THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Through Rainbow-Tinted Glasses: Machine Learning-Driven Modeling of Chromophores

Eric Lindgren

Department of Physics and Astronomy
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2026

Through Rainbow-Tinted Glasses:
Machine Learning-Driven Modeling of Chromophores
Eric Lindgren

ISBN 978-91-8103-433-2

Doktorsavhandlingar vid Chalmers tekniska högskola.
Ny serie nr 5890
ISSN 0346-718X
DOI https://doi.org/10.63959/chalmers.dt/5890

Department of Physics and Astronomy
Chalmers University of Technology
SE-412 96 Göteborg, Sweden
Telephone +46 (0)31 772 10 00

Cover: A perylene molecule viewed through the author's own pair of rainbow-tinted glasses, surrounded by a neural network, the DYNASOR brontosaurus, and the CALORINE salamander. Composite of a photo and rendered assets with Blender, photo taken and all 3D models made by the author. No "AI" tools were used in the making of this cover.

Typeset in XƎLaTeX using the Memoir class. Template by M. Petisme and J. M. Rahm.

Chalmers digitaltryck
Göteborg, Sweden 2026

Through Rainbow-Tinted Glasses:
Machine Learning-Driven Modeling of Chromophores

ERIC LINDGREN
Department of Physics and Astronomy
Chalmers University of Technology

Abstract

Chromophores are a class of molecules that give color to the world around us, from the chlorophyll in plants that enables photosynthesis to the retinal molecules in our eyes that allow us to see. Chromophores are also fundamental for developing a range of technologies crucial for a transition to a sustainable society, including photovoltaics and energy storage. While chromophores have been widely studied experimentally, we still lack a sufficient understanding of their structure and dynamics on the atomic scale. In particular, many chromophores are glass formers or form supramolecular aggregates due to intermolecular interactions, leading to a complicated landscape of interactions spanning many time and length scales. Addressing this limitation requires atomistic simulations capable of connecting microscopic dynamics to experimentally measurable quantities. In this spirit, this thesis presents a simulation framework that links electronic structure calculations via molecular dynamics simulations to experiments, with a focus on neutron scattering.

The key ingredient in this framework is machine-learned interatomic potentials, enabling simulations with the accuracy of quantum mechanical calculations for large systems of chromophores. Methodological developments focus on the neuroevolution potential framework implemented in the GPUMD package. A major contribution is the development of the CALORINE package, which is a companion software for GPUMD that interfaces with the broader scientific software ecosystem.

The framework is applied to three challenging applications: glass formation, optical response, and neutron scattering. First, glass formation occurs beyond timescales accessible to molecular dynamics simulations, and this limitation is circumvented by extrapolating relaxation processes from nanoseconds to experimental timescales using Bayesian regression. Second, optical properties are obtained by using the neuroevolution potential framework to predict tensorial properties such as the dipole moment, or spectral quantities such as the electronic dielectric function. Finally, instrument-specific inelastic neutron scattering signatures are predicted using electronic structure calculations, machine learning, and correlation functions. Together, these developments establish a framework for connecting atomistic simulations with experimental observables, enabling modeling of chromophores over multiple time and length scales. The framework is transferable and directly applicable to other systems.

Keywords: chromophores, machine learning, machine-learned interatomic potentials, molecular dynamics, neutron scattering, optical absorption, foundation models

LIST OF APPENDED PAPERS

This thesis consists of a background to the field and the following papers. It is partly based on the author's licentiate thesis (E. Lindgren, *Shedding light on liquid chromophores using machine learning* (2024)).

- **I** **Probing Glass Formation in Perylene Derivatives via Atomic Scale Simulations and Bayesian Regression**
  Eric Lindgren, Jan Swenson, Christian Müller, and Paul Erhart
  The Journal of Physical Chemistry B **129** (26), 6613–6619 (2025)
- **II** **GPUMD: A package for constructing accurate machine-learned potentials and performing highly efficient atomistic simulations**
  Zheyong Fan, Yanzhou Wang, Penghua Ying, Keke Song, Junjie Wang, Yong Wang, Zezhu Zeng, Ke Xu, **Eric Lindgren**, J. Magnus Rahm, Alexander J. Gabourie, Jiahui Liu, Haikuan Dong, Jianyang Wu, Yue Chen, Zheng Zhong, Jian Sun, Paul Erhart, Yanjing Su and Tapio Ala-Nissila
  The Journal of Chemical Physics **157** (11), 114801 (2022)
- **III** **GPUMD 4.0: A high-performance molecular dynamics package for versatile materials simulations with machine-learned potentials**
  Ke Xu, Hekai Bu, Shuning Pan, **Eric Lindgren**, Yongchao Wu, Yong Wang, Jiahui Liu, Keke Song, Bin Xu, Yifan Li, Tobias Hainer, Lucas Svensson, Julia Wiktor, Rui Zhao, Hongfu Huang, Cheng Qian, Shuo Zhang, Zezhu Zeng, Bohan Zhang, Benrui Tang, Yang Xiao, Zihan Yan, Jiuyang Shi, Zhixin Liang, Junjie Wang, Ting Liang, Shuo Cao, Yanzhou Wang, Penghua Ying, Nan Xu, Chengbing Chen, Yuwen Zhang, Zherui Chen, Xin Wu, Wenwu Jiang, Esmée Berger, Yanlong Li, Shunda Chen, Alexander J. Gabourie, Haikuan Dong, Shiyun Xiong, Ning Wei, Yue Chen, Jianbin Xu, Feng Ding, Zhimei Sun, Tapio Ala-Nissila, Ari Harju, Jincheng Zheng, Pengfei Guan, Paul Erhart, Jian Sun, Wengen Ouyang, Yanjing Su, and Zheyong Fan
  Materials Genome Engineering Advances **3** (3), e70028 (2025)
- **IV** **calorine: A Python package for constructing and sampling neuroevolution potential models**
  Eric Lindgren, Magnus Rahm, Erik Fransson, Fredrik Eriksson, Nicklas Österbacka, Zheyong Fan, and Paul Erhart
  Journal of Open Source Software **9** (2024)
- **V** **NEP89: Universal neuroevolution potential for inorganic and organic materials across 89 elements**
  Ting Liang, Ke Xu, **Eric Lindgren**, Zherui Chen, Rui Zhao, Jiahui Liu, Esmée Berger, Benrui Tang, Bohan Zhang, Yanzhou Wang, Keke Song, Penghua Ying, Nan Xu, Haikuan

Dong, Shunda Chen, Paul Erhart, Zheyong Fan, Tapio Ala-Nissila, and Jianbin Xu
In review, (2025)

VI **Tensorial Properties via the Neuroevolution Potential Framework: Fast Simulation of Infrared and Raman Spectra**
Nan Xu, Petter Rosander, Christian Schäfer, **Eric Lindgren**, Nicklas Österbacka, Mandi Fang, Wei Chen, Yi He, Zheyong Fan, and Paul Erhart
Journal of Chemical Theory and Computation **20** (8), 3273–3284 (2024)

VII **Optical response of silica across temperature, pressure, and disorder via pretrained atomistic representations**
Tobias Hainer, **Eric Lindgren**, Lucas Svensson, Julia Wiktor, and Paul Erhart
In manuscript

VIII **Predicting neutron experiments from first principles: A workflow powered by machine learning**
Eric Lindgren, Adam J. Jackson, Erik Fransson, Esmée Berger, Goran Škoro, Svemir Rudić, Rastislav Turanyi, Sanghamitra Mukhopadhyay, and Paul Erhart
Journal of Materials Chemistry A **13** (31), 25509–25520 (2025)

The author's contribution to the papers:

I I designed and performed the molecular dynamics simulations and Bayesian extrapolation, calculated and analyzed the correlation functions, wrote the first draft, and finalized the paper with my co-authors.

II I am one of the developers of the GPUMD package, which we present in this paper. My main contribution was parts of the initial version of the CALORINE package, a companion software presented together with GPUMD, and I assisted in manuscript preparation and revision.

III Here we present version 4 of the GPUMD package. I contributed, e.g., with the development of the active learning module and improvements to the MSD module, and I helped revise the manuscript.

IV I am the main developer of the CALORINE paper, and I wrote the first draft and finalized the paper with my co-authors.

V I performed fine-tuning for crystalline benzene and provided the simulated inelastic neutron scattering example to the paper. I also contributed to the final editing and revision process of the paper.

VI I implemented on-the-fly calculation of tensorial properties during molecular dynamics, enabling efficient large scale evaluation without saving the full trajectory, and helped revise the paper.

VII I primarily explored other model types, investigated the extension to molecular systems, performed a large part of the analysis in the Supplementary Information, and helped finalize the paper together with the main author.

VIII I performed DFT calculations and trained the NEP model for crystalline benzene, performed all molecular dynamics simulations, calculated and post-processed the dynamic structure factors, wrote the first draft, and finalized the paper with my co-authors.

PUBLICATIONS NOT INCLUDED IN THIS THESIS

General-purpose machine-learned potential for 16 elemental metals and their alloys
Keke Song, Rui Zhao, Jiahui Liu, Yanzhou Wang, **Eric Lindgren**, Yong Wang, Shunda Chen, Ke Xu, Ting Liang, Penghua Ying, Nan Xu, Zhiqiang Zhao, Jiuyang Shi, Junjie Wang, Shuang Lyu, Zezhu Zeng, Shirong Liang, Haikuan Dong, Ligang Sun, Yue Chen, Zhuhua Zhang, Wanlin Guo, Ping Qian, Jian Sun, Paul Erhart, Tapio Ala-Nissila, Yanjing Su, and Zheyong Fan
Nature Communications **15**, 10208 (2024)

Machine Learning for Polaritonic Chemistry: Accessing chemical kinetics
Christian Schäfer, Jakub Fojt, **Eric Lindgren**, and Paul Erhart
Journal of the American Chemical Society **146** (8), 5402–5413 (2024)

Dynasor 2: From simulation to experiment through correlation functions
Esmée Berger, Erik Fransson, Fredrik Eriksson, **Eric Lindgren**, Göran Wahnström, Thomas Holm Rod, and Paul Erhart
Computer Physics Communications **316**, 109759 (2025)

List of abbreviations

ACF autocorrelation function. 15, 16, 43, 44, 49

CPU central processing unit. 24

DFT density functional theory. 2, 21–25, 33, 35, 36, 39–41, 44, 45, 47, 48, 57, 61, 62

FF force field. 21–25, 40, 41, 57, 58, 60
FFT fast Fourier transform. 16

GNN graph neural network. 36, 47, 48
GPU graphics processing unit. 30, 36–38, 63

INS inelastic neutron scattering. 54–62
IR infrared. 15, 43–46, 59, 60

MD molecular dynamics. 2, 6, 8, 11, 12, 14–17, 19–24, 28, 33–38, 41–45, 47, 50, 53–56, 58–63
ML machine learning. 2, 20, 27, 46
MLIP machine-learned interatomic potential. 2, 20, 21, 24–27, 33, 38, 39, 41, 43, 46, 55, 59–63
MSD mean squared displacement. 16, 17, 40, 42

NACF normal-vector autocorrelation function. 17–19, 60
NEP neuroevolution potential. 2, 21, 24–31, 34–48, 55–63
NN neural network. 25–27, 29, 32, 43

OLED organic light emitting diode. 3–6
OSC organic solar cell. 3–6, 17, 62

PBC periodic boundary conditions. 12, 35
PCA principal component analysis. 29, 30

QENS quasi-elastic neutron scattering. 16, 54, 55, 61

RMSE root mean squared error. 31, 33, 35, 36, 40, 41

SNES separable natural evolution strategy. 31, 34

TNEP tensorial neuroevolution potential. 43–46

VFT Vogel–Fulcher–Tammann. 9, 17, 19

XC exchange–correlation. 23, 47, 48, 61

Contents

1

Introduction

There is something mesmerizing about colors. Their alluring appeal has fascinated mankind throughout history, including the modern materials scientist. This is not without reason. A colorful substance is not only pretty to look at but promises interesting optical properties that can be used in devices. One class of organic molecules that live up to this promise are chromophores, from the ancient Greek for "color-bearer", reflecting the role of the part of a molecule responsible for color. Naturally occurring examples include chlorophyll, enabling photosynthesis in plants, retinal molecules in our eyes facilitating vision, and β-carotene, responsible for the orange color of carrots and autumn leaves. The color arises from conjugated systems, π-bonded atoms over which electrons are delocalized [1, Chapter 7]. The extent of the system determines the wavelength of incoming light that is absorbed, acting as a mini-antenna.

Chromophores seldom come alone; they tend to aggregate [2–4], for example through π–π-stacking between neighboring molecules. These groupings are known as supramolecular aggregates, formed through intermolecular interactions *between* molecules. Intermolecular interactions are relatively weak compared to the covalent bonds within individual molecules (intramolecular bonds), and are mediated by electrostatic Coulomb interaction, hydrogen bonding, or van der Waals forces [5]. The combination of intra- and intermolecular interactions in chromophores leads to a veritable mosh pit of atoms [6], with a broad range of relevant time and length scales. These span from atomic vibration and torsion of molecules on the order of Ångström and picoseconds, to molecular aggregation and glass formation over micrometers and hours. Crucially, these processes affect the structure and optical properties of chromophores, which in turn influence device performance and the material's industrial processability. Chromophores play a key role in the transition to a sustainable society, with applications including organic semiconductors [7–9], notably solar cells [10–12] and organic light-emitting diodes [13–15], as well as solar thermal storage systems [16–21], all relying on tunable optical properties. Computer simulations can help shed light on the structure and dynamics of chromophores, which hold the key to controlling their optical properties.

Simulating these multiscale systems is computationally demanding and requires combining multiple techniques. Molecular dynamics (MD) simulations can capture multiscale dynamics up to microseconds, but their predictive power depends on force accuracy. Traditionally, forces are obtained either from empirical interatomic potentials with limited transferability, or from accurate but computationally expensive electronic-structure calculations such as density functional theory (DFT). Over the last two decades, machine-learned interatomic potentials (MLIPs) have emerged as a compelling solution [22–28], providing near-DFT accuracy at the cost of empirical interatomic potentials. Accurate MD simulations can be directly compared to experiment through correlation functions, linking, e.g., neutron scattering to the time-dependent particle density, and infrared and Raman spectroscopy to the dipole moment and polarizability. In this framework, machine learning (ML) bridges simulations and experiments, improving our understanding of the interplay between structure, dynamics, and optical properties in chromophores.

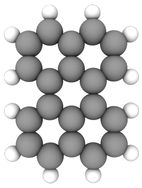

Perylene
$C_{20}H_{12}$

The aim of this thesis is to develop a simulation framework combining MD and MLIPs based on the neuroevolution potential (NEP) framework to study the structure and dynamics of chromophores, and connect these results to experimental observables. Perylene-based chromophores ($C_{20}H_{12}$) form an ideal model system, due to their tendency to aggregate, their relatively simple molecular structure allowing side-chains and substitutions, and their extensive experimental characterization. Three main research questions are in focus:

> What is the accuracy and reliability of NEP models for capturing intra- and intermolecular interactions in perylene derivatives?
>
> How effectively does the simulation protocol capture the glass-forming behavior of perylene derivatives?
>
> To what degree can the simulation framework predict experimental observables from neutron scattering and optical spectroscopy?

This thesis consists of five main chapters, and a summary and outlook. Chapter 2 discusses the relation between chromophore structure, dynamics, and device performance, followed by Chapter 3, which focuses on MD simulations and the link to experiments via correlation functions. Chapter 4 details the theory of MLIPs with a specific focus on the NEP framework, while Chapter 5 and Chapter 6 showcase two applications of the workflow, predicting optical signatures and neutron scattering experiments, respectively. The thesis concludes with a summary and outlook on further developments.

2

Applications of perylene-based chromophores and organic glasses

> As usual, computer simulations are a priori severely limited in such an endeavour.
>
> *Modern computational studies of the glass transition*
> *Berthier et al. Ref. [29]*

This chapter discusses the interplay between intermolecular interactions and optical performance in chromophores, viewed through the lens of organic solar cells (OSCs) and organic light emitting diodes (OLEDs) as example applications. Extra emphasis will be placed on how molecular aggregation and glassy dynamics jointly control optoelectronic performance and industrial processability through atomic-scale interactions (Fig. 2.1).

2.1 Chromophores for organic electronic devices

Perylene-based chromophores belong to the broader class of organic semiconductors, which in recent decades have garnered significant attention in electronics manufacturing due to their natural abundance, low-cost production through techniques such as roll-to-roll printing, and mechanical flexibility allowing, e.g., devices to be integrated into clothing [30–33]. Optoelectronic performance is of central importance in applications of organic semiconductors, and is governed by the interactions between molecules. Intermolecular interactions influence molecular aggregation, ultimately affecting how the system couples to and interacts with light. Two representative device classes that highlight this interplay between structure and optoelectronic properties are OSCs and

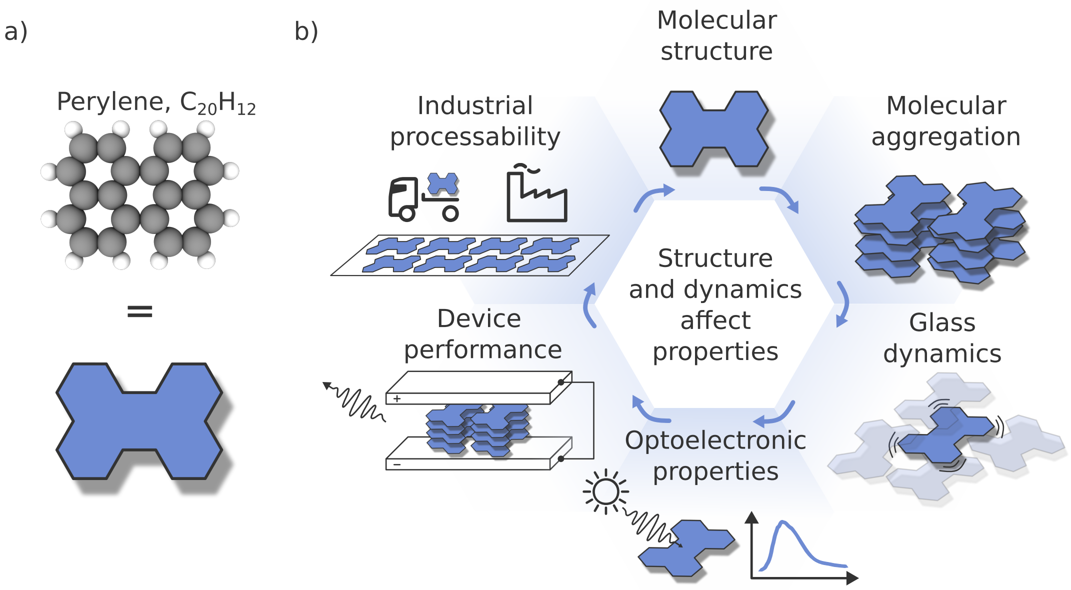

Figure 2.1: (a) Perylene ($C_{20}H_{12}$), the molecule considered as a prototype system in this thesis. (b) A schematic of the interplay between molecular structure, aggregation, and dynamics, and how these affect application performance in perylene-based organic semiconductors. These properties are interlinked; device performance and industrial processability set requirements on the molecular structure and dynamics, which in turn feed back into the optoelectronic properties. A simulation framework for chromophores must therefore integrate molecular structure, packing, dynamics, and optoelectronics in order to model the nuances of organic semiconductors.

OLEDs. The basic operating principles in these two types of devices are complementary, and can at a high level be viewed as each other's opposites; an OLED is an OSC run in reverse, with visible light emission instead of absorption.

The central component of an OSC is the active light-absorbing material, which often is a chromophore. The specific chromophore can vary, ranging from small molecules such as perylenediimide [12, 34–36], via larger molecules such as the Y6-family [37, 38], to polymers such as *p*-phenylene vinylene [39, 40] (Fig. 2.2a). The optical band gap determines which parts of the incoming sunlight are absorbed[1]; photons with an energy above the band gap will be absorbed, exciting an electron from the valence band to the conduction band, forming a localized electron-hole pair known as an exciton. A second semiconductor with a lower-lying conduction band than the light absorbing material is then introduced, acting as an electron acceptor that splits the exciton into a free electron and a free hole that generate a current. Examples of acceptor materials include fullerenes [31, 39], polymers, and molecules such as ITIC and Y6 [41] (Fig. 2.2b).

---

[1] The optical absorption coefficient is a frequency-dependent quantity determined by the electronic dielectric function, which we will return to in Chapter 5.

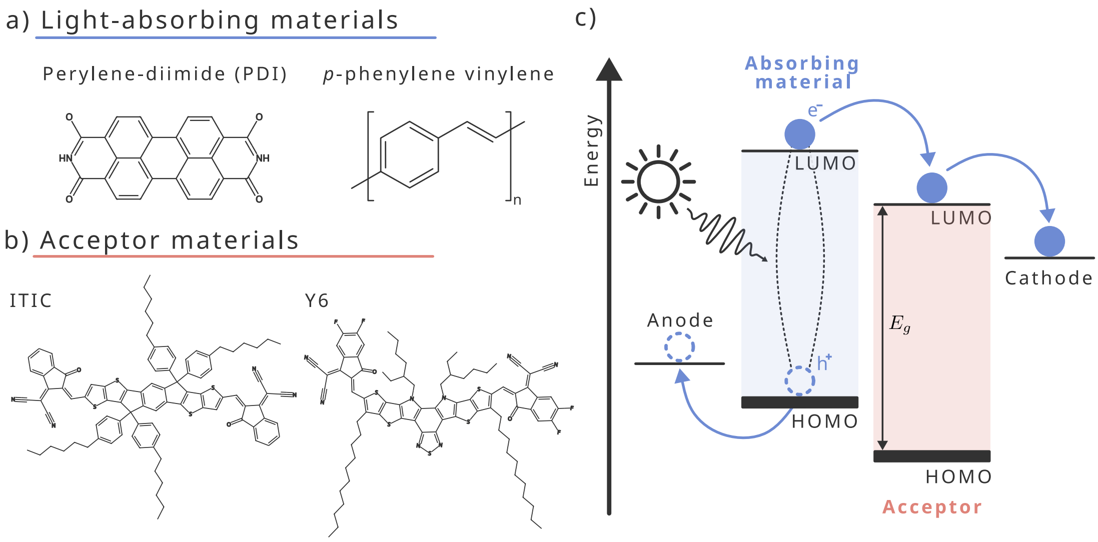

Figure 2.2: (a) Examples of light absorbing materials and (b) electron acceptor materials. (c) Schematic of the exciton splitting in OSCs, where the exciton formed by the electron-hole pair is split and the electron is transferred to the electron acceptor. E_g refers to the optical band gap.

The electron acceptor and light absorbing material are typically mixed to form a blend known as a bulk heterojunction, in order to facilitate efficient exciton splitting and thus high device performance. It is important to control the positions and orientations of the donor- and acceptor molecules in the blend, i.e., the morphology of the blend, for optimal solar cell performance [41, 42]. For example, perylene-based acceptors like perylene-diimide tend to self-assemble into large supramolecular aggregates [35, 43, 44], decreasing the solubility and the effective interface size between the acceptor and donor molecules, leading to a decrease in the exciton-splitting efficiency. One strategy for controlling self-assembly is side-group engineering, where side-groups are grafted onto the core molecule in order to reduce or enhance intermolecular interactions [4, 45]. However, side-group engineering might have adverse effects; in perylene-based chromophores this can lead to a twisting of the perylene core, in turn affecting the charge-transport properties of the material, which benefit from molecules being packed closer together [46]. This tradeoff between solubility and charge transfer can be observed in many electron acceptor/donor systems [31]. Controlling the supramolecular structure is thus key when designing improved OSCs, which in turn demand controlling the interactions between the molecules in the cell [42].

Controlling molecular morphology is also key in designing efficient OLEDs that only emit light in a narrow band around the desired frequency [47, 48]. In an OLED, an externally applied voltage injects electrons and holes from the cathode and anode, respectively, which will move through the device and form excitons in the photoactive emissive layer. These excitons then recombine, emitting a photon with a frequency determined by the band gap of the chosen chromophore. Supramolecular aggregation through π-π-stacking can decrease the efficiency of the OLED, and molecular rearrangement as well as intramolecular vibrations can lead to a broadening of the emitted spectra [48].

As we have seen, intermolecular interactions play a key role for controlling the optoelectronic performance of OLEDs and OSCs, as in many other applications of organic semiconductors. Semiconductor design therefore requires an atomic-scale understanding of structure and dynamics. In the next section, we will turn to how the interplay between electronic properties and molecular structure can be modeled.

2.2 Molecular structure affects electronic properties

The applications of chromophores in OLEDs and OSCs highlight the importance of controlling molecular aggregation to achieve high device performance. Accurate theoretical models describing how aggregation influences optical response are therefore needed. Already in the 1960s, Kasha proposed the theory of H–J aggregates, based on the Coulomb coupling between two molecules [49–52]. Compared to the single molecule, Coulomb coupling leads to a splitting of the S_1 excited state into $S_1^{\uparrow\uparrow}$ and $S_1^{\uparrow\downarrow}$ in a dimer. However, only one of these excited states will be optically visible, depending on the configuration of the two molecules. In a face-to-face configuration, known as an H-aggregate, the higher energy state with parallel dipole moments $S_1^{\uparrow\uparrow}$ is visible, leading to a blue shift to higher energies of the $S_0 \rightarrow S_1$ transition compared to the free monomer (Fig. 2.3b). Arranging the molecules head-to-tail results in a J-aggregate, in which the $S_0 \rightarrow S_1$ transition is red-shifted compared to the free monomer. It is possible to transition between these two extremes, by translating one of the molecules along the other (Fig. 2.3a).

Kasha's H-J theory can be extended to aggregates containing more than two molecules, as well as to include more excited states $S_2, S_3, \ldots$ [52]. However, although H-J theory has been successful in qualitatively describing the optical properties of chromophore aggregates, it fails to describe the dynamic, disordered, and vibronically coupled aggregates encountered in real devices. Specifically, the vibrations of the atoms within the molecule, wavefunction overlap, and charge transfer between neighboring molecules are neglected. A number of theoretical approaches have been developed to incorporate these effects, often focusing on either vibrations or electronic excitations, but a unified framework that simultaneously treats molecular structure, dynamics, and optical response with atomic resolution in large, disordered systems remains challenging [52, 57–61]. Atomistic modeling techniques such as molecular dynamics (MD) provide a detailed description of the structure and dynamics of molecular aggregation. When combined with electronic-structure models, and especially recent machine learning–based approaches, such simulations offer a promising route to connecting molecular-scale structure and dynamics with optical properties, as will be explored in Chapter 5.

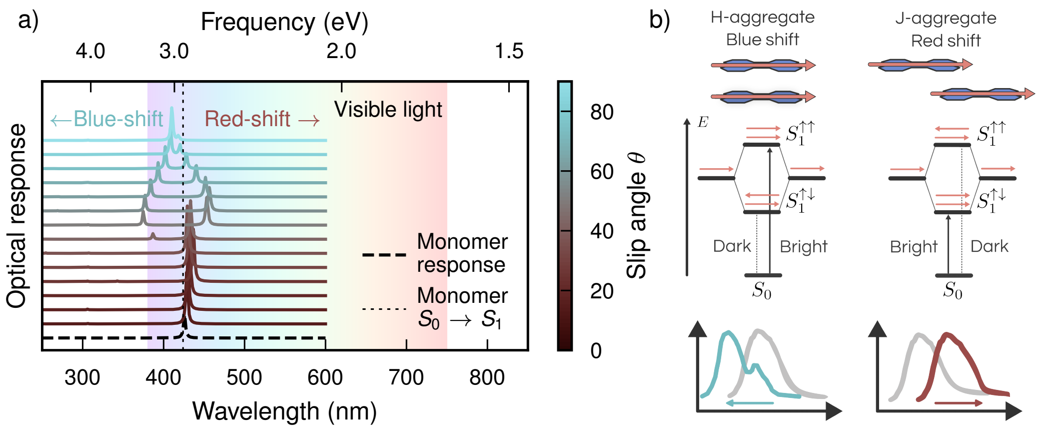

Figure 2.3: (a) Optical spectrum for a perylene dimer as a function of slip angle θ, from time-dependent density functional theory calculations performed with the VeloxChem package [53–55] using the CAM-B3LYP functional [56] and the DEF2-SV(P) basis set. The slip angle is defined to be 90° when the molecules are face-to-face, an H-aggregate, and zero when head-to-tail, a J-aggregate. The $S_0 \to S_1$ transition is red-shifted for the J-aggregates, and slightly blue-shifted for H-aggregates, as predicted by Kasha theory. Interestingly, the lower excited state predicted to be dark by Kasha theory is visible in the time-dependent density functional theory calculations, highlighting the limitations of Kasha theory. (b) Schematic of a perylene dimer in an H- and J-aggregate configuration and their corresponding spectral shifts.

2.3 Glass dynamics in chromophores — processability and stability

Optimizing and modeling the molecular nanostructure is required for understanding optoelectronic performance of chromophore-based semiconductors, but also for ensuring their ease of manufacture, their processability, and long-term device stability. Organic semiconductors are often designed to be in an amorphous glassy state lacking long-range order, where the molecular packing can be carefully controlled. However, the metastable nanostructure can be disturbed during processing or operation of the device. For example, chromophore-based organic semiconductors can benefit from cheap manufacturing techniques such as ink-jet and roll-to-roll printing, slot-die coating, or vapor deposition, allowing for large-scale production [30, 62, 63]. Directly after processing, the material is usually in a non-equilibrium state, and the molecular structure will relax over time to reach a lower energy configuration, which can affect the optoelectronic properties. Another challenge is crystallization, which drives the molecular structure toward equilibrium and can be detrimental for both optoelectronic device performance and industrial processability. Controlling the properties of organic glasses is therefore required to effectively manufacture organic semiconductor devices.

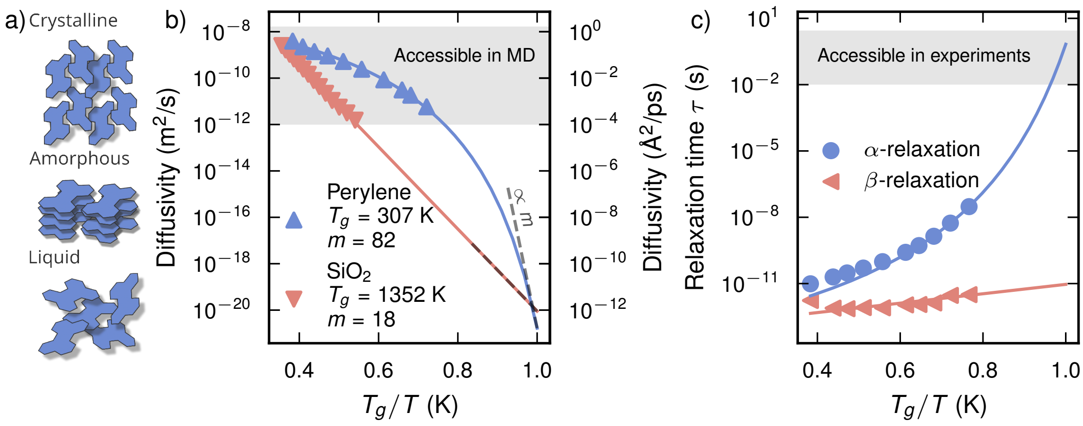

Figure 2.4: (a) Schematic of three different phases of organic glass formers. The crystalline phase has long-range order, whilst the amorphous phase consists of aggregates of molecules. Depending on the specifics of the material, it will either be in the crystalline or metastable amorphous phase at low temperatures, $T < T_g$. The system will be in a disordered, liquid phase at high temperatures, $T > T_g$. (b) Angell plot of the diffusivity as a function of inverse temperature for a strong glass former, SiO$_2$, and a weak glass former, perylene. Correspondingly, SiO$_2$ has a low kinetic fragility, and perylene has a large kinetic fragility. These results were obtained from MD simulations, which will be the focus of Chapter 3. (c) Examples of α- and β-relaxation in liquid perylene, simulated with MD. The α-relaxation is well described by a Vogel-Fulcher-Tammann equation, and the β-relaxation follows Arrhenius behavior. In this case, the α process corresponds to molecular reorientation, and the β process intramolecular atomic vibrations.

Two main properties characterize a glass former: the glass transition temperature T_g and the kinetic fragility m [64–66]. The glass transition temperature T_g is the temperature below which the conformation of the glass former is essentially fixed on experimental timescales. In chromophore systems, this means that molecular reorganization and polymer segment motion are effectively frozen below T_g. The characteristic timescale of these relaxation processes, often denoted τ, increases dramatically with temperature as T_g is approached from above, resulting in a slowdown of the dynamics. Consequently, in organic semiconductor devices, T_g sets an upper bound on the operational temperature, since above T_g the molecular structure tuned to achieve the desired optoelectronic properties can be disrupted when molecular mobility is restored [30].

The kinetic fragility m describes how the relaxation processes in the glass former slow down when approaching T_g [67–69]. The fragility is defined as the slope of the relaxation time τ, typically represented by the viscosity, of the material as T_g is approached from above [64],

$$m = \left.\frac{\partial \log_{10} \tau}{\partial (T_g/T)}\right|_{T=T_g}. \tag{2.1}$$

Above the melting temperature the relaxation times follow Arrhenius behavior,

$$\tau(T) = \tau_0 e^{-E_a/k_B T}, \tag{2.2}$$

where τ_0 is a proportionality constant, k_B is the Boltzmann constant, and E_a is the apparent activation energy of the process. However, when the temperature approaches T_g the relaxation time might grow faster than stipulated by the Arrhenius equation. The kinetic fragility corresponds to the degree of deviation from Arrhenius behavior, where a larger m corresponds to a larger deviation (Fig. 2.4b). Such glass formers are known as *fragile* glass formers. Conversely, materials with a low fragility are known as strong glass formers. Many organic molecules, including chromophores, polymers, and pharmaceutical molecules, are fragile glass formers, while strong glass formers include, for example, silica, SiO_2 [66, 70, 71]. Because of this deviation from Arrhenius behavior, the relaxation time in fragile glass forming systems is often described by the Vogel–Fulcher–Tammann (VFT) equation [64],

$$\tau(T) = \tau_0 e^{-B/k_B(T-T_{\mathrm{VF}})}, \tag{2.3}$$

where B is a pseudo-activation energy, and T_{VF} is a constant known as the Vogel–Fulcher temperature. Generally, the slowest relaxation process in a glassy system obeys the VFT equation, and is known as primary or α-relaxation. This is typically a local process, encompassing phenomena such as molecular rearrangement and relaxation of short polymer chain segments, though some polymer glasses may also exhibit slower relaxation of the entire polymer chain, known as normal mode relaxation [72]. Faster processes, such as intramolecular motions like atomic vibration, generally follow the Arrhenius equation, and do not freeze in at T_g. These are known as secondary or β-relaxation processes (Fig. 2.4b) [64, 73]. Note that the VFT equation is an empirical model that may have limited accuracy when extrapolating to low temperatures [74]. Glass dynamics can be modeled more explicitly using, e.g., mode coupling theory and Monte-Carlo methods [29, 75]. With this in mind, we will focus on the VFT equation due to its simplicity.

The kinetic fragility depends heavily on the interactions between molecules in the material, and thus the glass characteristics can vary dramatically between seemingly similar molecular systems. An example is the effect of mixing; perylene-based chromophores are known to be fragile glass formers in single-component systems, but upon mixing the fragility decreases dramatically, down to becoming a strong glass former for eight-component mixtures [76]. Mixing is thus a method that can be exploited to engineer the viscosity and related properties for processing and device stability. Another option is side-chain engineering, where alkyl side chains can be grafted onto the core chromophore molecule to improve solubility in organic solvents, required for, e.g., solution processing [30]. However, both of these options may affect the molecular aggregation behavior and the optoelectronic properties of the material. In practice, optoelectronic properties and industrial processability are interlinked, where optimizing one sets constraints on the other.

As we have seen in this chapter, intermolecular interactions determine molecular aggregation. Molecular aggregation in turn affects the optical properties, and glassy dynamics affects the stability and processability of the molecular aggregates. These three aspects therefore cannot be optimized independently of each other, and we thus need a simulation approach that encompasses atomic and molecular structure, optoelectronic properties, and glass dynamics. Furthermore, many of the relevant properties that we have encountered in this chapter, such as optical response, structure, and dynamics, are experimentally accessible through optical spectroscopy and neutron scattering. However, these experimental techniques provide ensemble-averaged information, making it difficult to disentangle the role of specific molecular arrangements. Atomistic simulations can bridge the gap by connecting experimental observables to molecular-level structure and dynamics, and will be the focus of the next chapter.

3

Simulating chromophores with molecular dynamics

Atoms go brrr

Nicklas, 2024

In this chapter we will acquaint ourselves with molecular dynamics (MD) simulations as a tool for studying chromophore structure and dynamics. The previous chapter highlighted the complex intermolecular dynamics of chromophores, and the need for an accurate description of these interactions to predict optical properties and glass-forming behavior. MD is well suited to studying several of these dynamic properties, including atomic vibrations, diffusion, and molecular aggregate formation. Glass dynamics, however, are challenging to model due to the limited timescales accessible in MD (Fig. 3.1a), and require special treatment.

We begin by describing MD with a focus on computation of observables such as correlation functions. Correlation functions selectively probe certain time and length scales, and offer a direct link to experiments. Next, we apply MD to simulate glass formation in perylene derivatives. The chapter concludes with a brief remark on the different methods available for extending the time scales in computer simulations beyond what is accessible in MD.

3.1 Modeling dynamics in an atomic system

MD is a simulation technique in which the position and velocity of each atom in the system is evolved in time using small increments, the time step, Δt. At each time step,

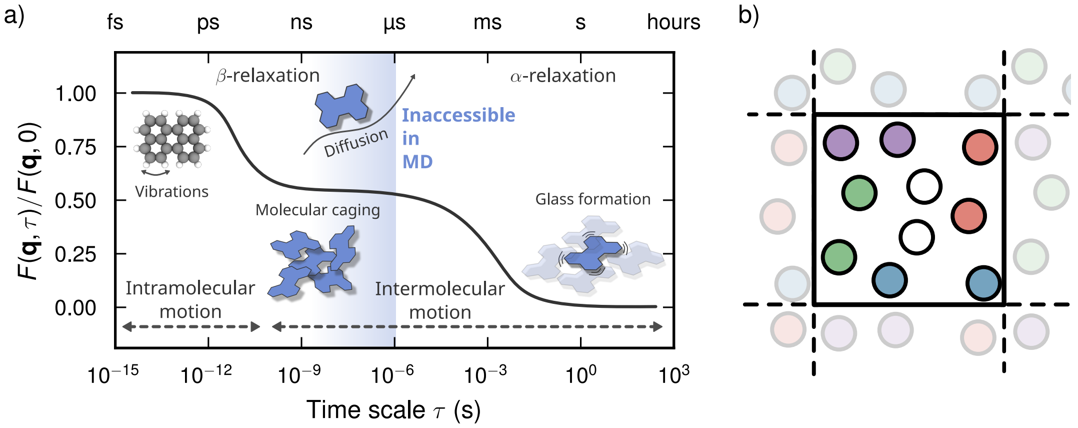

Figure 3.1: (a) Schematic of the different molecular interactions and relaxation processes associated with glass formation, shown relative to timescales accessible in MD. $F(\boldsymbol{q},t)$ denotes the intermediate scattering function, a correlation function that captures the dynamics of glass formation. α marks the slowest collective relaxation processes that drives glass formation in the system and spans many orders of magnitude in time, while β denote localized and weakly temperature-dependent processes. (b) Illustration of periodic boundary conditions (PBC). Atoms may cross the simulation box boundary (black square) and interact with periodic images of the atoms on the other side of the cell. As long as the box is sufficiently large to avoid self-interaction, this scheme effectively represents an infinite system.

the positions and velocities are updated according to the forces acting on each atom. Typical simulations span 1 ps to a few nanoseconds, with a time step on the order of 1 fs. Thermodynamic quantities such as temperature, pressure, and density can then be computed from the atomic motions. Because each atom is treated explicitly, simulations are typically limited to millions of atoms [77–79], which in turn limits the accessible length scales. Periodic boundary conditions (PBC) are therefore often used in MD simulations, such that atoms crossing the boundary re-enter from the other side, approximating an infinite system when the box is sufficiently large (Fig. 3.1b).

For computational reasons, the short time step Δt limits the time scales accessible in the simulation (Fig. 3.1a), with the absolute longest simulations reaching roughly tens of microseconds [80–82]. This is many orders of magnitude smaller than the 100 s time scale of experimental α-relaxation. Luckily, since α-relaxation spans a broad range of time scales and temperatures, aspects of it can be accessed in MD in some cases.

To be more precise, in a classical MD simulation the nuclei are treated explicitly, while electronic structure is implicitly included through the potential energy surface (Born-Oppenheimer approximation). The movements of the atoms are governed by Newton's equations of motion [83, Chapter 3],

$$m_i\ddot{\boldsymbol{r}}_i = \boldsymbol{f}_i \quad \text{and} \quad \boldsymbol{f}_i = -\nabla_{\boldsymbol{r}_i}\phi(\boldsymbol{r}_1,\dots,\boldsymbol{r}_i,\dots,\boldsymbol{r}_N), \tag{3.1}$$

where m_i is the mass of atom i, $\phi(\ldots)$ is the potential which is a function of the positions of the N atoms in the system, and $\boldsymbol{f}_i$ is the total force acting on atom i.

Equation (3.1) is a second-order differential equation, the solutions $\boldsymbol{r}_i(t)$ of which yield the trajectory for atom i. This equation can be solved using an iterative scheme known as the Velocity-Verlet algorithm [84], which updates positions and velocities using the acceleration $\boldsymbol{a}\,(t)$ at the current and next time steps,

$$\begin{aligned}\boldsymbol{v}\left(t+\frac{1}{2}\Delta t\right) &= \boldsymbol{v}(t)+\frac{\Delta t}{2}\boldsymbol{a}(t)\\ \boldsymbol{r}\,(t+\Delta t) &= \boldsymbol{r}(t)+\Delta t\boldsymbol{v}(t)+\frac{\Delta t^2}{2}\boldsymbol{a}(t)\\ \boldsymbol{v}\,(t+\Delta t) &= \boldsymbol{v}(t)+\Delta t\frac{\boldsymbol{a}(t)+\boldsymbol{a}(t+\Delta t)}{2}.\end{aligned} \tag{3.2}$$

The computationally expensive step in this scheme is the evaluation of forces, and in Chapter 4 we discuss how these are obtained from interatomic potentials.

3.2 Extracting information from molecular dynamics

We now turn to how we extract estimates of quantities of interest from the trajectory of positions $\boldsymbol{r}_i(t)$ and velocities $\boldsymbol{v}_i(t)$ that we obtain from the Velocity-Verlet scheme in Eq. (3.2). For a system of N atoms, the instantaneous state can be described by $3N$ position and $3N$ velocity components, which can be collected in a $6N$-dimensional vector $\boldsymbol{\Gamma}$. $\boldsymbol{\Gamma}$ is a point in the phase space of the system, with each point in phase space representing a different configuration of positions and velocities of the atoms. Any observable $A(t)$ of the system is a function of $\boldsymbol{\Gamma}$, and its equilibrium value is given by an ensemble average taken over the entire phase space of the system,

$$\langle A\rangle_{\text{ensemble}} = \int_{\boldsymbol{\Gamma}} A(\boldsymbol{\Gamma})\rho_{\text{ensemble}}(\boldsymbol{\Gamma})d\boldsymbol{\Gamma}. \tag{3.3}$$

The ensemble average is a weighted average over all configurations in phase space by the probability of observing that configuration, $\rho_{\text{ensemble}}(\boldsymbol{\Gamma})d\boldsymbol{\Gamma}$. Computing Eq. (3.3) explicitly for a realistic system is computationally difficult. Luckily, if the system is ergodic, meaning that an infinite simulation will visit every possible configuration $\boldsymbol{\Gamma}$ in phase space, we can replace the ensemble average with a time average over N_t time steps of a sufficiently long trajectory,

$$\begin{aligned}\langle A\rangle_{\text{ensemble}} &\overset{\text{ergodic}}{\Longleftrightarrow} \langle A\left(\boldsymbol{\Gamma}(t)\right)\rangle_t \Rightarrow\\ A_{\text{estimate}} &\approx \langle A\left(\boldsymbol{\Gamma}(t)\right)\rangle_t = \frac{1}{N_t}\sum_{i=1}^{N_t} A\left(\boldsymbol{\Gamma}(t_i)\right).\end{aligned} \tag{3.4}$$

The probability of observing a specific configuration over the trajectory is $\rho_{\text{ensemble}}(\mathbf{\Gamma})d\mathbf{\Gamma}$. $\rho_{\text{ensemble}}(\mathbf{\Gamma})$ is determined by which thermodynamic ensemble the simulation is conducted in, which in turn is governed by which macroscopic variables are kept constant during the simulation. So far, we have kept the number of atoms N and the volume V of the simulation box constant, and Newton's equations of motion fulfill conservation of energy E. This is known as the micro-canonical or NVE ensemble [83, Chapter 2], and the accessible states in phase space $(\mathbf{\Gamma})$ lie on an isosurface where N, V and E are all constant. One central assumption in statistical mechanics is that all accessible states are equally probable [85]. The probability distribution $\rho_{NVE}(\mathbf{\Gamma}) = 1/Z_{NVE}$ is thus a uniform distribution, where Z_{NVE} is called the partition function and is in this case equal to the number of states in $\mathbf{\Gamma}$ that keep N and V constant, and have energy E.

If one allows the total energy of the system to change by allowing the system to exchange energy with a heat bath at a constant temperature T, we arrive at the canonical or NVT ensemble. The accessible states $(\mathbf{\Gamma})$ may now have different energy $E_{\mathbf{\Gamma}}$. The probability distribution in this case, $\rho_{NVT}(\mathbf{\Gamma})$, is known as the Boltzmann distribution [85]. Similarly, by additionally allowing the simulation box volume V to change we obtain the isothermal-isobaric or NPT ensemble, where the pressure P is kept constant [83],

$$\rho_{NVT}(\mathbf{\Gamma}) = \frac{e^{-E_{\mathbf{\Gamma}}/k_BT}}{\sum_{\mathbf{\Gamma'}} e^{-E_{\mathbf{\Gamma'}}/k_BT}} \quad \text{and} \quad \rho_{NPT}(\mathbf{\Gamma}) = \frac{e^{-(E_{\mathbf{\Gamma}}+PV)/k_BT}}{\sum_{\mathbf{\Gamma'}} e^{-\left(E_{\mathbf{\Gamma'}}+PV\right)/k_BT}}. \tag{3.5}$$

The NVE, NVT and NPT ensembles are the most common in the context of MD simulations, and indeed in this thesis. The NVE ensemble is sampled by Newton's equations of motions, but we need to use thermostats and barostats to control the temperature and pressure in order to sample the NVT and NPT ensembles. While the NPT ensemble most closely mimics experimental conditions, these thermostats and barostats alter the dynamics which might introduce artifacts. Dynamic quantities such as correlation functions are therefore often computed in the NVE ensemble.

3.3 Correlation functions connect to experiments

In the previous section we concerned ourselves with the average of an observable A over the trajectory $\mathbf{\Gamma}(t)$. The instantaneous snapshots $A(t)$, however, fluctuate as the simulation samples different regions of phase space. These fluctuations are not noise; they are informative, and can be studied using correlation functions. Correlation functions describe how systems "wiggle" and evolve in space and time [86, Chapter 10], and how it relaxes from external stimuli. They thus provide insight into both the structure and dynamics of a system. In particular, the relaxation dynamics that drive glass formation can be studied using correlation functions, as we shall see in Sect. 3.4. Beyond theory, correlation functions are also central to experiments. Many experimental techniques directly probe system response, and hence correlation functions link simulations to experiments [87, Chapter 1]. In Chapter 5 and Chapter 6 we will encounter two examples of this connection, where MD simulations are connected to infrared (IR) spectroscopy and neutron scattering experiments via correlation functions of the dipole moment $\boldsymbol{\mu}(t)$ and the time-dependent particle density $\rho(\boldsymbol{r}, t)$, respectively.

A general correlation function of an observable $A(\boldsymbol{x}, t)$ takes the form

$$C(\boldsymbol{r}, \tau) = \langle A(\boldsymbol{x} + \boldsymbol{r}, t + \tau) A(\boldsymbol{x}, t) \rangle_{\boldsymbol{x},t}\,, \tag{3.6}$$

where the ensemble average is taken over $\boldsymbol{x}$ and t. Note that correlation functions are stationary, depending only on the time lag τ and not on the start time t. Two limits in space $\boldsymbol{r}$ and time τ are of special interest. The first is $C(\boldsymbol{r}, 0)$, the static correlation function, which describes how the observable varies spatially on average. The density-density correlation function $C(\boldsymbol{r}) = \langle \rho(\boldsymbol{x} + \boldsymbol{r}) \rho(\boldsymbol{x}) \rangle$ where $\rho(\boldsymbol{r}) = \sum_i \delta(r - \boldsymbol{r}_i)$ is a static correlation function that leads to the radial distribution function, which encodes the probability of finding another particle at a distance r from a particle,

$$g(r) = \frac{V}{N^2} \left\langle \sum_{i=1}^{N} \sum_{i \neq j}^{N} \delta(r - r_{ij}) \right\rangle \tag{3.7}$$

where N is the number of atoms, V is the system volume and r_{ij} is the pairwise distance between particles i and j. The radial distribution function will exhibit peaks where the probability of finding particles is high, such as for nearest-neighbors, and for a fluid in the long distance limit $g(r \to \infty)$ it goes to 1.

The second important limit is $C(0, \tau)$, the time autocorrelation function (ACF), which describes how an observable correlates with itself in time. An example of an ACF is the velocity autocorrelation function, defined as [87, Chapter 3]

$$C_v(\tau) = \frac{\langle \boldsymbol{v}(t + \tau) \cdot \boldsymbol{v}(t) \rangle}{\langle \boldsymbol{v}(t) \cdot \boldsymbol{v}(t) \rangle} \tag{3.8}$$

The denominator normalizes $C_v(\tau)$ such that $C_v(\tau) = 1$ as is customary for an ACF, and equals $\langle \boldsymbol{v}(t) \cdot \boldsymbol{v}(t) \rangle = 3k_BT/m$ for identical particles with mass m at temperature T. The vibrational density of states can be obtained by taking the Fourier transform of $C_v(\tau)$, and the diffusivity D by integrating $C_v(\tau)$,

$$D = \frac{k_BT}{m} \int_0^\infty C_v(\tau) d\tau. \tag{3.9}$$

The diffusivity is an important quantity when studying glass forming systems, describing how fast molecules move and reorganize, and can be measured experimentally, for

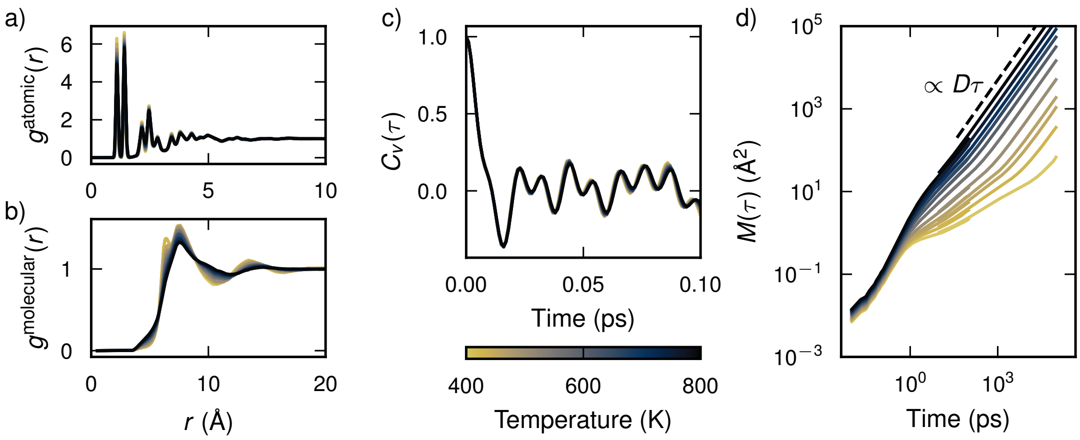

Figure 3.2: Examples of correlation functions probing different time and length scales in liquid perylene over a broad range of temperatures. (a) and (b) demonstrate the difference between the radial distribution functions for the atoms, $g^{\text{atomic}}(r)$, and molecules, $g^{\text{molecular}}(r)$, in the system. In this case, the velocity autocorrelation function in panel (c) probes short sub-picosecond timescales, while the mean squared displacement (MSD) for the molecules' center of mass in panel (d) captures the diffusive behavior of the molecules up to 100 ns. Note that the molecular MSD curves are made up of results from two individual simulations for each temperature, hence the slight mismatch around 100 ps.

example via quasi-elastic neutron scattering (QENS). The diffusivity can also be obtained as the asymptotic slope of the MSD $M(\tau)$ in the long-time limit [88, Chapter 13],

$$M(\tau) = \left\langle \|\boldsymbol{r}_i(t+\tau) - \boldsymbol{r}_i(t)\|_2^2 \right\rangle_{it} \qquad \lim_{\tau\to\infty} M(\tau) = 6D\tau, \tag{3.10}$$

with the ensemble average being taken over all starting times t and all particles i with position $\boldsymbol{r}_i$. $M(\tau)$ can be interpreted as the average distance squared traveled by particles after time τ. Note that the MSD is not formally a correlation function, but it can be recast in terms of an ACF [89]. The benefit of recasting the MSD is that although correlation functions are computationally challenging to compute in real space, they can be efficiently computed in Fourier space. The Wiener-Khinchin theorem links the Fourier power spectrum of a signal to its ACF, allowing for efficient evaluation [90, Chapter 15],

$$C(\tau) = \lim_{T\to\infty} \frac{1}{T}\frac{1}{2\pi} \int_{-T}^{T} |A(\omega)|^2 e^{i\omega\tau} d\omega. \tag{3.11}$$

The interpretation of Eq. (3.11) is that we can obtain the ACF of $A(t)$ by measuring it over a long time period T, computing its power spectrum $|A(\omega)|^2$, and then inverse Fourier transforming to obtain $C(\tau)$. Fourier transforms using fast Fourier transform (FFT) algorithms are computationally efficient, making this approach highly practical when analyzing MD simulations.

3.4 Accessing glass dynamics with MD

Equipped with the principles of MD, we now turn to how we can probe the glass-forming behavior in perylenes. One of the main challenges when studying glass formation with MD is that the dramatic changes in timescales of the dynamics close to the glass transition lead to a situation where the system appears frozen in the simulation, i.e., on timescales in the range of 1 ns to 100 ns, while remaining liquid on experimental timescales 1 s to 100 s. Trying to directly estimate key properties such as the glass transition temperature T_g from a MD simulation is thus challenging. For example, an approach often used in the literature for estimating T_g is to mimic experiments through simulated annealing, where the temperature is slowly decreased until the system enters a glassy state. The glass transition temperature is taken as the point at which the density changes slope as a function of temperature. Estimates of T_g from such annealing simulations agree with experiments in metal-organic frameworks [91], phosphates [92], disordered molecular solids [93], and organic solar cell (OSC) materials [94, 95], but overestimate T_g in perylene derivatives (Fig. 3.3b). The underlying assumption for the simulated annealing method is that the α-relaxation process is accessible in the MD simulation, such that the trend in density with simulation temperature can be extrapolated to infinitely slow cooling rates from the extremely high cooling rates in MD simulations; even a slow rate by MD standards of $1\,\mathrm{K\,ns^{-1}}$ corresponds to $10^9\,\mathrm{K\,s^{-1}}$. This assumption is violated in perylene derivatives: although molecular diffusion appears to cease at 400 K in the simulation, leading to a change of slope in the density, the experimental $T_g^{\exp}$ is approximately 250 K. The discrepancy arises because individual perylene molecules continue to reorient and exhibit librational motion far below 400 K, however, indicating a slower α-relaxation process that appears frozen only due to the short timescales accessible in the MD simulation.

In paper I, we present an approach for extrapolating the changes in dynamics as a function of temperature far above T_g, circumventing the problems with directly simulating the glass transition (Fig. 3.3c). Using a combination of correlation functions and Bayesian regression we are able to obtain estimates of T_g and the kinetic fragility m that are significantly closer to the experimental values than what can be obtained from simulated annealing (Fig. 3.3e and f). We compute the diffusivity D and the NACF $J(\tau)$ as functions of temperature for disordered liquid systems of five different perylene derivatives. The diffusivity is obtained from the long-time limit of the MSD (Eq. (3.10)), and is then fitted to the Vogel–Fulcher–Tammann (VFT) equation (Eq. (2.3)),

$$D(T) = D_0 e^{-B/(T-T_{\mathrm{VF}})}. \tag{3.12}$$

Machine learning aficionados might identify Eq. (3.12) as a *sloppy model* [96–98], where the model is highly sensitive to changes in B and T_{VF} due to the exponential, but less so to D_0 (Fig. 3.3d). To handle the sloppy model, we turn to Bayesian regression. Bayes'

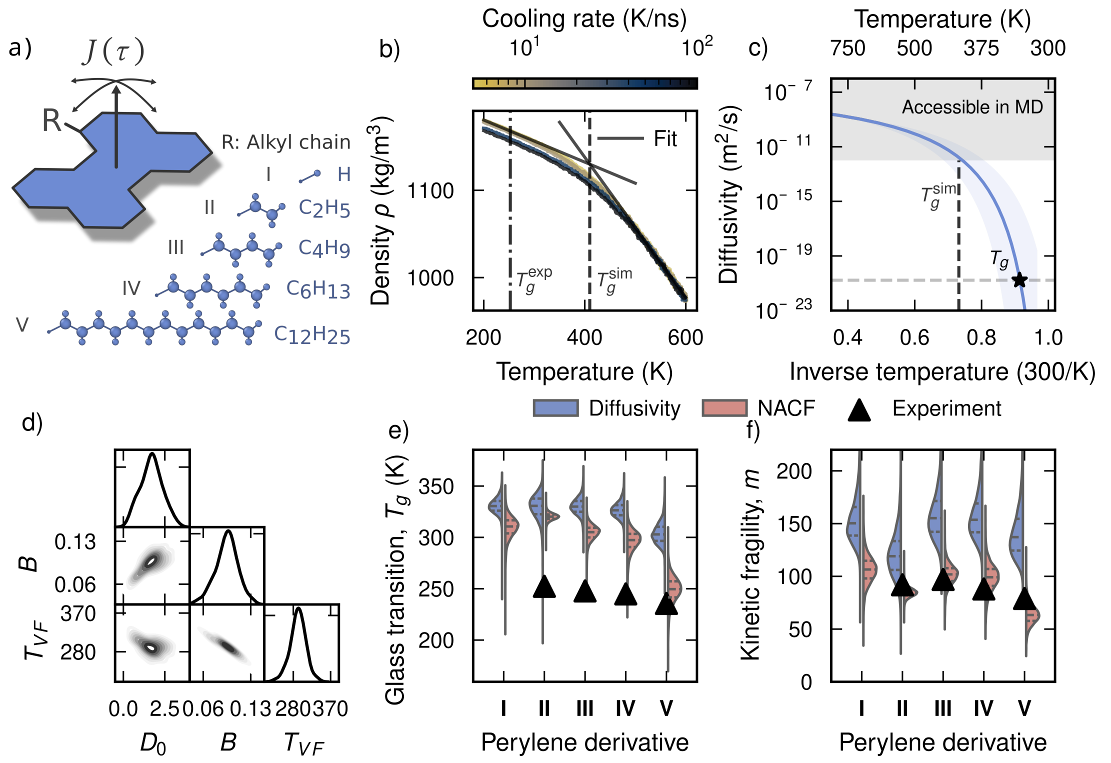

Figure 3.3: (a) The five different perylene derivatives in Paper 1, with the normal-vector autocorrelation function (NACF) $J(\tau)$ schematically drawn from the perylene core. (b) Glass transition temperature from simulated annealing for derivative II. T_g^{sim} overestimates the experimental T_g^{exp} by over 150 K. (c) Bayesian extrapolation of mean diffusivity $D(t)$ with one standard deviation posterior error band, extrapolated to T_g. T_g^{sim} from simulated annealing shown for reference. (d) Corner plot of the posterior $p(\boldsymbol{\theta}|\mathscr{D}, I)$. The banana-shaped marginal $p(B, T_{VF}|\mathscr{D}, I)$ reflects strong B–T_{VF} correlation, indicative of a sloppy model. (e) Glass transition T_g from Bayesian extrapolation of diffusivity and normal-vector autocorrelation function (NACF) via Bayesian regression. (f) Kinetic fragility m obtained analogously to panel (e). The predictions agree semi-quantitatively with experiments, where T_g is overestimated by about 50 K. Note that neat perylene (I) crystallizes under experimental conditions, and hence no experimental data for T_g or m are available.

theorem relates the posterior distribution of solutions $p(\boldsymbol{\theta}|\mathscr{D}, I)$ for the parameters $\boldsymbol{\theta} = \{D_0, B, T_{VF}\}$ to data and prior information,

$$p(\boldsymbol{\theta}|\mathscr{D}, I) = \frac{p(\mathscr{D}|\boldsymbol{\theta}, I)\,p(\boldsymbol{\theta}|I)}{p(\mathscr{D}|I)}. \tag{3.13}$$

where $\mathscr{D}$ is a set of data, $\mathscr{D} = \{(\boldsymbol{x_i}, y_i)\}_{i=1}^{N_d}$, and I denotes information we might have about the system *a priori*. The prior distribution $p(\boldsymbol{\theta}|I)$ is key in ensuring the stability of fitting Eq. (3.12), alleviating the problems with sloppy models by constraining the parameters $\boldsymbol{\theta} = \{D_0, B, T_{VF}\}$ to physical values. The denominator $p(\mathscr{D}|I)$ is called the evidence and acts as a normalization constant. The likelihood $p(\mathscr{D}|\boldsymbol{\theta}, I)$ can be interpreted as the probability of observing the data $\mathscr{D}$ given the parameters $\boldsymbol{\theta}$.

The resulting fit of the diffusivity as a function of temperature $D(T)$ can then be readily extrapolated until the diffusivity is so low that the molecules are not moving on experimental timescales, marking the glass transition temperature T_g (Fig. 3.3c). By sampling solutions from the posterior $p(\boldsymbol{\theta}|\mathscr{D}, I)$ we obtain a full distribution for T_g, allowing us to quantify the uncertainty in T_g stemming from the fit to the VFT equation.

We also predict T_g from the normal-vector autocorrelation function (NACF) in order to investigate if the diffusivity is actually capturing the α-relaxation process, the slowest relaxation process in the system. The NACF $J(\tau)$ is a home-grown correlation function based on the normal-vector $\hat{n}_i(t)$ defined against the plane of the perylene core of molecule i (Fig. 3.3a),

$$J(\tau) = \langle \hat{\boldsymbol{n}}_i(t) \cdot \hat{\boldsymbol{n}}_i(t+\tau) \rangle_{it}\,, \tag{3.14}$$

where the ensemble average is taken over all molecules i and all times t. $J(\tau)$ can efficiently be evaluated using the Wiener-Khinchin theorem, and is well described by a triple exponential. The fitted exponentials span timescales from femtosecond vibrational motion to nanosecond molecular reorientation. The NACF represents how long a molecule remains oriented in the same direction, and directly allows probing the cooperative reorganization of molecules associated with relaxation in a glass former. The diffusivity and the NACF thus probe two related but distinct relaxation processes, that freeze in at slightly different temperatures. Both molecular diffusion and reorientation slow down close to the glass transition, but the molecules will continue to reorient even when the molecules have stopped diffusing, as evident by the lower T_g estimate obtained from the NACF (Fig. 3.3e). This indicates that the processes captured by the NACF might be more closely related to α-relaxation in these perylene derivatives, and demonstrates the importance of choosing a suitable correlation function.

Using Bayesian regression is one approach that in some cases might help to circumvent the limited timescales accessible in MD. However, Bayesian regression is not magic; the uncertainty does not include systematic errors, including issues with the phenomenological VFT equation. Furthermore, the approach relies on the MD simulation sampling the actual α-relaxation process in order to extrapolate the temperature dependence. Some systems such as mixtures of chromophores can, for example, undergo a liquid-liquid phase transition before the glass transition, leading to a change in dynamics not easily detectable on the order of 100s of nanoseconds [76]. This raises the question: can signatures of such slow processes be identified already on MD timescales? Further research is needed to answer this question, but Bayesian regression remains a promising tool for when α-relaxation is not directly accessible in MD simulations.

3.5 Extending the accessible timescales in MD

The fundamental limitation in MD simulations is the short time step, which makes accessing the long timescales required for studying glass formation computationally demanding, or impossible. Several techniques exist to address this, though not all tackle the timescale problem directly. Adaptive biasing methods such as metadynamics [99], parallel simulations [100], and replica exchange [101, 102] broaden the configurational space explored [103, 104], but do not extend the timescales accessible to a single simulation. Coarse-graining takes a more direct approach, reducing the number of degrees of freedom by grouping atoms into effective particles, which increases computational efficiency and extends the accessible timescales [105–107]. However, coarse grained potentials are often limited to the thermodynamic state point at which they were designed, although transferable potentials are being developed [108, 109].

Beyond MD, lattice glass models use Monte Carlo methods to rearrange particles on a grid, though long relaxation times in the supercooled regime remains challenging [29, 110, 111]. Mode coupling theory takes a different approach, predicting relaxation dynamics from static observables of the structure, such as the structure factor $S(k)$ (revisited in Chapter 6), the Fourier transform of the radial distribution function [75]. The predicted relaxation dynamics are qualitatively in agreement with simulations [112, 113] and experiments [114], but tend to overestimate the glass transition temperature [75].

Finally, a particularly interesting development in recent years are machine learning (ML) methods based on deep generative models [111], such as Boltzmann generators for direct sampling of structural conformations from the desired thermodynamic ensemble [115]. Models can also learn the dynamics in the system to take longer steps than the MD time step, with techniques such as score dynamics [116] or the implicit-transfer operator approach [117]. However, these methods can be difficult to train, are trained for a specific system and temperature, and are often limited to comparatively small systems [118], although transferability is an active area of development [119].

Compared with these techniques, atomistic MD remains a transferable and versatile simulation method, and has thus been the method of choice throughout this thesis. Its reliability, however, depends on the availability of accurate interatomic potentials, as force calculations are both the most demanding and arguably most important step in a MD simulation. Machine-learned interatomic potentials (MLIPs) offer a pathway to more accurate and efficient force evaluation, and will be the focus of the next chapter.

4

Molecular dynamics with machine-learned potentials

> For potentials I'm now convinced that equivariance is useless.
>
> *Michele Ceriotti,*
> *CECAM workshop, 2024*

In the previous chapters, we have seen how the intermolecular interactions in chromophores leads to aggregation and glass formation, its effect on device performance, and how molecular dynamics (MD) can be used to simulate these processes. These large-scale processes are ultimately rooted in quantum mechanical interactions between atoms and molecules. Capturing both aspects simultaneously is the central computational challenge: electronic structure methods such as density functional theory (DFT) provide the required quantum mechanical accuracy, but are limited to systems of hundreds of atoms, far short of the millions needed to model chromophore assemblies. Classical force fields (FFs), on the other hand, scale readily to large systems but sacrifice accuracy, possibly not fully describing the subtle intermolecular interactions in chromophores.

Machine-learned interatomic potential (MLIP) have in recent years emerged as a way to resolve this tradeoff, offering accuracy approaching DFT at a computational cost closer to classical FFs. In this chapter, we introduce the key ideas behind MLIPs, before focusing on the neuroevolution potential (NEP) framework as a particularly efficient implementation that is well-suited to large-scale chromophore simulations. We also discuss how to estimate when a MLIP is likely to be wrong, and how recently developed foundational models can be adapted to new systems through fine-tuning.

4.1 Electronic-structure and classical force fields

The goal of a FF is to provide a set of forces for a given configuration of atoms. Atomic systems behave according to the laws of quantum mechanics. The nuclei and electrons in the system obey the many-body Schrödinger equation [120, Chapter 3],

$$\begin{aligned}\hat{H}\Psi &= E\Psi \\ \hat{H} &= -\frac{1}{2}\sum_i \nabla_i^2 - \sum_I \frac{1}{2M_I}\nabla_I^2 - \sum_{i,I}\frac{Z_I}{|\boldsymbol{r}_i - \boldsymbol{R}_I|} \\ &+ \frac{1}{2}\sum_{i\neq j}\frac{1}{|\boldsymbol{r}_i - \boldsymbol{r}_j|} + \frac{1}{2}\sum_{I\neq J}\frac{Z_I Z_J}{|\boldsymbol{R}_I - \boldsymbol{R}_J|},\end{aligned} \tag{4.1}$$

where we have used atomic units ($m_e = e = 4\pi\epsilon_0 = \hbar = 1$) for brevity. Lower-case indices run over electrons and uppercase indices run over nuclei, which have mass M_I and atomic number Z_I. $\boldsymbol{r}_i$ denotes the position of electron i and $\boldsymbol{R}_J$ is the position of nucleus J. Solving Eq. (4.1) would give a complete description of the system at the current moment in time, including all the forces between the atoms, which we need in order to propagate a MD simulation. However, the many-body Schrödinger equation is only analytically solvable in the simplest of cases involving a single electron, and numerically only for moderately sized or high-symmetry systems [121]. One major hurdle in solving Eq. (4.1) are the terms involving pairs of particles, denoted by the sums over $i \neq j$, i, I, and $I \neq J$, which scale exponentially in complexity as the number of particles in the system increases. A first step we can take to simplify this problem is to assume the slow-moving nuclei to be stationary relative to the electrons, which is known as the Born-Oppenheimer approximation [122]. DFT then proposes to reformulate the Schrödinger equation in terms of the electron density through the Hohenberg-Kohn theorem, replacing the fully-interacting electrons with N fictitious non-interacting electrons that still have the same ground-state density $n(\boldsymbol{r})$ through the Kohn-Sham equations [121, 123],

$$\begin{aligned}\left(-\frac{\nabla^2}{2} + v_s[n](\boldsymbol{r})\right)\varphi_j(\boldsymbol{r}) &= \epsilon_j\varphi_j(\boldsymbol{r}) \\ v_s[n](\boldsymbol{r}) &= v(\boldsymbol{r}) + \int d^3r' \frac{n(\boldsymbol{r'})}{|\boldsymbol{r} - \boldsymbol{r'}|} + v_{xc}[n](\boldsymbol{r}) \\ v_{xc}[n](\boldsymbol{r}) &= \frac{\delta E_{xc}[n]}{\delta n(\boldsymbol{r})} \\ n(\boldsymbol{r}) &= \sum_{j=1}^{N}|\varphi_j(\boldsymbol{r})|^2.\end{aligned} \tag{4.2}$$

$\varphi_j(\boldsymbol{r})$ is the Kohn-Sham orbital for state j of the non-interacting system, with associated energy ϵ_j, and $[n]$ denotes that a quantity is a functional of the electron density $n(\boldsymbol{r})$. A feature of the Kohn-Sham equations is that the ground state density $n(\boldsymbol{r})$ is a function of $\{\varphi_j(\boldsymbol{r})\}_j$, which in turn also depends on the density. The equations thus have to be solved self-consistently, by iteratively solving for and updating the density $n(\boldsymbol{r})$ until it has converged. Physical properties such as the forces between the nuclei can then be computed from the ground state density $n(\boldsymbol{r})$. Note that the external potential $v(\boldsymbol{r})$ includes the Coulomb interaction from the nuclei acting on the electrons.

The Kohn-Sham equations are formally exact, and could be solved if the exchange–correlation (XC) energy functional $E_{xc}[n]$ and its corresponding potential $v_{xc}[n](\boldsymbol{r})$ was known. Unfortunately, $E_{xc}[n]$ is in general unknown, leaving users of DFT with no other choice but to rely on approximations for $E_{xc}[n]$. There exists a plethora of DFT functionals, and depending on the choice of functional a DFT calculation can yield different results. As such, DFT calculations are dependent on the choice of functional.

The Kohn-Sham equations in Eq. (4.2) are significantly easier to solve computationally than the many-body Schrödinger equation in Eq. (4.1), but they still scale poorly with the number of electrons in the system. In practice, DFT is often not used for systems with more than hundreds or possibly thousands of atoms, which is short of the hundreds of thousands or millions in MD simulations. Classical FFs take a more computationally efficient but potentially less accurate approach by foregoing the Schrödinger equation entirely. One approach is to expand the system energy in terms depending on the positions $\boldsymbol{r}_i^N$ of the atoms in the system [83],

$$V(\boldsymbol{r}_i) = \sum_i V_1(\boldsymbol{r}_i) + \sum_i \sum_{j>i} V_2(\boldsymbol{r}_i, \boldsymbol{r}_j) + \sum_i \sum_{j>i} \sum_{k>j>i} V_2(\boldsymbol{r}_i, \boldsymbol{r}_j, \boldsymbol{r}_k) + \dots \tag{4.3}$$

where the first term, V_1, corresponds to interactions of all N atoms with an external potential, the second term with V_2 corresponds to interactions between pairs of particles, the third sum to interactions between triplets, and so forth. Note that this is one possible expansion of the potential. Including more terms in Eq. (4.3) gives in general a more accurate but less computationally efficient model. Truncating the expansion at the V_2 term gives a class of potentials known as pair potentials, the most well-known of which is the Lennard-Jones potential [83],

$$V^{LJ}(\boldsymbol{r}_i) = \sum_i \sum_{j>i} 4\epsilon \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \left(\frac{\sigma}{r_{ij}} \right)^{6} \right]. \tag{4.4}$$

The Lennard-Jones potential involves two parameters, ϵ and σ, controlling the energy and length scale of the potential, which are fitted to reproduce experimentally observed properties, such as the density, in MD simulations. This is a common strategy for classical FFs, where different potentials mainly differ in the interactions considered, the number of free parameters, and the materials they have been optimized for. Once optimized, classical FFs can be readily evaluated for a large number of atoms, but at the cost of not being as accurate as DFT calculations.

4.2 Machine-learned interatomic potentials

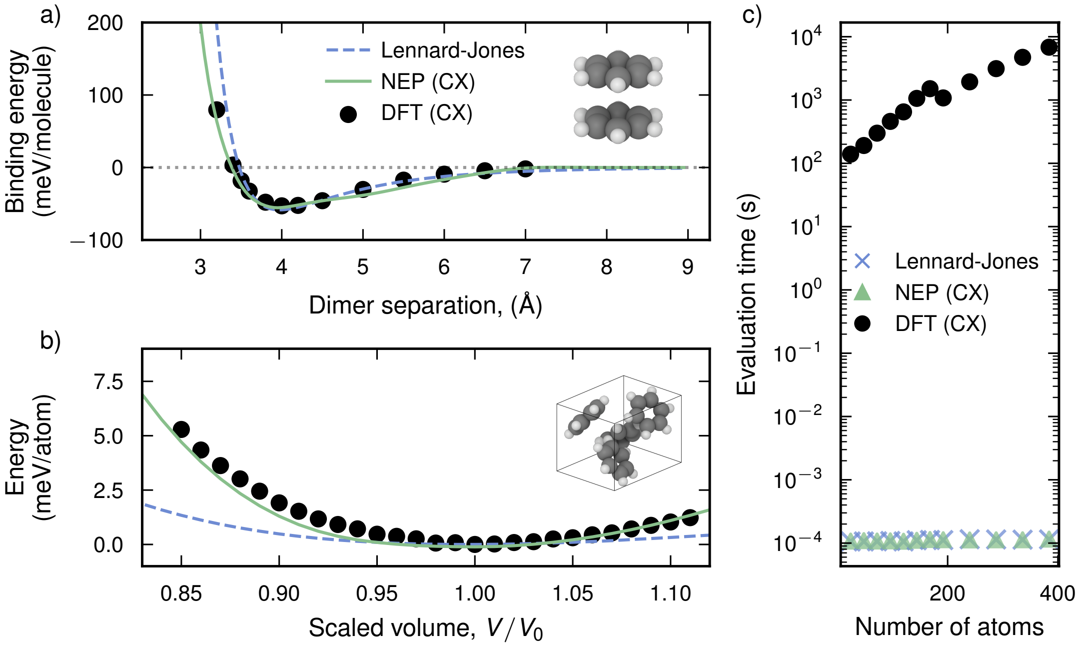

Figure 4.1: (a) Binding energy between two benzene molecules as a function of distance, and (b) the change in potential energy as the unit cell of crystalline benzene is scaled, using three different methods; a Lennard-Jones potential, a MLIP based on the NEP formalism trained on vdW-DF-cx DFT data [124], and DFT using the same functional. (c) Comparison of the computational cost of evaluating the energies with each method. NEP accurately captures the target curves from DFT, while still having a computational cost comparable to the Lennard-Jones potential.

We can demonstrate the tradeoff in accuracy and performance between DFT and FFs for benzene, as the simplest possible representative of the chromophores studied in this thesis. A Lennard-Jones potential is able to reproduce the binding energy between two benzene molecules as a function of separation (Fig. 4.1a) compared to DFT, but fails for the more complex case of the change in potential energy as the axes of the unit cell of crystalline benzene are scaled (Fig. 4.1b). DFT, on the other hand, scales prohibitively compared to the Lennard-Jones potential as a function of the number of atoms in the system, taking over 32 central processing unit (CPU) core hours to evaluate the forces for a system of $\sim$300 atoms compared to less than a millisecond for the Lennard-Jones potential (Fig. 4.1c). Performing large-scale MD simulations for tens of thousands of atoms is thus computationally intractable for DFT.

Fortunately, MLIP bridge the gap between FF and DFT by combining the computational efficiency of FFs with the accuracy of DFT, achieved by training a flexible parametric model to reproduce DFT energies and forces. The promise of MLIPs is fulfilled for our benzene example; a MLIP based on the NEP framework accurately reproduces both the binding energy curve and energy-volume curve from DFT, while retaining the high computational performance of the Lennard-Jones potential. As for classical FFs, there exists a large number of different types of MLIPs based on different functional forms. We will highlight two of these here, namely kernel-based methods and neural-network based methods.

4.2.1 Kernel-based methods

Kernel-based methods typically employ a probabilistic approach, modeling the potential energy surface as a Gaussian process [125, Chapter 2],

$$\begin{aligned} E(\boldsymbol{q}) &\sim \mathcal{GP}\left(m(\boldsymbol{q}), k(\boldsymbol{q},\boldsymbol{q}')\right) \\ m(\boldsymbol{q}) &= \mathbb{E}\left[E(\boldsymbol{q})\right] \\ k(\boldsymbol{q},\boldsymbol{q}') &= \mathbb{E}\left[\left(E(\boldsymbol{q}) - m(\boldsymbol{q})\right)\left(E(\boldsymbol{q}') - m(\boldsymbol{q}')\right)\right]. \end{aligned} \tag{4.5}$$

$\boldsymbol{q} \in \mathbb{R}^{3N}$ represents a specific configuration of the N atoms in the system, and the covariance function $k(\boldsymbol{q},\boldsymbol{q}')$ encodes the similarity between two configurations $\boldsymbol{q}$ and $\boldsymbol{q}'$. By modeling the potential energy as a Gaussian process, the potential energy surface $E(\boldsymbol{q})$ is constrained to be smoothly varying as the positions of the atoms represented by the vector $\boldsymbol{q}$ are changed, with each configuration having an associated predicted mean energy $E(\boldsymbol{q})$ with standard deviation $\sigma_E(\boldsymbol{q})$. This leads us to the main benefit of Gaussian process-based methods: the uncertainty in the predictions can readily be extracted as the predicted variance $\sigma_E(\boldsymbol{q})$, which, for example, enables augmenting the training set through active learning (Sect. 4.3.3). The disadvantage of kernel-based methods is that they are computationally expensive to evaluate, with the cost of evaluating the model scaling as $\mathcal{O}(n^3)$, where n is the number of data points in the training dataset, though some strategies exist for partially mitigating this limitation [126].

Examples for Gaussian process-based MLIPs include the Gradient-Domain Machine-Learning model (GDML, [127]) and the Gaussian Approximation Potential (GAP) [128]. sGDML is an extension of GDML in which relevant symmetries are incorporated to improve the efficiency of the model [129].

4.2.2 Neural network-based methods

Popularized by Behler and Parrinello [130, 131], neural network (NN)-based potentials have dominated the MLIP field over the last decade. A NN potential denoted U with weights $\boldsymbol{w}$ predicts the energies for each atom i in a system of N atoms as follows,

$$\begin{aligned} E_i &= U(\boldsymbol{w}, \boldsymbol{q}(\{\boldsymbol{r}_{ij}\})) \\ \rightarrow \boldsymbol{F}_i &= -\nabla_{\boldsymbol{r}_i} U(\boldsymbol{w}, \boldsymbol{q}(\{\boldsymbol{r}_{ij}\})). \end{aligned} \tag{4.6}$$

E_i and $\boldsymbol{F}_i$ are the per-atom potential energy and the forces acting on atom i respectively. The per-atom potential energies are summed up to yield $E = \sum_i^N E_i$, as only the total potential energy for a structure is defined, and E can be seen as a model for the potential energy surface of the system. The forces $\boldsymbol{F}_i$ are obtained as the gradient of U with respect to the coordinates of atom i, $\boldsymbol{r}_i$.

Uncertainty estimates cannot be extracted as easily from NN-based MLIPs as from Gaussian process-based, but they are more computationally efficient. One way to estimate the uncertainty for NN MLIPs is to train an ensemble of models and compute the forces with each model, providing an uncertainty estimate at a higher cost [132].

In addition to the Behler-Parrinello network, examples of neural network-based methods include Deep Potential (DP) [133, 134], Embedded Atom Neural Network (EANN) [135], Recursively embedded atom neural network (REANN) [136, 137], ANI-1 [138], SchNet [139], Nequip [140], Allegro [141], MACE [142], and NEP which we will return to in Sect. 4.3. There also exists families of models that predict the energy of forces through regression, but forego NNs, in the form of linear basis function models. These include Moment Tensor Potentials (MTP) [143, 144], Spectral Neighbor Analysis Potential (SNAP) [145], and Atomic Cluster Expansions (ACE) [146–148].

4.2.3 Using descriptors to represent atomic structures

Both kernel-based and neural network-based MLIPs do not typically have the Cartesian coordinates of a set of atoms, $\{\boldsymbol{r}_i\}$, as input, but rather a so-called descriptor vector $\boldsymbol{q}(\{\boldsymbol{r}_{ij}\})$. This is a function of the relative positions of atom i and all neighboring atoms j, and can be thought of as a chemical fingerprint describing the environment around atom i. This seemingly cumbersome descriptor vector ensures that the model fulfills certain symmetries dictated by physics.

First of all, there are a set of invariances, transformations of the atomic system that should not change the predicted energies of forces. These include invariance under translations and rotations of the system, and permutations of atoms of the same element [149]. Recently, equivariances have additionally started to be incorporated in MLIPs [142, 150]. An equivariant transformation can, for instance, be a rotation; when the input structure is rotated, the output forces should be rotated accordingly. Specifically, both descriptors and the model output are designed to transform equivariantly under SO(3) rotations of the input. These are often implemented using equivariant graph neural networks [140, 142, 151]. However, as we shall see later in Sect. 4.3.4, equivariant transformations often come with a significant computational cost. This has lead to some research groups opting to train non-equivariant models, instead relying on extensive training sets to ensure that the model learns these equivariances, such as the PET-MAD family of models [152].

Understanding the descriptor vector is key to understanding the success of MLIPs in recent years [23]. A machine learning (ML) model in general, and a NN in particular, is simply a "dumb" mathematical function where the parameters are adjusted during training such that it best mimics the examples of inputs and corresponding target outputs in the training data set. Crucially, this means that a MLIP does not "know" anything about physics. Furthermore, as ML models typically involve many parameters interacting non-linearly, it is difficult to know beforehand what the output will be for a specific input, a problem that becomes exponentially more difficult as the size of the model increases. Taken together, this means that a ML model can produce unexpected outputs when presented with an input that is different from the examples in the training data set, a problem that is known as out-of-distribution prediction and is researched in the ML literature [153]. A MLIP can to some extent be protected from the issues of out-of-distribution prediction by incorporating a descriptor vector that ensures that the input to the MLIP is physically meaningful. However, even with a descriptor vector a MLIP still suffers from the issues of out-of-distribution predictions, for instance if structures encountered during simulations are vastly different from the ones in the training data set. Like with any other ML model, MLIPs can be hardened against this problem by ensuring that the training data set samples the chemical space of interest well, for example by using entropy-maximized datasets [154, 155], but in practice one can never guarantee the robustness of the predictions of the MLIP.

The points discussed in this section are the general considerations that go into crafting a physically accurate descriptor, the exact implementation of which can change for different MLIPs. Examples of often-used descriptors are atom-centered symmetry functions (ACSF) [130, 131], smooth overlap of atomic orbitals (SOAP) [156], spherical harmonics [157] and the many-body tensor representation [158]. In the next section, we will take a look at a specific implementation of an NN-based MLIP, and discuss its associated descriptor vector in detail.

4.3 Neuroevolution potentials

We will now turn to a specific type of neural network-based MLIP, namely the neuroevolution potential (NEP) developed and implemented in the GPUMD package to which I am a contributing developer, presented in papers II and III. In this section, we will describe NEP models in detail, starting with the formalism.

4.3.1 The NEP formalism

Similarly to the Behler-Parrinello NN-MLIPs described in Eq. (4.6), in NEP the energy of atom i, E_i, is predicted as a function of a descriptor vector with N_{des} components, denoted $\boldsymbol{q}_i$. The NN consists of a single fully connected hidden layer, and the predicted energy takes the form

$$E_i = U(\boldsymbol{w}, \boldsymbol{q}) = \sum_{\mu=1}^{N_\mathrm{neu}} w_\mu^{(1)} \tanh\left(\sum_{\nu=1}^{N_\mathrm{des}} w_{\mu\nu}^{(0)} q_\nu^i - b_\mu^{(0)}\right) - b^{(1)}. \tag{4.7}$$

$\boldsymbol{w}^{(0)}$ and $\boldsymbol{b}^{(0)}$ are the weight matrix from the input descriptor vector to the hidden layer, and $\boldsymbol{w}^{(1)}$ and $b^{(1)}$ are the weights and bias term from the hidden layer to the single output neuron. The activation function for the hidden layer is the hyperbolic tangent.

Note that the parameters for the model in Eq. (4.7), which for historical reasons is known as a NEP3 model, are shared between all atoms in the system. This has the crucial benefit that the model does not increase in size as the number of atoms in the simulated system increases, which could otherwise lead to a model that is computationally impossible to evaluate for the millions of atoms in a MD simulation. However, sharing weights between all atoms in a system can lead to an insufficiently flexible model, especially in systems with many different atomic species and thus a potentially large input space of possible descriptor vectors $\boldsymbol{q}$. NEP4 increases the flexibility of NEP models by having an individual network for each atomic species α in the system, $U(\boldsymbol{w}^\alpha, \boldsymbol{q})$ [159].

The descriptor vector $\boldsymbol{q}$ takes the same shape for both NEP3 and NEP4, and is comprised of a radial part and an angular part (Fig. 4.2a). The radial descriptors encode pairwise atomic interactions, and the angular part captures three-body, four-body, and five-body interactions. The radial part has $n_\mathrm{max}^\mathrm{R} + 1$ components and is defined as

$$q_n^i = \sum_{j \neq i} g_n(r_{ij}). \tag{4.8}$$

The summation runs over all neighboring atoms j to atom i, where r_{ij} is the distance between them. The contribution from each neighbor, $g_n(r_{ij})$, is in turn computed from $N_\mathrm{bas}^\mathrm{R} + 1$ basis functions,

$$\begin{aligned} g_n(r_{ij}) &= \sum_{k=0}^{N_\mathrm{bas}^\mathrm{R}} c_{nk}^{ij} f_k(r_{ij}) \\ f_k(r_{ij}) &= \frac{1}{2}\left[T_k\left(2\left(\frac{r_{ij}}{r_c^\mathrm{R} - 1}\right)^2 - 1\right) + 1\right] f_c(r_{ij}) \end{aligned} \tag{4.9}$$

where $T_k(\ldots)$ is the kth-order Chebyshev polynomial of the first kind. $f_c(r_{ij})$ is a cutoff function that ensures that the contribution $g_n(r_{ij})$ from atom j decreases smoothly to zero as the distance r_{ij} approaches the radial cutoff distance r_c^R, and is defined as

$$f_c(r_{ij}) = \begin{cases} \frac{1}{2}\left[1 + \cos\left(\pi \frac{r_{ij}}{r_c^\mathrm{R}}\right)\right], & r_{ij} \leq r_c^\mathrm{R} \\ 0, & r_{ij} > r_c^\mathrm{R}. \end{cases} \tag{4.10}$$

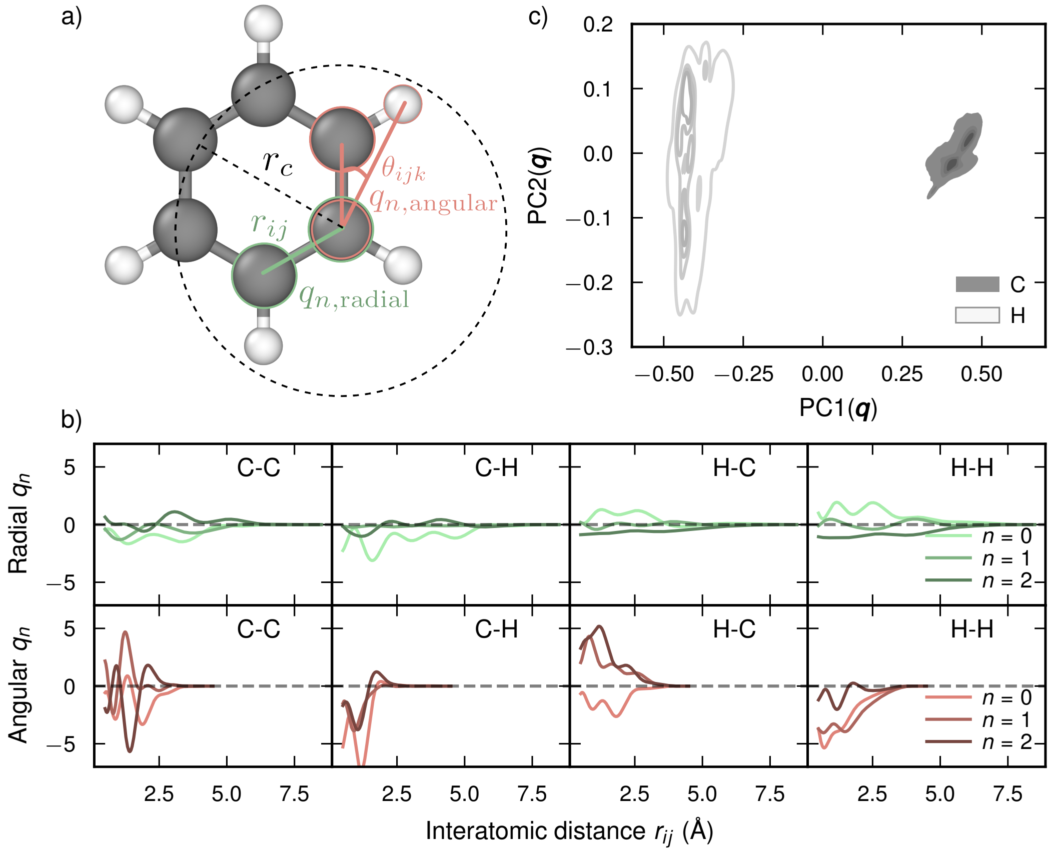

Figure 4.2: (a) Schematic of radial (green) and angular (red) descriptor components. (b) Visualization of three radial and angular descriptor components. Note that the peak position of descriptor component n varies. (c) Principal component analysis (PCA) of the two principal descriptor components for about 900 structures of crystalline benzene, colored by atomic species.

$n^{\mathrm{R}}_{\mathrm{max}}$, $N^{\mathrm{R}}_{\mathrm{bas}}$, and r^{R}_c are hyperparameters that are set before training. A key feature of the NEP formalism is that the coefficients c^{ij}_{nk} in the radial basis expansion are free parameters that are optimized jointly with the weights $\boldsymbol{w}$ and biases $\boldsymbol{b}$ of the NN. These coefficients depend on the species of atom i and j, which allows NEP to tailor the message $g_n(r_{ij})$ from each neighbor, increasing the flexibility of the model.

The angular descriptor vector are similarly defined as

$$q^i_{nl} = \frac{2l+1}{4\pi} \sum_{j\neq i} \sum_{k\neq i} g_n(r_{ij}) g_n(r_{ik}) P_l\left(\cos\theta_{ijk}\right) \tag{4.11}$$

with $0 < n < n^A_{\mathrm{max}}$ and $1 \leq l \leq l^{\mathrm{3b}}_{\mathrm{max}}$ as hyperparameters that control the size of the basis expansion. $P_l(...)$ is the Legendre polynomial of order l, and θ_{ijk} is the angle formed between the two pairs of atoms, ij and ik. This expression is a three-body descriptor as it involves three atoms; the central atom i and two neighboring atoms j and k. Higher order terms, such as four-body or five-body interactions, can be included in the NEP formalism, but we will not describe those in detail here.

The different radial and angular components as a function of interatomic distance r_{ij} exhibit peaks at different positions, which can be interpreted as the different components represents probing different regions of the chemical environment (Fig. 4.2b). Typically, a NEP model encounters a large number of chemical environments during inference, each with their own descriptor vectors. Using PCA to visualize the descriptors of about 900 structures of crystalline benzene shows that atoms of the same species cluster in similar regions, yet vary dramatically within them (Fig. 4.2c).

The descriptor vector fulfills the invariance requirements we discussed in the previous section. Invariance under translation and rotation of the system is respected as the descriptors only depend on the relative distance r_{ij} between pairs of atoms, as well as the angle θ_{ijk} between triplets of atoms. Furthermore, invariance under permutations of atoms is guaranteed by the summation over neighbors in Eq. (4.8) and Eq. (4.11).

The partial force acting on atom i can be computed using the chain rule,

$$\begin{aligned}\frac{\partial E_i}{\partial \boldsymbol{r}_{ij}} &= \sum_{n=0}^{n_{\max}^{\mathrm{R}}} \frac{\partial E_i}{\partial q_n^i}\frac{\partial q_n^i}{\partial \boldsymbol{r}_{ij}} + \sum_{n=0}^{n_{\max}^{\mathrm{A}}}\sum_{l=1}^{l_{\max}^{\mathrm{3b}}} \frac{\partial E_i}{\partial q_{nl}^i}\frac{\partial q_{nl}^i}{\partial \boldsymbol{r}_{ij}} \\ &+ \sum_{n=0}^{n_{\max}^{\mathrm{A}}}\sum_{l=1}^{l_{\max}^{\mathrm{4b}}} \frac{\partial E_i}{\partial q_{nlll}^i}\frac{\partial q_{nlll}^i}{\partial \boldsymbol{r}_{ij}} + \sum_{n=0}^{n_{\max}^{\mathrm{A}}}\sum_{l=1}^{l_{\max}^{\mathrm{5b}}} \frac{\partial E_i}{\partial q_{nllll}^i}\frac{\partial q_{nllll}^i}{\partial \boldsymbol{r}_{ij}},\end{aligned} \tag{4.12}$$

where we have additionally included the four- and five-body angular descriptors. Note that the derivative is taken with regard to the distance vector between atoms i and j, $\boldsymbol{r}_{ij}$. We can now construct the force acting on atom i due to atom j to respect Newton's third law, $\boldsymbol{F}_{ij} = -\boldsymbol{F}_{ji}$, as $\boldsymbol{F}_{ij} = \partial E_i/\partial \boldsymbol{r}_{ij} - \partial E_j/\partial \boldsymbol{r}_{ji}$. The total force acting on atom i from all neighboring atoms can be obtained by direct summation,

$$\boldsymbol{F}_i = \sum_{i\neq j} \boldsymbol{F}_{ij}. \tag{4.13}$$

The per-atom virial, from which properties such as stress and heat-current can be derived, can also be defined in terms of the partial force,

$$\boldsymbol{W}_i = \sum_{j\neq i} \boldsymbol{r}_{ij} \otimes \frac{\partial U_j}{\partial \boldsymbol{r}_{ji}}. \tag{4.14}$$

These analytical expressions for the energies, forces, and virials are computationally cheap to evaluate, and since Eq. (4.7), Eq. (4.13), and Eq. (4.14) can be evaluated for all atoms in the system in parallel the NEP formalism maps efficiently onto graphics processing units (GPUs).

4.3.2 Training NEP models

NEPs are trained by minimizing the following loss function, where the first three terms are the root mean squared error (RMSE) loss with regards to energy, forces, and virials, respectively,

$$\begin{aligned}
L(\boldsymbol{z}) = {} & \lambda_e \left(\frac{1}{N_{\text{str}}} \sum_{n=1}^{N_{\text{str}}} \left(E^{\text{NEP}}(n, \boldsymbol{z}) - E^{\text{tar}} \right)^2 \right)^{1/2} \\
& + \lambda_f \left(\frac{1}{3N} \sum_{i=1}^{N} \left(F_i^{\text{NEP}}(\boldsymbol{z}) - F_i^{\text{tar}} \right)^2 \right)^{1/2} \\
& + \lambda_v \left(\frac{1}{6N_{\text{str}}} \sum_{n=1}^{N_{\text{str}}} \sum_{\mu\nu} \left(W_{\mu\nu}^{\text{NEP}}(n, \boldsymbol{z}) - W_{\mu\nu}^{\text{tar}} \right)^2 \right)^{1/2} \\
& + \lambda_1 \frac{1}{N_{\text{par}}} \sum_{n=1}^{N_{\text{par}}} |z_n| + \lambda_2 \left(\frac{1}{N_{\text{par}}} \sum_{n=1}^{N_{\text{par}}} z_n^2 \right)^{1/2}
\end{aligned} \tag{4.15}$$

with $\boldsymbol{z}$ denoting the trainable parameters of the model, N_{str} the number of structures in the current batch, and N the total number of atoms in the current batch. Superscripts NEP and tar represent predicted and target values, respectively. The last two terms, weighted by factors λ_1 and λ_2, are $L1$ and $L2$-regularization terms, which makes Eq. (4.15) an elastic-net loss [160]. An elastic net combines the benefits of L1 (Lasso) and L2 (ridge) regression, and yields a sparse model.

The loss in Eq. (4.15) is minimized using a separable natural evolution strategy (SNES) [161], which is an evolutionary optimization algorithm. The general idea of SNES is to optimize a distribution for each parameter, instead of a single value as in most other optimization techniques. This scheme is implemented as follows [162]. Let the parameters $\boldsymbol{z}$ be distributed according to a joint N_{par}-dimensional Gaussian distribution, $\boldsymbol{z} \sim \mathcal{N}(\boldsymbol{m}, \boldsymbol{s})$, where $\boldsymbol{m}$ and $\boldsymbol{s}$ are the mean and standard deviation vectors respectively. This parameter distribution is iteratively updated according to the natural gradient of the fitness $J(\boldsymbol{z})$,

$$J(\boldsymbol{z}) = \mathbb{E}[-L(\boldsymbol{z})] = -\int L(\boldsymbol{z}) p(\boldsymbol{z}|\boldsymbol{m}, \boldsymbol{s}) d\boldsymbol{z}, \tag{4.16}$$

which is the expected value of the loss function under the search parameter distribution, $p(\boldsymbol{z}|\boldsymbol{m}, \boldsymbol{s})$. First, N_{pop} samples $\boldsymbol{z}_k$ are drawn from the distribution, $\boldsymbol{z}_k = \boldsymbol{m} + \boldsymbol{s} \odot \boldsymbol{r}_k$ where $\boldsymbol{r}_k \sim \mathcal{N}(0, 1)$, which each can be seen as an instance of the NEP model. The symbol $\odot$ denotes the Hadamard (element-wise) product. Second, $L(\boldsymbol{z}_k)$ is evaluated for each of the N_{pop} models in the current generation, the models are sorted in ascending order of

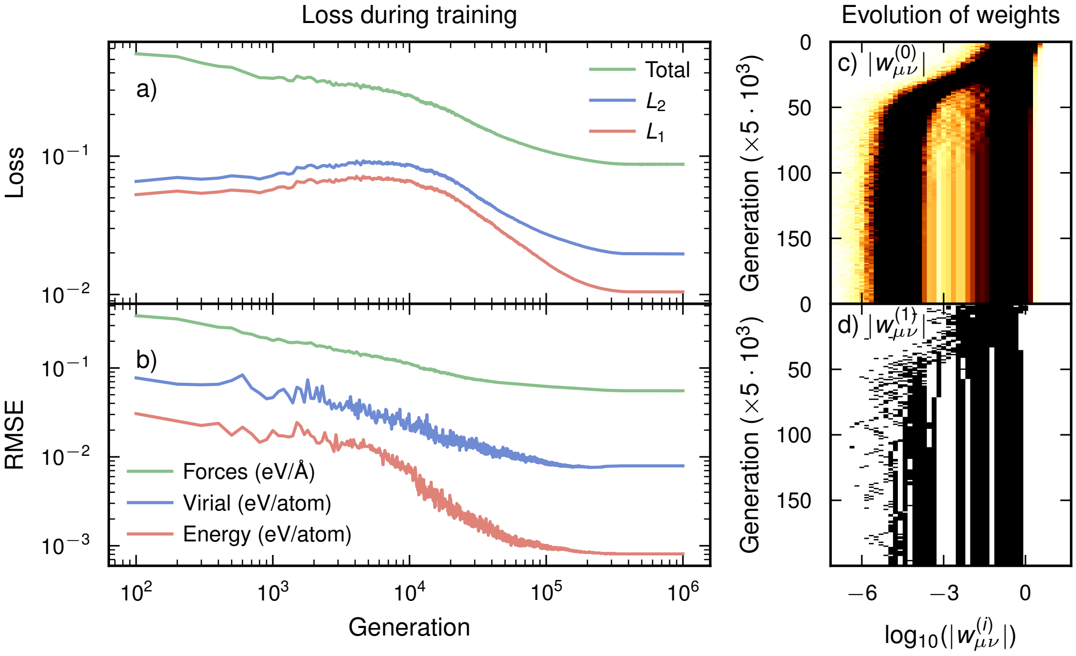

Figure 4.3: An example of how the loss function in Eq. (4.15) (a, b) and the magnitude of the weights in the hidden and output layers of the NN change (c, d) as a function of training generation. The effect of regularization can clearly be observed in many of the model weights decreasing in magnitude as the training progresses, starting around 50 000 generations.

the loss score, and each of the models is assigned a value u_k according to its rank (see [163] for explicit values of u_k). Third, the natural gradients of the fitness with respect to $\boldsymbol{m}$ and $\boldsymbol{s}$ are computed,

$$\begin{aligned} \nabla_{\boldsymbol{m}} J(\boldsymbol{z}) &= \sum_{k=1}^{N_\text{pop}} u_k \boldsymbol{r}_k \\ \nabla_{\boldsymbol{s}} J(\boldsymbol{z}) &= \sum_{k=1}^{N_\text{pop}} u_k (\boldsymbol{r}_k \odot \boldsymbol{r}_k - 1), \end{aligned} \tag{4.17}$$

which, finally, are used to update the mean and standard deviations of the parameter distribution,

$$\begin{aligned} \boldsymbol{m} &\leftarrow \boldsymbol{m} + \eta_{\boldsymbol{m}} (\boldsymbol{s} \odot \nabla_{\boldsymbol{m}} J(\boldsymbol{z})) \\ \boldsymbol{s} &\leftarrow \boldsymbol{s} \odot \exp\left(\frac{\eta_{\boldsymbol{s}}}{2} \nabla_{\boldsymbol{s}} J(\boldsymbol{z})\right). \end{aligned} \tag{4.18}$$

$\eta_{\boldsymbol{m}}$ and $\eta_{\boldsymbol{s}}$ are the equivalent of learning rates, and are set to $\eta_{\boldsymbol{m}} = 1$ and $\eta_{\boldsymbol{s}} = (3 + \ln N_\text{par})/(5\sqrt{N_\text{par}})$ following Refs. [163, 164].

Optimizing the loss function using the natural gradient instead of, e.g., the Euclidian gradient as in regular steepest descent optimization is beneficial, as the natural gradient takes the curvature of the loss landscape into consideration. This feature, in conjunction with the elastic-net regularization and the genetic form of the algorithm, leads to an efficient optimization scheme that yields both an accurate and sparse final model with few parameters contributing (Fig. 4.3). Sparsity is desirable as it decreases the effect of overfitting, where the parameters are adjusted too tightly to the training set. Overfitting can lead to the model predicting unphysical forces when it encounters an out-of-distribution atomic configuration, affecting the accuracy of the MD trajectory.

4.3.3 Active learning and uncertainty estimates

A comprehensive training set that covers the most relevant parts of configuration space is the most efficient method for minimizing the risk of the model extrapolating to unknown structures. To achieve this, the training dataset $\mathscr{D}_i$ is augmented using active learning [165, 166]. Typically, an ensemble of at least $N_{\text{ens}} = 5$ models is trained on various random subsets of the training dataset $\mathscr{D}_i$ through a process known as bootstrapping and aggregation, or bagging [167]. Then, a short MD simulation is run with one of the models, and for each atomic configuration the forces are predicted with each of the N_{ens} models. The uncertainty of each structure is then estimated as the force on atom i with the maximum standard deviation σ^i_{Force} over the ensemble models. After the simulation, the structures with the largest uncertainties are selected and target energies, forces and virials are computed using DFT, after which they are added to the dataset to yield an augmented dataset $\mathscr{D}_{i+1}$. Selecting structures with too large uncertainty might negatively impact the performance of the model, as the structures might be unphysical and subsequently have extremely large forces. This procedure is then repeated around 3-10 times, until the uncertainty of the structures encountered during MD fall beneath the force RMSE over the entire dataset, which can be seen as the achievable accuracy of the model. The initial dataset, $\mathscr{D}_0$, is typically comprised of various rattled and strained structures, starting from a reference structure via DFT.

Care must be taken when constructing the ensemble for active learning to ensure that the estimate of the prediction uncertainty is not underestimated. The model uncertainty will vary greatly depending on if one varies the data set, model hyperparameters, or trains with different random seeds, when constructing the models in the ensemble (Fig. 4.4). These different ensembles will sample increasingly different local minima in the loss landscape, where a more varied ensemble leads to a more representative estimate of the model uncertainty. Note, however, that the ensemble only probes the model uncertainty part of the overall prediction uncertainty. The reliability of the reference DFT calculations stemming from choice of functional or an incomplete training data set are examples of systematic sources of error, which ideally also should be investigated and modelled when training a MLIP.

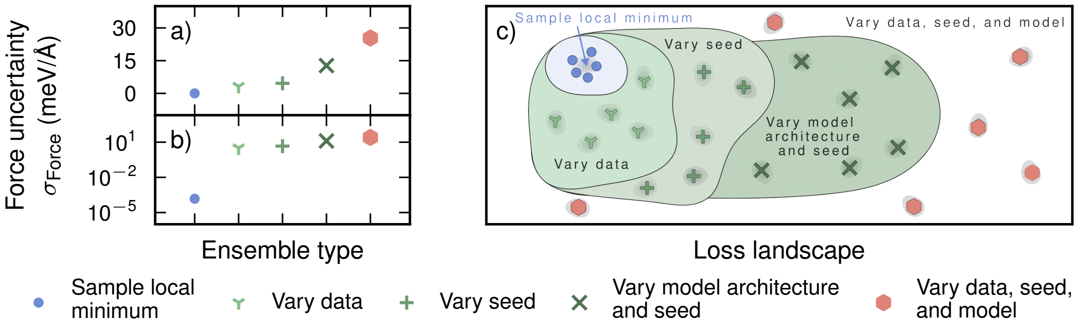

Figure 4.4: (a, b) Estimated force uncertainty σ_Force evaluated for five model ensembles on a SiO_2 dataset. The uncertainty magnitude depends on ensemble construction. The lowest uncertainty is obtained by sampling the SNES parameter distributions in Eq. (4.18), which only samples a single local minimum. Introducing variation via different training subsets (bagging; "Vary data") or optimizer initializations ("Vary seed") increases uncertainty by several orders of magnitude. Further increases occur when model architecture is also varied, with the combination of data, seed, and architecture variation producing the largest uncertainty. (c) These ensembles correspond to sampling regions of different extent in the loss landscape: more diverse ensembles explore more widely separated minima and better estimate the true model uncertainty.

By using active learning and ensemble uncertainties naïvely there is thus a risk of underestimating the prediction uncertainty. However, even an ensemble consisting of different model architectures, trained on different data and with different initialization, might underestimate the true prediction uncertainty. We can demonstrate this by forcing a MD simulation to sample configurations that are outside of the distribution covered by the training set, structures that are out-of-distribution for the NEP model, by running a high temperature and pressure simulation. The descriptor vector alone is an imperfect indicator of whether the model can distinguish two atomic environments: two structures may have different descriptors yet identical latent space representations, in which case the model predicts identical energies regardless. The latent space representation $\boldsymbol{\gamma}^i$ for each atom, corresponding to Eq. (4.7) before summing to yield a predicted energy, is therefore a better probe [168, 169],

$$\boldsymbol{\gamma}^i = w_\mu^{(1)} \tanh\left(\sum_{\nu=1}^{N_\text{des}} w_{\mu\nu}^{(0)} q_\nu^i - b_\mu^{(0)}\right). \tag{4.19}$$

To identify out-of-distribution configurations, we fit a Gaussian Mixture Model M with K Gaussians to the latent space configurations in the training set $\mathscr{D}$, and evaluate the probability of observing each latent space representation encountered during the MD simulation,

$$p(\boldsymbol{\gamma}^i|\mathscr{D}, M) = \sum_{k=1}^{K} \phi_k \mathcal{N}(\boldsymbol{\gamma}^i|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \tag{4.20}$$

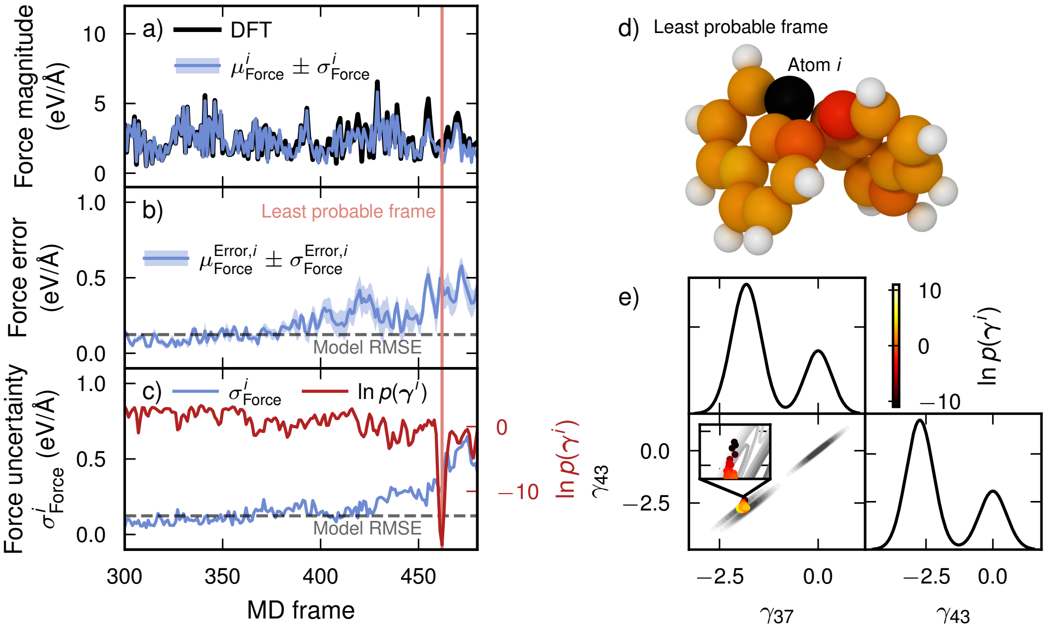

Figure 4.5: A high-temperature (1000 K), high-pressure (5 GPa) MD simulation of a perylene molecule under periodic boundary conditions (PBC) for 5 ps, that samples configurations beyond the training set. Atomic environment probabilities are computed via Eq. (4.20); the least probable atom i is highlighted in black (d). (a) Force magnitude for atom i from DFT and NEP, with mean μ^i_{Force} and $\sigma^i_{\mathrm{Force}}$. (b) Force error across the ensemble relative to DFT. Note that $\sigma^{\mathrm{Error,\,i}}_{\mathrm{Force}} \neq \sigma^i_{\mathrm{Force}}$, although they are similar; the former is the standard deviation in the *errors*, whilst the latter is the standard deviation of the predictions. The dashed line indicates model RMSE. (c) Force uncertainty and log-probability $\ln p(\boldsymbol{\gamma}^i|\mathscr{D}, M)$ over time; uncertainty increases and probability decreases near the least probable frame (462). (d) Atomic configuration at this frame, colored by log-probability (colormap in e); atom i in black. (e) Kernel density estimate of latent representations for $\mathscr{D}$ (black), with the trajectory of atom i colored by log-probability.

where $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_k$ are the mean vector and covariance matrix of Gaussian k, respectively, and the mixture weights ϕ_k sum to 1. The probability $p(\boldsymbol{\gamma}^i|\mathscr{D}, M)$ can then be interpreted as the out-of-distributionness for a latent space vector $\boldsymbol{\gamma}^i$; a low probability corresponds to a less probable chemical environment around atom i, while a high probability means that the configuration is firmly within the training dataset. The least probable frame and atom in that frame can then be identified as the configuration which has the lowest probability for any atom over the MD simulation (Fig. 4.5d).

By studying the predictions for the least probable atom i, we can see that our ensemble uncertainty estimate falls short. Already a few hundred MD steps before the least probable configuration, the latent space probability begins to decline and the force errors relative to DFT start to grow (Fig. 4.5a–c). Yet the standard deviation in the force prediction for atom i, σ^i_{Force}, remains close to the model RMSE throughout, and does not reflect the actual error (Fig. 4.5b and c). The ensemble is *confidently wrong*.

The probability dips sharply for the least probable atomic configuration (Fig. 4.5c), and then recovers but trends downwards for the subsequent frames while σ^i_{Force} continues to increase. This analysis suggests that $\ln p(\boldsymbol{\gamma}^i|\mathscr{D}, M)$ can be used to identify out-of-distribution configurations during a MD simulation, perhaps even as a replacement to using an ensemble to compute an uncertainty estimate. Training a single model instead of an ensemble of $N_{\text{ens}} = 5$ NEP models would dramatically save compute resources, as training a NEP model typically requires on the order of 100-400 GPUh[1]. The full ensemble must be refit for each generation of active learning, leading to a considerable computational effort.

As an aside, the out-of-distribution analysis also yields an heuristic for how to select structures from active learning. The point where σ^i_{Force} starts to increase roughly coincides with the increase in force error compared to DFT (Fig. 4.5b and c), which is precisely what we aim to minimize through active learning. A good threshold for the uncertainty at which to select structures during active learning is thus when σ^i_{Force} is greater than the model RMSE in the NEP loss (Eq. (4.15)), i.e., when the uncertainty is greater than the prediction error.

4.3.4 The NEP framework in practice

The formalism, optimization, and training procedure outlined in the three previous subsections are key for the success of the NEP approach. To demonstrate this in practice, extensive benchmarking comparing NEP to five other state-of-the-art models has been performed. This includes MACE [142], NequIP [140], Allegro [141], and linear and non-linear models based on the ACE framework [146–148]; all models have been tested for a range of different hyperparameters. The models were all trained and evaluated on the extensive test for SiO_2 from Leimeroth et al. [170]. NEP achieves competitive accuracy on both energies and forces, although it is outperformed by the graph neural network (GNN)–based message passing models MACE and NequIP (Fig. 4.6a and b). Message passing propagates information along the molecular graph, effectively increasing the receptive field beyond the cutoff radius r_c of the model. Local models, such as NEP, have a receptive field limited by the cutoff radius r_c, and can be loosely viewed as a GNNs with a single message passing step. However, message passing negatively impacts computational performance, and this is where NEP excels. For a small cell of 6144 atoms, NEP is almost two orders of magnitude faster than the most accurate MACE and NequIP models, and roughly one order faster than the least accurate but fastest models. Compared to linear and non-linear ACE, models that architecturally are similar to NEP, NEP is

---

[1] For example, training a model on a dataset with about 1000 structures and 100 000 atoms takes around 50 h using four NVIDIA A100 GPUs.

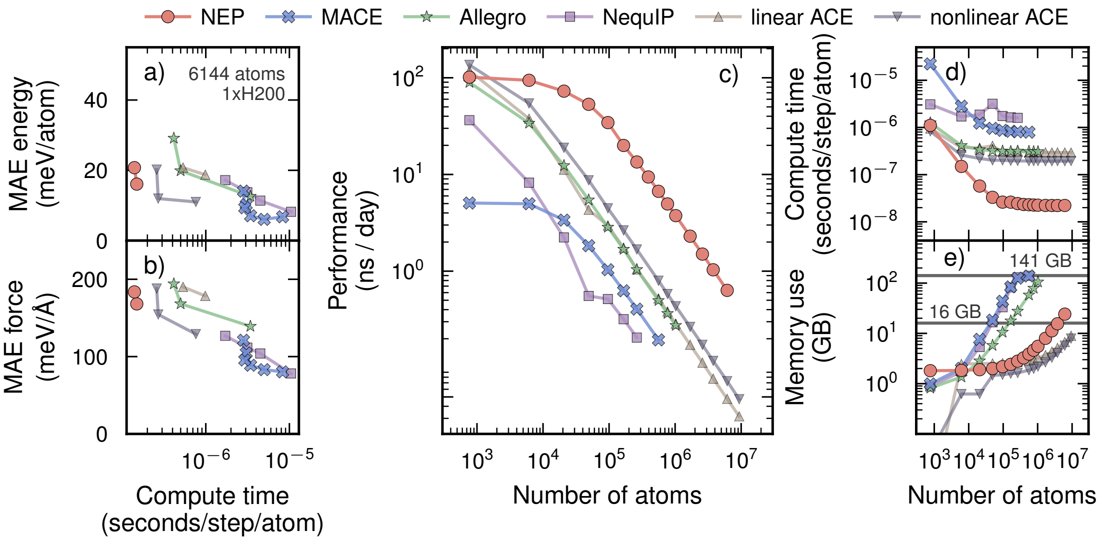

Figure 4.6: Benchmarks comparing NEP to various other state-of-the-art models as of 2026 for an amorphous SiO_2 system from Leimeroth et al. [170]. The MACE and ACE models where taken from Ref. [170], while the NequIP and Allegro models were retrained to benefit from recent optimizations using model compilation for these models. The models used in panels (c–e) are the fastest models from panel (a) and (b). The simulations for NEP were performed using GPUMD, while the rest were run with Kokkos in LAMMPS, all using a time step of 1 fs and run on a single NVIDIA H200 GPU.

almost a factor of two faster. Interestingly, the fastest Allegro models achieve worse performance than NEP, and the model that outperforms NEP in accuracy has approximately the same computational performance as MACE. Allegro is a local-only version of NequIP without any message passing steps, just like NEP, and the fact that a complex model is required to beat the accuracy of NEP demonstrates the efficiency of the NEP approach. The MD simulations for the timing results were performed with GPUMD for NEP, and with LAMMPS for the other models [171, 172]. All LAMMPS simulations were performed using the Kokkos package, and model compilation was enabled where supported [173–176]. ACE models also used the PACE implementation in LAMMPS [177].

NEP scales well to large systems, and is approximately five times faster than the fastest competing model, a small basis set linear ACE model, for 10^6 atoms (Fig. 4.6d). With NEP, you can perform 80 ns of MD simulation in a day with a 1 fs time step for 10^4 atoms on an NVIDIA H200 GPU, and 3 ns for 10^6 atoms (Fig. 4.6c). This is consistently almost an order of magnitude more than what can be achieved with the fastest Allegro and MACE models for the same system sizes.

Memory use is the primary driver for what hardware is needed in order to run simulations. A high-end consumer GPU like the NVIDIA RTX 5080 that can be found in a workstation has 16 GB of graphics memory. This limits simulations to tens of thousands

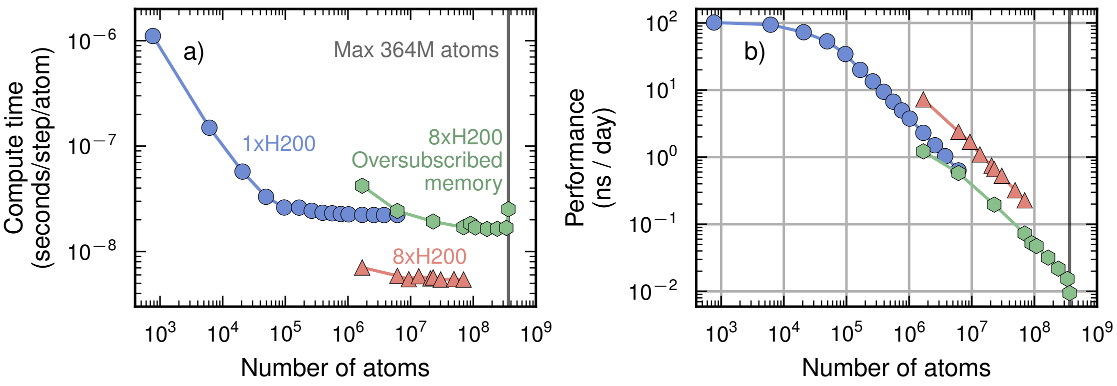

Figure 4.7: Pushing the limits of the NEP approach on a single node, by performing large MD simulations for SiO_2. All simulations were run with GPUMD, with a time step of 1 fs, on NVIDIA H200 GPUs. NEP scales efficiently to eight GPUs, achieving a speedup of five, and reaches a maximum system size of 364 million atoms using memory oversubscription from GPU to host memory.

of atoms for MACE, and hundreds of thousands for Allegro, while NEP and ACE scale into the millions (Fig. 4.6e). The scaling results are for the small and thus fastest models for MACE, NequIP, and Allegro, with few layers and small basis sets. Memory usage grows sharply with model size for these architectures, where the most accurate MACE model with 256 message passing channels exhausts the H200's 141 GB of video memory around 2×10^4 atoms, while NEP and ACE remain roughly constant as the model size is changed. Note also that all models here are trained specifically for SiO_2; the memory footprint of foundation models such as MACE-MP-0 is considerably larger.

NEP also parallelizes efficiently to multiple GPUs, achieving a speedup of about a factor of five when using eight H200 GPUs (Fig. 4.7). A speedup of five out of a possible eight indicates reasonably efficient scaling, despite the communication overhead that grows with GPU count. By enabling oversubscribing of GPU memory to using the host memory, which for a full node with 8 H200 GPUs leads to a combined 1.1 TB of video memory and 2.1 TB of system memory, we can simulate up to 364 million atoms on a single node. Using the host memory is significantly slower than the video memory on the GPUs and is currently an experimental feature in GPUMD, which explains why the performance when oversubscribing memory is on par with that of a single GPU, albeit at much larger system sizes.

These benchmarks highlight the design principle of NEP: high computational efficiency without sacrificing accuracy, enabled by GPUMD's native C++/CUDA implementation. NEP in GPUMD thus directly addresses the need of simulating large systems of chromophores for long time scales identified in the previous chapter, and has been the MLIP of choice throughout the thesis. Raw performance alone is not sufficient, however.

Features, ecosystem integration, and usability are equally important for a practical simulation engine. To this end, I have contributed several features to GPUMD, including the active learning module, and developed CALORINE, presented in paper IV. CALORINE is a Python toolbox for GPUMD simulations and NEP construction that streamlines the user experience. It also makes NEP-models transferable to other Python-based workflows through integration with the Atomic Simulation Environment (ASE) [178], making NEPs available to the broader materials research community.

4.4 Fine-tuning foundation models

So far in this chapter we have focused on MLIPs trained for a specific system, which we can call custom models. This means that a new model needs to be trained for each system or material of interest, which requires considerable computational effort and time. A recent development is the concept of foundational MLIPs, models trained on large parts of the periodic table. These models boast competitive performance on a broad range of inorganic and organic systems, at the cost of extensive computational effort in terms of training data and training of the actual model, as well as increased model complexity. Examples of foundation models include ones based on the MACE (e.g. MACE-MP0 [179], MACE-OFF23 [180]), and NequIP frameworks [181], as well as M3GNet [182], CHGNet [183], and PET-MAD [152]. These are highly accurate equivariant models, but scale poorly to large systems. In particular, foundational models based on these frameworks tend to scale even worse than the custom models based on the NequIP and MACE frameworks that we considered in Sect. 4.3.4. In paper V we present the NEP89 foundational model, which can handle systems containing up to 89 different chemical elements while preserving the simple architecture and high computational efficiency of the NEP framework. This is enabled by the NEP4 framework [184], which ensures descriptors scale only quadratically and the neural network linearly in the number of chemical species. NEP89 scales to systems containing tens of millions of atoms and beyond, and is trained for crystalline, amorphous, and molecular systems.

Foundation models follow the old adage "jack-of-all-trades, master of none". Although they are performant on average for a broad range of systems, in practice they fall short of the accuracy required for many applications. In these cases, foundational models can be fine-tuned by training on small amounts of data for the system of interest. We can demonstrate this for the NEP89 model by fine-tuning on 100 structures from the extensive dataset for SiO_2 from Leimeroth et al. [170] that we used in Sect. 4.3.4. The crystalline phases of SiO_2 are quite close in energy, and NEP89 gives the wrong phase order compared with DFT, where α-quartz is predicted as the second lowest phase in energy. The order of phases change during fine-tuning, with the biggest change occurring at around 600 fine-tuning generations, at which point α-quartz and β-quartz are the lowest in energy (Fig. 4.8a and b). Although the losses continue to decrease as

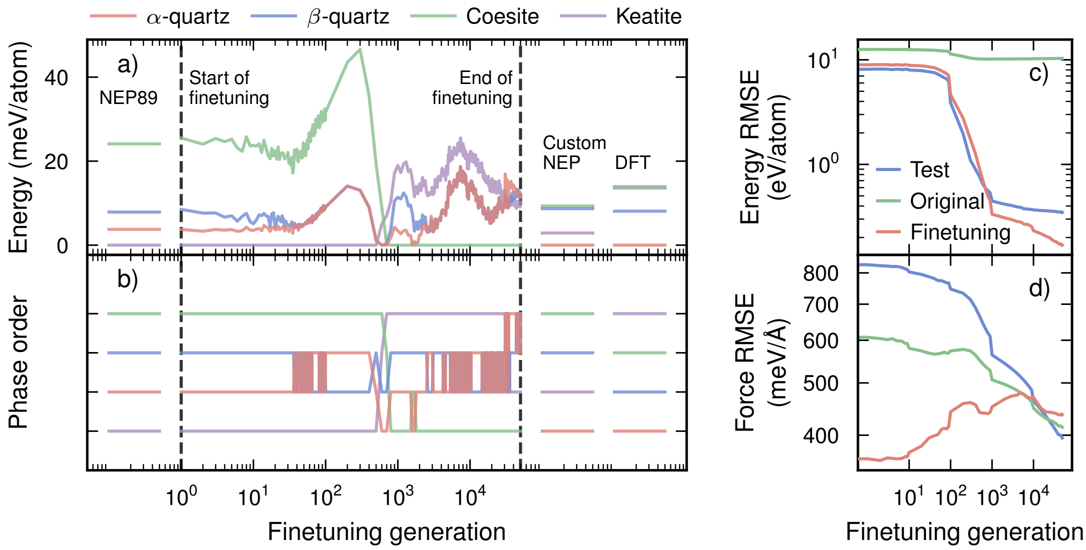

Figure 4.8: Finetuning on 100 structures from Leimeroth et al. [170] for SiO_2 over 50 000 generations. (a) Predicted phase energy and (b) order for each generation of finetuning, compared to NEP89, a custom NEP, and DFT with the SCAN exchange correlation functional. (c) Energy and (d) force RMSE evaluated on the test set from Leimeroth et al. (Test), the part of the original NEP89 training set containing SiO_2 (Original), and the finetuning dataset (Finetuning).

finetuning continues, the phase order changes dramatically, with the final fine-tuned model giving a vastly different phase order compared to a custom NEP model and DFT, with coesite as the lowest lying phase. This suggests that the model is starting to overfit to the finetuning data, leading to a catastrophic “forgetting” of the original training data that worsens performance, which might not be obvious from the monitored losses alone (Fig. 4.8c and d). It is thus important to carefully monitor performance on the property of interest during fine-tuning, in addition to the original dataset.

Molecular systems, such as perylene, pose additional challenges due to the importance of weak intermolecular interactions. We can illustrate this issue by fine-tuning NEP89 on 42 perylene structures. Fine-tuning improves the accuracy of the predicted inter- and intramolecular forces on an external test set compared to the original NEP89 model, albeit not to the same level as a custom NEP (Fig. 4.9). However, all NEP models perform significantly better on the stronger intramolecular forces than on the weaker intermolecular forces (Fig. 4.9). This trend is particularly pronounced for OPLS, a classical FF for organic systems [185] that explicitly treats intermolecular van der Waals-forces and electrostatic Coulomb interaction. Interestingly, OPLS reproduces the mean squared displacement (MSD) and density of the fine-tuned and custom NEP models reasonably well (Fig. 4.10b and c), yet its molecular centroid-to-centroid pair distribution function differs markedly, exhibiting fewer peaks indicating a different supramolecu-

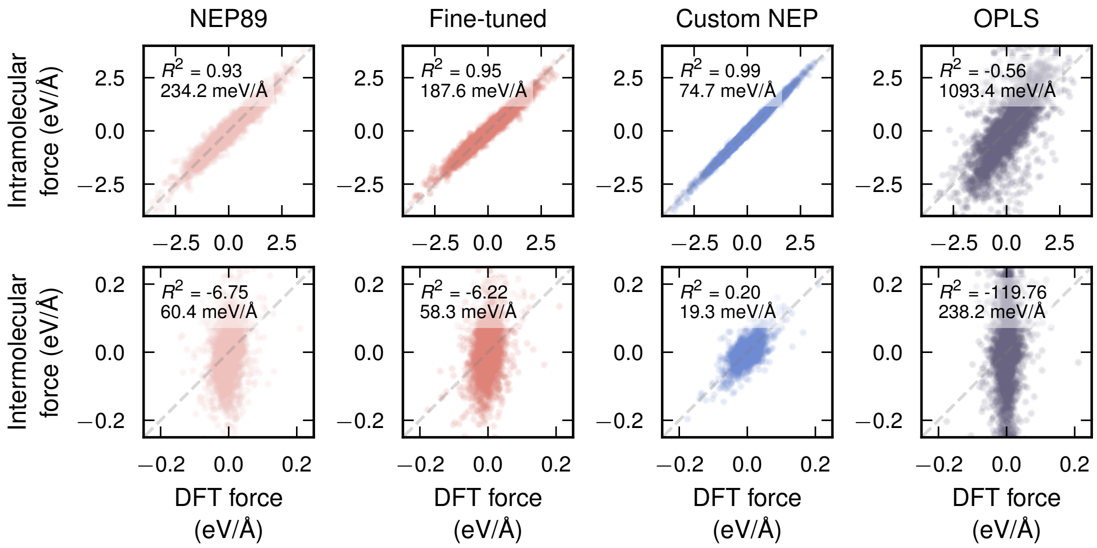

Figure 4.9: Predicted intramolecular and intermolecular forces for NEP89, fine-tuned NEP89 (42 perylene structures), a custom NEP, and the FF OPLS, with RMSE and R^2 insets. DFT reference forces were computed with the vdW-DF-cx functional [124]. All NEP models reproduce the intramolecular forces well but struggle with the intermolecular forces, which are roughly an order of magnitude smaller. OPLS performs worse than all NEP models.

lar structure (Fig. 4.10a). This inconsistency highlights the subtlety of intermolecular interactions, where an inaccurate treatment can go undetected in some properties while giving a clear signature in others.

Intermolecular interactions from, e.g, van der Waals-forces and electrostatic interactions, are typically not explicitly modeled in MLIPs, including NEP. The expectation is instead that the MLIP will learn these effects implicitly from the training data. In practice, this implicit learning is hampered by intermolecular forces contributing weakly to the loss function relative to the stronger intramolecular forces. A further, more fundamental challenge is the inherently local nature of cutoff-based MLIPs, with cutoff radii on the order of 6 Å. While this is enough for many atomic and crystalline systems, intermolecular interactions in molecular system can extend over longer distances. Message-passing architectures such as MACE and NequIP increase the receptive field of the model and possibly the accuracy in reproducing these forces, but they do not address this fundamental limitation.

Aggregation and glass dynamics in chromophores rely on an accurate treatment of intermolecular interactions, and remains an outstanding challenge for the workflow developed in this thesis. This can be addressed in several ways: Grimme D3-dispersion [186] can be added during MD simulations, the loss function can be reweighted to penalize errors on the weak intermolecular forces separately from the intramolecular ones,

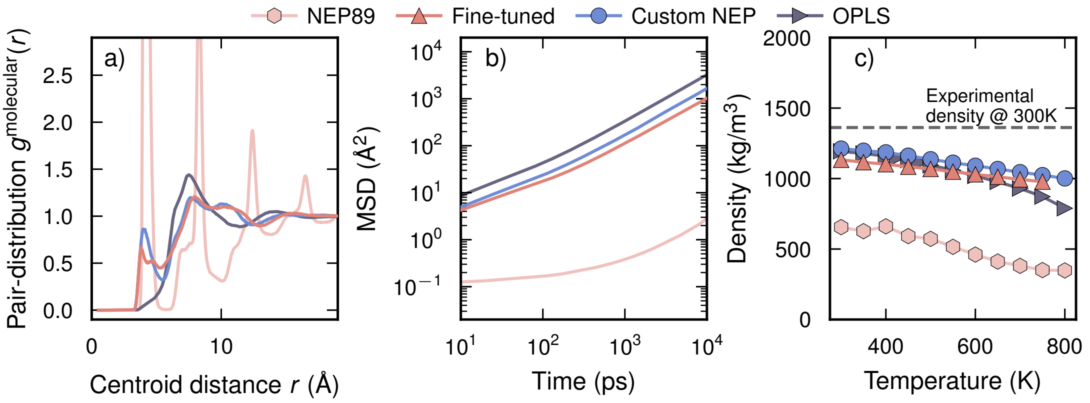

Figure 4.10: Simulated properties from large-scale MD simulations of 2000 perylene molecules with NEP89, a fine-tuned NEP89 model trained on 42 structures, a custom NEP model for perylene, and the classical force field OPLS. (a) Molecular pair distribution function at 600 K. (b) Molecular MSD based on the position of the centroid in each molecule, at 600 K. (c) Density as a function of temperature. Note that the density is missing for the fine-tuned NEP89 model at 800 K, since the simulation failed.

and coarse-grained representations could extend the effective cutoff range. Most promising is qNEP [187], which predicts atomic charges alongside energies and forces, enabling long-range Coulomb interactions and directly addressing the locality limitation.

Despite this challenge, NEP offers exceptional computational performance and competitive accuracy, precisely the properties demanded by the large-scale MD simulations in this thesis.

5

Connecting to experiments I — Beyond forces

> I will primarily use unnecessarily complicated and highly inefficient but really huge neural network models to burn a lot of computer time. The goal of this project is to max out the power consumption of as many GPUs as possible. The approach is based on large language models that have been shown to be super duper good in creating heat.
>
> *Paul, 2024*

Equipped with accurate and efficient machine-learned interatomic potentials (MLIPs) based on the neuroevolution potential (NEP) framework, we now have everything needed for accurate and large-scale molecular dynamics (MD) simulations. Connecting these simulations to experiments requires correlation functions of specific observables depending on the experiment, as discussed in Chapter 3. However, many of these observables require additional electronic structure calculations, negating the computational gains of using a MLIP. This raises the question: can we extend the NEP framework to predict such observables directly? In this chapter, we explore two extensions in this direction. The first, the tensorial neuroevolution potential (TNEP) framework, extends NEP to tensorial properties, such as dipoles and polarizabilities, establishing a direct link to infrared (IR) and Raman spectroscopy via their autocorrelation functions (ACFs). The second extension takes the idea further by using the NEP descriptors as input for a more expressive neural network (NN), capable of predicting the frequency-dependent complex-valued dielectric function, from which observables such as absorption spectra and optical band gaps can be obtained.

5.1 Tensorial properties with the NEP framework

In paper VI we present the TNEP framework, which builds tensorial predictions on two related quantities: the partial force $\partial U_i/\partial r_{ij}^{\nu}$, and the rank-2 virial tensor $W^{\mu\nu}$,

$$W^{\mu\nu} = -\sum_{i}^{N}\sum_{j\neq i} r_{ij}^{\mu}\frac{\partial U_i}{\partial r_{ij}^{\nu}}, \tag{5.1}$$

where r_{ij}^{ν} is the ν-component of the interatomic distance vector $\boldsymbol{r}_{ij} = \boldsymbol{r}_j - \boldsymbol{r}_i$. Although rank-1 tensors such as the molecular dipole moment $\boldsymbol{\mu}$ or the electric polarization $\boldsymbol{P}$ might seem directly predictable from the partial forces, this fails since the partial forces sum to zero by construction. Instead, the TNEP framework defines rank-1 tensors through a contraction of the virial tensor with a vector,

$$\mu^{\nu} = -\sum_{i}^{N}\sum_{j\neq i}\sum_{\mu} r_{ij}^{\mu}\left(r_{ij}^{\mu}\frac{\partial U_i}{\partial r_{ij}^{\nu}}\right) = -\sum_{i}^{N}\sum_{j\neq i} r_{ij}^{2}\frac{\partial U_i}{\partial r_{ij}^{\nu}}. \tag{5.2}$$

For rank-2 tensors such as the molecular polarizability $\alpha^{\mu\nu}$ or electric susceptibility $\chi^{\mu\nu}$, the virial tensor is slightly adapted for numerical stability,

$$\alpha^{\mu\nu} = \sum_{i}^{N} U_i\delta^{\mu\nu} - \sum_{i}^{N}\sum_{j\neq i} r_{ij}^{\mu}\frac{\partial U_i}{\partial r_{ij}^{\nu}}, \tag{5.3}$$

where the Kronecker delta $\delta^{\mu\nu}$ stabilizes the predictions, since the diagonal elements tend to dominate over the off-diagonal elements.

In paper VI we apply the TNEP framework to the molecular dipole, electric polarization, molecular polarizability, and the electric susceptibility[1] to enable simulation of IR and Raman spectra, two widely used, nondestructive spectroscopic techniques. These spectra can be simulated at finite temperatures using *ab-initio* MD, where the relevant tensorial quantities are computed at every step using density functional theory (DFT), which is accurate but limited in accessible system sizes and time scales. The TNEP approach predicts both forces and tensorial properties, inheriting the efficiency of the NEP framework (Fig. 5.1a). For IR spectroscopy, the relevant quantity is the dipole $\boldsymbol{\mu}$ or polarization $\boldsymbol{P}$, the ACF of which yields the isotropic IR absorption cross section [192],

$$\sigma(\omega) \propto \frac{\omega^2}{2\pi}\int_{-\infty}^{\infty}\langle|\boldsymbol{\mu}(t)||\boldsymbol{\mu}(t+\tau)|\rangle_t \, d\tau. \tag{5.4}$$

Other machine-learned approaches for the dipole moment and related spectroscopic quantities have been proposed [193–195], including the recent MACE4IRMOL foundation model [196], though TNEP stands out for its favourable computational scaling.

---

[1] The distinction between molecular and electronic properties refers to the ambiguity in defining polarization for periodic systems, handled via the *modern theory of polarization* [188–191].

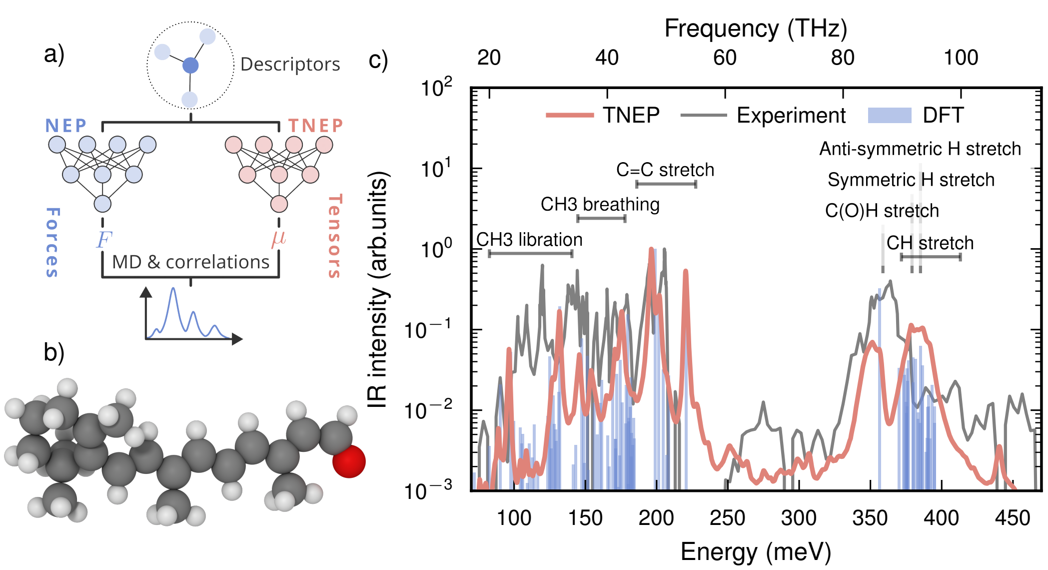

Figure 5.1: (a) Schematic of how a NEP and TNEP model evaluate forces and the molecular dipole moment $\boldsymbol{\mu}$ at each time step during MD to predict IR spectra via correlation functions. (b) The all-*trans* isomer of retinal, $C_{20}H_{28}O$. (c) IR spectrum for all-*trans*-retinal from TNEP obtained at 300 K, compared with experiments [198] and the 0 K IR spectrum from DFT calculated using VeloxChem [53–55] with the B3LYP functional [199] and the DEF2-SV(P) basis set.

As an example, consider the IR spectrum of the chromophore retinal (Fig. 5.1b). Retinal is a vitamin A derivative responsible for light absorption in eyes, enabling vision among vertebrates [197]. It is a preferred model system over perylene for TNEP, since the permanent dipole moment in retinal makes it a straightforward quantity to learn. Perylene, on the other hand, is non-polar, and any instantaneous molecular dipole arises from small variations in the molecular shape due to thermal fluctuations, which are difficult to learn due to a small signal-to-noise ratio in the reference dipole moments.

By constructing a NEP and a TNEP model for retinal on about 800 structures, we predict both forces and dipole moments over a NVE MD simulation with equilibration in NVT at 300 K, from which the IR spectrum can be computed using Eq. (5.4). The TNEP spectrum is consistent with the 0 K DFT reference, though both overestimate vibrational frequencies, most notably in the C–H stretch region around 325 meV to 400 meV, relative to Fourier-Transform IR (FTIR) measurements [198] (Fig. 5.1c). The intensities of the softer librational modes of the methyl groups are also underestimated by both methods, which indicates an underestimation of the associated dipole moment by DFT. Overall, however, the predicted spectra from TNEP and DFT are both in semi-quantitative agreement with experiment, but the NEP and TNEP models can be evaluated at a fraction of the cost of DFT. NEP and TNEP additionally gives access to temperature dependence and anharmonic effects, which are not included in the 0 K DFT spectrum.

5.2 Dielectric functions from NEP descriptors

While the TNEP approach captures tensorial properties, they are by construction evaluated as static quantities at each time step. The frequency dependence of dynamical observables such as an IR spectrum can be obtained via correlation functions, but electronic properties such as the dielectric function or the electronic density of states are intrinsically electronic and cannot be resolved from the ionic dynamics alone. This raises the question of whether the full frequency-dependent spectrum of such a property can be predicted directly from atomistic descriptors using machine learning (ML). The electronic density of states is an example that has attracted attention in the literature, with several models trained to predict it across broad materials classes, following the philosophy of universal MLIPs [200–202]. However, universal models require large training datasets containing tens of thousands to millions of structures, and although they yield on average acceptable performance they might fail for individual systems. System-specific models can in principle achieve comparable or better accuracy with far less data.

Another, more challenging, target that depends on pairs of electronic states is the electronic dielectric function $\boldsymbol{\epsilon}(\omega)$. The electronic dielectric function is a complex-valued, frequency dependent rank 2-tensor that relates the polarization $\boldsymbol{P}$, the induced total dipole moment per unit volume in the system, to an external field $\boldsymbol{E}$ [192],

$$\boldsymbol{P} = \boldsymbol{\chi}(\omega)\boldsymbol{E} = \frac{1}{4\pi}\left(\boldsymbol{\epsilon}(\omega) - 1\right)\boldsymbol{E} \tag{5.5}$$

where $\boldsymbol{\chi}(\omega)$ is the electric susceptibility. A key property of $\boldsymbol{\epsilon}(\omega)$ is that the real and imaginary part of the dielectric function, $\boldsymbol{\epsilon}(\omega) = \boldsymbol{\epsilon}^{\mathrm{Re}}(\omega) + i\boldsymbol{\epsilon}^{\mathrm{Im}}(\omega)$, are not independent. They are related to each other via the Kramers-Kronig relation,

$$\epsilon_{ij}^{\mathrm{Re}}(\omega) = 1 + \frac{2}{\pi}\mathscr{P}\int_0^{\infty} \epsilon_{ij}^{\mathrm{Im}}(\omega')\frac{\omega'}{\omega'^2 - \omega^2}d\omega', \tag{5.6}$$

where i and j are Cartesian directions and $\mathscr{P}$ denotes the principal value of the integral. Predicting either $\boldsymbol{\epsilon}^{\mathrm{Re}}(\omega)$ or $\boldsymbol{\epsilon}^{\mathrm{Im}}(\omega)$ with a ML model is therefore sufficient to recover the other. Evaluating the loss over both representations of the dielectric function effectively adds a regularization constraint that ensures physically meaningful predictions.

In paper VII we present a system-specific model for predicting the electronic dielectric function $\boldsymbol{\epsilon}(\omega)$ directly from pretrained NEP descriptors for crystalline, amorphous, and liquid SiO_2. Compared to the broad chemical range targeted by current universal models for the density of states, we target the dependence of $\boldsymbol{\epsilon}(\omega)$ on temperature, pressure, phase, and local structure, which requires considerably higher accuracy. We focus on a single system and predict the isotropic average, $\epsilon(\omega) = \epsilon^{\mathrm{Re}}(\omega) + i\epsilon^{\mathrm{Im}}(\omega)$, rather than the full tensor. To ensure that the spectrum is physically well-behaved, $\epsilon(\omega)$ is parametrized as a sum of Lorentzian functions by the model (Fig. 5.2), from which both the real and imaginary parts are constructed analytically.

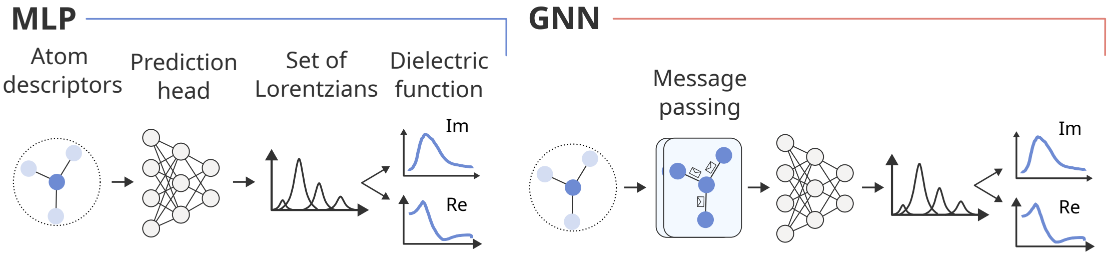

Figure 5.2: The model classes, multi-layer perceptron and graph neural network (GNN), used for predicting the dielectric function from atom descriptors in paper VII. From an atomic descriptor a set of Lorentzians are predicted, typically 60, from which the real and imaginary part of the pseudo-dielectric function for that atom is constructed. The dielectric function for the entire system is then obtained by averaging the pseudo-dielectric functions.

Pretrained NEP descriptors make the approach very data efficient, accurately predicting the dielectric function with only 200 reference calculations (Fig. 5.3a-c). We trained two model classes, based on a multi-layer perceptron and a two-layer GNN, both of which reproduce the DFT reference equally well. This indicates that the NEP descriptors account for most of the prediction accuracy, rather than the model architecture.

The dielectric function gives direct access to the optical absorption coefficient [203],

$$\eta(\omega) \propto \omega\sqrt{|\epsilon(\omega)| - \epsilon^{\mathrm{Re}}(\omega)}. \tag{5.7}$$

The optical band gap can be estimated from the absorption coefficient via the Tauc gap [204], conventionally obtained by extrapolating the tail of the absorption coefficient to zero. Following Ref. [205], we instead used a less conventional approach, where the Tauc gap is estimated as the energy at which the absorption coefficient exceeds a threshold value of $5 \times 10^3\,\mathrm{cm}^{-1}$ (Fig. 5.3c). We evaluated the Tauc gap as a function of temperature using large-scale MD simulations, averaging over an ensemble of multi-layer perceptron and GNN models to obtain prediction uncertainty estimates (Fig. 5.3d). The ensemble reproduces the experimentally observed temperature slope, but with a constant offset. This offset is a consequence of the band gap underestimation in the PBE exchange–correlation (XC) functional [206, 207], which was corrected for by shifting the dielectric functions to match the experimental band gap of the α-quartz phase.

A conceptual challenge in this approach is that the dielectric function is a global property of the system, while the NEP descriptors are local to individual atoms. Rather than aggregating local features at an intermediate stage of the model, we instead predict a pseudo-dielectric function for each atom individually and take the mean to produce the system-level dielectric function. This choice is physically motivated by mixed-phase systems, in which the predicted dielectric function is the average of the contributions of each phase. A single global prediction can in these cases lead to non-physical behavior, as mixed-phase structures were not included in the training set. The approach does

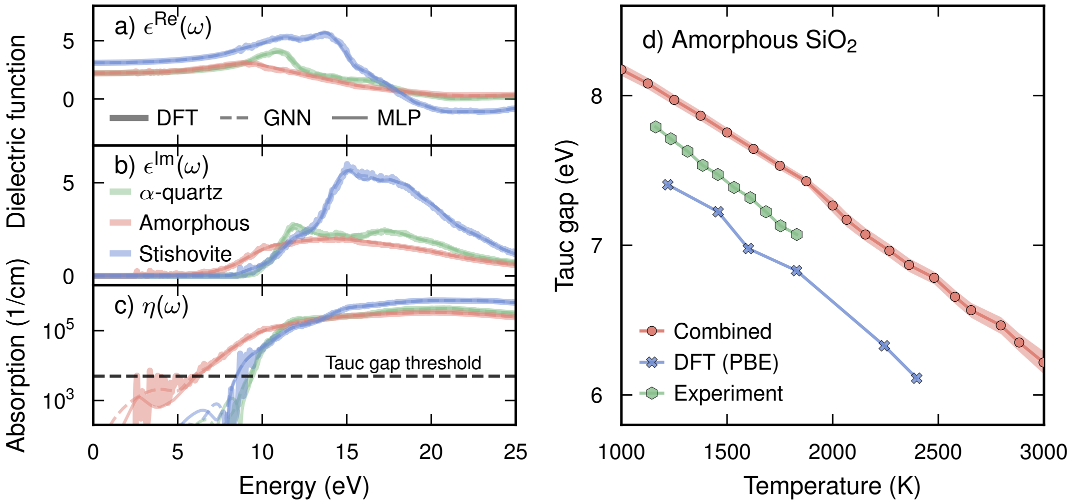

Figure 5.3: Predicted real (a) ϵ^{Re} and imaginary (b) ϵ^{Im} part of the dielectric function, as well as absorption coefficient $\eta(\omega)$ (c), for both model classes and compared to a DFT reference. All models reproduce the DFT reference. Panel (c) also shows the threshold used for estimating the Tauc-gap, following Ref. [205]. (d) The temperature dependence of the Tauc-gap in amorphous $\mathrm{SiO_2}$ for a combined model ensemble consisting of a total of 12 multi-layer perceptron and GNN models trained on different splits of the training data, with the error band corresponding to a 95% confidence interval. Reference DFT calculations used the PBE XC functional [206] and are from [203], and experimental values are from [205].

carry limitations: storing the full pseudo-dielectric function per atom is memory intensive, and the simple mean implicitly assumes independent contributions from each atom. A global prediction, on the other hand, is more memory efficient and sufficient for single-phase systems. Whether an aggregation scheme that considers correlations between pseudo-dielectric functions would improve predictions in mixed-phase systems is an open question and a natural direction for future work.

Extending this approach to chromophores is considerably more demanding. Both intramolecular and intermolecular interactions shape the electronic structure and optical response, producing a dielectric function that is highly sensitive to the relative orientation of the molecules. For example, small structural changes drive the spectral shifts characteristic of H- and J-aggregates. Accurately reproducing these shifts is an essential benchmark for a unified framework connecting chromophore structure, dynamics, and optical properties. This remains out of reach with the current NEP framework, where intermolecular interactions are not yet accurately captured (Sect. 4.4). Improving the description of these interactions is therefore a prerequisite for extending the dielectric function approach to molecular systems, and a natural next step toward a fully predictive simulation framework for chromophores.

6

Connecting to experiments II — Simulating neutron scattering

> I'm afraid neutrons will not be of any use to any one.
>
> *Sir James Chadwick,*
> *discoverer of the neutron*

Neutron scattering is one of the most important experimental techniques for studying the structure and dynamics of chromophores, as neutrons are excellent probes for organic matter. In neutron scattering, the neutron scatters directly off the atomic nuclei, with scattering cross-sections that vary greatly and non-monotonically across the periodic table. Hydrogen, in particular, scatters neutrons strongly, making the technique especially sensitive to dynamics in hydrogen-rich organic matter such as chromophores. Neutrons additionally transfer relatively small amounts of kinetic energy to the sample under study, typically in the meV range, which is comparable to the energy scales associated with molecular motion and other processes relevant to the dynamics of chromophores. Unfortunately, producing neutrons requires large facilities, such as nuclear reactors or spallation sources, and, hence, these are only available in a handful of locations globally (Fig. 6.1). Simulations play a key role in optimally utilizing these scarce neutrons, not only in interpreting experimental results, but also serving as a guide for which experiments to conduct.

Just as we saw in Chapter 5 for IR and Raman spectroscopy, correlation functions also link simulations to neutron scattering. The dynamic structure factor $S(\boldsymbol{q}, \omega)$ is the central quantity in neutron scattering, obtained as an autocorrelation function (ACF) of the time-dependent particle density, and it is a remarkable correlation function that describes the structure and dynamics in a system [87, Chapter 5].

(a) TU Delft, the Netherlands **(b)** SINQ, Switzerland

(c) J-PARC, Japan **(d)** ESS (under construction), Sweden

Figure 6.1: Four of the handful of neutron scattering facilities globally. The TU Delft experimental reactor in the Netherlands, the SINQ beamline in Switzerland, the J-PARC beamline in Japan, and the ESS beamline in Lund, Sweden. Photos taken by me during visits to various beamlines.

In this chapter we will follow Squires [208, Chapters 1, 2 and 4] in deriving the dynamic structure factor, and demonstrate how it can be computed from the time-dependent particle density obtained via molecular dynamics (MD) simulations. We then turn to how the dynamic structure factor can be weighted and broadened by the resolution function of particular neutron scattering instruments, enabling direct comparison between simulated and experimental signatures.

6.1 Basic neutron scattering theory

The quantity measured in neutron scattering is the scattered intensity from each measured neutron, $I(\boldsymbol{q}, \omega)$, where $\boldsymbol{q}$ is the momentum transferred to the sample, and ω is the angular frequency of the excitation gained or lost by the sample. Conceptually, this

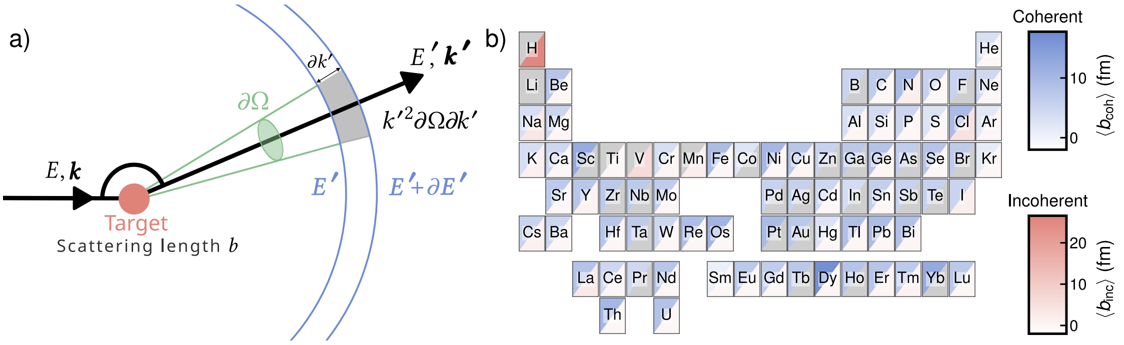

Figure 6.2: (a) Scattering of a neutron with incoming wave vector $\boldsymbol{k}$ and energy E into solid angle $\partial\Omega$ with momentum change $\partial k'$ when hitting a target with scattering length b, with the final neutron having wave vector $\boldsymbol{k'}$ and energy between E' and $E' + \partial E'$. (b) Coherent and incoherent scattering lengths from [209], which squared yield the cross sections, across the periodic table. Note that some species have negative scattering lengths; these are marked in gray. Species for which no data is available are skipped.

can be reformulated as $I(\theta, \partial E')$, where θ is the change in angle of the incoming neutron wave and ∂E is the energy transferred to the sample from the neutron, which leads to the interpretation of $I(\theta, \partial E')$ as the measured intensity of neutrons that change direction by an angle θ and transfer energy $\partial E'$ to the sample. If $\partial E' = 0$, the process is called elastic neutron scattering, and if $\partial E' \neq 0$ it is called inelastic neutron scattering. Broadly speaking, neutron scattering techniques that utilize elastic scattering are used to study the structure of materials, while inelastic techniques are used to study their dynamics.

Formally, the intensity measured in a neutron scattering experiment is directly related to the partial differential cross section,

$$\frac{\partial^2\sigma}{\partial\Omega\partial E'} = \frac{1}{\Phi}\frac{\begin{array}{c}\text{No. neutrons scattered per time into}\\ \text{solid angle } \partial\Omega \text{ with final energy between } E' \text{ and } E' + \partial E'\end{array}}{\partial\Omega\partial E'} \tag{6.1}$$

with Φ being the flux of the incoming neutron beam, i.e., the number of neutrons per unit area and second, and σ denoting the cross section. Let the incoming neutron have wave vector $\boldsymbol{k}$ and energy E, and wave vector $\boldsymbol{k'}$ and energy E' after the scattering event (Fig. 6.2a). With these definitions one can express Eq. (6.1) explicitly,

$$\frac{\partial^2\sigma}{\partial\Omega\partial E'} = \frac{|k'|}{|k|}\frac{1}{2\pi\hbar}\sum_{jj'} b_j b_{j'} \int_{-\infty}^{\infty} \left\langle e^{-i\boldsymbol{q}\cdot\boldsymbol{R}_{j'}(0)} e^{i\boldsymbol{q}\cdot\boldsymbol{R}_j(t)} \right\rangle e^{-i\omega t} dt \tag{6.2}$$

where the double sum runs over all pairs of scattering nuclei with position $\boldsymbol{R}_j$ and scattering length b_j. The scattering length is different for different isotopes and is in general a complex number, with the imaginary part representing neutron absorption. For most nuclei the imaginary part is small, and hence b has been taken as real in Eq. (6.2).

6.2 Coherent and incoherent neutron scattering

We split the sum over j, j' in Eq. (6.2) into two terms, corresponding to $j \neq j'$ and $j = j'$,

$$\begin{aligned}\frac{\partial^2\sigma}{\partial\Omega\partial E'} = \frac{|k'|}{|k|}\frac{1}{2\pi\hbar}\Bigg(&\sum_{j\neq j'} b_j b_j' \int_{-\infty}^{\infty} \left\langle e^{-i\boldsymbol{q}\cdot\boldsymbol{R}_{j'}(0)} e^{i\boldsymbol{q}\cdot\boldsymbol{R}_j(t)}\right\rangle e^{-i\omega t} dt \\ &+ \sum_j b_j^2 \int_{-\infty}^{\infty} \left\langle e^{-i\boldsymbol{q}\cdot\boldsymbol{R}_j(0)} e^{i\boldsymbol{q}\cdot\boldsymbol{R}_j(t)}\right\rangle e^{-i\omega t} dt\Bigg).\end{aligned} \tag{6.3}$$

In principle, each nucleus in the sample can have a different scattering length b_j that occurs with abundance f_j. However, a macroscopic sample contains a large number of nuclei, and thus we can replace the factors b_j^2 and $b_j b_{j'}$ in Eq. (6.3) with their averages, $\overline{b_j^2}$ and $\overline{b_j b_{j'}}$. The average scattering lengths are

$$\overline{b_j^2} = \sum_j f_j b_j^2 \equiv \overline{b^2} \quad \text{and} \quad \overline{b_j b_{j'}} = \sum_j f_j b_j \sum_{j'} f_{j'} b_{j'} \equiv \overline{b}^2. \tag{6.4}$$

By adding and subtracting the missing term for $j = j'$ in the first term in Eq. (6.3) we arrive at

$$\begin{aligned}\frac{\partial^2\sigma}{\partial\Omega\partial E'} =& \frac{|k'|}{|k|}\frac{1}{2\pi\hbar}\overline{b}^2 \sum_{jj'} \int_{-\infty}^{\infty} \left\langle e^{-i\boldsymbol{q}\cdot\boldsymbol{R}_{j'}(0)} e^{i\boldsymbol{q}\cdot\boldsymbol{R}_j(t)}\right\rangle e^{-i\omega t} dt \\ &+ \frac{|k'|}{|k|}\frac{1}{2\pi\hbar}\left(\overline{b^2} - \overline{b}^2\right) \sum_j \int_{-\infty}^{\infty} \left\langle e^{-i\boldsymbol{q}\cdot\boldsymbol{R}_j(0)} e^{i\boldsymbol{q}\cdot\boldsymbol{R}_j(t)}\right\rangle e^{-i\omega t} dt.\end{aligned} \tag{6.5}$$

By introducing $\sigma_{\text{coh}} = 4\pi\overline{b}^2$ and $\sigma_{\text{inc}} = 4\pi\left(\overline{b^2} - \overline{b}^2\right)$ we identify the two terms

$$\begin{aligned}\left(\frac{\partial^2\sigma}{\partial\Omega\partial E'}\right)_{\text{coh}} =& \frac{\sigma_{\text{coh}}}{4\pi}\frac{|k'|}{|k|}\frac{1}{2\pi\hbar} \sum_{jj'} \int_{-\infty}^{\infty} \left\langle e^{-i\boldsymbol{q}\cdot\boldsymbol{R}_{j'}(0)} e^{i\boldsymbol{q}\cdot\boldsymbol{R}_j(t)}\right\rangle e^{-i\omega t} dt \\ \left(\frac{\partial^2\sigma}{\partial\Omega\partial E'}\right)_{\text{inc}} =& \frac{\sigma_{\text{inc}}}{4\pi}\frac{|k'|}{|k|}\frac{1}{2\pi\hbar} \sum_{j} \int_{-\infty}^{\infty} \left\langle e^{-i\boldsymbol{q}\cdot\boldsymbol{R}_j(0)} e^{i\boldsymbol{q}\cdot\boldsymbol{R}_j(t)}\right\rangle e^{-i\omega t} dt,\end{aligned} \tag{6.6}$$

as the coherent and the incoherent partial neutron cross sections, respectively. The sum in the incoherent cross section runs over each atom j, and, thus, the incoherent scattering contains the scattered intensity from the individual nuclei. The coherent cross section, on the other hand, describes the scattering contributions from all *pairs* of nuclei. Since the scattered neutrons behave like waves, the coherent scattering can exhibit interference effects, which manifest themselves as the peaks seen in a typical neutron diffraction experiment.

The coherent and incoherent scattering lengths vary dramatically over the periodic table (Fig. 6.2b), which is a key advantage of neutron scattering for studying chromophores. Unlike X-ray scattering, where cross sections scale with the number of electrons, leaving light elements like hydrogen poorly resolved, the incoherent neutron scattering length for hydrogen is relatively large, yielding a strong signal. However, this strong signal can be problematic, manifesting as an incoherent background signal in diffraction experiments. In these cases, the scattering from hydrogen can be reduced by replacing it with the isotope deuterium, which has a smaller incoherent scattering length. This scheme is exploited in contrast matching, where hydrogen is selectively substituted with deuterium, allowing isolated study of parts of the system [210].

6.3 Simulating neutron scattering via correlation functions

We can further manipulate the coherent and incoherent partial cross sections in Eq. (6.6) into a form that makes them easily relatable to observables in MD simulations. Starting from Eq. (6.6), we can move the sums inside the integral sign and define them as the intermediate scattering function $F(\boldsymbol{q}, t)$ for the coherent cross section, and the self intermediate scattering function $F_s(\boldsymbol{q}, t)$ for the incoherent cross section,

$$F(\boldsymbol{q},t) = \frac{1}{N}\sum_{jj'}\left\langle e^{-i\boldsymbol{q}\cdot\boldsymbol{R}_{j'}(0)} e^{i\boldsymbol{q}\cdot\boldsymbol{R}_j(t)}\right\rangle \quad \text{and} \quad F_s(\boldsymbol{q},t) = \frac{1}{N}\sum_{j}\left\langle e^{-i\boldsymbol{q}\cdot\boldsymbol{R}_{j}(0)} e^{i\boldsymbol{q}\cdot\boldsymbol{R}_j(t)}\right\rangle, \tag{6.7}$$

with N being the number of nuclei in the system.

From $F(\boldsymbol{q}, t)$, we can further define the time-dependent pair-correlation function $G(\boldsymbol{r}, t)$ and the dynamic structure factor $S(\boldsymbol{q}, \omega)$,

$$G(\boldsymbol{r},t) = \frac{1}{(2\pi)^3}\int F(\boldsymbol{q},t)e^{-i\boldsymbol{q}\cdot\boldsymbol{r}}d\boldsymbol{q} \quad \text{and} \quad S(\boldsymbol{q},\omega) = \frac{1}{2\pi\hbar}\int F(\boldsymbol{q},t)e^{-i\omega t}dt. \tag{6.8}$$

Similarly, from $F_s(\boldsymbol{q}, t)$ we get the self time-dependent pair-correlation function $G_s(\boldsymbol{r}, t)$ and the incoherent dynamic structure factor $S_i(\boldsymbol{q}, \omega)$. We can now rewrite the coherent and incoherent cross sections in terms of $S(\boldsymbol{q}, \omega)$ and $S_i(\boldsymbol{q}, \omega)$,

$$\begin{aligned}\left(\frac{\partial^2\sigma}{\partial\Omega\partial E'}\right)_{\text{coh}} &= \frac{\sigma_{\text{coh}}}{4\pi}\frac{|k'|}{|k|}NS(\boldsymbol{q},\omega)\\ \left(\frac{\partial^2\sigma}{\partial\Omega\partial E'}\right)_{\text{inc}} &= \frac{\sigma_{\text{inc}}}{4\pi}\frac{|k'|}{|k|}NS_i(\boldsymbol{q},\omega).\end{aligned} \tag{6.9}$$

Equation (6.9) is quite remarkable. If we can calculate $S(\boldsymbol{q},\omega)$ and $S_i(\boldsymbol{q},\omega)$ from an MD simulation, then we can estimate the partial differential cross section, and by extension the intensity that one would measure in a neutron scattering experiment. The remaining question is, how do we calculate $S(\boldsymbol{q},\omega)$? Let $\rho(\boldsymbol{r},t)$ be the time-dependent particle density,

$$\rho(\boldsymbol{r},t) = \sum_j \delta(\boldsymbol{r} - \boldsymbol{R}_j(t)), \tag{6.10}$$

which we can Fourier transform in space to obtain $F(\boldsymbol{q},t)$ (Eq. (6.7)),

$$\begin{aligned} \rho(\boldsymbol{q},t) &= \sum_j e^{-i\boldsymbol{q}\cdot\boldsymbol{R}_j(t)} \\ \Rightarrow F(\boldsymbol{q},t) &= \frac{1}{N}\left\langle \rho(\boldsymbol{q},0)\rho(-\boldsymbol{q},t)\right\rangle \end{aligned} \tag{6.11}$$

We can get the self intermediate scattering function, $F_s(\boldsymbol{q},t)$, by only considering the terms in which $j = j'$ in Eq. (6.11). From Eq. (6.8) we know that we can obtain $S(\boldsymbol{q},\omega)$ and $S_i(\boldsymbol{q},\omega)$ by Fourier transforming $F(\boldsymbol{q},t)$ and $F_s(\boldsymbol{q},t)$ respectively in time. In a MD simulation we know the positions of all atoms, and, thus, we can record $\rho(\boldsymbol{r},t)$ throughout the simulation. We can then after the fact compute the double Fourier transform in time and space over this trajectory, and via Eq. (6.11) and Eq. (6.8) compute $S(\boldsymbol{q},\omega)$ and $S_i(\boldsymbol{q},\omega)$, and by extension the simulated intensity for a neutron scattering experiment.

6.4 Connecting simulated dynamic structure factors to experimentally measured spectra

In order to compare the dynamic structure factor from MD with experiments, we need to weigh it with the resolution function and kinematic constraint of the specific instrument. Neutron scattering instruments have their unique combination of sample position, detector angles, and incoming and outgoing energy ranges, and thus probes different parts of $(\boldsymbol{q},\omega)$–space with varying resolution and weight. Different neutron scattering techniques exploit this to target specific regions of $(\boldsymbol{q},\omega)$–space (Fig. 6.3). The two most relevant here are inelastic neutron scattering (INS) and quasi-elastic neutron scattering (QENS), which focus on vibrational dynamics and slower dynamics such as diffusion, respectively. Temperature also plays a role; INS experiments in particular are typically conducted at cryogenic temperatures, since the coherent scattered intensity is attenuated by thermal motion via the Debye-Waller factor [208, Chapter 3],

$$\left(\frac{\partial^2\sigma}{\partial\Omega\partial E'}\right)_{\text{coh}} \propto \exp\left(-q^2\left\langle u^2\right\rangle/3\right), \tag{6.12}$$

where $\langle u^2\rangle$ is the mean-squared thermal displacement.

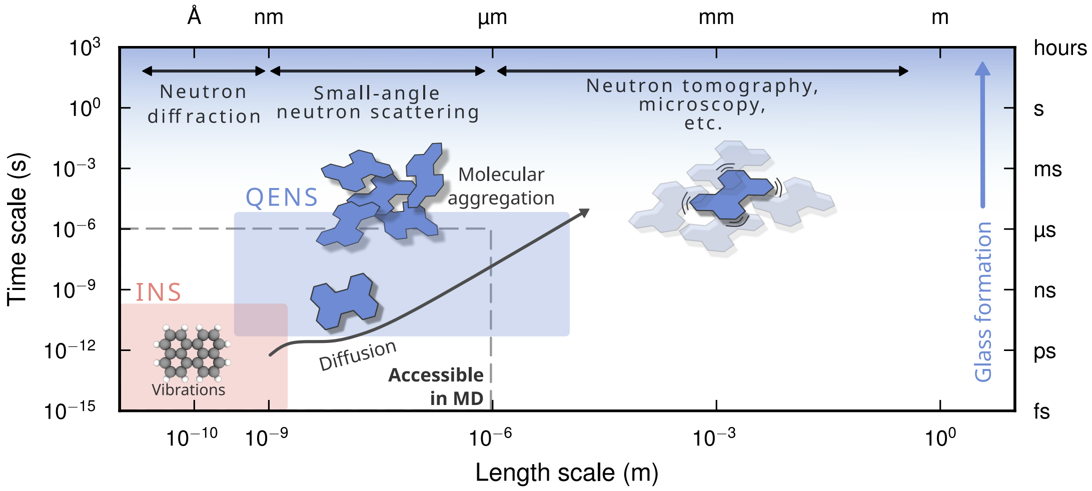

Figure 6.3: A schematic diagram of a few different neutron scattering techniques and the spatial and temporal ranges they probe. Neutron diffraction, neutron small angle scattering, and neutron tomography primarily probe material structure from a few Ångström up to macroscopic scales. Energy-resolved techniques such as inelastic neutron scattering (INS) and quasi-elastic neutron scattering (QENS) probe intermolecular vibrations and longer-time processes such as molecular diffusion and organization, respectively.

Unlike experiments, the $(\boldsymbol{q}, \omega)$–ranges sampled in a MD simulation are limited only by the time step and the size of the simulation cell, with all regions equally weighted. We can account for the instrumental resolution function and kinematic constraint by weighting the simulated dynamic structure factor $S_{\text{simulation}}(\boldsymbol{q}, \omega)$ by a combined resolution function and kinematic constraint function that we denote $f(\boldsymbol{q}, \omega)$,

$$S_{\text{experiment}}(\boldsymbol{q}, \omega) = f(\boldsymbol{q}, \omega) \odot S_{\text{simulation}}(\boldsymbol{q}, \omega), \qquad (6.13)$$

where $\odot$ denotes element-wise multiplication. $f(\boldsymbol{q}, \omega)$ can be obtained from several sources: from standard neutron analysis software packages such as MANTID [211] or FULLPROF [212, 213], by simulating the entire instrument using Monte-Carlo ray-tracing with packages such as MCSTAS [214–216] or VITESSE [217, 218], or by asking the instrument scientist nicely. For simulating INS weighted with $f(\boldsymbol{q}, \omega)$ specifically, dedicated tools include OCLIMAX [219, 220] for MD, as well as EUPHONIC [221] and the ABINS [222] routine in MANTID for phonon-based INS simulations.

In paper VIII we present a fully Python-based workflow for simulating INS, combining machine-learned interatomic potentials (MLIPs) based on the NEP framework with MD in GPUMD, correlation functions computed with DYNASOR [223, 224], and kinematic constraint and resolution function weighting via RESINS. NEP enables larger system sizes and longer simulations than other MLIP-based INS simulations [225]. Beyond

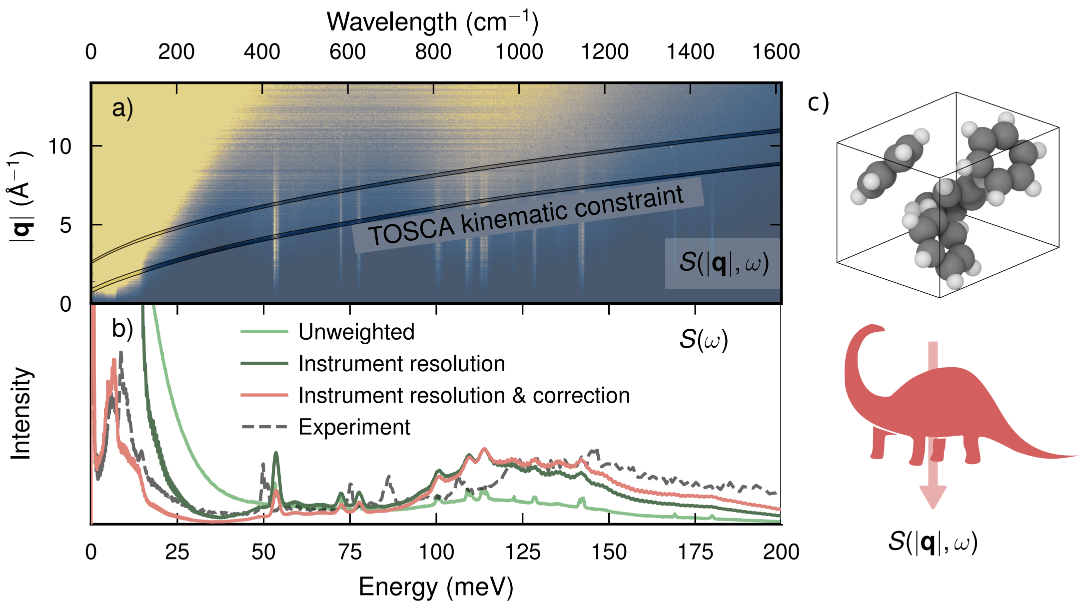

Figure 6.4: (a) Simulated $S(|\boldsymbol{q}|, \omega)$ for crystalline benzene at 127 K from MD overlaid with the TOSCA kinematic constraint (ISIS Neutron and Muon Source, UK). (b) Simulated INS spectrum $S(\omega)$ for crystalline benzene with instrument resolution including the kinematic constraint in (a), and a quantum correction for the classical statistics in MD. The red-shift of the simulated spectra compared to experiments is attributed to the vdW-DF-cx functional [124] and the neuroevolution potential (NEP) model not fully capturing intermolecular interactions. All data is from paper VIII. (c) Phase-*I* crystalline benzene, and how $S(|\boldsymbol{q}|, \omega)$ is obtained using DYNASOR.

neutron scattering lengths, DYNASOR also supports X-ray form factors and electronic scattering factors for comparison with other scattering techniques.

We apply the workflow to elemental Si, crystalline benzene C_6H_6, and hydrogenated Sc-doped $BaTiO_3$ in paper VIII to demonstrate its versatility, where benzene can be viewed as a prototype system for the chromophores studied in this thesis. The experimental INS spectrum $S(\omega) = \int S(q, \omega) dq$ for benzene was collected at the TOSCA neutron spectrometer [226, 227] at the ISIS Neutron and Muon Source, UK, and applying the instrument resolution and kinematic constraint function $f(q, \omega)$ to the simulated $S(q, \omega)$ yields a marked improvement in agreement with experiment (Fig. 6.4). Agreement improves further by correcting for classical phonon statistics in MD using a first-order correction from Stokes-Raman scattering[1] [228],

$$S(q, \omega)_{\text{corrected}} = S(q, \omega) \frac{\beta \hbar \omega}{1 - \exp(-\beta \hbar \omega)}. \tag{6.14}$$

---

[1] This correction was also applied for Raman scattering in paper VI.

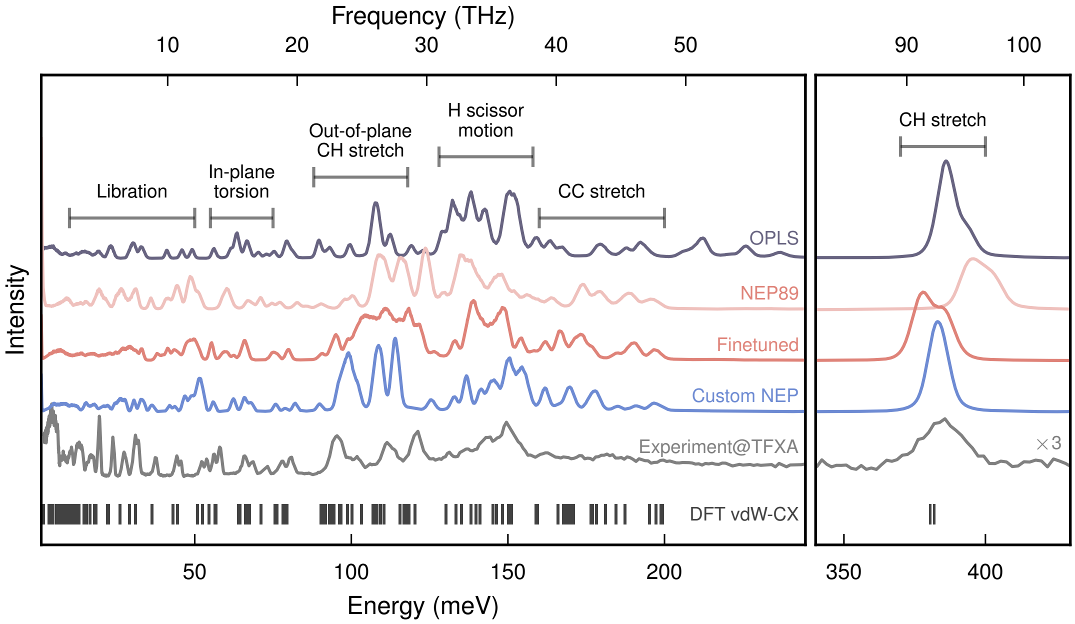

Figure 6.5: Simulated INS spectra using the TOSCA resolution function compared with experimental data from TFXA. Note that no direct quantitative comparison is possible since the instruments differ. The experimental C–H stretch peak is scaled by a factor of 3 for visibility. The models shown are the same as in Sect. 4.4: a custom NEP, a fine-tuned NEP89 model for perylene, the foundation model NEP89, and the classical force field (FF) OPLS. The training data for the custom and fine-tuned NEP models used the vdW-DF-cx functional [124], and harmonic DFT modes using this functional are shown for reference.

The corrected spectrum still exhibits a red shift of approximately 20 meV relative to the experimental spectrum, which is due to the vdW-DF-cx functional [124] used to generate the density functional theory (DFT) training data, as well as difficulties for the NEP model to accurately capture the weak intermolecular interactions.

We can also apply the workflow to crystalline perylene at 20 K, comparing against literature INS data [229] measured at the TFXA spectrometer [230] at ISIS (Fig. 6.5). The resolution function and kinematic constraint of TFXA are currently unavailable in RESINS, and hence the same settings as for TOSCA in paper VIII were used as a proxy. This is a reasonable approximation since TFXA is the predecessor to TOSCA and has a very similar resolution function, but note that this means that no direct quantitative comparison between experiment and simulation can be performed in this case. We compare the same four models we considered in Sect. 4.4: the foundation model NEP89, a fine-tuned NEP89 model for perylene, a custom NEP model, and the classical FF OPLS. All NEP models reproduce the experimental spectrum qualitatively, with fine-tuning improving over NEP89 most notably for the high frequency C–H stretch band. The custom NEP outperforms the fine-tuned model, however, which shows broader features particularly for the out-of-plane C–H stretch.

The classical FF OPLS yields worse agreement with experiment than the NEP models, and overestimates the C–C bond stretch region with features that are absent in both experiments and the other models. Note that the OPLS simulation was performed with GROMACS [231] in the NVT ensemble, and required using a thermostat with a rather strong coupling (time constant of 100 fs), rather than the NVE ensemble as for the NEP models. GROMACS is designed for room temperature simulations of biological systems, and for technical reasons, energy conservation in the NVE ensemble is violated. NVT affects the dynamics in the simulation, and it is thus plausible that some of the features in the INS spectrum for OPLS originate from the thermostat.

No model captures the librational modes around 30 meV to 50 meV well, consistent with the difficulty all models showed in reproducing the intermolecular interactions in Sect. 4.4. Disregarding this low-energy region, the good overall agreement highlights the predictive power of the workflow. Simulated INS spectra for a specific instrument could thus be obtained before or during an experiment. The examples here took several hours to compute, but a less converged spectrum could be produced in minutes by running a shorter MD simulation and sampling fewer $\boldsymbol{q}$-points. This prospect is particularly interesting for NEP89, which requires no system-specific training and could therefore be deployed in an automated pipeline at neutron facilities to help instrument scientists and users select optimal settings and interpret data before and during beamtime.

7

Conclusions and outlook

You could look it up, but it's way more fun to speculate!

Esmée, 2024

This thesis set out to develop a simulation framework capable of modeling structure and dynamics in chromophores, and connect the results to experiments. The framework has taken shape in the form of three interlocking components: molecular dynamics (MD) simulations, machine-learned interatomic potentials (MLIPs) using the neuroevolution potential (NEP) framework, and correlation functions as a bridge to experiments. In this final chapter, we evaluate the workflow, discuss its current limitations, and outline the most promising directions for future work.

7.1 Evaluation of the framework

Three research questions were posed in the introduction of this thesis, against which we will evaluate the simulation framework.

> What is the accuracy and reliability of NEP models for capturing intra- and intermolecular interactions in perylene derivatives?

NEP models captures intramolecular interactions well for both perylene and retinal, as demonstrated by the reproduction of their infrared (IR) and inelastic neutron scattering (INS) spectra, and they generally perform well for broad materials classes as demonstrated in papers II, III, and V. Intermolecular interactions present a more significant challenge: NEP fails to fully reproduce the weak intermolecular forces in perylenes, a shortcoming evident both in the fine-tuning analysis in Sect. 4.4 and in the low-frequency features of the vibrational spectra. Since intermolecular interactions govern molecular aggregation and glass formation, and thereby the optoelectronic properties, industrial processability, and stability of chromophore-based devices, this represents the primary limitation of the workflow. Incorporating long-range interactions in NEP by separating out intermolecular forces in the loss function or explicit architecture extensions is an active area of development, and remains an open challenge for MLIPs more broadly [25, 232–234]. Improving this aspect of the NEP framework is therefore the most consequential next step for the workflow as a whole.

> How effectively does the simulation protocol capture the glass-forming behavior of perylene derivatives?

Throughout the thesis, we have focused on glass forming behavior through the lens of predicting key quantities, specifically the glass transition temperature T_g and the kinetic fragility m. The Bayesian extrapolation scheme introduced in paper I substantially improves the glass transition temperature estimates for perylene derivatives over standard simulated annealing. Simulated annealing dramatically overestimates T_g since the molecular structure appears frozen on MD timescales, even though individual molecules continue to reorient at temperatures far below the estimated transition. By constructing correlation functions that isolate this molecular reorientation, the Bayesian scheme bypasses the timescale limitation, yielding estimates of T_g and the kinetic fragility that better agree with experiments.

Quantitative agreement with experiment remains partial, however: both T_g and the kinetic fragility are overestimated, which can be explained by two compounding limitations. First, the classical OPLS force field (FF) used in paper I fails to capture intermolecular interactions in perylenes, performing worse than NEP as demonstrated in Sect. 4.4. Second, the chosen normal-vector autocorrelation function (NACF) might not isolate the true α-relaxation process, introducing a systematic bias in the extrapolation independently of force field quality. Assuming the NACF is an appropriate choice, improving the description of intermolecular interactions in the NEP framework is therefore expected to improve predictions of T_g and kinetic fragility. Overall, the Bayesian protocol enables semi-quantitative prediction of glass-forming behavior in perylene derivatives.

> To what degree can the simulation framework predict experimental observables from neutron scattering and optical spectroscopy?

The workflow successfully predicts IR and INS spectra, as demonstrated in Chapter 5 and Chapter 6. Agreement with experiment is good across mid-to-high frequency features but breaks down at low frequencies, consistent with the inadequate description of intermolecular interactions in NEP. The same limitation affects the prediction of the electronic dielectric function, where extending the approach from bulk systems like SiO_2 to molecular systems remains challenging.

Incorporating other experimental techniques in the workflow is in principle straightforward, by utilizing appropriate correlation functions. Two examples of relevant techniques are quasi-elastic neutron scattering (QENS) and small-angle neutron scattering, capturing molecular diffusion and aggregation, respectively. These techniques probe, however, time and length scales at the edge of what is accessible even with MLIP-based MD simulations, and computational cost is therefore a major consideration [235].

In summary, the workflow addresses all three research questions. Intramolecular interactions are well captured in perylene derivatives, but the weak intermolecular interactions remain challenging. The Bayesian extrapolation protocol enables prediction of glass-forming behavior in perylenes, with remaining discrepancies compared to experiments largely attributable to the shortcoming in describing intermolecular interactions. Correlation functions connect MD simulations with a range of spectroscopic and scattering techniques, with agreement limited predominantly by the same intermolecular interactions. Taken together, these results establish the workflow presented in this thesis as a viable framework for studying structure, dynamics, and experimental observables in chromophores, and the predictive power of this workflow is expected to improve substantially as the description of long-range interactions in MLIPs in general and NEP in particular matures.

7.2 Limitations

The NEP models that form the basis of the workflow are trained on density functional theory (DFT) data and inherit its limitations, making the choice of exchange–correlation (XC) functional consequential. The vdW-DF-cx functional used throughout this thesis targets a balance between computational efficiency and accuracy for both bulk and molecular systems, with explicit treatment of van der Waals dispersion [124]. This functional performs well across most cases encountered in this thesis, though minor discrepancies remain, such as in paper VIII[1] where a slight softening of the intramolecular vibrational frequencies relative to INS data can be observed. Errors such as these are directly inherited by NEP models through the training data. Higher-level methods such as coupled cluster theory [236–238] offer a route to more accurate reference data, at substantially greater computational cost. Fine-tuning addresses this tradeoff directly: a NEP model pretrained on a large DFT dataset can be refined on a small number of high-level calculations, concentrating the computational expense only where accuracy matters most. Fine-tuning is therefore a data-efficient approach for obtaining more accurate NEP models. Long-range intermolecular interactions would still be challenging due to their small magnitude and the locality of the model, but fine-tuning on high-level data could potentially improve the short-ranged part of the intermolecular force description of NEP.

---

[1] Figure S16 in the Supplementary Information of Paper VIII.

Another limitation is that the workflow operates entirely on the ground-state potential energy surface. In chromophores, excited-state interactions directly influence molecular structure and electronic properties, affecting, for example, the efficiency of perylene-based organic solar cells (OSCs) [239]. Excited states can be accessed via time-dependent DFT [121], where MLIPs for excited-state potential energy surfaces have been demonstrated [193, 240, 241], including for the NEP framework [242]. Extending the present workflow in this direction is a natural next step.

A final limitation is the restricted time and length scales accessible to MD, discussed in Sect. 3.5. This constrains direct simulation of molecular aggregation and glass formation, necessitating indirect approaches such as the Bayesian extrapolation scheme. A particular example of this that is not covered in this thesis is the liquid–liquid phase transition in multi-component perylene mixtures during cooling, which gives rise to ultralow fragility and strong glass forming behavior [76]. This phenomenon remains out of reach for standard MD, despite significant computational effort. Extending the framework to, for example, coarse-grained MD could potentially enable access to the liquid-liquid transition, but care must in that case be taken to construct a coarse-grained model that is transferable over a broad range of temperatures.

7.3 Future work

In addition to addressing the limitations mentioned so far, there are several natural directions for extending the workflow presented in this thesis. Among these, uncertainty quantification for MLIPs stands out as particularly important. Current approaches typically quantify uncertainty in model predictions of energies and forces, but do not propagate these uncertainties to the derived properties obtained from MD simulations. In particular, obtaining formal error estimates stemming from model uncertainty in properties from correlation functions, such as experimental spectra, without resorting to ensemble averages would be highly useful in predicting experiments, and is to the best of my knowledge still an open challenge.

Propagating uncertainty to experimental observables would also facilitate direct validation of MLIPs against experiment, a need that has already been recognized in the community [25], complementing the current standard of benchmarking against DFT or quantum chemical calculations. This could be achieved by creating, for example, reference databases of INS spectra spanning a wide range of organic and inorganic compounds. Such a database could be developed collaboratively between neutron scattering facilities and computational materials modeling groups, starting from well-characterized compounds with existing measurements. Simulation workflows capable of reproducing instrument-specific signatures would then be key for meaningful comparison between simulated and measured spectra.

Beyond validation, experimental observables could also be used to actively guide MD simulations. By computing the error between a predicted observable, such as an absorption spectrum or pair distribution function, and an experimental reference, corrective forces can be back-propagated to the atomic coordinates. This would enable identification of structural configurations responsible for specific experimental signatures, which can be difficult to determine from experiments alone. The idea has been explored in the context of X-ray diffraction [243], though not, to my knowledge, with MLIPs. A particularly compelling application in the context of this thesis would be driving simulations toward the molecular aggregates formed during the liquid-liquid phase transition in perylene mixtures [76], which would give direct structural insight into how mixing controls glass-forming behavior.

A key theme throughout this thesis has been the computational efficiency of NEP, enabling access to accurate large-scale MD simulations. The time and length-scales accessible to computer simulations are set to improve in the coming years, through increasing computational performance of supercomputing hardware and graphics processing units (GPUs). Beyond raw performance, however, the more consequential implication may be one of accessibility. Large-scale simulations of hundreds of thousands of atoms with NEP can be performed on consumer-grade GPUs in standard desktop computers. Supercomputing infrastructure is expensive and unevenly distributed; as of November 2025, the vast majority of the world's top 500 supercomputers are located in developed countries [244]. Simulation frameworks that run efficiently on modest hardware therefore have the potential to broaden participation in computational materials research beyond groups with access to large computer time allocations at supercomputers.

Fitting with the theme of accessibility is the role of free and open-source software in computational research. The workflow presented in this thesis is built on open-source foundations: GPUMD, CALORINE, and DYNASOR are all freely available, and depend on community infrastructure including NUMPY [245], SCIPY [246], PANDAS [247], and the Atomic Simulation Environment (ASE) [178], which provides a consistent interface across much of the computational materials modeling community. Open-source software is similarly foundational to the modern Internet, yet much of it is sustained by voluntary effort from a small number of contributors.[2] If any of these tools have served you too, contributing back through code, documentation, or bug reports is one of the most direct ways to strengthen the research infrastructure the community relies on. I firmly believe that this kind of collaboration is the way forward, towards addressing the remaining challenges outlined in this thesis, and beyond. It might just be my perspective through my rainbow-tinted glasses, but I see a complete understanding of chromophores to be well within sight.

---

[2] `https://xkcd.com/2347/`

Acknowledgments

Finally, here comes the part when I get to thank all of you who have made this thesis possible. First and foremost, to my supervisor: thank you, Paul. All the way back in 2021 when we were all at home "enjoying" social distancing, you took the time to have several Zoom calls with an uncertain master thesis student who didn't know if a PhD was right for him. Thank you for your patience and convincing me to finally take the plunge[3], and for your support throughout these past five years, both academic and through affirming me in my opinion that music truly was better before the 2000s.

To my co-supervisors, Jan Swenson, Christian Müller, Sanghamitra Mukhopadhyay, and Thomas Holm-Rod, as well as my scientific collaborator Adam Jackson, thank you for your guidance on everything that moves slowly, is fragile, and scatters neutrons. To my examiner, Henrik Grönbeck, thank you for your support and prompt feedback, and for crediting me courses in Ladok without asking too many questions.

To all the Neutron Nerds in SwedNESS-II: thank you for all the fun times we've had at conferences and during hands-on neutron training. I couldn't think of a better gang for coming up with weird things to put in the tomography beamline in Delft.

I'd also like to thank the SwedNESS graduate school for funding me, and the supercomputing centers at Kungliga Tekniska Högskolan (PDC), Linköpings Tekniska Universitet (NSC), Uppsala Universitet (UPPMAX), and Chalmers (C3SE) for providing the computational resources I need for my research, and that I love to tinker with.

A special thank you also goes out to each and every one of my colleagues at the Condensed Matter and Materials Theory division at Chalmers. To my current and my ex-office mate, Felix Uddén and Erik Fransson, thank you for making each day at the office a blast and a learning experience, regardless of whether it's about correlation functions or how to achieve optimal fluffiness when baking sourdough bread.

Till sist, till mina vänner och min familj, tack för all er kärlek och ert stöd. Ett särskilt stort tack till min fästmö Amanda, för att du står ut med mig. Du hade rätt, det här var bara att göra.

Ha det gôtt, vi ses!
Eric Lindgren, maj 2026

---

[3] `https://youtu.be/cU2AvkKA4kM?si=hWDplvMox1oDcnl3&t=118`; "Ten Meter Tower".

Bibliography

[1] J. Clayden, N. Greeves, and S. G. Warren, *Organic Chemistry* (Oxford University Press, 2012). ISBN 978-0-19-166621-6.

[2] A. Laiho, B. M. Smarsly, C. F. J. Faul, and O. Ikkala, *Macroscopically Aligned Ionic Self-Assembled Perylene-Surfactant Complexes within a Polymer Matrix,* Advanced Functional Materials **18**, 1890 (2008). doi:10.1002/adfm.200701496.

[3] X. Li, L. E. Sinks, B. Rybtchinski, and M. R. Wasielewski, *Ultrafast Aggregate-to-Aggregate Energy Transfer within Self-assembled Light-Harvesting Columns of Zinc Phthalocyanine Tetrakis(Perylenediimide),* Journal of the American Chemical Society **126**, 10810 (2004). doi:10.1021/ja047176b.

[4] S. Herbst, B. Soberats, P. Leowanawat, M. Lehmann, and F. Würthner, *A Columnar Liquid-Crystal Phase Formed by Hydrogen-Bonded Perylene Bisimide J-Aggregates,* Angewandte Chemie International Edition **56**, 2162 (2017). doi:10.1002/anie.201612047.

[5] F. J. M. Hoeben, P. Jonkheijm, E. W. Meijer, and A. P. H. J. Schenning, *About Supramolecular Assemblies of π-Conjugated Systems,* Chemical Reviews **105**, 1491 (2005). doi:10.1021/cr030070z.

[6] J. L. Silverberg, M. Bierbaum, J. P. Sethna, and I. Cohen, *Collective Motion of Humans in Mosh and Circle Pits at Heavy Metal Concerts,* Physical Review Letters **110**, 228701 (2013). doi:10.1103/PhysRevLett.110.228701.

[7] S. Wang, T.-P. Ruoko, G. Wang, S. Riera-Galindo, S. Hultmark, Y. Puttisong, F. Moro, H. Yan, W. M. Chen, M. Berggren, C. Müller, and S. Fabiano, *Sequential Doping of Ladder-Type Conjugated Polymers for Thermally Stable n-Type Organic Conductors,* ACS Applied Materials & Interfaces **12**, 53003 (2020). doi:10.1021/acsami.0c16254.

[8] M. E. Gemayel, K. Börjesson, M. Herder, D. T. Duong, J. A. Hutchison, C. Ruzié, G. Schweicher, A. Salleo, Y. Geerts, S. Hecht, E. Orgiu, and P. Samorì, *Optically Switchable Transistors by Simple Incorporation of Photochromic Systems into Small-Molecule Semiconducting Matrices,* Nature Communications **6**, 6330 (2015). doi:10.1038/ncomms7330.

[9] M. Gsänger, D. Bialas, L. Huang, M. Stolte, and F. Würthner, *Organic Semiconductors Based on Dyes and Color Pigments,* Advanced Materials **28**, 3615 (2016). doi:10.1002/adma.201505440.

[10] L. Yu, D. Qian, S. Marina, F. A. A. Nugroho, A. Sharma, S. Hultmark, A. I. Hofmann, R. Kroon, J. Benduhn, D.-M. Smilgies, K. Vandewal, M. R. Andersson, C. Langhammer, J. Martín, F. Gao, and C. Müller, *Diffusion-Limited Crystallization: A Rationale for the Thermal Stability of Non-Fullerene Solar Cells,* ACS Applied Materials & Interfaces **11**, 21766 (2019). doi:10.1021/acsami.9b04554.

[11] A. Diacon, O. Krupka, and P. Hudhomme, *Fullerene-Perylenediimide (C60-PDI) Based Systems: An Overview and Synthesis of a Versatile Platform for Their Anchor Engineering,* Molecules **27**, 6522 (2022). doi:10.3390/molecules27196522.

[12] J. Zhang, Y. Li, J. Huang, H. Hu, G. Zhang, T. Ma, P. C. Y. Chow, H. Ade, D. Pan, and H. Yan, *Ring-Fusion of Perylene Diimide Acceptor Enabling Efficient Nonfullerene Organic Solar Cells with a Small Voltage Loss,* Journal of the American Chemical Society **139**, 16092 (2017). doi:10.1021/jacs.7b09998.

[13] N. Kobayashi, T. Kasahara, T. Edura, J. Oshima, R. Ishimatsu, M. Tsuwaki, T. Imato, S. Shoji, and J. Mizuno, *Microfluidic White Organic Light-Emitting Diode Based on Integrated Patterns of Greenish-Blue and Yellow Solvent-Free Liquid Emitters,* Scientific Reports **5**, 14822 (2015). doi:10.1038/srep14822.

[14] L. Yao, S. Zhang, R. Wang, W. Li, F. Shen, B. Yang, and Y. Ma, *Highly Efficient Near-Infrared Organic Light-Emitting Diode Based on a Butterfly-Shaped Donor–Acceptor Chromophore with Strong Solid-State Fluorescence and a Large Proportion of Radiative Excitons,* Angewandte Chemie **126**, 2151 (2014). doi:10.1002/ange.201308486.

[15] V. Venunath Patil, K. Hyung Lee, and J. Yeob Lee, *Isomeric Fused Benzocarbazole as a Chromophore for Blue Fluorescent Organic Light-Emitting Diodes,* Journal of Materials Chemistry C **8**, 8320 (2020). doi:10.1039/D0TC01268H.

[16] K. Börjesson, D. Dzebo, B. Albinsson, and K. Moth-Poulsen, *Photon Upconversion Facilitated Molecular Solar Energy Storage,* Journal of Materials Chemistry A **1**, 8521 (2013). doi:10.1039/C3TA12002C.

[17] K. Stranius and K. Börjesson, *Determining the Photoisomerization Quantum Yield of Photoswitchable Molecules in Solution and in the Solid State,* Scientific Reports **7**, 41145 (2017). doi:10.1038/srep41145.

[18] M. Jevric, A. U. Petersen, M. Mansø, S. Kumar Singh, Z. Wang, A. Dreos, C. Sumby, M. B. Nielsen, K. Börjesson, P. Erhart, and K. Moth-Poulsen, *Norbornadiene-Based Photoswitches with Exceptional Combination of Solar Spectrum Match and Long-Term Energy Storage,* Chemistry – A European Journal **24**, 12767 (2018). doi:10.1002/chem.201802932.

[19] V. Gray, A. Dreos, P. Erhart, B. Albinsson, K. Moth-Poulsen, and M. Abrahamsson, *Loss Channels in Triplet–Triplet Annihilation Photon Upconversion: Importance of Annihilator Singlet and Triplet Surface Shapes,* Physical Chemistry Chemical Physics **19**, 10931 (2017). doi:10.1039/C7CP01368J.

[20] M. Quant, A. Hamrin, A. Lennartson, P. Erhart, and K. Moth-Poulsen, *Solvent Effects on the Absorption Profile, Kinetic Stability, and Photoisomerization Process of the Norbornadiene–Quadricyclanes System,* The Journal of Physical Chemistry C **123**, 7081 (2019). doi:10.1021/acs.jpcc.9b02111.

[21] G. Charalambidis, E. Georgilis, M. K. Panda, C. E. Anson, A. K. Powell, S. Doyle, D. Moss, T. Jochum, P. N. Horton, S. J. Coles, M. Linares, D. Beljonne, J.-V. Naubron, J. Conradt, H. Kalt, A. Mitraki, A. G. Coutsolelos, and T. S. Balaban, *A Switchable Self-Assembling and Disassembling Chiral System Based on a Porphyrin-Substituted Phenylalanine–Phenylalanine Motif,* Nature Communications **7**, 12657 (2016). doi:10.1038/ncomms12657.

[22] O. T. Unke, S. Chmiela, H. E. Sauceda, M. Gastegger, I. Poltavsky, K. T. Schütt, A. Tkatchenko, and K.-R. Müller, *Machine Learning Force Fields,* Chemical Reviews **121**, 10142 (2021). doi:10.1021/acs.chemrev.0c01111.

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

[24] J. Behler and G. Csányi, *Machine Learning Potentials for Extended Systems: A Perspective,* The European Physical Journal B **94**, 142 (2021). doi:10.1140/epjb/s10051-021-00156-1.

[25] D. M. Anstine and O. Isayev, *Machine Learning Interatomic Potentials and Long-Range Physics,* The Journal of Physical Chemistry A **127**, 2417 (2023). doi:10.1021/acs.jpca.2c06778.

[26] B. Mortazavi, X. Zhuang, T. Rabczuk, and A. V. Shapeev, *Atomistic Modeling of the Mechanical Properties: The Rise of Machine Learning Interatomic Potentials,* Materials Horizons **10**, 1956 (2023). doi:10.1039/D3MH00125C.

[27] Y.-W. Zhang, V. Sorkin, Z. H. Aitken, A. Politano, J. Behler, A. P Thompson, T. W. Ko, S. P. Ong, O. Chalykh, D. Korogod, E. Podryabinkin, A. Shapeev, J. Li, Y. Mishin, Z. Pei, X. Liu, J. Kim, Y. Park, S. Hwang, S. Han, K. Sheriff, Y. Cao, and R. Freitas, *Roadmap for the Development of Machine Learning-Based Interatomic Potentials,* Modelling and Simulation in Materials Science and Engineering **33**, 023301 (2025). doi:10.1088/1361-651X/ad9d63.

[28] B. Kalita, H. Gokcan, and O. Isayev, *Machine Learning Interatomic Potentials at the Centennial Crossroads of Quantum Mechanics,* Nature Computational Science **5**, 1120 (2025). doi:10.1038/s43588-025-00930-6.

[29] L. Berthier and D. R. Reichman, *Modern Computational Studies of the Glass Transition,* Nature Reviews Physics **5**, 102 (2023). doi:10.1038/s42254-022-00548-x.

[30] C. Müller, *On the Glass Transition of Polymer Semiconductors and Its Impact on Polymer Solar Cell Stability,* Chemistry of Materials **27**, 2740 (2015). doi:10.1021/acs.chemmater.5b00024.

[31] G. Zhang, F. R. Lin, F. Qi, T. Heumüller, A. Distler, H.-J. Egelhaaf, N. Li, P. C. Y. Chow, C. J. Brabec, A. K.-Y. Jen, and H.-L. Yip, *Renewed Prospects for Organic Photovoltaics,* Chemical Reviews **122**, 14180 (2022). doi:10.1021/acs.chemrev.1c00955.

[32] C. Liao, Y. Xiong, Y. Fu, X. Chen, and L. G. Occhipinti, *Organic Semiconductors Based Wearable Bioelectronics,* Wearable Electronics **2**, 23 (2025). doi:10.1016/j.wees.2024.12.003.

[33] T. A. Dela Peña, W. Su, Y. Li, Q. Fan, Y. Zhang, C. Yan, K. Feng, A. Uddin, Y. M. Lam, R. Ma, and G. Li, *Organic Semiconductors: A Versatile next-Generation Photovoltaics Technology,* Information & Functional Materials **2**, 155 (2025). doi:10.1002/ifm2.70000.

[34] Q. Yan, Y. Zhou, Y.-Q. Zheng, J. Pei, and D. Zhao, *Towards Rational Design of Organic Electron Acceptors for Photovoltaics: A Study Based on Perylenediimide Derivatives,* Chemical Science **4**, 4389 (2013). doi:10.1039/C3SC51841H.

[35] P.-O. Schwartz, L. Biniek, E. Zaborova, B. Heinrich, M. Brinkmann, N. Leclerc, and S. Méry, *Perylenediimide-Based Donor–Acceptor Dyads and Triads: Impact of Molecular Architecture on Self-Assembling Properties,* Journal of the American Chemical Society **136**, 5981 (2014). doi:10.1021/ja4129108.

[36] A. G. Macedo, L. P. Christopholi, A. E. X. Gavim, J. F. de Deus, M. A. M. Teridi, A. R. b. M. Yusoff, and W. J. da Silva, *Perylene Derivatives for Solar Cells and Energy Harvesting: A Review of Materials, Challenges and Advances,* Journal of Materials Science: Materials in Electronics **30**, 15803 (2019). doi:10.1007/s10854-019-02019-z.

[37] L. Zhu, M. Zhang, J. Xu, C. Li, J. Yan, G. Zhou, W. Zhong, T. Hao, J. Song, X. Xue, Z. Zhou, R. Zeng, H. Zhu, C.-C. Chen, R. C. I. MacKenzie, Y. Zou, J. Nelson, Y. Zhang, Y. Sun, and F. Liu, *Single-Junction Organic Solar Cells with over 19% Efficiency Enabled by a Refined Double-Fibril Network Morphology,* Nature Materials **21**, 656 (2022). doi:10.1038/s41563-022-01244-y.

[38] Y. Ding, W. A. Memon, S. Xiong, S. Gong, M. Li, Z. Deng, H. Liu, Y. Liu, X. Chen, N. Zheng, and F. He, *Molecular Design of Dimeric Acceptor Enables Binary Organic Solar Cells with 19.78% Efficiency and Enhanced Stability,* Advanced Materials **37**, 2501671 (2025). doi:10.1002/adma.202501671.

[39] G. Yu, J. Gao, J. C. Hummelen, F. Wudl, and A. J. Heeger, *Polymer Photovoltaic Cells: Enhanced Efficiencies via a Network of Internal Donor-Acceptor Heterojunctions,* Science **270**, 1789 (1995). doi:10.1126/science.270.5243.1789.

[40] J. J. M. Halls, C. A. Walsh, N. C. Greenham, E. A. Marseglia, R. H. Friend, S. C. Moratti, and A. B. Holmes, *Efficient Photodiodes from Interpenetrating Polymer Networks,* Nature **376**, 498 (1995). doi:10.1038/376498a0.

[41] J. Yi, G. Zhang, H. Yu, and H. Yan, *Advantages, Challenges and Molecular Design of Different Material Types Used in Organic Solar Cells,* Nature Reviews Materials **9**, 46 (2024). doi:10.1038/s41578-023-00618-1.

[42] J. Xin, Z. He, Z. Liu, X. Liu, H. Zhu, Z. Zhang, C. Song, X. Yin, Q. Liang, and J. Liu, *Recent Advances in Polymorphism of Organic Solar Cells,* Small **21**, 2409411 (2025). doi:10.1002/smll.202409411.

[43] Z. Chen, V. Stepanenko, V. Dehm, P. Prins, L. D. A. Siebbeles, J. Seibt, P. Marquetand, V. Engel, and F. Würthner, *Photoluminescence and Conductivity of Self-Assembled π–π Stacks of Perylene Bisimide Dyes,* Chemistry – A European Journal **13**, 436 (2007). doi:10.1002/chem.200600889.

[44] S. Chen, P. Slattum, C. Wang, and L. Zang, *Self-Assembly of Perylene Imide Molecules into 1D Nanostructures: Methods, Morphologies, and Applications,* Chemical Reviews **115**, 11967 (2015). doi:10.1021/acs.chemrev.5b00312.

[45] M. Zhang, Y. Bai, C. Sun, L. Xue, H. Wang, and Z.-G. Zhang, *Perylene-Diimide Derived Organic Photovoltaic Materials,* Science China Chemistry **65**, 462 (2022). doi:10.1007/s11426-021-1171-4.

[46] R. Singh, E. Giussani, M. M. Mróz, F. Di Fonzo, D. Fazzi, J. Cabanillas-González, L. Oldridge, N. Vaenas, A. G. Kontos, P. Falaras, A. C. Grimsdale, J. Jacob, K. Müllen, and P. E. Keivanidis, *On the Role of Aggregation Effects in the Performance of Perylene-Diimide Based Solar Cells,* Organic Electronics **15**, 1347 (2014). doi:10.1016/j.orgel.2014.03.044.

[47] G. Méhes, A. S. D. Sandanayaka, J.-C. Ribierre, and K. Goushi, *Physics and Design Principles of OLED Devices,* in *Handbook of Organic Light-Emitting Diodes* (Springer, Tokyo, 2020), p. 1. doi:10.1007/978-4-431-55761-6_49-1.

[48] J. M. Ha, S. H. Hur, A. Pathak, J.-E. Jeong, and H. Y. Woo, *Recent Advances in Organic Luminescent Materials with Narrowband Emission,* NPG Asia Materials **13**, 1 (2021). doi:10.1038/s41427-021-00318-8.

[49] M. Kasha, *Relation between Exciton Bands and Conduction Bands in Molecular Lamellar Systems,* Reviews of Modern Physics **31**, 162 (1959). doi:10.1103/RevModPhys.31.162.

[50] M. Kasha, *Energy Transfer Mechanisms and the Molecular Exciton Model for Molecular Aggregates,* Radiation Research **20**, 55 (1963). doi:10.2307/3571331.

[51] M. Kasha, H. R. Rawls, and M. Ashraf El-Bayoumi, *The Exciton Model in Molecular Spectroscopy,* Pure and Applied Chemistry **11**, 371 (1965). doi:10.1351/pac196511030371.

[52] N. J. Hestand and F. C. Spano, *Expanded Theory of H- and J-Molecular Aggregates: The Effects of Vibronic Coupling and Intermolecular Charge Transfer,* Chemical Reviews **118**, 7069 (2018). doi:10.1021/acs.chemrev.7b00581.

[53] X. Li, M. Linares, and P. Norman, *VeloxChem: GPU-Accelerated Fock Matrix Construction Enabling Complex Polarization Propagator Simulations of Circular Dichroism Spectra of G-Quadruplexes,* The Journal of Physical Chemistry A **129**, 633 (2025). doi:10.1021/acs.jpca.4c07510.

[54] S. Lehtola, C. Steigemann, M. J. T. Oliveira, and M. A. L. Marques, *Recent Developments in Libxc — A Comprehensive Library of Functionals for Density Functional Theory,* SoftwareX **7**, 1 (2018). doi:10.1016/j.softx.2017.11.002.

[55] J. Kussmann, H. Laqua, and C. Ochsenfeld, *Highly Efficient Resolution-of-Identity Density Functional Theory Calculations on Central and Graphics Processing Units,* Journal of Chemical Theory and Computation **17**, 1512 (2021). doi:10.1021/acs.jctc.0c01252.

[56] T. Yanai, D. P. Tew, and N. C. Handy, *A New Hybrid Exchange–Correlation Functional Using the Coulomb-attenuating Method (CAM-B3LYP),* Chemical Physics Letters **393**, 51 (2004). doi:10.1016/j.cplett.2004.06.011.

[57] S. Tretiak and S. Mukamel, *Density Matrix Analysis and Simulation of Electronic Excitations in Conjugated and Aggregated Molecules,* Chemical Reviews **102**, 3171 (2002). doi:10.1021/cr0101252.

[58] D. Beljonne, J. Cornil, R. Silbey, P. Millié, and J. L. Brédas, *Interchain Interactions in Conjugated Materials: The Exciton Model versus the Supermolecular Approach,* The Journal of Chemical Physics **112**, 4749 (2000). doi:10.1063/1.481031.

[59] J. Cornil, D. A. dos Santos, X. Crispin, R. Silbey, and J. L. Brédas, *Influence of Interchain Interactions on the Absorption and Luminescence of Conjugated Oligomers and Polymers: A Quantum-Chemical Characterization,* Journal of the American Chemical Society **120**, 1289 (1998). doi:10.1021/ja973761j.

[60] M. Hoffmann, K. Schmidt, T. Fritz, T. Hasche, V. M. Agranovich, and K. Leo, *The Lowest Energy Frenkel and Charge-Transfer Excitons in Quasi-One-Dimensional Structures: Application to MePTCDI and PTCDA Crystals,* Chemical Physics **258**, 73 (2000). doi:10.1016/S0301-0104(00)00157-9.

[61] G. Ritschel, D. Suess, S. Möbius, W. T. Strunz, and A. Eisfeld, *Non-Markovian Quantum State Diffusion for Temperature-Dependent Linear Spectra of Light Harvesting Aggregates,* The Journal of Chemical Physics **142**, 034115 (2015). doi:10.1063/1.4905327.

[62] C. R. Snyder and D. M. DeLongchamp, *Glassy Phases in Organic Semiconductors,* Current Opinion in Solid State and Materials Science **22**, 41 (2018). doi:10.1016/j.cossms.2018.03.001.

[63] M. D. Ediger, *Perspective: Highly Stable Vapor-Deposited Glasses,* The Journal of Chemical Physics **147**, 210901 (2017). doi:10.1063/1.5006265.

[64] J. Rault, *Origin of the Vogel–Fulcher–Tammann Law in Glass-Forming Materials: The α–β Bifurcation,* Journal of Non-Crystalline Solids **271**, 177 (2000). doi:10.1016/S0022-3093(00)00099-5.

[65] P. G. Debenedetti and F. H. Stillinger, *Supercooled Liquids and the Glass Transition,* Nature **410**, 259 (2001). doi:10.1038/35065704.

[66] P. Chakravarty, K. Pandya, and K. Nagapudi, *Determination of Fragility in Organic Small Molecular Glass Forming Liquids: Comparison of Calorimetric and Spectroscopic Data and Commentary on Pharmaceutical Importance,* Molecular Pharmaceutics **15**, 1248 (2018). doi:10.1021/acs.molpharmaceut.7b01068.

[67] C. A. Angell, *Relaxation in Liquids, Polymers and Plastic Crystals — Strong/Fragile Patterns and Problems,* Journal of Non-Crystalline Solids **131–133**, 13 (1991). doi:10.1016/0022-3093(91)90266-9.

[68] C. A. Angell, *Formation of Glasses from Liquids and Biopolymers,* Science **267**, 1924 (1995). doi:10.1126/science.267.5206.1924.

[69] R. Böhmer, K. L. Ngai, C. A. Angell, and D. J. Plazek, *Nonexponential Relaxations in Strong and Fragile Glass Formers,* The Journal of Chemical Physics **99**, 4201 (1993). doi:10.1063/1.466117.

[70] V. N. Novikov and A. P. Sokolov, *Poisson's Ratio and the Fragility of Glass-Forming Liquids,* Nature **431**, 961 (2004). doi:10.1038/nature02947.

[71] P. Lunkenheimer, F. Humann, A. Loidl, and K. Samwer, *Universal Correlations between the Fragility and Interparticle Repulsion of Glass-Forming Liquids,* The Journal of Chemical Physics **153**, 124507 (2020). doi:10.1063/5.0014457.

[72] A. Schoenhals, *Relation between Main and Normal Mode Relaxations for Polyisoprene Studied by Dielectric Spectroscopy,* Macromolecules **26**, 1309 (1993). doi:10.1021/ma00058a019.

[73] G. D. Smith and D. Bedrov, *Relationship between the α- and β-Relaxation Processes in Amorphous Polymers: Insight from Atomistic Molecular Dynamics Simulations of 1,4-Polybutadiene Melts and Blends,* Journal of Polymer Science Part B: Polymer Physics **45**, 627 (2007). doi:10.1002/polb.21064.

[74] J. C. Mauro, Y. Yue, A. J. Ellison, P. K. Gupta, and D. C. Allan, *Viscosity of Glass-Forming Liquids,* Proceedings of the National Academy of Sciences **106**, 19780 (2009). doi:10.1073/pnas.0911705106.

[75] L. M. C. Janssen, *Mode-Coupling Theory of the Glass Transition: A Primer,* Frontiers in Physics **6**, (2018). doi:10.3389/fphy.2018.00097.

[76] S. Hultmark, A. Cravcenco, K. Kushwaha, S. Mallick, P. Erhart, K. Börjesson, and C. Müller, *Vitrification of Octonary Perylene Mixtures with Ultralow Fragility,* Science Advances **7**, eabi4659 (2021). doi:10.1126/sciadv.abi4659.

[77] E. M. Bringa, K. Rosolankova, R. E. Rudd, B. A. Remington, J. S. Wark, M. Duchaineau, D. H. Kalantar, J. Hawreliak, and J. Belak, *Shock Deformation of Face-Centred-Cubic Metals on Subnanosecond Timescales,* Nature Materials **5**, 805 (2006). doi:10.1038/nmat1735.

[78] L. A. Zepeda-Ruiz, A. Stukowski, T. Oppelstrup, and V. V. Bulatov, *Probing the Limits of Metal Plasticity with Molecular Dynamics Simulations,* Nature **550**, 492 (2017). doi:10.1038/nature23472.

[79] J. Jung, W. Nishima, M. Daniels, G. Bascom, C. Kobayashi, A. Adedoyin, M. Wall, A. Lappala, D. Phillips, W. Fischer, C.-S. Tung, T. Schlick, Y. Sugita, and K. Y. Sanbonmatsu, *Scaling Molecular Dynamics beyond 100,000 Processor Cores for Large-Scale Biophysical Simulations,* Journal of Computational Chemistry **40**, 1919 (2019). doi:10.1002/jcc.25840.

[80] L. Freddolino, F. Liu, M. Gruebele, and K. Schulten, *Ten-Microsecond Molecular Dynamics Simulation of a Fast-Folding WW Domain,* Biophysical Journal **94**, L75 (2008). doi:10.1529/biophysj.108.131565.

[81] L. C. Pierce, R. Salomon-Ferrer, C. Augusto F. de Oliveira, J. A. McCammon, and R. C. Walker, *Routine Access to Millisecond Time Scale Events with Accelerated Molecular Dynamics,* Journal of Chemical Theory and Computation **8**, 2997 (2012). doi:10.1021/ct300284c.

[82] D. E. Shaw, P. J. Adams, A. Azaria, J. A. Bank, B. Batson, A. Bell, M. Bergdorf, J. Bhatt, J. A. Butts, T. Correia, R. M. Dirks, R. O. Dror, M. P. Eastwood, B. Edwards, A. Even, P. Feldmann, M. Fenn, C. H. Fenton, A. Forte, J. Gagliardo, G. Gill, M. Gorlatova, B. Greskamp, J. Grossman, J. Gullingsrud, A. Harper, W. Hasenplaugh, M. Heily, B. C. Heshmat, J. Hunt, D. J. Ierardi, L. Iserovich, B. L. Jackson, N. P. Johnson, M. M. Kirk, J. L. Klepeis, J. S. Kuskin, K. M. Mackenzie, R. J. Mader, R. McGowen, A. McLaughlin, M. A. Moraes, M. H. Nasr, L. J. Nociolo, L. O'Donnell, A. Parker, J. L. Peticolas, G. Pocina, C. Predescu, T. Quan, J. K. Salmon, C. Schwink, K. S. Shim, N. Siddique, J. Spengler, T. Szalay, R. Tabladillo, R. Tartler, A. G. Taube, M. Theobald, B. Towles, W. Vick, S. C. Wang, M. Wazlowski, M. J. Weingarten, J. M. Williams, and K. A. Yuh, *Anton 3: Twenty Microseconds of Molecular Dynamics Simulation before Lunch,* in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, (St. Louis Missouri), 1, ACM, 2021. doi:10.1145/3458817.3487397.

[83] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Oxford University Press, 1987).

[84] W. C. Swope, H. C. Andersen, P. H. Berens, and K. R. Wilson, *A Computer Simulation Method for the Calculation of Equilibrium Constants for the Formation of Physical Clusters of Molecules: Application to Small Water Clusters,* The Journal of Chemical Physics **76**, 637 (1982). doi:10.1063/1.442716.

[85] R. Bowley, M. Sanchez, R. Bowley, and M. Sanchez, *Introductory Statistical Mechanics* (Oxford, New York: Oxford University Press, 1999). ISBN 978-0-19-850576-1.

[86] J. P. Sethna, *Entropy, Order Parameters, and Complexity* (Oxford: Clarendon Press, 2021).

[87] J. P. Boon and S. Yip, *Molecular Hydrodynamics* (McGraw-Hill Education, 1980).

[88] M. E. Tuckerman, *Statistical Mechanics: Theory and Molecular Simulation* (New York, NY: Oxford University Press, 2010). ISBN 978-0-19-852526-4.

[89] V. Calandrini, E. Pellegrini, P. Calligari, K. Hinsen, and G. Kneller, *nMoldyn - Interfacing Spectroscopic Experiments, Molecular Dynamics Simulations and Models for Time Correlation Functions,* École thématique de la Société Française de la Neutronique **12**, 201 (2011). doi:10.1051/sfn/201112010.

[90] F. Reif, *Fundamentals of Statistical and Thermal Physics* (Singapore: McGraw-Hill, 1985).

[91] N. Farzi and M. Ebrahim, *Mechanical Properties and Glass Transition Temperature of Metal-Organic Framework-Filled Epoxy Resin: A Molecular Dynamics Study,* Materials Chemistry and Physics **314**, 128874 (2024). doi:10.1016/j.matchemphys.2023.128874.

[92] N. Marchin, S. Urata, and J. Du, *Effect of Three-Body Interaction on Structural Features of Phosphate Glasses from Molecular Dynamics Simulations,* The Journal of Chemical Physics **161**, 154507 (2024). doi:10.1063/5.0225188.

[93] K.-H. Lin, L. Paterson, F. May, and D. Andrienko, *Glass Transition Temperature Prediction of Disordered Molecular Solids,* npj Computational Materials **7**, 1 (2021). doi:10.1038/s41524-021-00647-w.

[94] S. E. Root, S. Savagatrup, C. J. Pais, G. Arya, and D. J. Lipomi, *Predicting the Mechanical Properties of Organic Semiconductors Using Coarse-Grained Molecular Dynamics Simulations,* Macromolecules **49**, 2886 (2016). doi:10.1021/acs.macromol.6b00204.

[95] C. P. Callaway, J. H. Bombile, W. Mask, S. M. Ryno, and C. Risko, *Thermomechanical Enhancement of DPP-4T through Purposeful π-Conjugation Disruption,* Journal of Polymer Science **60**, 559 (2022). doi:10.1002/pol.20210494.

[96] J. J. Waterfall, F. P. Casey, R. N. Gutenkunst, K. S. Brown, C. R. Myers, P. W. Brouwer, V. Elser, and J. P. Sethna, *Sloppy-Model Universality Class and the Vandermonde Matrix,* Physical Review Letters **97**, 150601 (2006). `doi:10.1103/PhysRevLett.97.150601`.

[97] R. N. Gutenkunst, J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers, and J. P. Sethna, *Universally Sloppy Parameter Sensitivities in Systems Biology Models,* PLOS Computational Biology **3**, e189 (2007). `doi:10.1371/journal.pcbi.0030189`.

[98] M. K. Transtrum, B. B. Machta, and J. P. Sethna, *Geometry of Nonlinear Least Squares with Applications to Sloppy Models and Optimization,* Physical Review E **83**, 036701 (2011). `doi:10.1103/PhysRevE.83.036701`.

[99] A. Barducci, G. Bussi, and M. Parrinello, *Well-Tempered Metadynamics: A Smoothly Converging and Tunable Free-Energy Method,* Physical Review Letters **100**, 020603 (2008). `doi:10.1103/PhysRevLett.100.020603`.

[100] A. F. Voter, *Hyperdynamics: Accelerated Molecular Dynamics of Infrequent Events,* Physical Review Letters **78**, 3908 (1997). `doi:10.1103/PhysRevLett.78.3908`.

[101] A. Mitsutake, Y. Sugita, and Y. Okamoto, *Generalized-Ensemble Algorithms for Molecular Simulations of Biopolymers,* Peptide Science **60**, 96 (2001). `doi:10.1002/1097-0282(2001)60:2<96::AID-BIP1007>3.0.CO;2-F`.

[102] J. D. Chodera and M. R. Shirts, *Replica Exchange and Expanded Ensemble Simulations as Gibbs Sampling: Simple Improvements for Enhanced Mixing,* The Journal of Chemical Physics **135**, 194110 (2011). `doi:10.1063/1.3660669`.

[103] R. C. Bernardi, M. C. R. Melo, and K. Schulten, *Enhanced Sampling Techniques in Molecular Dynamics Simulations of Biological Systems,* Biochimica et Biophysica Acta (BBA) - General Subjects **1850**, 872 (2015). `doi:10.1016/j.bbagen.2014.10.019`.

[104] J. Hénin, T. Lelièvre, M. R. Shirts, O. Valsson, and L. Delemotte, *Enhanced Sampling Methods for Molecular Dynamics Simulations,* Living Journal of Computational Molecular Science **4**, 1583 (2022). `doi:10.33011/livecoms.4.1.1583`.

[105] S. Kmiecik, D. Gront, M. Kolinski, L. Wieteska, A. E. Dawid, and A. Kolinski, *Coarse-Grained Protein Models and Their Applications,* Chemical Reviews **116**, 7898 (2016). `doi:10.1021/acs.chemrev.6b00163`.

[106] L. J. Lewis and G. Wahnström, *Molecular-Dynamics Study of Supercooled Ortho-Terphenyl,* Physical Review E **50**, 3865 (1994). `doi:10.1103/PhysRevE.50.3865`.

[107] A. Lin, K. K. Huguenin-Dumittan, Y.-C. Cho, J. Nigam, and R. K. Cersonsky, *Expanding Density-Correlation Machine Learning Representations for Anisotropic Coarse-Grained Particles,* The Journal of Chemical Physics **161**, 074112 (2024). `doi:10.1063/5.0210910`.

[108] W. Xia, J. Song, C. Jeong, D. D. Hsu, F. R. J. Phelan, J. F. Douglas, and S. Keten, *Energy-Renormalization for Achieving Temperature Transferable Coarse-Graining of Polymer Dynamics,* Macromolecules **50**, 8787 (2017). `doi:10.1021/acs.macromol.7b01717`.

[109] D. Rosenberger and N. F. A. van der Vegt, *Addressing the Temperature Transferability of Structure Based Coarse Graining Models,* Physical Chemistry Chemical Physics **20**, 6617 (2018). `doi:10.1039/C7CP08246K`.

[110] G. Biroli and M. Mézard, *Lattice Glass Models,* Physical Review Letters **88**, 025501 (2001). `doi:10.1103/PhysRevLett.88.025501`.

[111] K. Zhu, E. Trizio, J. Zhang, R. Hu, L. Jiang, T. Hou, and L. Bonati, *Enhanced Sampling in the Age of Machine Learning: Algorithms and Applications,* Chemical Reviews **126**, 671 (2026). `doi:10.1021/acs.chemrev.5c00700`.

[112] W. Kob, M. Nauroth, and F. Sciortino, *Quantitative Tests of Mode-Coupling Theory for Fragile and Strong Glass Formers,* Journal of Non-Crystalline Solids **307–310**, 181 (2002). `doi:10.1016/S0022-3093(02)01457-6`.

[113] F. Weysser, A. M. Puertas, M. Fuchs, and Th. Voigtmann, *Structural Relaxation of Polydisperse Hard Spheres: Comparison of the Mode-Coupling Theory to a Langevin Dynamics Simulation,* Physical Review E **82**, 011504 (2010). doi:10.1103/PhysRevE.82.011504.

[114] W. Götze, *Recent Tests of the Mode-Coupling Theory for Glassy Dynamics,* Journal of Physics: Condensed Matter **11**, A1 (1999). doi:10.1088/0953-8984/11/10A/002.

[115] F. Noé, S. Olsson, J. Köhler, and H. Wu, *Boltzmann Generators: Sampling Equilibrium States of Many-Body Systems with Deep Learning,* Science **365**, eaaw1147 (2019). doi:10.1126/science.aaw1147.

[116] T. Hsu, B. Sadigh, V. Bulatov, and F. Zhou, *Score Dynamics: Scaling Molecular Dynamics with Picoseconds Time Steps via Conditional Diffusion Model,* Journal of Chemical Theory and Computation **20**, 2335 (2024). doi:10.1021/acs.jctc.3c01361.

[117] M. Schreiner, O. Winther, and S. Olsson, *Implicit Transfer Operator Learning: Multiple Time-Resolution Models for Molecular Dynamics,* Advances in Neural Information Processing Systems **36**, 36449 (2023).

[118] S. Olsson, *Generative Molecular Dynamics,* Current Opinion in Structural Biology **96**, 103213 (2026). doi:10.1016/j.sbi.2025.103213.

[119] L. Klein, A. Krämer, and F. Noé, *Equivariant Flow Matching,* 2023. doi:10.48550/arXiv.2306.15030.

[120] R. M. Martin, *Electronic Structure: Basic Theory and Practical Methods* (Cambridge: Cambridge University Press, 2004). doi:10.1017/CBO9780511805769.

[121] C. A. Ullrich, *Time-Dependent Density-Functional Theory: Concepts and Applications* (Oxford University Press, 2011). ISBN 978-0-19-177513-0. doi:10.1093/acprof:oso/9780199563029.001.0001.

[122] M. Born and R. Oppenheimer, *Zur Quantentheorie Der Molekeln,* Annalen der Physik **389**, 457 (1927). doi:10.1002/andp.19273892002.

[123] W. Kohn and L. J. Sham, *Self-Consistent Equations Including Exchange and Correlation Effects,* Physical Review **140**, A1133 (1965). doi:10.1103/PhysRev.140.A1133.

[124] K. Berland and P. Hyldgaard, *Exchange Functional That Tests the Robustness of the Plasmon Description of the van Der Waals Density Functional,* Physical Review B **89**, 035412 (2014). doi:10.1103/PhysRevB.89.035412.

[125] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (Cambridge, Mass: MIT Press, 2006). ISBN 978-0-262-18253-9.

[126] E. Gilboa, Y. Saatçi, and J. P. Cunningham, *Scaling Multidimensional Inference for Structured Gaussian Processes,* IEEE Transactions on Pattern Analysis and Machine Intelligence **37**, 424 (2015). doi:10.1109/TPAMI.2013.192.

[127] S. Chmiela, A. Tkatchenko, H. E. Sauceda, I. Poltavsky, K. T. Schütt, and K.-R. Müller, *Machine Learning of Accurate Energy-Conserving Molecular Force Fields,* Science Advances **3**, e1603015 (2017). doi:10.1126/sciadv.1603015.

[128] V. L. Deringer, A. P. Bartók, N. Bernstein, D. M. Wilkins, M. Ceriotti, and G. Csányi, *Gaussian Process Regression for Materials and Molecules,* Chemical Reviews **121**, 10073 (2021). doi:10.1021/acs.chemrev.1c00022.

[129] S. Chmiela, H. E. Sauceda, K.-R. Müller, and A. Tkatchenko, *Towards Exact Molecular Dynamics Simulations with Machine-Learned Force Fields,* Nature Communications **9**, 3887 (2018). doi:10.1038/s41467-018-06169-2.

[130] J. Behler and M. Parrinello, *Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces,* Physical Review Letters **98**, 146401 (2007). doi:10.1103/PhysRevLett.98.146401.

[131] J. Behler, *Atom-Centered Symmetry Functions for Constructing High-Dimensional Neural Network Potentials,* The Journal of Chemical Physics **134**, 074106 (2011). doi:10.1063/1.3553717.

[132] M. Gastegger, J. Behler, and P. Marquetand, *Machine Learning Molecular Dynamics for the Simulation of Infrared Spectra,* Chemical Science **8**, 6924 (2017). doi:10.1039/C7SC02267K.

[133] H. Wang, L. Zhang, J. Han, and W. E, *DeePMD-kit: A Deep Learning Package for Many-Body Potential Energy Representation and Molecular Dynamics,* Computer Physics Communications **228**, 178 (2018). doi:10.1016/j.cpc.2018.03.016.

[134] L. Zhang, J. Han, H. Wang, W. Saidi, R. Car, and W. E, *End-to-End Symmetry Preserving Inter-atomic Potential Energy Model for Finite and Extended Systems,* in *Advances in Neural Information Processing Systems,* vol. 31, Curran Associates, Inc., 2018.

[135] Y. Zhang, C. Hu, and B. Jiang, *Embedded Atom Neural Network Potentials: Efficient and Accurate Machine Learning with a Physically Inspired Representation,* The Journal of Physical Chemistry Letters **10**, 4962 (2019). doi:10.1021/acs.jpclett.9b02037.

[136] Y. Zhang, J. Xia, and B. Jiang, *REANN: A PyTorch-based End-to-End Multi-Functional Deep Neural Network Package for Molecular, Reactive, and Periodic Systems,* The Journal of Chemical Physics **156**, 114801 (2022). doi:10.1063/5.0080766.

[137] Y. Zhang, J. Xia, and B. Jiang, *Physically Motivated Recursively Embedded Atom Neural Networks: Incorporating Local Completeness and Nonlocality,* Physical Review Letters **127**, 156002 (2021). doi:10.1103/PhysRevLett.127.156002.

[138] J. S. Smith, O. Isayev, and A. E. Roitberg, *ANI-1: An Extensible Neural Network Potential with DFT Accuracy at Force Field Computational Cost,* Chemical Science **8**, 3192 (2017). doi:10.1039/C6SC05720A.

[139] K. T. Schütt, H. E. Sauceda, P.-J. Kindermans, A. Tkatchenko, and K.-R. Müller, *SchNet – A Deep Learning Architecture for Molecules and Materials,* The Journal of Chemical Physics **148**, 241722 (2018). doi:10.1063/1.5019779.

[140] S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt, and B. Kozinsky, *E(3)-Equivariant Graph Neural Networks for Data-Efficient and Accurate Interatomic Potentials,* Nature Communications **13**, 2453 (2022). doi:10.1038/s41467-022-29939-5.

[141] A. Musaelian, S. Batzner, A. Johansson, L. Sun, C. J. Owen, M. Kornbluth, and B. Kozinsky, *Learning Local Equivariant Representations for Large-Scale Atomistic Dynamics,* Nature Communications **14**, 579 (2023). doi:10.1038/s41467-023-36329-y.

[142] I. Batatia, D. P. Kovacs, G. Simm, C. Ortner, and G. Csanyi, *MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields,* Advances in Neural Information Processing Systems **35**, 11423 (2022).

[143] A. V. Shapeev, *Moment Tensor Potentials: A Class of Systematically Improvable Interatomic Potentials,* Multiscale Modeling & Simulation **14**, 1153 (2016). doi:10.1137/15M1054183.

[144] I. S. Novikov, K. Gubaev, E. V. Podryabinkin, and A. V. Shapeev, *The MLIP Package: Moment Tensor Potentials with MPI and Active Learning,* Machine Learning: Science and Technology **2**, 025002 (2020). doi:10.1088/2632-2153/abc9fe.

[145] A. P. Thompson, L. P. Swiler, C. R. Trott, S. M. Foiles, and G. J. Tucker, *Spectral Neighbor Analysis Method for Automated Generation of Quantum-Accurate Interatomic Potentials,* Journal of Computational Physics **285**, 316 (2015). doi:10.1016/j.jcp.2014.12.018.

[146] R. Drautz, *Atomic Cluster Expansion for Accurate and Transferable Interatomic Potentials,* Physical Review B **99**, 014104 (2019). doi:10.1103/PhysRevB.99.014104.

[147] D. P. Kovács, C. van der Oord, J. Kucera, A. E. A. Allen, D. J. Cole, C. Ortner, and G. Csányi, *Linear Atomic Cluster Expansion Force Fields for Organic Molecules: Beyond RMSE,* Journal of Chemical Theory and Computation **17**, 7696 (2021). doi:10.1021/acs.jctc.1c00647.

[148] G. Dusson, M. Bachmayr, G. Csányi, R. Drautz, S. Etter, C. van der Oord, and C. Ortner, *Atomic Cluster Expansion: Completeness, Efficiency and Stability,* Journal of Computational Physics **454**, 110946 (2022). doi:10.1016/j.jcp.2022.110946.

[149] J. Behler, *Constructing High-Dimensional Neural Network Potentials: A Tutorial Review,* International Journal of Quantum Chemistry **115**, 1032 (2015). doi:10.1002/qua.24890.

[150] M. Thürlemann, L. Böselt, and S. Riniker, *Learning Atomic Multipoles: Prediction of the Electrostatic Potential with Equivariant Graph Neural Networks,* Journal of Chemical Theory and Computation **18**, 1701 (2022). doi:10.1021/acs.jctc.1c01021.

[151] V. G. Satorras, E. Hoogeboom, and M. Welling, *E(n) Equivariant Graph Neural Networks,* 2022. doi:10.48550/arXiv.2102.09844.

[152] A. Mazitov, F. Bigi, M. Kellner, P. Pegolo, D. Tisi, G. Fraux, S. Pozdnyakov, P. Loche, and M. Ceriotti, *PET-MAD as a Lightweight Universal Interatomic Potential for Advanced Materials Modeling,* Nature Communications **16**, 10653 (2025). doi:10.1038/s41467-025-65662-7.

[153] J. Liu, Z. Shen, Y. He, X. Zhang, R. Xu, H. Yu, and P. Cui, *Towards Out-Of-Distribution Generalization: A Survey,* 2023. doi:10.48550/arXiv.2108.13624.

[154] M. Karabin and D. Perez, *An Entropy-Maximization Approach to Automated Training Set Generation for Interatomic Potentials,* The Journal of Chemical Physics **153**, 094110 (2020). doi:10.1063/5.0013059.

[155] D. Schwalbe-Koda, S. Hamel, B. Sadigh, F. Zhou, and V. Lordi, *Model-Free Estimation of Completeness, Uncertainties, and Outliers in Atomistic Machine Learning Using Information Theory,* Nature Communications **16**, 4014 (2025). doi:10.1038/s41467-025-59232-0.

[156] A. P. Bartók, R. Kondor, and G. Csányi, *On Representing Chemical Environments,* Physical Review B **87**, 184115 (2013). doi:10.1103/PhysRevB.87.184115.

[157] S. Jindal, S. Chiriki, and S. S. Bulusu, *Spherical Harmonics Based Descriptor for Neural Network Potentials: Structure and Dynamics of Au147 Nanocluster,* The Journal of Chemical Physics **146**, 204301 (2017). doi:10.1063/1.4983392.

[158] H. Huo and M. Rupp, *Unified Representation of Molecules and Crystals for Machine Learning,* Machine Learning: Science and Technology **3**, 045017 (2022). doi:10.1088/2632-2153/aca005.

[159] K. Song, R. Zhao, J. Liu, Y. Wang, E. Lindgren, Y. Wang, S. Chen, K. Xu, T. Liang, P. Ying, N. Xu, Z. Zhao, J. Shi, J. Wang, S. Lyu, Z. Zeng, S. Liang, H. Dong, L. Sun, Y. Chen, Z. Zhang, W. Guo, P. Qian, J. Sun, P. Erhart, T. Ala-Nissila, Y. Su, and Z. Fan, *General-Purpose Machine-Learned Potential for 16 Elemental Metals and Their Alloys,* 2023. doi:10.48550/arXiv.2311.04732.

[160] H. Zou and T. Hastie, *Regularization and Variable Selection Via the Elastic Net,* Journal of the Royal Statistical Society Series B: Statistical Methodology **67**, 301 (2005). doi:10.1111/j.1467-9868.2005.00503.x.

[161] Y. Sun, F. Gomez, T. Schaul, and J. Schmidhuber, *A Linear Time Natural Evolution Strategy for Non-Separable Functions,* 2011. doi:10.48550/arXiv.1106.1998.

[162] Z. Fan, Z. Zeng, C. Zhang, Y. Wang, K. Song, H. Dong, Y. Chen, and T. Ala-Nissila, *Neuroevolution Machine Learning Potentials: Combining High Accuracy and Low Cost in Atomistic Simulations and Application to Heat Transport,* Physical Review B **104**, 104309 (2021). doi:10.1103/PhysRevB.104.104309.

[163] T. Glasmachers, T. Schaul, S. Yi, D. Wierstra, and J. Schmidhuber, *Exponential Natural Evolution Strategies,* in *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation - GECCO '10*, (Portland, Oregon, USA), 393, ACM Press, 2010. doi:10.1145/1830483.1830557.

[164] Z. Fan, *Improving the Accuracy of the Neuroevolution Machine Learning Potential for Multi-Component Systems,* Journal of Physics: Condensed Matter **34**, 125902 (2022). doi:10.1088/1361-648X/ac462b.

[165] D. A. Cohn, Z. Ghahramani, and M. I. Jordan, *Active Learning with Statistical Models,* Journal of Artificial Intelligence Research **4**, 129 (1996). doi:10.1613/jair.295.

[166] H. Liu, Y.-S. Ong, and J. Cai, *A Survey of Adaptive Sampling for Global Metamodeling in Support of Simulation-Based Complex Engineering Design,* Structural and Multidisciplinary Optimization **57**, 393 (2018). doi:10.1007/s00158-017-1739-8.

[167] L. Breiman, *Bagging Predictors,* Machine Learning **24**, 123 (1996). doi:10.1007/BF00058655.

[168] J. Paul Janet, C. Duan, T. Yang, A. Nandy, and H. J. Kulik, *A Quantitative Uncertainty Metric Controls Error in Neural Network-Driven Chemical Discovery,* Chemical Science **10**, 7913 (2019). doi:10.1039/C9SC02298H.

[169] J. A. Vita, A. Samanta, F. Zhou, and V. Lordi, *LTAU-FF: Loss Trajectory Analysis for Uncertainty in Atomistic Force Fields,* Machine Learning: Science and Technology **6**, 015048 (2025). doi:10.1088/2632-2153/adb4b9.

[170] N. Leimeroth, L. C. Erhard, K. Albe, and J. Rohrer, *Machine-Learning Interatomic Potentials from a Users Perspective: A Comparison of Accuracy, Speed and Data Efficiency,* Modelling and Simulation in Materials Science and Engineering **33**, 065012 (2025). doi:10.1088/1361-651X/adf56d.

[171] 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**, 108171 (2022). doi:10.1016/j.cpc.2021.108171.

[172] S. J. Plimpton, A. Kohlmeyer, A. P. Thompson, S. G. Moore, and R. Berger, *LAMMPS: Large-scale Atomic/Molecular Massively Parallel Simulator,* Zenodo, 2023. doi:10.5281/zenodo.10806852.

[173] H. Carter Edwards, C. R. Trott, and D. Sunderland, *Kokkos: Enabling Manycore Performance Portability through Polymorphic Memory Access Patterns,* Journal of Parallel and Distributed Computing **74**, 3202 (2014). doi:10.1016/j.jpdc.2014.07.003.

[174] C. Trott, L. Berger-Vergiat, D. Poliakoff, S. Rajamanickam, D. Lebrun-Grandie, J. Madsen, N. Al Awar, M. Gligoric, G. Shipman, and G. Womeldorff, *The Kokkos EcoSystem: Comprehensive Performance Portability for High Performance Computing,* Computing in Science & Engineering **23**, 10 (2021). doi:10.1109/MCSE.2021.3098509.

[175] C. R. Trott, D. Lebrun-Grandié, D. Arndt, J. Ciesko, V. Dang, N. Ellingwood, R. Gayatri, E. Harvey, D. S. Hollman, D. Ibanez, N. Liber, J. Madsen, J. Miles, D. Poliakoff, A. Powell, S. Rajamanickam, M. Simberg, D. Sunderland, B. Turcksin, and J. Wilke, *Kokkos 3: Programming Model Extensions for the Exascale Era,* IEEE Transactions on Parallel and Distributed Systems **33**, 805 (2022). doi:10.1109/TPDS.2021.3097283.

[176] C. W. Tan, M. L. Descoteaux, M. Kotak, G. d. M. Nascimento, S. R. Kavanagh, L. Zichi, M. Wang, A. Saluja, Y. R. Hu, T. Smidt, A. Johansson, W. C. Witt, B. Kozinsky, and A. Musaelian, *High-Performance Training and Inference for Deep Equivariant Interatomic Potentials,* 2025. doi:10.48550/arXiv.2504.16068.

[177] Y. Lysogorskiy, C. van der Oord, A. Bochkarev, S. Menon, M. Rinaldi, T. Hammerschmidt, M. Mrovec, A. Thompson, G. Csányi, C. Ortner, and R. Drautz, *Performant Implementation of the Atomic Cluster Expansion (PACE) and Application to Copper and Silicon,* npj Computational Materials **7**, 97 (2021). doi:10.1038/s41524-021-00559-9.

[178] 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**, 273002 (2017). doi:10.1088/1361-648X/aa680e.

[179] I. Batatia, P. Benner, Y. Chiang, A. M. Elena, D. P. Kovács, J. Riebesell, X. R. Advincula, M. Asta, M. Avaylon, W. J. Baldwin, F. Berger, N. Bernstein, A. Bhowmik, S. M. Blau, V. Cărare, J. P. Darby, S. De, F. D. Pia, V. L. Deringer, R. Elijošius, Z. El-Machachi, F. Falcioni, E. Fako, A. C. Ferrari, A. Genreith-Schriever, J. George, R. E. A. Goodall, C. P. Grey, P. Grigorev, S. Han, W. Handley, H. H. Heenen, K. Hermansson, C. Holm, J. Jaafar, S. Hofmann, K. S. Jakob, H. Jung, V. Kapil, A. D.

Kaplan, N. Karimitari, J. R. Kermode, N. Kroupa, J. Kullgren, M. C. Kuner, D. Kuryla, G. Liepuoniute, J. T. Margraf, I.-B. Magdău, A. Michaelides, J. H. Moore, A. A. Naik, S. P. Niblett, S. W. Norwood, N. O'Neill, C. Ortner, K. A. Persson, K. Reuter, A. S. Rosen, L. L. Schaaf, C. Schran, B. X. Shi, E. Sivonxay, T. K. Stenczel, V. Svahn, C. Sutton, T. D. Swinburne, J. Tilly, C. van der Oord, E. Varga-Umbrich, T. Vegge, M. Vondrák, Y. Wang, W. C. Witt, F. Zills, and G. Csányi, *A Foundation Model for Atomistic Materials Chemistry,* 2024. doi:10.48550/arXiv.2401.00096.

[180] D. P. Kovács, J. H. Moore, N. J. Browning, I. Batatia, J. T. Horton, Y. Pu, V. Kapil, W. C. Witt, I.-B. Magdău, D. J. Cole, and G. Csányi, *MACE-OFF: Short-Range Transferable Machine Learning Force Fields for Organic Molecules,* Journal of the American Chemical Society **147**, 17598 (2025). doi:10.1021/jacs.4c07099.

[181] S. R. Kavanagh and M. G. . Harvard, *NequIP & Allegro Foundation Potentials,* Zenodo, 2026. doi:10.5281/zenodo.18775904.

[182] C. Chen and S. P. Ong, *A Universal Graph Deep Learning Interatomic Potential for the Periodic Table,* Nature Computational Science **2**, 718 (2022). doi:10.1038/s43588-022-00349-3.

[183] 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**, 1031 (2023). doi:10.1038/s42256-023-00716-3.

[184] K. Song, R. Zhao, J. Liu, Y. Wang, E. Lindgren, Y. Wang, S. Chen, K. Xu, T. Liang, P. Ying, N. Xu, Z. Zhao, J. Shi, J. Wang, S. Lyu, Z. Zeng, S. Liang, H. Dong, L. Sun, Y. Chen, Z. Zhang, W. Guo, P. Qian, J. Sun, P. Erhart, T. Ala-Nissila, Y. Su, and Z. Fan, *General-Purpose Machine-Learned Potential for 16 Elemental Metals and Their Alloys,* Nature Communications **15**, 10208 (2024). doi:10.1038/s41467-024-54554-x.

[185] W. L. Jorgensen and J. Tirado-Rives, *The OPLS [Optimized Potentials for Liquid Simulations] Potential Functions for Proteins, Energy Minimizations for Crystals of Cyclic Peptides and Crambin,* Journal of the American Chemical Society **110**, 1657 (1988). doi:10.1021/ja00214a001.

[186] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, *A Consistent and Accurate Ab Initio Parametrization of Density Functional Dispersion Correction (DFT-D) for the 94 Elements H-Pu,* The Journal of Chemical Physics **132**, 154104 (2010). doi:10.1063/1.3382344.

[187] Z. Fan, B. Tang, E. Berger, E. Berger, E. Fransson, K. Xu, Z. Yan, Z. Liu, Z. Song, H. Dong, S. Chen, L. Li, Z. Wang, Y. Zhu, J. Wiktor, and P. Erhart, *qNEP: A Highly Efficient Neuroevolution Potential with Dynamic Charges for Large-Scale Atomistic Simulations,* Journal of Chemical Theory and Computation , (2026). doi:10.1021/acs.jctc.6c00146.

[188] N. A. Spaldin, *A Beginner's Guide to the Modern Theory of Polarization,* Journal of Solid State Chemistry **195**, 2 (2012). doi:10.1016/j.jssc.2012.05.010.

[189] R. D. King-Smith and D. Vanderbilt, *Theory of Polarization of Crystalline Solids,* Physical Review B **47**, 1651 (1993). doi:10.1103/PhysRevB.47.1651.

[190] R. Resta, *Macroscopic Electric Polarization as a Geometric Quantum Phase,* Europhysics Letters **22**, 133 (1993). doi:10.1209/0295-5075/22/2/010.

[191] D. Vanderbilt and R. D. King-Smith, *Electric Polarization as a Bulk Quantity and Its Relation to Surface Charge,* Physical Review B **48**, 4442 (1993). doi:10.1103/PhysRevB.48.4442.

[192] D. A. McQuarrie, *Statistical Mechanics* (New York: Harper Collins, 1976). ISBN 978-1-891389-15-3.

[193] J. Westermayr and P. Marquetand, *Machine Learning Spectroscopy to Advance Computation and Analysis,* Chemical Science **16**, 21660 (2025). doi:10.1039/d5sc05628d.

[194] L. Zhao, J. Zhang, Y. Zhang, S. Ye, G. Zhang, X. Chen, B. Jiang, and J. Jiang, *Accurate Machine Learning Prediction of Protein Circular Dichroism Spectra with Embedded Density Descriptors,* JACS Au **1**, 2377 (2021). doi:10.1021/jacsau.1c00449.

[195] H. Ren, H. Li, Q. Zhang, L. Liang, W. Guo, F. Huang, Y. Luo, and J. Jiang, *A Machine Learning Vibrational Spectroscopy Protocol for Spectrum Prediction and Spectrum-Based Structure Recognition,* Fundamental Research **1**, 488 (2021). doi:10.1016/j.fmre.2021.05.005.

[196] N. Bhatia, O. Krejci, S. Botti, P. Rinke, and M. A. L. Marques, *MACE4IRmol: An Uncertainty-Aware Foundation Model for Molecular Infrared Spectroscopy,* 2026. doi:10.48550/arXiv.2508.19118.

[197] R. D. Fernald, *Casting a Genetic Light on the Evolution of Eyes,* Science **313**, 1914 (2006). doi:10.1126/science.1127889.

[198] T. Lóránd, P. Molnár, J. Deli, and G. Tóth, *FT-IR Study of Some Seco- and Apocarotenoids,* Journal of Biochemical and Biophysical Methods **53**, 251 (2002). doi:10.1016/S0165-022X(02)00113-6.

[199] P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch, *Ab Initio Calculation of Vibrational Absorption and Circular Dichroism Spectra Using Density Functional Force Fields,* The Journal of Physical Chemistry **98**, 11623 (1994). doi:10.1021/j100096a001.

[200] W. B. How, P. Febrer, S. Chong, A. Mazitov, F. Bigi, M. Kellner, S. Pozdnyakov, and M. Ceriotti, *A Universal Machine Learning Model for the Electronic Density of States,* 2026. doi:10.48550/arXiv.2508.17418.

[201] P. R. Kaundinya, K. Choudhary, and S. R. Kalidindi, *Prediction of the Electron Density of States for Crystalline Compounds with Atomistic Line Graph Neural Networks (ALIGNN),* JOM **74**, 1395 (2022). doi:10.1007/s11837-022-05199-y.

[202] N. Lee, H. Noh, S. Kim, D. Hyun, G. S. Na, and C. Park, *Density of States Prediction of Crystalline Materials via Prompt-guided Multi-Modal Transformer,* 2023. doi:10.48550/arXiv.2311.12856.

[203] B. Sadigh, P. Erhart, D. Åberg, A. Trave, E. Schwegler, and J. Bude, *First-Principles Calculations of the Urbach Tail in the Optical Absorption Spectra of Silica Glass,* Physical Review Letters **106**, 027401 (2011). doi:10.1103/PhysRevLett.106.027401.

[204] J. Tauc, *Optical Properties and Electronic Structure of Amorphous Ge and Si,* Materials Research Bulletin **3**, 37 (1968). doi:10.1016/0025-5408(68)90023-8.

[205] K. Saito and A. J. Ikushima, *Absorption Edge in Silica Glass,* Physical Review B **62**, 8584 (2000). doi:10.1103/PhysRevB.62.8584.

[206] J. P. Perdew, K. Burke, and M. Ernzerhof, *Generalized Gradient Approximation Made Simple,* Physical Review Letters **77**, 3865 (1996). doi:10.1103/PhysRevLett.77.3865.

[207] Á. Morales-García, R. Valero, and F. Illas, *An Empirical, yet Practical Way To Predict the Band Gap in Solids by Using Density Functional Band Structure Calculations,* The Journal of Physical Chemistry C **121**, 18862 (2017). doi:10.1021/acs.jpcc.7b07421.

[208] G. L. Squires, *Introduction to the Theory of Thermal Neutron Scattering* (Cambridge: Cambridge University Press, 2012). ISBN 978-1-107-64406-9. doi:10.1017/CBO9781139107808.

[209] V. F. Sears, *Neutron Scattering Lengths and Cross Sections,* Neutron News **3**, 26 (1992). doi:10.1080/10448639208218770.

[210] P. Schurtenberger and J. Oberdisse, *Contrast and Contrast Variation in Neutron, X-ray, and Light Scattering,* in *Neutrons, X-rays, and Light* (Elsevier, 2025), p. 151. doi:10.1016/B978-0-443-29116-6.00022-9.

[211] M. Almakki, R. Applin, M. Backman, R. Baust, J. Borreguero, R. Boston, A. Bridger, J. Clarke, A. Diaz-Alvarez, R. Farooq, C. Finn, S. Foxley, D. Ganyushin, J. Haigh, T. Hampson, D. Ioannide, A. J. Jackson, W. P. Jayasundara Abeykoon Wickramasingha, D. Le, M. Lewis, Z. Morgan, M. Patrou, G. Pereira, P. F. Peterson, K. Qianli Ma, A. Savici, S. Schomann, C. Sears, K. TacTac, K. Travis, R. Waite, M. Walsh, R. Whitfield, J. Yusuf, C. Zhang, and Y. Zhang, *Mantid 6.12.0: Manipulation and Analysis Toolkit for Instrument Data.,* Mantid Project, 2025. doi:10.5286/SOFTWARE/MANTID6.12.

[212] J. Rodríguez-Carvajal, *Recent Advances in Magnetic Structure Determination by Neutron Powder Diffraction,* Physica B: Condensed Matter **192**, 55 (1993). doi:10.1016/0921-4526(93)90108-I.

[213] J. Rodriguez-Carvajal, J. Gonzalez-Platas, and N. A. Katcho, *Magnetic Structure Determination and Refinement Using FullProf,* Acta Crystallographica Section B: Structural Science, Crystal Engineering and Materials **81**, 302 (2025). doi:10.1107/S2052520625003944.

[214] K. Lefmann and K. Nielsen, *McStas, a General Software Package for Neutron Ray-Tracing Simulations,* Neutron News **10**, 20 (1999). doi:10.1080/10448639908233684.

[215] P. K. Willendrup and K. Lefmann, *McStas (i): Introduction, Use, and Basic Principles for Ray-Tracing Simulations,* Journal of Neutron Research **22**, 1 (2020). doi:10.3233/JNR-190108.

[216] P. K. Willendrup and K. Lefmann, *McStas (Ii): An Overview of Components, Their Use, and Advice for User Contributions,* Journal of Neutron Research **23**, 7 (2021). doi:10.3233/JNR-200186.

[217] D. Wechsler, G. Zsigmond, F. Streffer, and F. Mezei, *VITESS: Virtual Instrumentation Tool for Pulsed and Continuous Sources,* Neutron News **11**, 25 (2000). doi:10.1080/10448630008233764.

[218] G. Zsigmond, K. Lieutenant, and F. Mezei, *Monte Carlo Simulations of Neutron Scattering Instruments by VITESS: Virtual Instrumentation Tool for ESS,* Neutron News **13**, 11 (2002). doi:10.1080/10448630208218488.

[219] Y. Q. Cheng, L. L. Daemen, A. I. Kolesnikov, and A. J. Ramirez-Cuesta, *Simulation of Inelastic Neutron Scattering Spectra Using OCLIMAX,* Journal of Chemical Theory and Computation **15**, 1974 (2019). doi:10.1021/acs.jctc.8b01250.

[220] Y. Q. Cheng and A. J. Ramirez-Cuesta, *Calculation of the Thermal Neutron Scattering Cross-Section of Solids Using OCLIMAX,* Journal of Chemical Theory and Computation **16**, 5212 (2020). doi:10.1021/acs.jctc.0c00569.

[221] R. Fair, A. Jackson, D. Voneshen, D. Jochym, D. Le, K. Refson, and T. Perring, *Euphonic: Inelastic Neutron Scattering Simulations from Force Constants and Visualization Tools for Phonon Properties,* Journal of Applied Crystallography **55**, 1689 (2022). doi:10.1107/S1600576722009256.

[222] K. Dymkowski, S. F. Parker, F. Fernandez-Alonso, and S. Mukhopadhyay, *AbINS: The Modern Software for INS Interpretation,* Physica B: Condensed Matter **551**, 443 (2018). doi:10.1016/j.physb.2018.02.034.

[223] E. Fransson, M. Slabanja, P. Erhart, and G. Wahnström, *Dynasor—A Tool for Extracting Dynamical Structure Factors and Current Correlation Functions from Molecular Dynamics Simulations,* Advanced Theory and Simulations **4**, 2000240 (2021). doi:10.1002/adts.202000240.

[224] E. Berger, E. Fransson, F. Eriksson, E. Lindgren, G. Wahnström, T. H. Rod, and P. Erhart, *Dynasor 2: From Simulation to Experiment through Correlation Functions,* Computer Physics Communications **316**, 109759 (2025). doi:10.1016/j.cpc.2025.109759.

[225] T. M. Linker, A. Krishnamoorthy, L. L. Daemen, A. J. Ramirez-Cuesta, K. Nomura, A. Nakano, Y. Q. Cheng, W. R. Hicks, A. I. Kolesnikov, and P. D. Vashishta, *Neutron Scattering and Neural-Network Quantum Molecular Dynamics Investigation of the Vibrations of Ammonia along the Solid-to-Liquid Transition,* Nature Communications **15**, 3911 (2024). doi:10.1038/s41467-024-48246-9.

[226] S. F. Parker, F. Fernandez-Alonso, A. J. Ramirez-Cuesta, J. Tomkinson, S. Rudic, R. S. Pinna, G. Gorini, and J. Fernández Castañon, *Recent and Future Developments on TOSCA at ISIS,* Journal of Physics: Conference Series **554**, 012003 (2014). doi:10.1088/1742-6596/554/1/012003.

[227] R. S. Pinna, S. Rudić, S. F. Parker, J. Armstrong, M. Zanetti, G. Škoro, S. P. Waller, D. Zacek, C. A. Smith, M. J. Capstick, D. J. McPhail, D. E. Pooley, G. D. Howells, G. Gorini, and F. Fernandez-Alonso, *The Neutron Guide Upgrade of the TOSCA Spectrometer,* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **896**, 68 (2018). doi:10.1016/j.nima.2018.04.009.

[228] M. Cardona and G. Güntherodt, eds., *Light Scattering in Solids II* (Berlin, Heidelberg: Springer, 1982). ISBN 978-3-540-11380-5 978-3-540-39075-6. doi:10.1007/3-540-11380-0.

[229] F. Fillaux, R. Papoular, A. Lautié, and J. Tomkinson, *Inelastic Neutron-Scattering Study of the Proton Dynamics in Coals,* Fuel **74**, 865 (1995). doi:10.1016/0016-2361(95)00017-Y.

[230] J. Penfold, J. Tomkinson, and C. U. Rutherford Appleton Lab., *The ISIS Time Focussed Crystal Analyser Spectrometer, TFXA,* Tech. Rep. RAL–86-019, Rutherford Appleton Lab., Chilton (UK), 1986.

[231] M. J. Abraham, T. Murtola, R. Schulz, S. Páll, J. C. Smith, B. Hess, and E. Lindahl, *GROMACS: High Performance Molecular Simulations through Multi-Level Parallelism from Laptops to Supercomputers,* SoftwareX **1–2**, 19 (2015). doi:10.1016/j.softx.2015.06.001.

[232] T. W. Ko, J. A. Finkler, S. Goedecker, and J. Behler, *A Fourth-Generation High-Dimensional Neural Network Potential with Accurate Electrostatics Including Non-Local Charge Transfer,* Nature Communications **12**, 398 (2021). doi:10.1038/s41467-020-20427-2.

[233] B. Gurlek, S. Sharma, P. Lazzaroni, A. Rubio, and M. Rossi, *Accurate Machine Learning Interatomic Potentials for Polyacene Molecular Crystals: Application to Single Molecule Host-Guest Systems,* npj Computational Materials **11**, 318 (2025). doi:10.1038/s41524-025-01825-w.

[234] Y. Ji, J. Liang, and Z. Xu, *Machine-Learning Interatomic Potentials for Long-Range Systems,* Physical Review Letters **135**, 178001 (2025). doi:10.1103/ssp9-7s81.

[235] R. Lavén, E. Fransson, P. Erhart, F. Juranyi, G. E. Granroth, and M. Karlsson, *Unraveling the Nature of Vibrational Dynamics in CsPbI3 by Inelastic Neutron Scattering and Molecular Dynamics Simulations,* The Journal of Physical Chemistry Letters **16**, 4812 (2025). doi:10.1021/acs.jpclett.5c00778.

[236] I. Y. Zhang and A. Grüneis, *Coupled Cluster Theory in Materials Science,* Frontiers in Materials **6**, (2019). doi:10.3389/fmats.2019.00123.

[237] F. Coester, *Bound States of a Many-Particle System,* Nuclear Physics **7**, 421 (1958). doi:10.1016/0029-5582(58)90280-3.

[238] F. Coester and H. Kümmel, *Short-Range Correlations in Nuclear Wave Functions,* Nuclear Physics **17**, 477 (1960). doi:10.1016/0029-5582(60)90140-1.

[239] C. Ye, V. Gray, K. Kushwaha, S. Kumar Singh, P. Erhart, and K. Börjesson, *Optimizing Photon Upconversion by Decoupling Excimer Formation and Triplet Triplet Annihilation,* Physical Chemistry Chemical Physics **22**, 1715 (2020). doi:10.1039/C9CP06561J.

[240] R. Barrett, C. Ortner, and J. Westermayr, *Transferable Machine Learning Potential X-MACE for Excited States Using Integrated DeepSets,* 2025. doi:10.48550/arXiv.2502.12870.

[241] S. Mausenberger, C. Müller, A. Tkatchenko, P. Marquetand, L. González, and J. Westermayr, *SPAINN: Equivariant Message Passing for Excited-State Nonadiabatic Molecular Dynamics,* Chemical Science **15**, 15880 (2024). doi:10.1039/D4SC04164J.

[242] C. Linderälv, N. Österbacka, J. Wiktor, and P. Erhart, *Optical Line Shapes of Color Centers in Solids from Classical Autocorrelation Functions,* npj Computational Materials **11**, 101 (2025). doi:10.1038/s41524-025-01565-x.

[243] D. J. Hsu, D. Leshchev, I. Kosheleva, K. L. Kohlstedt, and L. X. Chen, *Integrating Solvation Shell Structure in Experimentally Driven Molecular Dynamics Using X-Ray Solution Scattering Data,* The Journal of Chemical Physics **152**, 204115 (2020). doi:10.1063/5.0007158.

[244] *List Statistics | TOP500,* https://www.top500.org/statistics/list/.

[245] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, *Array Programming with NumPy,* Nature **585**, 357 (2020). doi:10.1038/s41586-020-2649-2.

[246] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. Van-

derPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, and P. van Mulbregt, *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,* Nature Methods **17**, 261 (2020). doi:10.1038/s41592-019-0686-2.

[247] W. McKinney, *Data Structures for Statistical Computing in Python,* in *Proceedings of the 9th Python in Science Conference*, 56, 2010. doi:10.25080/Majora-92bf1922-00a.