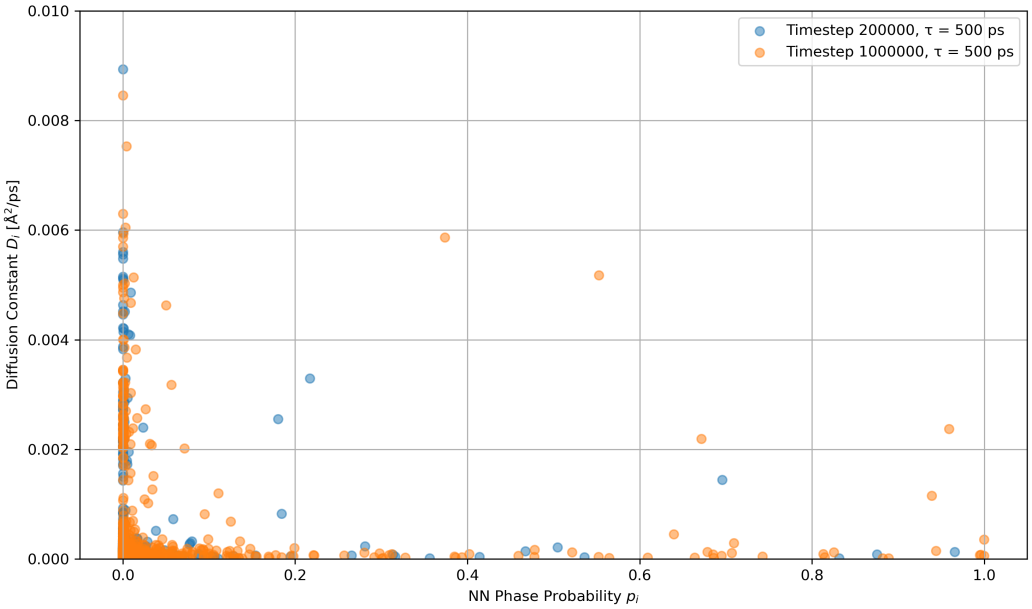

Figure 4.5: Atomwise Diffusion Constant vs NN Output Probability (low-chalcocite phase)

High-Temperature Regime At timestep 5.5m, which corresponds to the fully developed high-temperature phase, the distribution of atoms in the diffusion-probability space exhibits a notable asymmetry. As demonstrated in Figure 4.6, the majority of copper atoms in the high chalcocite phase are assigned a neural network phase probability $p_i \approx 1$. Nevertheless, the diffusion constants D_i demonstrate a considerable range, extending from almost zero up to $0.15 \text{Å}^2/ps$.

This observation challenges the simplistic expectation that high structural phase probability should always correlate with high mobility. Instead, the data suggest that, even within a high-phase environment, a significant proportion of copper atoms remain relatively immobile.

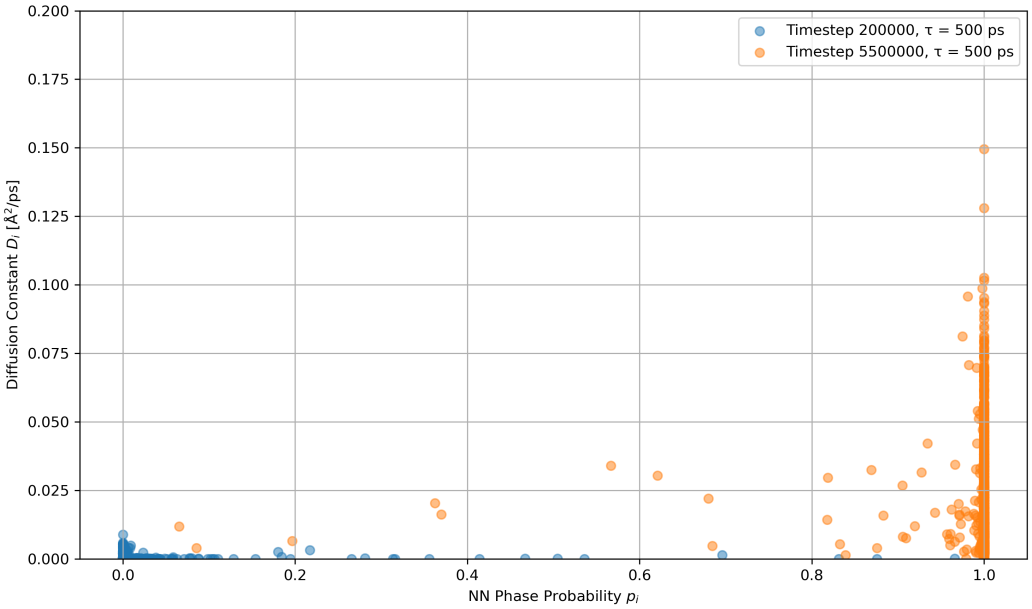

Figure 4.6: Atomwise Diffusion Constant vs NN Output Probability (low-vs.high chalcocite phase)

Nucleation Regime At time step 4.2m 4.7, which is well within the transition window, the distribution of NN phase probabilities p_i shows two dominant groups: one centred near $p_i \approx 0$ and one near $p_i \approx 1$, with relatively few atoms occupying intermediate values. This reflects the coexistence of atoms in low- and high-phase-like local environments. At this stage, the diffusion constants D_i are generally low in both groups, indicating that large-scale copper mobility has not yet developed.

By timestep 4.7m 4.8, the relative populations of the $p_i \approx 0$ and $p_i \approx 1$ groups have barely changed, indicating that the structural partition between low- and high-phase-like atoms has already been established. However, the diffusion constants, particularly for the $p_i \approx 1$ group, are higher than at timestep 4.2m. This suggests that the high-phase-like regions are becoming more dynamically active, despite their structural extent not having significantly expanded.

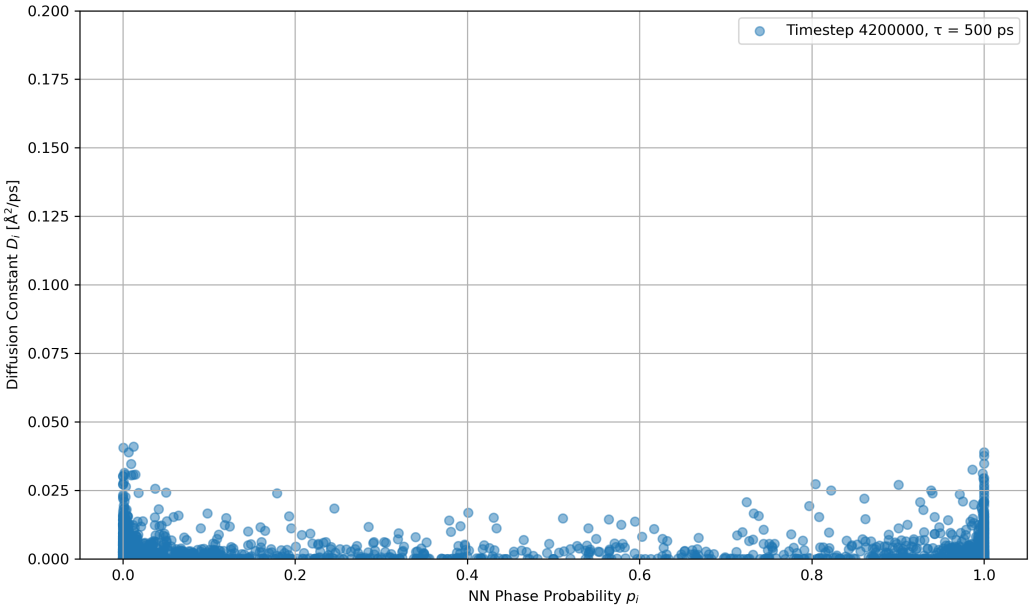

Figure 4.7: Atomwise Diffusion Constant vs NN Output Probability (during phase transition)

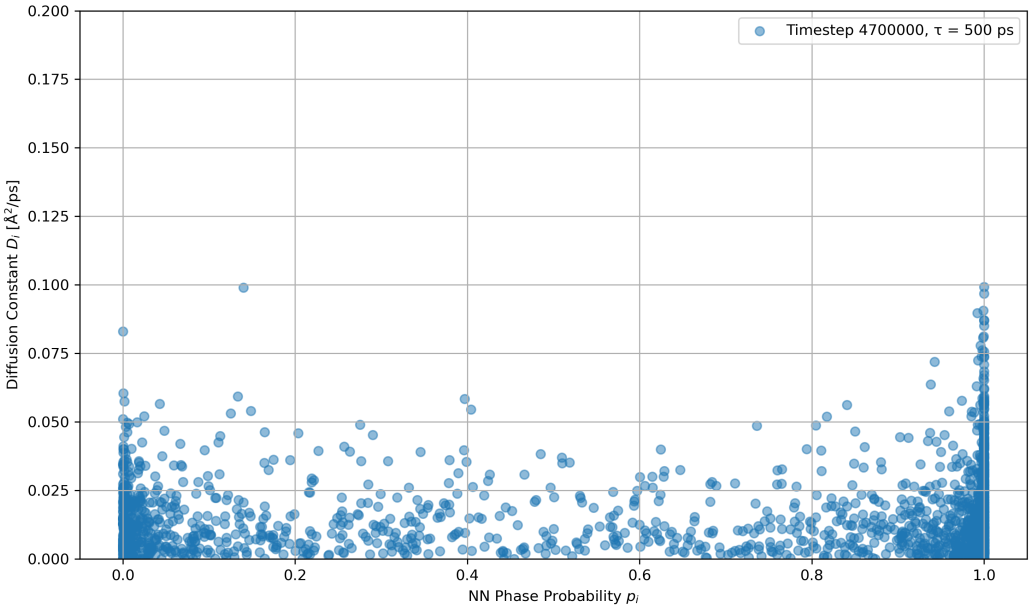

Figure 4.8: Atomwise Diffusion Constant vs NN Output Probability (during phase transition)

4.6 D_i by Phase Probability

In timesteps ranging from 104k to 2m MD steps (see Figures 4.9a - 4.9e, the copper sublattice is predominantly classified as low-chalcocite-like ($LC, p_i < 0.2$), with hardly any atoms meeting the high-chalcocite ($HC, p_i > 0.8$) criterion. The LC histograms in this regime have a sharp peak at $D_i \approx 0\text{Å}^2/ps$, and the vast majority of atoms have diffusion constants well below $10^{-3}\text{Å}^2/ps$. There is only a very short tail in the distributions towards higher D_i, and the frequency of atoms beyond $2 \times 10^{-3}\text{Å}^2/ps$ are probably due to fluctuations.

This pattern reflects a structurally ordered and kinetically frozen low-chalcocite phase, in which the copper atoms remain confined to their lattice positions throughout the 500ps observation period. The absence of HC-classified atoms suggests that no significant high-chalcocite-like local environments have formed at these temperatures, and that there is no detectable mobility contrast between phases.

At timestep 2.5-3.2m 4.9f - 4.9h, the LC population still constitutes the majority, but a significant HC fraction emerges. These HC-classified atoms span a broader D_i range than LC atoms, with a modest but noticeable tail extending beyond $2 \times 10^{-3}\text{Å}^2/ps$. Although the LC remains strongly peaked at very low diffusion constants, its tail also begins to elongate, indicating local mobility enhancement near HC domains.

By timestep 3.4-3.8m 4.9i - 4.9k, the HC population has increased substantially. The HC histogram becomes markedly broader and flatter, with a significant number of atoms exhibiting D_i values approaching $1 \times 10^{-2}\text{Å}^2/ps$. LC atoms now exhibit a much more pronounced tail, and there is greater overlap between the mobility distributions of LC and HC, reflecting spatial heterogeneity. Immobile LC islands coexist with LC-classified atoms in more mobile environments, likely at LC–HC interfaces.

At around timesteps 4.0-4.4m 4.9l - 4.9n, the LC and HC distribution overlap the whole D_i range. The majority of HC atoms exhibit mobilities that far exceed those observed in the initial stage. The LC fraction continues to decline and its distribution broadens further, losing its sharp confinement to the near-zero D_i regime. This indicates that the phase transition is ongoing.

By approximately 4.6 to 4.8m 4.9o - 4.9p, the system demonstrates a predominantly HC-like copper sublattice. The HC histogram exhibits a high degree of uniformity across the range of plotted D_i values, suggesting that copper dynamics resemble those of a fluid. The presence of LC atoms is observed to be limited in number and distributed over a broad range of D_i values. A subset of these atoms exhibit minimal mobility, while others demonstrate mobilities that are comparable to those of HC atoms.

In the final stage of the trajectory, the system exhibits a fully transformed and highly mobile copper sublattice. The histograms demonstrate that the vast majority of copper atoms are classified as high-chalcocite-like ($HC, p_i > 0.8$), with the low-chalcocite (LC) population ($p_i < 0.2$) effectively vanishing.

The HC D_i distribution is extensive and relatively uniform across the entire range up to $10^{-2}\text{Å}^2/ps$, with no significant peak at low values. This is indicative of the absence of a rigid lattice and the presence of liquid-like copper dynamics, where atoms undergo significant displacements over the 500ps observation window. The near-uniform population of bins across the D_i range indicates that atoms explore a wide spectrum of mobilities, from moderate diffusion to very rapid motion, consistent with a disordered, high-temperature state.

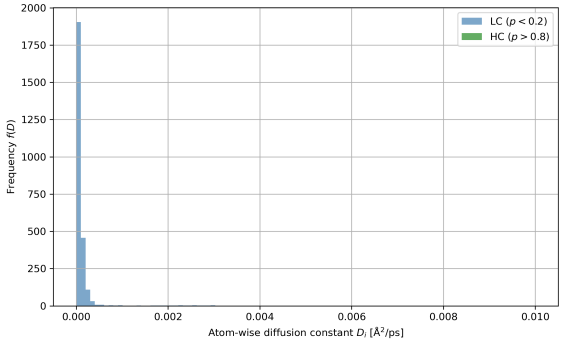

(a) Timestep 104k

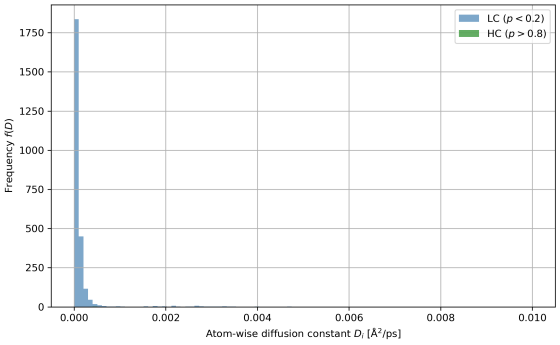

(b) Timestep 500k

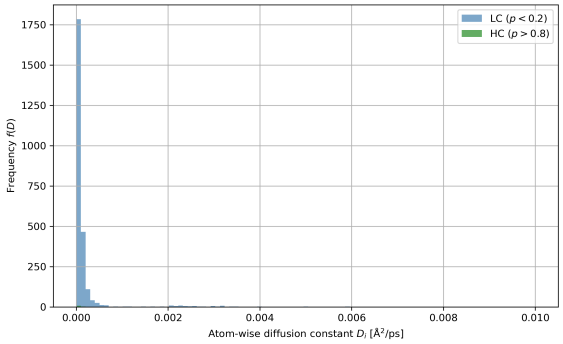

(c) Timestep 1m

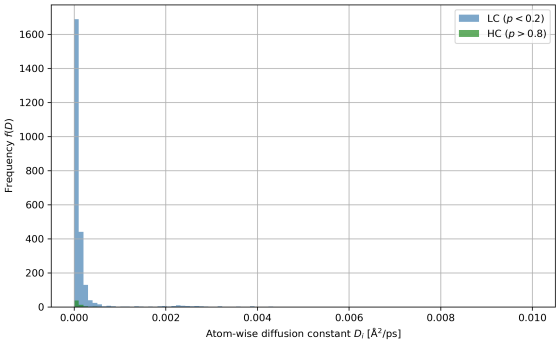

(d) Timestep 1.5m

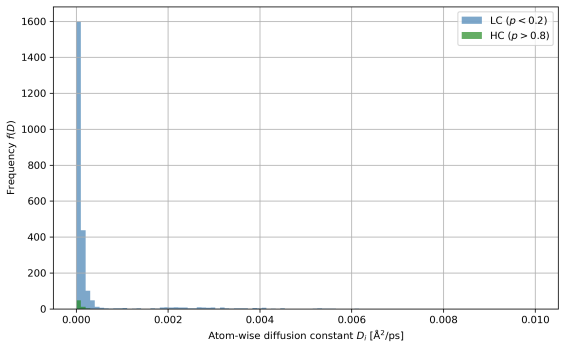

(e) Timestep 2m

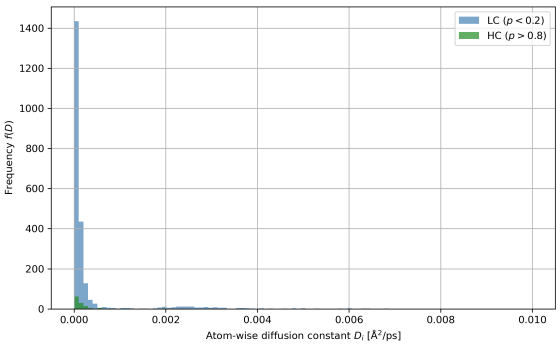

(f) Timestep 2.5m

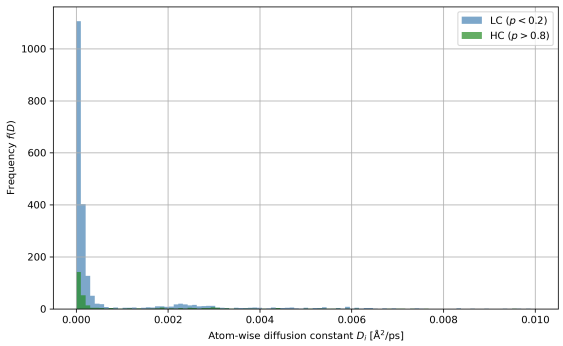

(g) Timestep 3m

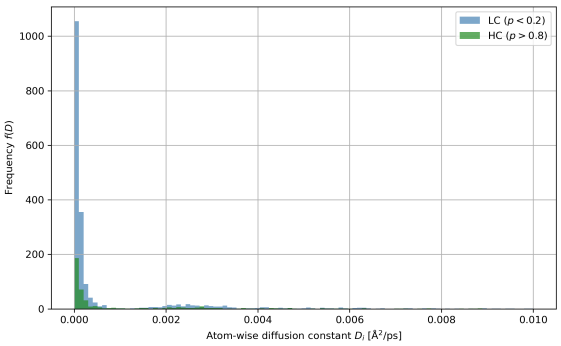

(h) Timestep 3.2m

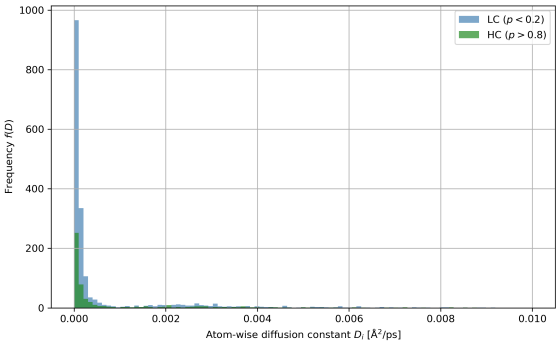

(i) Timestep 3.4m

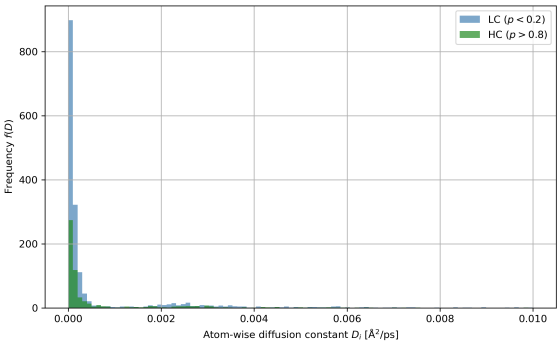

(j) Timestep 3.6m

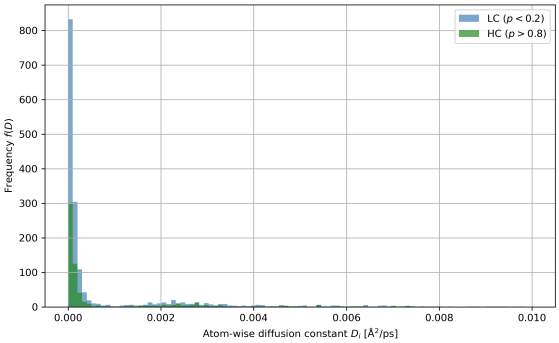

(k) Timestep 3.8m

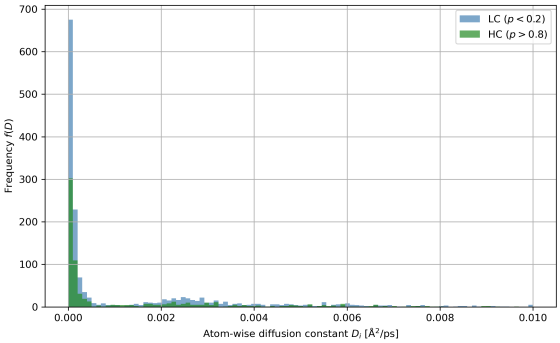

(l) Timestep 4m

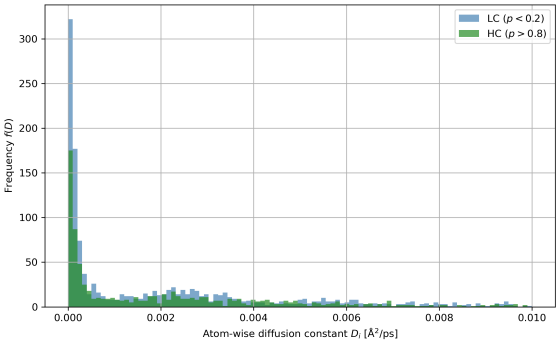

(m) Timestep 4.2m

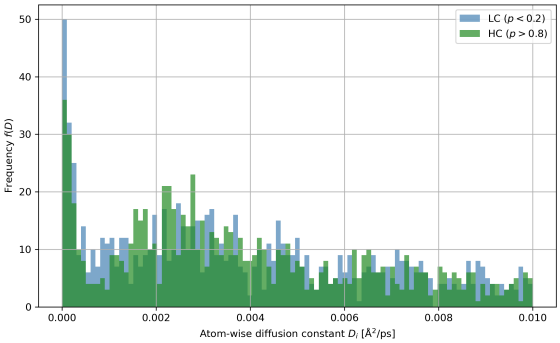

(n) Timestep 4.4m

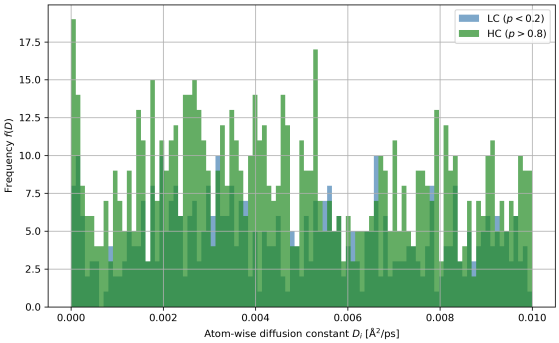

(o) Timestep 4.6m

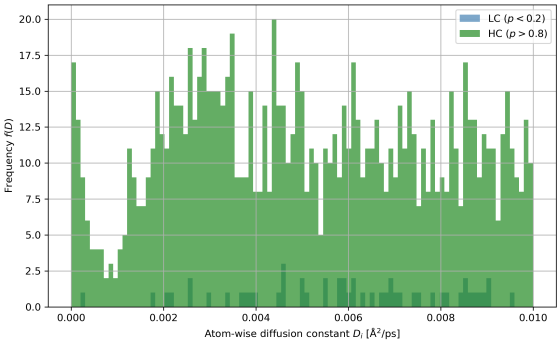

(p) Timestep 4.8m

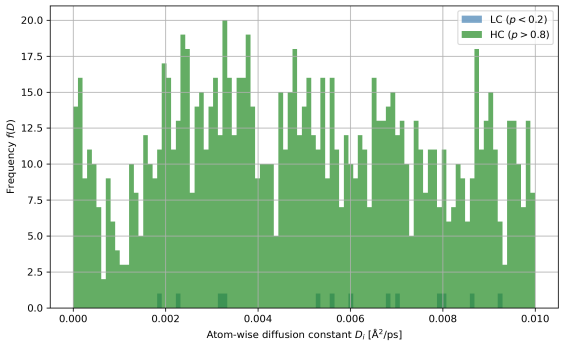

(q) Timestep 5m

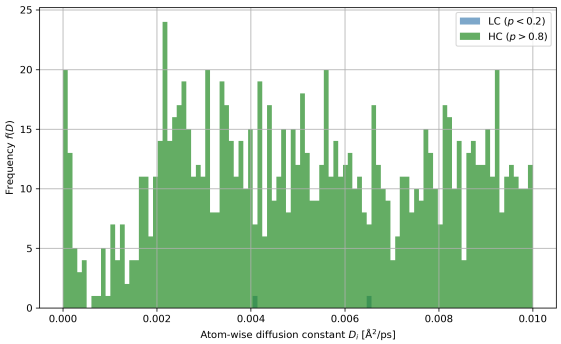

(r) Timestep 5.5m

Figure 4.9: D_i by phase probability

5 Discussion

The PaCMAP analysis, conducted across a wide range of simulation timesteps, provides an unexpected perspective on the microscopic mechanisms underlying the superionic phase transition in Cu2S. Contrary to the prevailing expectation of a classical nucleation process, where a distinct fluid-like phase gradually emerges and grows from localized nucleation sites, the dimensionality-reduced atomic environments reveal a fundamentally different dynamic.

In the initial time periods corresponding to timesteps 200k-500k, the DBSCAN algorithm identifies over 80 distinct clusters, reflecting a rich local environment diversity in the low-chalcocite phase. It is argued that a key reason for the large number of distinct clusters observed in the low-temperature regime could be the phenomenon of symmetry breaking. The low-chalcocite phase of Cu_2S is characterised by a highly ordered structure, exhibiting a restricted number of crystallographically inequivalent copper sites. However, due to thermal motion even at relatively low temperatures, atoms are subject to constant fluctuations around their equilibrium positions. These fluctuations are not perfectly symmetrical across the entire crystal structure, thereby disrupting the ideal symmetry of the crystal on a local scale. The Atom-Centered Symmetry Functions (ACSFs) utilised in this study are characterised by a high degree of sensitivity to such local perturbations. It has been demonstrated that even minor variations in the spatial configuration of neighboring atoms can result in substantial disparities in the ACSF descriptors. Consequently, atoms occupying nominally equivalent sites are assigned to different points in descriptor space and may be separated into distinct clusters by the PaCMAP + DBSCAN pipeline.

The clustering outcome in the ACSF-PaCMAP analysis is intrinsically linked to the design and parameterisation of the underlying symmetry functions. As ACSF encodes local geometry with adjustable resolution, the number, sharpness and spacing of the radial and angular components can either emphasise or mask physically meaningful differences. Therefore, careful tuning is essential in order to strike a balance between sensitivity to structural changes and robustness against thermal noise.

Despite this fragmentation, a clear distinction between the sulfur and copper environments can be seen from the earliest stages. This aligns with the underlying physical structure, whereby the sulfur atoms form a comparatively rigid hexagonal close-packed lattice (Wang[2]), while the copper atoms randomly occupy Wyckoff sites. Over time, a second copper cluster gradually emerges and grows in size. This evolution becomes particularly apparent from timestep 3.8m onwards and continues into the fully developed

high-temperature regime beyond 4.4m, where two dominant copper clusters coexist: one compact and dense and the other more diffuse and varied. This bifurcation strongly suggests a phase transition.

What is striking is the absence of a clear nucleation event, i.e., a well-defined, spatially localized formation of the high-temperature phase embedded in a surrounding low-temperature matrix. Instead, the transition appears to occur collectively and delocalized, with many copper atoms simultaneously shifting their environments across the system. Further trials were conducted with Smooth Overlap of Atomic Positions (SOAP) descriptors and UMAP dimensionality reduction to assess whether alternative local environment encodings could better resolve transient intermediate states. In this instance, SOAP+UMAP and SOAP+PaCMAP did not yield separation as clear or physically interpretable as that obtained with ACSF+PaCMAP. The performance of the clustering process was hindered in many ways. This outcome underscores the critical role of descriptor selection and sensitivity calibration in the detection of the symmetry-breaking and parallel-domain transition signatures observed here.

The combined analyses of the high-chalcocite phase fraction (Section 4.3), copper diffusion evolution (Section 4.4), atomwise diffusion vs. neural network phase probability (Section 4.5, and D_i distributions by phase probability (Section 4.6) present a coherent picture of the phase-transition dynamics in Cu_2S that diverges from a simple, spatially localized nucleation-and-growth model. Instead, the results obtained point towards a delocalised, partially synchronised transformation mechanism.

The temporal evolution of the high-chalcocite fraction reveals a relatively smooth, gradual increase across the transition window, with no single abrupt jump that would correspond to the appearance and growth of a localized nucleus. This observation is consistent with the findings of Zheng Rivest et al. [13] that, in confined systems, high- and low-chalcocite domains can fluctuate dynamically before stabilising, with transition pathways being sensitive to defects and interface energetics. However, while their in situ TEM study resolved nucleation fronts initiating at surfaces or cores, our simulation data suggest that in bulk-like periodic conditions, many copper atoms shift their local environments nearly simultaneously, limiting the detectability of a classical nucleation front.

The copper diffusion analysis provides further support for this theory. The increase in the mean diffusion coefficient $\langle D \rangle$ during the transition does not affect only a small subset of atoms, but instead spans a significant portion of the Cu sublattice. This suggests that there is correlated mobility enhancement rather than localised melting-like behaviour. This differs from the nanorod case studied by Zheng and Rivest [13], in which propagation fronts could be spatially resolved. However, it aligns with their observation

that multiple small domains can fluctuate in parallel.

Joint analysis of atomwise diffusion and phase probability strengthens the link between structural and dynamical changes. While atoms with a high predicted probability of forming high-chalcocite tend to exhibit elevated diffusion constants, this correlation is not absolute. Interestingly, some atoms have low D_i values despite having a high phase probability. This suggests that structural rearrangement towards high-chalcocite coordination does not necessarily require, or immediately lead to high mobility for all sites. This heterogeneity likely reflects transient trapping in energetically favourable local configurations, similar to the defect-pinned domains described by Zheng et al. [13].

Finally, the D_i distributions stratified by phase probability manifest a bimodal structure in the high-probability regime, with one peak at elevated diffusivities and another at low diffusivities. This finding indicates that during the transition, two distinct copper populations coexist within the same structural phase label, one that is dynamically active and one that is relatively immobile. This suggests that mobility activation may lag behind structural transformation for a portion of the sublattice.

The diffusion histograms resolved by neural-network phase probability reveal a transient, phase-selective mobility scale that is naturally interpreted within a CTRW framework. In a CTRW, copper motion is described as a sequence of spatially localised jumps separated by waiting times. If the mean waiting time is $\langle\tau\rangle$ and the mean-square jump length is $\langle\ell^2\rangle$, the self-diffusion coefficient in three dimensions is

$$D = \frac{\langle\ell^2\rangle}{6\langle\tau\rangle} \tag{5.1}$$

At timestep 4.4m a pronounced shoulder/peak appears near $3 \times 10^{-3}\text{Å}^2/ps$ predominantly in the HC-classified subpopulation. With the analysis lag fixed at $\tau = 500ps$, the position of this feature corresponds to a characteristic displacement

$$\ell^\star \;=\; \sqrt{6\,\tau\,D^\star} = \sqrt{6 * 500 * 0.003}\,\text{Å} \approx 3\text{Å} \tag{5.2}$$

.

This identifies the shoulder as a one-hop mode: during the transition onset, a sizeable fraction of copper atoms performs exactly one site-to-site jump of ~ 3Å within $500ps$. The absence of an analogous enhancement around 1×10^{-3}Å$^2/ps$ indicates that $\sim 1-2$Å is not a preferred site-to-site spacing, sub-hop cage vibrations populate a smooth low-D background rather than a peak. The extracted hop length aligns with the Cu–Cu first-neighbour shell (~ 3Å) (see Singraber Figure 7 [1]), supporting the interpretation that the sulfur framework defines the accessible copper sites while kinetics controls whether a full intersite jump occurs within the window.

6 Conclusion

In summary, combined PaCMAP, neural network classification and diffusion analyses reveal that the superionic phase transition in Cu_2S occurs through the delocalised, partially synchronised transformation of the copper sub-lattice, rather than via the classical nucleation and growth process. The evolution of local atomic environments shows that many copper atoms undergo structural reorganisation almost simultaneously, while the sulfur framework remains largely immobile. This behaviour is consistent with a first-order transition, but lacks a spatially localised nucleation front.

The observed heterogeneity in mobility, whereby structurally high-phase copper atoms can remain dynamically immobile, suggests a lag between structural transformation and the activation of full mobility. The extracted characteristic jump length of ~ 3Å corresponds to Cu–Cu nearest-neighbour distances, which supports the idea of discrete site-to-site hopping within the sulfur lattice during the transition.

While this study uses machine-learned potentials to provide a detailed microscopic picture, further investigation is necessary to fully quantify spatial correlations, defect effects and the role of local energetics. In principle, direct ab initio molecular dynamics would offer a more fundamental description, however, the required system sizes and simulation timescales make such calculations too expensive using current computational resources. Instead, future work should explore hybrid approaches, such as combining HDNNPs with targeted ab initio validation, to improve our understanding of superionic transport mechanisms in copper sulfide and related materials.

List of Figures

Bibliography

[1] A.Singraber, T.Morawietz, J.Behler and C.Dellago *Parallel Multistream Training of High-Dimensional Neural Network Potentials*, Journal of Chemical Theory and Computation **2019** 15 (5), 3075-3092 DOI: 10.1021/acs.jctc.8b01092

[2] L.-Wang Wang, *High Chalcocite Cu2S: A Solid-Liquid Hybrid Phase*, Physical Review Letters **2012** 108, 085703 DOI: 10.1103/PhysRevLett.108.085703

[3] J.Behler, *Four Generations of High-Dimensional Neural Network Potentials*, Chemical Reviews **2021** 121 (16), 10037-10072 DOI: 10.1021/acs.chemrev.0c00868

[4] L.Kývala, P.M de Hijes and C.Dellago *Unsupervised identification of crystal defects from atomistic potential descriptors*, **2024** arXiv:2405.01320

[5] D.Chmielowski et al. *Are Binary Copper Sulfides/Selenides Really New and Promising Thermoelectric Materials?*, **2014** DOI: 10.1002/aenm.201301581

[6] C. Coughlan et al. *Compound Copper Chalcogenide Nanocrystals*, **2017** 117 (9), 5865-6109 DOI: 10.1021/acs.chemrev.6b00376

[7] Yingfang Wang et al. *Understanding How Dimension Reduction Tools Work: An Empirical Approach to Deciphering t-SNE, UMAP, TriMAP, and PaCMAP for Data Visualization*, **2021** arXiv:2012.04456

[8] Jörg Behler and Michele Parrinello *Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces*, Physical Review Letters **2007** DOI: 10.1103/PhysRevLett.98.146401

[9] Howard T. Evans, Jr. *Djurleite ($Cu_{1.94}S$) and Low Chalcocite (Cu_2S): New Crystal Structure Studies*, Science Vol. 203, No. 4378 pp. 356- 358 **1979** DOI: 10.1126/science.203.4378.356

[10] Howard T. Evans, Jr. *The crystal structures of low chalcocite and djurleite**, Zeitschrift für Kristallographie 150, 299-320 **1979**

[11] W.Kohn and L.J.Sham *Self-Consistent Equations Including Exchange and Correlation Effects**, Physical Review Vol. 140, No. 4A **1965**

[12] A.Singraber, J.Behler and C.Dellago *Library-Based LAMMPS Implementation of High-Dimensional Neural Network Potentials*, Journal of Chemical Theory and Computation **2019** 15 (3), 1827–1840 DOI: 10.1021/acs.jctc.8b00770

[13] Haimei Zheng et al. *Observation of Transient Structural-Transformation Dynamics in a* Cu_2S *Nanorod*, Science **2011** 333,206-209 DOI:10.1126/science.1204713

[14] Thomas B. Blank et al. *Neural network models of potential energy surfaces*, The Journal of Chemical Physics **1995** 103, 4129–4137 DOI:10.1063/1.469597

[15] Bobby G. Sumpter and Donald W. Noid *Potential energy surfaces for macromolecules. A neural network technique*, Chemical Physics Letters **1992** 192, 455-462 DOI:10.1016/0009-2614(92)85498-Y

[16] Steven Hobday, Roger Smith and Joe Belbruno *Applications of neural networks to fitting interatomic potential functions*, Modelling and Simulation in Materials Science and Engineering **1999** 7, 397412 DOI:10.1088/0965-0393/7/3/308

[17] Steven Hobday, Roger Smith and Joe Belbruno *Applications of genetic algorithms and neural networks to interatomic potentials*, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms **1999** 153, 247263 DOI:10.1016/S0168-583X(99)00057-9