Investigating amorphous graphene and graphene oxide using machine-learned potentials

Zakariya El-Machachi

St Edmund Hall

A thesis submitted to the University of Oxford
for the degree of Doctor of Philosophy
Supervisor: Prof. Volker L. Deringer

Inorganic Chemistry Laboratory

University of Oxford

Declaration

The work presented in this thesis was carried out between October 2021 and July 2025 in the Inorganic Chemistry Laboratory, University of Oxford under the supervision of Prof. Volker Deringer. This dissertation is the result of my own original work, and where it draws on the work of others, this is acknowledged at appropriate points in the text. This dissertation has not been submitted in whole or in part for a degree at this or any other institution.

Zakariya El-Machachi

Abstract

The atomic-level structure of amorphous solids, which lack long-range order, has been a long-standing challenge at the crossroads between chemistry, physics and materials science for over a century. This thesis addresses this fundamental question by employing state-of-the-art atomistic simulation techniques to investigate topological disorder in two-dimensional carbonaceous materials: graphene and graphene oxide (GO).

A central theme is the application of machine learning interatomic potentials (MLIPs) to explore complex potential energy surfaces that are inaccessible to conventional methods. First, a novel approach is introduced where machine-learned atomic energies are used to drive bond-switching Monte–Carlo simulations. This method successfully navigates the configurational space of graphene, providing feasible computational models of both continuous random network and paracrystalline (cybotactic) structures and offering a unique energetic fingerprint to distinguish between them where traditional structural descriptors, such as the radial distribution function, fall short.

Next, this thesis details the development of a bespoke, highly accurate MLIP for graphene oxide (`GO-MACE-23`) using an iterative, on-the-fly training workflow. This potential is used to simulate the thermal reduction of a large-scale GO structural model, revealing the dynamic transformation into reduced graphene oxide (rGO). The resulting rGO structure, which is cybotactic in nature and features pores and embedded oxygen functional groups, shows excellent agreement with experimental observations, a conclusion supported by simulated X-ray photoelectron spectroscopy (XPS) that matches experimental spectra.

Finally, the `GO-MACE-23` potential is applied to investigate the structure-property relationships governing the mechanical behaviour of GO and rGO. Uniaxial strain simulations reveal that the mechanical response is highly dependent on the type and distribution of oxygen functional groups. Structures rich in epoxide groups exhibit a more plastic response, while hydroxyl-rich structures are more brittle, providing key insights for designing functionalised graphene materials with tailored mechanical properties for applications in flexible electronics and composite reinforcement.

Collectively, this work demonstrates the power of machine learning in advancing our fundamental understanding of amorphous materials, providing not only new structural insights but also robust computational tools for the predictive design of next-generation carbon-based technologies.

Acknowledgements

I'm currently perched outside 101 coffee shop in Jericho, Oxford – it's 10 AM on a comfortable Friday morning as the sun swarms over a slumbering Walton Street. Will has gifted me an espresso outshining any I have had in the past four years, so that should lay the foundations for my attempt at writing the ostensibly trivial acknowledgements. Indeed, it is far from trivial. For those of you who know me well, you will be familiar with my tendency to extend beyond verbosity into the realm of the excessive – this section alone could be a thesis in itself – a testament to the many people who have supported me throughout this journey. I apologise for just how long this is going to be.

It's difficult to comprehend this moment, the culmination of years of hard work, late nights, and countless cups of coffee.[*] I am beyond overwhelmed with gratitude, to the point of near paralysis, but there is one person I must start with: my supervisor, mentor, and friend, Volker.

As it turns out, we both started at Oxford back in 2019, you had just started in your new faculty position and I was beginning my Master's course, but it wasn't until the summer of 2020 that our paths would first cross. I joined as a summer student during *that pandemic*, and the group was tiny at that point. I vividly remember us all on those early video calls, back when your long hair resembled that of a certain Abrahamic prophet, discussing projects and where they placed on this abstract looking *triangle*, and of course all SSH'ing into one of those poor Linux desktop machines.[†] But one thing stood out to me then as it does now: your infectious enthusiasm and unwavering belief in our potential as researchers. There have been moments of self-doubt, times when I questioned my abilities and external pressures that made me feel like giving up, but you were always there to patiently guide me, to remind me of my ability, and give me space and time to learn, which has no doubt played a role in shaping the man I am today. Of course, there are the memories from GAP D&U 2022, DPG 2022, EuChemS 2024 and many more conferences that I will cherish, drinking beers at Fat Lizard and listening to your football and champions league analogies, even though you never watched football! Today, I am beyond proud to witness the group grow into the size of a small, semi-autonomous state, and there is no doubt that the group, led by your vision, will go on to achieve great things. I will forever be indebted to you and am truly grateful for your support, guidance, and friendship throughout this journey. From the bottom of my heart, thank you Volker.

Of course, I must also thank Mark, the director of the TMCS CDT which I was a part of, and also co-author on my first publication. Your support and guidance during my time in the CDT were invaluable, and I am grateful for the opportunity to have worked with you. I also want to extend my gratitude to the TMCS CDT team, including all the staff and students who made my time there so enjoyable. I was also fortunate to spend 1 month at BAM visiting Janine and the group, I learnt

---

[*] I have to acknowledge the group coffee machine, that thing is a tank.

[†] Faujasite – gone but never forgotten (2019–???)

a lot about automation and bonding analysis during my overall 3 months in Berlin. Thank you all for your support and encouragement, as well as the wider group.

Similarly I want to thank another mentor, co-author and mad scientist; Miguel for hosting me at Aalto. This was a great experience allowing me to really think creatively about SOAP descriptors. I had an awesome time in Helsinki, where I also got to sing Livin' La Vida Loca with the elder Helsinki locals at Tennka karaoke bar, introduced to me by my good friend, Tigany. Thank you both for making it an unforgettable visit.

I also want to extend my gratitude to all the teaching staff at the University of Southampton, in particular Graeme who introduced me to interatomic potentials on a 2nd year physical chemistry course at Southampton back in 2017. Your course captured my imagination and set me on this trajectory after hosting me for a 3 month summer project working on crystal structure prediction, with Chris who I must also thank for his patience, in the summer of 2017.

The group has evolved over the years, and I am grateful to all past and present colleagues for their support and friendship. I want to thank all of the CG7 crew. In particular, John and Zoe, who have been there since the beginning and have provided much needed trilateral emotional support during the early days. I won't forget that morning when you both looked sheepish as you walked into the office together and gave me that look, confirming what I had begun to suspect! Daniel, you have been a constant, unending source of conversation and you never fail to make me laugh. I have loads of amazing memories including showing you around Berlin and our sim racing evening just yesterday. A big hug goes out to Tom, my older sibling in the group. Although our intersection of interests in music came from opposite ends of the spectrum, you always kept me interested in your music and vice versa. Huge congratulations on the post-doc, looking forward to catching up over a cold beer in Belgium soon.

Moving away from CG7 to CG9 now, I have to start with the 3-point sharpshooter, Yuxing! You mentored me back in 2020 looking at the liquid–liquid phosphours transition, and since then you have been a constant support and a good friend. There are two distinct memories that stand out: the first is when you cooked the most unreal beef dish and brought me some in the office, and the second is when I hit you with my pump fake and layed up an easy basket. Of course, what I never said was that you won by at least double digits... Congratulations on your new fellowship! Shixuan, my neighbour to my left, I'll never forget the hotpot we had at your place, and your extensive knowledge of classic UK comedy shows. Natascia, thank you for always listening to my nonsense, honestly I don't know how you put up with me. Beyond your dodgy lunch choices, the rants and laughs over the years have kept me sane, so thank you for that. Hope you find your West Country voice soon! Louise, where do I start? I couldn't have wished for a better neighbour. I think we would win the heptathlon of procrastination, especially when it comes to guessing how old and where a picture is taken. Thank you for listening to me ramble on about my life, I'm going to miss your seriously good cakes and nattering over a pack of strawberry laces. Chiheb, khoya, your classic cynical takes followed by your trademark self-aware smile always left me laughing when we would chat about all things from Tunisian football to Swiss import taxes – you still made it funny effortlessly. I've even started waving my hands in the same way you do and I think I'm stuck with it now. I'll miss our walks to various food spots around Oxford at lunch time. We have unfinished business in Salford, lmallem gha f call wa7ed.

Over the years I have witnessed the rich, social tapestry of Oxford evolve from the colourful characters that compose it, some of which I have had the privilege of friendship with. I want to start where it all began, a thousand hugs to Harry 'Swanny',[*] Dave 'The Beige', and Jagoda 'Jags' – you took a chance on a random guy off SpareRooms back in 2019 and the rest is history. There are way too many memories to list but I will never forget that dance we had in Bernwood forest to Confusion the Waitress – Underworld and when Jags and I were stuck during lockdown going insane together. At least we had a lovely swimming pool that went a bit green.

Big ups to the 178a Kingston Road crew, Maria for the endless supply of Barry's tea, guitar serenades and good craic; Felix for genuinely helping me speed run the mechanical properties paper (I really don't think I would have finished it without you), the best movie tips and tabouleh for days; and of course, my favourite kombucha brothers Charlie 'Chuck Strangelove' and Lil' Cam. Chuck, you have single-handedly kept the organ musically relevant today with your titanically deep house productions. I am so honoured to have spent the last few years learning music production tips from you, bouncing ideas off each other and sharing meaningful moments on dance floors of all types. Just got to work on your ball control sunshine. Lil' Cam, tricksy and nimble, always willing to provoke the big man before scurrying off only to go and stack it. Man, what can I say. I've taken a lot of inspiration from your wide collection of incredibly headsy tunes – Krautrock to the Proggers, you got it all in the locker. A lot of poignant moments centred around our shared love of music, I'm honoured to call you a mate and I can't wait to see you sending it on the boat. Jess, got to shout you out big time. We got a lot in common you and I. You've been a big part of my life and I want you to know I am grateful for you and your wonky humour and boundless wisdom.

A huge congratulations to the newlyweds, Harrison and Gabriella! I often reflect on that chance encounter at 101, and Harry asked me about my headphones. Soon enough, we were mixing records in my room and my entire perspective on music shifted seismically. I eventually met you Lella and before you knew it, we're all in Morocco getting a good ol' shakedown in the Hamam. Harry, thank you for being a constant source of inspiration, whether its in music to the way you approach life, I have learnt a lot from you. Lella, your kindness and warmth have been a blessing, and I am grateful to have you both in my life. I look forward to many more adventures together, whether it's in the UK or Morocco.

Nathan 'Nay-Nay', pleaaaase! It's been a long and beautiful time together. If only you knew how much I appreciate you. From our days in Southampton to now, we have grown a lot together and I have so many hilarious memories I'm laughing hard typing right now. Probably best left unsaid. Our brotherhood runs deep and

---

[*] Happy Birthday for tomorrow! (26/07)

I cannot imagine doing this DPhil without you. I am so proud of who you have grown into and what you have achieved, I look forward to drinking some more baby Guinness at Florian's.

Bob, my favourite goblin, you are an inspiration to me in everything that you do in your life. It is hard to find people who are as genuine and carry themselves with conviction and intent. Our adventures and projects over the years have fulfilled me with joy and perhaps most importantly, a purpose. My trip to visit you in Japan changed my life and I will never forget the enlightenment and freedom which followed. Not to mention all the gear you have given me and your tips on jamming live have been instrumental. Let's keep making noise, 23 and beyond.

Selina, where do I start with you? From when we first met in the Teddy Hall bar in our first week back in 2019 to now, it's been a journey of one of the deepest bonds I have ever had with anyone. That time we spent in Berlin together at the start of 2024 to now has been one of the most meaningful and transformative experiences of my life. Words fail to describe how much you mean to me, and I am beyond lucky to share this connection with you. My dance partner, music aficionado, clip board and hi-vis organising, fashionista, and most importantly, my best friend and more. I am so grateful for your love and support throughout this journey, and I look forward to many more adventures in Berlin together. Just go easy on the boiled eggs.

I want to acknowledge some in my family who have provided me with unwavering love and support throughout my life. To my two aunties, Leila and Zara 'Agga'; you have always been by my side, always taking care of me and sheltering me from the storms of life. Leila, you've always been the cool aunt, taking us out to Arsenal games and getting tattooed together. You have always looked after me and made sure I was happy. Agga, you've been more than a second mother to me, always cooking the most unreal food and making sure I was comfortable and happy. Your salmon is still unrivalled to this day. To my beautiful grandma, Ghailana 'Mit', I always love hearing your voice and your smile when I kiss your cheek. Your food was always the best, especially your couscous. To my little brother, Hesh, you're the best little brother I could ask for. I think we must of spent thousands and thousands of hours playing each other at various games and you would never win, especially at Tekken. Eventually my old age caught up with me and you took advantage. You have always been chill keeping me grounded and reminding me to take it easy. I am so proud of you and everything you have achieved, and I look forward to seeing you continue to grow and thrive.

Mum, I don't know where to start. You single handedly played the role of both parents, working tirelessly to provide for both me and Hesh, and give us all the opportunities possible, sacrificing so much to ensure we had the chance of a better life. I remember we used to dream of times like these – well here we are. Your love defeated all odds, motivated and inspired me to punch well above my weight, and I am so grateful for everything you have done for me. You are the strongest person I know, and I am so proud to be your son. I may be only half of you by blood, but my life is full because of your love. I love you mum, this is for you.

This thesis is dedicated to the memory of Mohammed 'Ba' Chaba.
To me he was more than my grandfather – an anchor of love, strength and perseverance. A deeply inspirational man who sacrificed everything at such a young age to provide a better life and opportunity for his family – your wisdom will always guide me.
Wherever you may be, I know you have a cat on your lap and a bird on your shoulder.
1946-2020

Publications

The articles below are all first-author publications conducted during the DPhil and form part of this integrated thesis as standalone chapters. They are listed in the order they appear in this thesis:

Chapter 3 Zakariya El-Machachi, Mark Wilson, and Volker L. Deringer, "Exploring the configurational space of amorphous graphene with machine-learned atomic energies", *Chemical Science*, 2022, **13**, 13720.

Chapter 4 Zakariya El-Machachi, Damyan Frantzov, A. Nijamudheen, Tigany Zarrouk, Miguel A. Caro and Volker L. Deringer, "Accelerated First-Principles Exploration of Structure and Reactivity in Graphene Oxide", *Angewandte Chemie International Edition*, 2024, **63**, e202410088.

Chapter 5 Zakariya El-Machachi, Bowen Cheng and Volker L. Deringer, "Mechanical properties of graphene oxide from machine-learning-driven simulations", *Chemical Communications*, 2025, **61**, 11405.

The remaining papers are co-authored publications:

- Tamás K. Stenczel, Zakariya El-Machachi, Guoda Liepuoniute, Joe D. Morrow, Albert P. Bartók, Matt I.J. Probert, Gábor Csányi, Volker L. Deringer, "Machine-learned acceleration for molecular dynamics in CASTEP", *Journal of Chemical Physics* 2023, **159**, 044803
- Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M. Elena, Dávid P. Kovács, ..., Zakariya El-Machachi, ..., Gábor Csányi, "A foundation model for atomistic chemistry", 2023, ArXiv:2401.00096 [physics.chem-ph]
- Chiheb Ben Mahmoud, Zakariya El-Machachi, Krystian A Gierczak, John L. A. Gardner, Volker L. Deringer, "Assessing zero-shot generalisation behaviour in graph-neural-network interatomic potentials", *Digital Discovery* 2025, Accepted Manuscript.

Abbreviations

a-C amorphous Carbon

a-Si amorphous Silicon

a-SiO$_2$ amorphous Silica

ABA Bernal stacking

ABC Rhombohedral stacking

ABOP Analytic bond-order potential

ACE Atomic Cluster Expansion

ACSF Atom-Centred Symmetry Functions

AFM Atomic Force Microscopy

AIMD *Ab Initio* Molecular Dynamics

AIREBO Adaptive Intermolecular Reactive Empirical Bond Order

B3LYP Becke, three-parameter, Lee-Yang-Parr

BO Born-Oppenheimer

C 1s Carbon 1s core level

CC Coupled Cluster

CCSD(T) Coupled Cluster (single and double excitations and a perturbative correction for triples)

CG-MD Coarse-Grained Molecular Dynamics

CI Configuration Interaction

CISD Configuration Interaction (single and double excitations)

COMB Charge Optimised Many Body

CRN Continuous Random Network

DFT Density Functional Theory

DSM Dynamic Structural Model

EDIP Environment-Dependent Interatomic Potential

GAP Gaussian Approximation Potential

GGA Generalised Gradient Approximation

GNN Graph Neural Network

GO Graphene Oxide

GPR Gaussian Process Regression

HF Hartree-Fock

(HR)TEM (High Resolution) Transmission Electron Microscopy

IR Infrared

KS Kohn-Sham

LAMMPS Large-scale Atomic/Molecular Massively Parallel Simulator

LBFGS Limited-memory Broyden–Fletcher–Goldfarb–Shanno

LCBOP Long-range Carbon Bond Order Potential

LDA Local-Density Approximation

LJ Lennard-Jones

MAC Monolayer Amorphous Carbon

MACE Message-passing Atomic Cluster Expansion

MC Monte-Carlo

MD Molecular Dynamics

MEAM Modified Embedded Atom Method

ML Machine Learning

ML-MD Machine Learning Molecular Dynamics

MLIP Machine Learning Interatomic Potential

MP Møller–Plesset

MPNN Message-Passing Neural Network

NAMD Non-Adiabatic Molecular Dynamics

NMR Nuclear Magnetic Resonance

NN Nearest Neighbours

NNN Next-Nearest Neighbours

NNP Neural Network Potential

NNs Neural Networks

NPT Isothermal-isobaric ensemble

NVE Microcanonical ensemble

NVT Canonical ensemble

O/C Oxygen/Carbon ratio

OTF On-The-Fly

PBE Perdew-Burke-Ernzerhof

PES Potential Energy Surface

PVDC Polyvinylidene Chloride

QM Quantum Mechanical

RDF Radial Distribution Function

ReaxFF Reactive Force Field

REBO Reactive Empirical Bond Order

RF Random Forests

rGO reduced Graphene Oxide

RMSE Root Mean Square Error

RSS Reactive State Summation

$\boldsymbol{S(Q)}$ Structure Factor

SCF Self-Consistent Field

SED-AIREBO Screened Environment Dependent AIREBO

SIE Self-Interaction Error

SOAP Smooth Overlap of Atomic Positions

SW Stone-Wales

SW Stillinger-Weber

$\boldsymbol{T_g}$ Glass transition temperature

TEM Transmission Electron Microscopy

XC Exchange-Correlation

XPS X-ray Photoelectron Spectroscopy

XRD X-ray Diffraction

Contents

List of Figures

List of Tables

Chapter 1

Introduction

> *"It must be frankly admitted that we know practically nothing about the atomic arrangement in glasses."* – W. H. Zachariasen (1932)[1]

Order and disorder are fundamental concepts in the study of complex systems. Such systems often exhibit emergent properties, where the collective behaviour of many individual components gives rise to phenomena that cannot be fully understood by examining the components in isolation. For example, in traditional Amazigh folk music (called Gnawa) the cadence of time signatures, vocal scales, percussion and dancing are constantly and dynamically evolving, creating a rich, complex tapestry of sound that is entrancing beyond what each individual component can provide. It is evident that simultaneously, order and disorder emerge from empirical rules and interactions operating at a localised level. "*The whole is different from the sum of its parts*". But what does this *actually* mean? In the context of the physical and natural sciences, we can turn a proverbial (and in some cases a literal) microscope on complex systems by examining the interplay between the micro and macro – a complex dialectic that reveals how local interactions govern the macroscopic features, and in a circular manner, responses to macroscopic perturbations can feed back to influence the local interactions.

This is particularly relevant in the study of disorder in materials, where the absence of long-range order at the microscopic scale does not imply the absence of structure altogether. Instead, emergent properties like mechanical strength, electronic conductivity, or ionic mobility are not entirely dependent on the (semi) local arrangement of atoms, which in turn, determines the extent of long-range order that may or may not be present. Disorder, then, is not just noise or randomness; it is a breaking of symmetry between the micro and macro, where the local arrangement of atoms can lead to a variety of macroscopic properties that are not immediately obvious at the atomic scale.

Amorphous materials (defined as being "non-crystalline") exhibit such behaviour. Yet, the very term *amorphous*[*] lacks a precise definition in the context of materials, where its ambiguity reflects the structural complexity it seeks to describe. In the context of this thesis, it is important to address the difficulty of pinning down what qualifies as amorphous, and to understand the various types of disorder that fall under this umbrella term. Much like the materials themselves, the concept of "amorphicity" defies simple categorisation and invites nuance. All amorphous materials possess a degree of randomness, but the extent to which this occurs and the form it takes span a wide range of possibilities. For example, randomness can

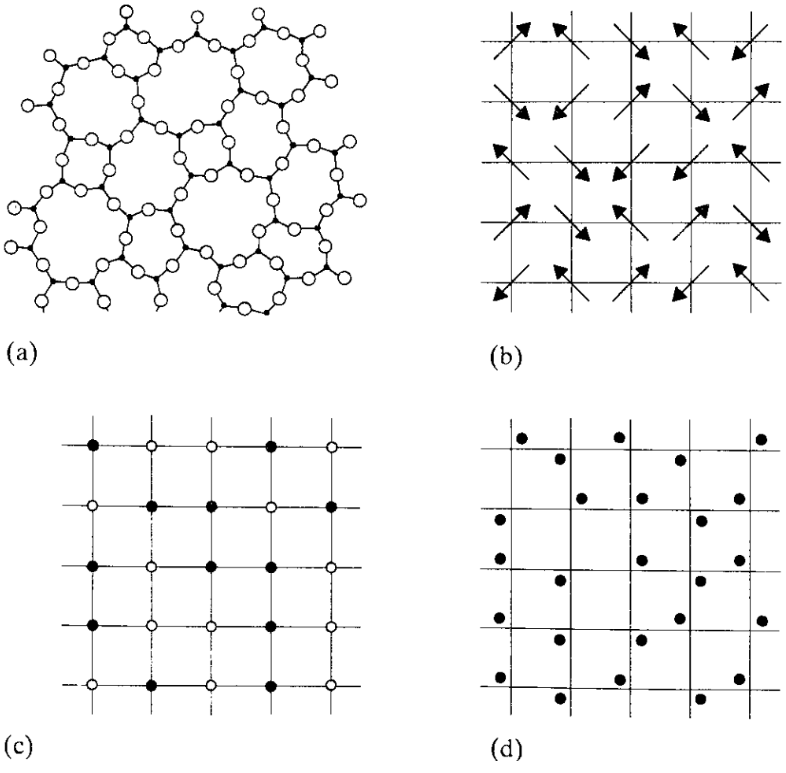

Figure 1.1: Types of disorder. **a** Topological disorder (adapted version from Zachariasen's figure)[1]; **b** spin disorder; **c** substitutional disorder and **d** vibrational disorder. Reproduced from Ref 2 with permission from Elsevier.

---

[*] Derived from Ancient Greek, *a-* (privative alpha) + *morphē* ("form") + *-os*. It can be interpreted as: (1): misshapen, ugly; (2) formless, shapeless. We will focus on the latter, however amorphous gems such as opal and obsidian are revered for their beauty.

be observed in topology, vibrations, spin, substitution (in the case of multi-species systems), defects and so on, and are illustrated in Figure 1.1. More often than not, amorphous materials are disordered in more than one motif. As such, disorder is not unique within itself, but instead must be compared to an ideal form for reference, which is typically a perfect crystalline phase for that material.

Crystalline materials are highly ordered with characteristic symmetries known as space groups, which are constrained due to the periodic repetition of atoms within a unit cell. A perfect crystal can be defined as:

> "*atoms (or groups of atoms or 'motifs') [that] are arranged in a pattern that repeats periodically in three dimensions to an infinite extent*".[2]

With a reference in place, we can begin to frame our understanding of the most important type of disorder for this thesis, topological disorder. Topological disorder is a breaking of translational symmetry, where some amorphous materials display considerable short-range order but lack long-range order. This can be seen in Fig-

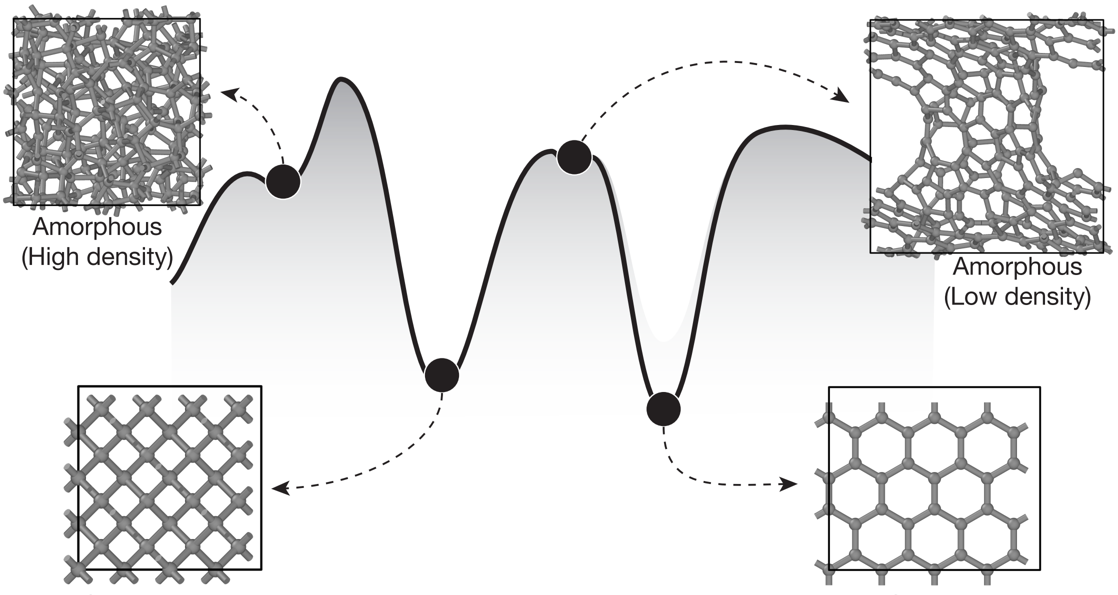

Figure 1.2: Schematic of potential energy surface (PES) for carbon with two local minima showing metastable amorphous states. The system can be trapped in either minimum, leading to different structural arrangements.

ure 1.2 where various carbon structures are shown, two of which are topologically disordered but display short-range ordering as seen by the local bond order corresponding to sp^2–like (low density) and sp^3-like (high density) environments. What is also illustrated is metastability, where graphite is shown to be the most stable form of carbon, as opposed to diamond which is slightly higher in energy. It is the kinetics of the diamond-to-graphite transition which are extremely sluggish, and if left at ambient conditions, will take longer than the age of the universe![*]

And so a question naturally arises, *why are some materials amorphous?* The phase a material takes depends on many factors: thermodynamics, kinetics and bonding, to name a few. According to the third law of thermodynamics, a perfect crystal has zero entropy S at 0 K, and the system must occupy its minimum-energy microstate. In practice this means that at absolute zero, the Gibbs free energy $G = H - TS$ reduces to the enthalpy H (or internal energy E) of the phase as $TS \to 0$. In other words, at 0 K the thermodynamically stable phase is the one with the lowest enthalpy (internal energy). Crystalline order usually maximises favourable bonding and minimises energy, so the lowest-enthalpy crystalline polymorph is typically the ground state.

In contrast, an amorphous solid generally has higher internal energy (and often residual configurational entropy), making it metastable relative to the crystal under equilibrium conditions. There are notable exceptions to the third law of thermodynamics, such as in the case of quantum liquids such as ^{4}He, which does not crystallise at 0 K due to low mass and weak interatomic forces resulting in considerable zero-point motion due to pronounced quantum effects.[4]

However, above absolute zero, statistical mechanics dictates that there is always a non-zero probability of atoms arranging themselves in such a way that the system becomes trapped in a local energy minimum – a metastable state. These configu-

---

[*] It has been estimated from experimentally determined activation energies for {110} faces of diamond that for a diamond of 15 μm ($\approx 1 \times 10^{-11}$g),[3] it would take 10 million years to graphitise at 900 K. Diamonds are not always forever!

rations may lack long-range order, resulting in amorphous structures that retain a degree of local order. For example, in SiO_2 glass,[*] the covalent bonding rules are preserved, but the tetrahedral symmetry is broken. An interesting example where the amorphous phase is thermodynamically more stable than the crystalline constituents is found in certain polymer-derived silicon oxycarbide and related glasses. These materials form a network of Si−O−C bonds with embedded nano-domains of graphene-like carbon and SiO_2-like clusters. Calorimetric and first-principles studies have shown that these amorphous Si−O−C networks have lower enthalpy than the combination of their crystalline constituents (silicon carbide, cristobalite).[5, 6]

The form topological disorder takes is not always a simple matter of energy minimisation, but also depends on the kinetics of the system. For example, if a material is cooled too quickly, it may not have enough time to arrange itself into a crystalline structure, resulting in a "supercooled" liquid for temperatures below the melting point, becoming more viscous as temperature reduces. This is known as the glass transition. The glass transition temperature T_g is loosely defined as the temperature at which the material becomes rigid and loses its ability to flow, and is typically below the melting point of the crystalline phase. It is usually characterised by a gradual change in volume as opposed to instantaneous change in volume as seen in a first-order phase transition during crystallisation. The region where the change of slope occurs is known as T_g, hence the loose definition.[2]

But perhaps more relevant to this thesis is the *atomic* arrangement in amorphous materials. Disorder is an inherently complex phenomenon that can manifest in many different ways, and the atomic arrangement in amorphous materials is no exception. Disorder can be thought of as a spectrum ranging from a completely random arrangement of atoms to highly ordered structures with only slight deviations from perfect order. It is unlikely to find a material with a totally random arrangement of atoms as this would not be energetically favourable owing to under- and

---

[*] A glass (also referred to as vitreous) is defined as an amorphous solid which exhibits a glass transition.[2]

over-coordinated atoms, which would result in an extremely high internal energy for the majority of systems. Instead, most amorphous materials exhibit some degree of short-range order, where atoms are arranged in a way that preserves local bonding rules owing to valency requirements, but crucially, lacks long-range order.

This short-range order is often described in terms of the local coordination environment, bond lengths, and bond angles. While the absence of long-range periodicity means that we cannot use the powerful tools of crystallography to fully describe the structure, we can gain insight through statistical measures. The most fundamental of these is the radial distribution function (RDF), $g(r)$, which describes the probability of finding another atom at a distance r from a reference atom. Experimental techniques such as X-ray or neutron diffraction are used to measure the structure factor $S(Q)$, from which $g(r)$ can be obtained via a Fourier transform. The peaks in $g(r)$ provide information about the average distances to nearest, next-nearest, and subsequent neighbours, confirming the existence of a well-defined local structure.

Beyond the first coordination shell, the structure becomes more complex to describe. The arrangement of local motifs, such as tetrahedra in amorphous silicon (a$-$Si) or silica (a$-$SiO$_2$), constitutes *medium-range order* when considering correlations beyond the 1st neighbour shell. Medium-range order describes correlations on a length scale of approximately 5–20 Å and is crucial for understanding many physical properties. For example, a$-$Si has a larger band gap than crystalline Si due to dangling bonds and coordination defects, which could in theory be tuned if the domain size of medium-range order could be controlled, making it useful for solar cell heterojunctions and thin film transistors.[7–9]

Yet it remains one of the most challenging aspects of characterising amorphous networks. Different models, such as the continuous random network (CRN) model (which will be discussed in the next section), have been proposed to capture the essence of amorphous structures, but a complete and universal description of atomic arrangement in the amorphous state remains an active goal of research.

Understanding the fundamental nature of bonding in amorphous solids has an interesting historical context and originated from a practical curiosity, rooted in the artisanal craft of glassblowing dating back to the ancients who produced glass for ceremonial and artistic expression.[10] What began as a materials challenge in craftsmanship gradually evolved, particularly in the 19th century, into a deeper scientific interest focused on producing optical glass with consistent properties.[11] It wasn't until around 1800 that Guinand and Fraunhofer developed special techniques for adding lead to glass, which dispersed light differently compared to lime-containing glass. Using a spectrometer, they observed the resulting changes in refractive index, effectively kickstarting the scientific exploration of glass as a material.[10] This transformation mirrored the broader shift in science from alchemical and protoscientific traditions to modern atomic theory, as researchers began to move beyond merely making glass to probing its underlying structure and bonding.[12] In this way, the study of amorphous solids transitioned from an empirical art to a fundamental problem at the crossroads between chemistry, physics and materials.

1.1 A Mirror of Division: Competing Theories of Glass Structure in East and West

The history of glass is ancient, that of its scientific exploration is comparatively young. In these formative years at the start of the 19th century, scientists such as Michael Faraday and Moritz Ludwig Frankenheim were speculating on the microscopic structure of various inorganic glasses, but it wasn't until the turn of the 20th century that brought about new technologies such as X-ray crystallography, which enabled a deeper understanding of bonding and structure at the atomic scale. The first X-ray diffraction (XRD) photograph for a glass was taken in 1917,[13] where the diffuse interference pattern was shown for quartz (SiO_2) glass compared with the sharper, more distinct crystalline quartz. The ability to indirectly "see" atomic structure kickstarted what would become a long-standing debate over the next century, with two competing schools of thought emerging: the crystallite theory and the random network theory, the former primarily developed in Soviet Russia, while the latter was championed by Western scientists. This subject has been extensively covered in Ref. 12, and so we will only briefly summarise the key points here.

1.1.1 The Crystallite Theory

Soon after the discovery of X-ray crystallograpy, theories of atomic arrangement in glasses began to emerge during the early 1920s led by A. A. Lebedev, and was based on the notion that glass was comprised of tiny crystals.[14] This theory became known as the "crystallite theory" and was coined to explain discontinuities in the property curves of glass systems. Specifically, Lebedev observed sudden changes in the properties, such as the refractive index, of silicate glasses at temperatures close to the $\alpha \rightleftharpoons \beta$ transition of quartz. However, subsequent XRD data (Figure 1.4) showed no characteristic peaks for both $\alpha-$ and $\beta-$quartz, and as we know today, this is instead associated with the glass transition range found at similar temperatures.[15, 16]

In 1924, Lebedev reverted on his belief that silica glass consisted of tiny embedded quartz crystals.[17] Instead, he suggested a polymorphic transformation between two phases: a high-temperature $\beta-$phase, representing the supercooled liquid, and a low-temperature $\alpha-$phase, which he associated with an ideal glass formed un-

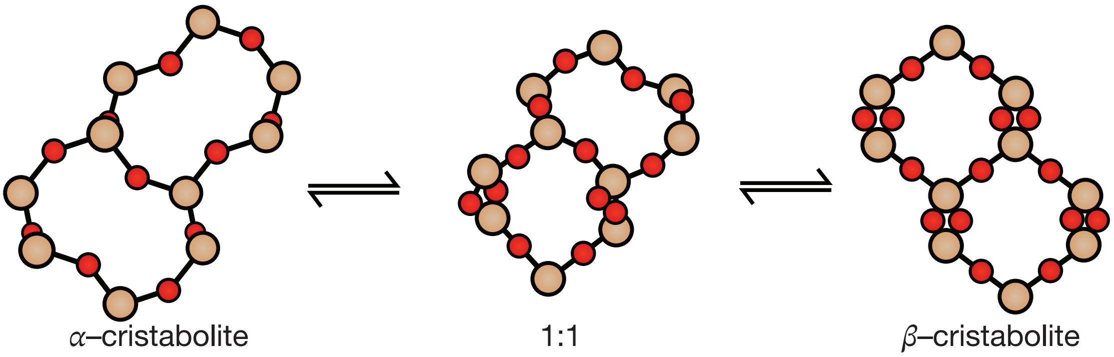

Figure 1.3: Transition of SiO_2 from $\alpha-$cristobalite to $\beta-$cristobalite. Red atoms are O with orange atoms as Si.

der infinitely slow cooling. The extent to which the $\beta-$phase transitioned into the $\alpha-$phase upon quenching was governed by the quench rate and also annealing.

Further XRD data of glassy silica suggested a significant presence of cristobalite crystals (Figure 1.3) of ≈ 15 Å and comprised roughly 80% of the total volume.[18]* This shift in interpretation was prompted by the mismatch between the first broad diffraction peak of glassy silica and the expected position of the corresponding peak in α-quartz (Figure 1.4). The broadness of such diffraction patterns was attributed to the surface having a greater degree of crystallinity compared to the inner bulk, giving a characteristic broader ring with distinct lines at the periphery according to crystalline theory.[19] However, this was difficult to prove conclusively at the time, leading some researchers to reject the crystalline theory outright.[12]

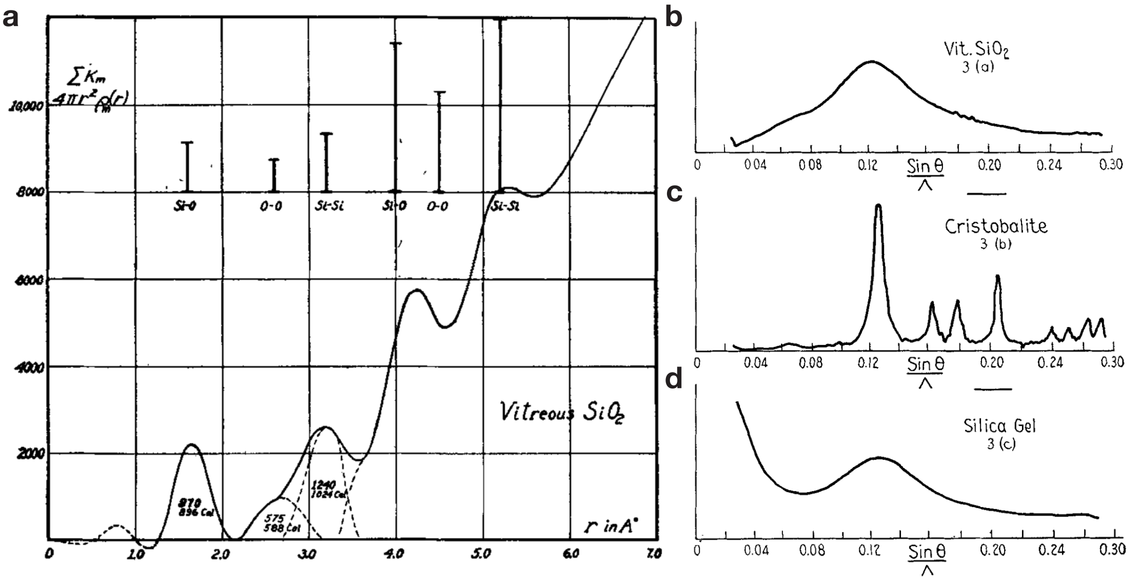

Figure 1.4: X-ray data of SiO_2. **a**) RDF curve of amorphous ("vitreous") SiO_2. **b)** XRD pattern of amorphous ("vitreous") SiO_2, **c)** XRD pattern of $\alpha-$cristobolite. **d)** XRD pattern of dried commercial SiO_2 gel. Reproduced from Ref. 20 with permission from John Wiley and Sons.

---

[*] It is highly likely this was an experimental artefact.[2]

1.1.2 The Continuous Random Network (CRN) Theory

A competing school of thought independently arose during the mid-1920s, with a highly prominent "continuous random network" (CRN) theory proposed by Walter Rosenhain in 1927.[21][*] This model suggested that the atomic arrangement in glasses form a random network of covalently bonded atoms, with no long-range order. W. H. Zachariasen, a prolific American glass scientist, proposed a set of "rules" in 1932 that described the structural criteria for forming oxide glass networks composed of atoms A and O in various stoichiometries:[1]

- **Rule 1:** "an oxygen atom is linked to not more than two atoms A".
- **Rule 2:** "the number of oxygen atoms surrounding atoms A must be small".
- **Rule 3:** "the oxygen polyhedra share corners with each other, not edges or faces".
- **Rule 4:** 'if we require the network to be three-dimensional, at least three corners in each oxygen polyhedron must be shared".

Following these rules, Zachariasen proposed his famous two-dimensional CRN for an AO_2 oxide glass such as SiO_2, as illustrated in Figure 1.5. The structure demonstrates that while atoms obey covalent bonding constraints, the bond angles and ring geometries vary. This enforces the presence of short-range order, at the expense of medium- and long-range order. Thus, the random network theory was a significant departure from the crystallite theory, as it proposed that the atomic arrangement in glasses was not comprised of tiny crystallites, but rather a large lattice which was similar in size to that of a crystal, but lacked symmetry.

A series of subsequent X-ray studies conducted by Bertram Warren Eugene on the structure of various glasses, including SiO_2,[22, 23] GeO_2,[23] BeF_2[24] and

---

[*] The initial proposal of "random network theory" is typically attributed to W. H. Zachariasen in 1932, however, Rosenhain described the features of such a model without explicitly calling it "random network theory" in 1927. It was actually Bertram Eugene Warren who introduced the term in 1933 whilst working on glassy silica.[22]

Na_2O-SiO_2[25] appeared to confirm the theory, offering empirical support for a random network. Warren argued that the position of the first XRD peak for glassy silica suggested a volume expansion of approximately 6% compared to $\alpha-$cristobalite,[26] a crystalline polymorph of silica, which was suggested at the time to make up glassy silica in tiny crystals (*see Section 1.1.1*). In addition, glassy silica displayed no instantaneous volume change between 200 and 300 °C where the cristobalite $\alpha \rightleftharpoons \beta$ transition (Figure 1.3) occurs, which normally results in a volume change of $\approx 3-4\%$.[26, 27] Finally, Warren argued that if a glass was held at high temperature, the tiny cristobalite crystals should grow, resulting in the narrowing of the broad diffraction peak. What was observed instead was an abrupt change to a sharp and narrow diffraction peak indicating $\alpha-$cristobalite.[26]

A major breakthrough came with the application of Fourier transform techniques to X-ray scattering data, enabling the extraction of real-space structural information from reciprocal space measurements. This enabled the projection of the weighted average distribution of neighbouring atoms expressed as a function of interatomic distance, also known as the RDF. The RDF became (and still is) a powerful tool for characterising the local structure of disordered materials, as it provides information

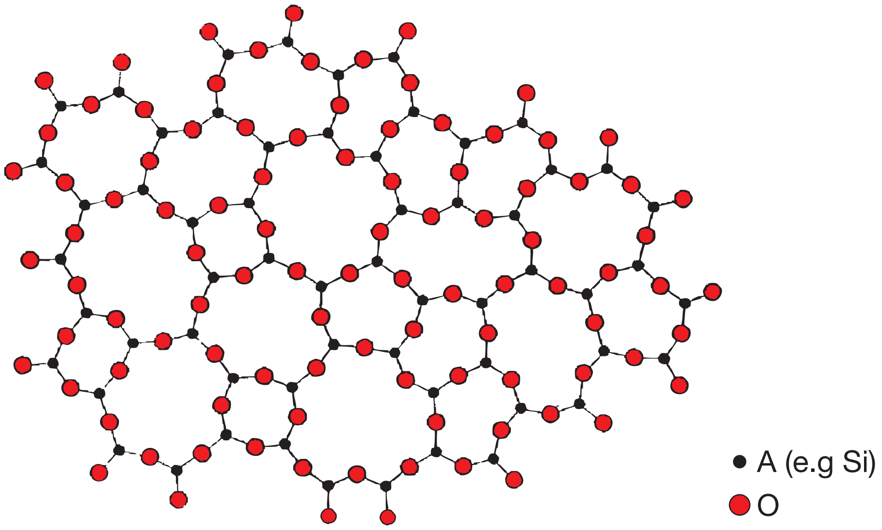

Figure 1.5: Zachariasen's two-dimensional CRN for an AO_2 oxide glass. A atoms are black, O atoms are red. The image has been adapted from Ref. 1.

about the arrangement of atoms in the first coordination shell and beyond, giving signature peaks corresponding to the distances between atoms which are nearest neighbours, next-nearest neighbours and so on. RDFs which are sharp with few peaks are indicative of long-range order, whereas broad peaks with a few distinct features are indicative of short- and medium-range order. Warren *et al.* found the RDF of SiO_2 and B_2O_3 glasses to have broad peaks at shorter distances with a few distinct features before eventually plateauing, which was consistent with the random network theory.[28]

1.1.3 Cybotactic Theory

In 1936, the crystallite theory was revisited in the face of developments from XRD studies at the time. N. N. Valenkov and E. A. Porai-Koshits, two prominent Soviet scientists who were influenced by Lebedev, performed XRD studies on silica and soda-silica (Na_2O-SiO_2) glasses, and observed a gradual change of the diffraction pattern on transition from amorphous to the crystalline phase, rather than a sudden change as proposed by random network theory,[19] further reinforcing the crystallite theory. This was because the crystallite theory predicted a continuous growth of crystallites at the crystallisation temperature, which would lead to a gradual change in the diffraction pattern. In contrast, the random network theory predicted a sudden change in the diffraction pattern at the crystallisation temperature, which was not observed.

Valenkov and Porai-Koshits proposed that every part of the glass is a tiny crystal, but one deformed by internal strain. In their 1936 paper they explicitly wrote:

> "*A hypothesis is made, according to which the crystallites are consisting of an inner part with regularly deformed (stretched or compressed) space lattice, and of outer portions with strong and arbitrary distortion*".[19]

In other words, glass was viewed as essentially entirely crystalline in core regions, with highly distorted grain-boundaries between nanocrystals. With the interpretation of the crystallite theory shifting to larger crystallites, the theory posited by Valenkov, Porai-Koshits and contemporaries suggested that these larger crystallites were embedded in an amorphous matrix which would connect the crystallites. And thus, the crystallite theory split into two schools of thought: the first was still focused on the initial proposal of tiny crystallites, while the second proposed that the crystallites were much larger. It is the latter which later became associated with the "crystallite theory", however it should be noted that there is a distinct difference between the two, and thus we will refer to the latter as the "cybotactic theory".[12]

The cybotactic theory proposed by Valenkov and Porai-Koshits did not oppose the random network theory, but rather suggested that the two theories were not mutually exclusive:

> "*It is evident that the theory of non-crystalline networks is not diametrically opposed to the crystalline theory, since on dissecting this non-crystalline network into separate portions, we obtain an aggregate of regular groups of atoms (ions), each one of these groups presenting an arbitrary distortion of the space lattice. . . . we will be sure that both theories do not differ substantially from one another.*"[19]

The evidence of the diffraction patterns, particularly the presence of broad diffuse halos without sharp Bragg peaks, was interpreted by Valenkov and Porai-Koshits as consistent with a structure composed of highly distorted, nanoscale crystalline regions embedded within a continuous, non-periodic network, thereby supporting a view that integrated aspects of both the cybotactic and random network theories.

The development of the cybotactic theory marked an important step in reconciling the apparent dichotomy between order and disorder in glassy materials. By suggesting that amorphous solids might contain locally ordered regions embedded within a disordered network, this theory provided a conceptual bridge between the statistical randomness of the CRN model and the structural motifs characteristic of crystalline phases. This integrative viewpoint was especially influential for oxide glasses, where experimental evidence from XRD and, later, NMR spectroscopy began to indicate the presence of medium-range order. Nevertheless, significant questions persisted regarding the nature and size of these ordered domains and their influence on bulk properties such as viscosity, mechanical strength, and the glass transition.

1.1.4 So what is the reality?

Despite these developments, the cybotactic theory did not gain as much traction in the West as it did in the Soviet Union. The prevailing view in Western science remained focused on the CRN model,[12] which was seen as a more general and versatile framework for understanding the structure of amorphous materials. By the late 1980s, amorphous materials were widely accepted to be CRNs,[2] but the debate persisted with novel imaging techniques and advanced simulations suggesting the presence of paracrystallites (small regions of localised order resembling the crystalline state), reinforcing the cybotactic theory.[29–31] Recent advanced, large scale simulations of amorphous silicon[9, 32] and synthesis of amorphous diamond[33, 34] demonstrated that changing conditions such as quench rate and pressure could lead to either a CRN or cybotactic structure, thus appearing to end the long-standing debate between the two theories, as both could be observed depending on method of synthesis and processing conditions.

However, the amorphous materials studied throughout these debates had been bulk, 3D materials – the nature of amorphous 2D materials remained an open research question. Investigating 2D materials was appealing as they could be directly imaged using electron microscope techniques as opposed to bulk materials. In addition, 2D networks are conceptually easier to model and simulate, as they are lower dimensional, and thus the number of degrees of freedom is reduced, making them more amenable to computational methods.

As experimental techniques advanced, including the advent of electron microscopy and the rise of computational modelling, research into the structural complexity of

2D amorphous materials accelerated. The following section traces the progression of these ideas and examines their impact on our current understanding of amorphous materials, with a particular focus on carbon, a material that has remained at the forefront of scientific investigation since the origins of modern materials science.

1.2 Disorder in Carbon

As we have seen, the study of amorphous and glassy materials has a rich history, with competing theories and models that have evolved over time. Whilst the theoretical underpinnings were focused on silica and borate glasses, much attention was paid to carbon as well, due to its unique bonding properties, the variety of allotropes it can form, but above all, carbon (in the form of coal) was the fuel by which the Industrial Revolution was powered and the birth of early capitalism emerged. Entire cities, states and even ideologies emerged from this inherently political material, and thus it is no surprise that carbon had been the subject of intense scientific scrutiny since the dawn of "modern science". Whilst the synergistic relationship between carbon technology and political power has always been in a positive feedback loop,[*] today we find ourselves in a world where carbon has become associated as being a "dirty" word, in the context of the current climate crisis. There is an inherently fascinating story to be told about the role of carbon in political and ideological discourse – unfortunately, it is important to note that this is outside of the scope of this thesis. We shall instead focus on the scientific and technological aspects of carbon, particularly in the context of disorder and amorphous carbon, which has been a subject of intense research for over a century.

---

[*] This loop was violently broken during the late 1970s with the implementation of "Monetarism", a neo-liberal economic theory implemented by Margaret Thatcher and Ronald Reagan resulting in the abrupt end of industrial power in the West, leading to widespread de-industrialisation, the decline of trade unionism and significant social unrest, a legacy which bears its scars today.

1.2.1 A Brief History of Amorphous Carbon

The earliest investigations into amorphous carbon date back to 1917 with a notable paper titled "*Über die Konstitution von Graphit und amorpher Kohle*" (translated: "On the Constitution of Graphite and amorphous Coal")[35] when Debye and Scherrer, renowned for their foundational work in XRD and crystallography, studied carbon black, a common form of amorphous carbon used as a pigment and reinforcing agent in rubber, and graphite, a layered honeycomb lattice of carbon atoms which will be further discussed later. The original papers are in German, however these conclusions were translated by Gibson, Holohan and Riley in 1946. The original quote from their paper is as follows:

> "*Debye and Scherrer (Physikal. Z., 1917, **18**, 291) found that such carbons, prepared by a variety of methods, all gave essentially the same type of X-ray powder photograph with the maxima of the diffuse diffractions in the same positions: "Amorphe Kohle ist also Graphit in einer so feinen Verteilung, wie dieselbe durch mechanische Mittel niemals erreicht werden kann"*[36]

(translated: "Amorphous coal is therefore graphite in such a fine distribution that can never be achieved by mechanical means"). Whilst the conclusion is aligned with the crystallite theory, it is important to note that Debye and Scherrer were not aware of the cybotactic theory at the time, and thus their conclusion was based on the notion of tiny crystallites, which was the prevailing view at the time.

This crystallite view of amorphous carbon was further supported by the work of Biscoe and Warren in 1942, who studied the XRD patterns of various carbon materials, including carbon black and graphite.[37] They found that the XRD patterns of these materials were similar to those of graphite, with broad diffuse peaks at the same positions as those of graphite, and if the carbon black samples were heated to high temperature (around 2000°C), the diffraction patterns became sharper, indicating an increase in crystallinity. This process would later become known as graphitisation, where amorphous carbon transforms into large domains of crystalline graphite upon heating, and a material which could transform under such a process would become known as a “graphitising carbon”. As a result, Biscoe and Warren suggested that the carbon black samples had a “turbostratic” structure, which is a type of disordered structure where the graphene layers are misaligned and stacked on top of each other, as illustrated in Figure 1.6.[*] This appears to tend towards the cybotactic theory, except the key difference here in this structural model is the absence of an amorphous matrix connecting these crystallites, which is a key feature of the cybotactic theory.

This work was further extended by Gibson, Holohan and Riley in 1946, who began their paper with the following statement:

> “*... further work has shown that truly amorphous carbon does exist.*”[36]

They thermally annealed various starting materials containing predominantly carbon, with varying compositions of other elements, such as oxygen, nitrogen and

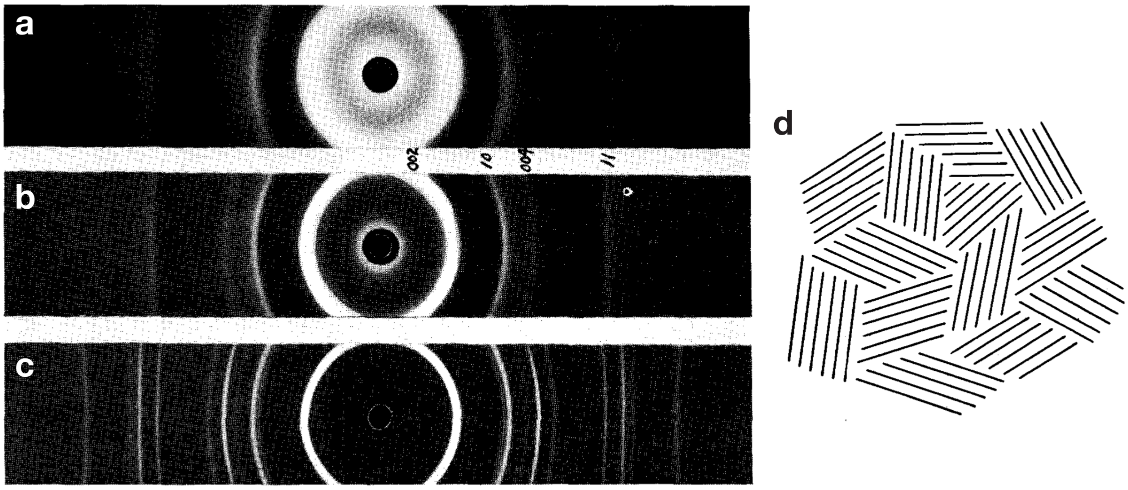

Figure 1.6: XRD patterns. **a** Unheated carbon black, **b** heated carbon black (2 hours at 2000°C), **c** graphite, **d** illustration of turbostratic structural model of carbon black. Reproduced from Ref. 37 with permission from AIP Publishing.

---

[*] Biscoe and Warren suggested that: “*A simple model of a turbostratic structure is readily obtained by shaking and rotating a bag full of penny-sized disks which have been coated with an adhesive. The agitation is suddenly stopped and the adhesive allowed to set. Examination will then show that the disks tend to be arranged in parallel groups with an average of about five parallel disks to a group.*”[37] Sadly, I ran out of coins to test this hypothesis, but I am sure it is a good analogy.

hydrogen, to investigate the structural changes that occurred during the process. What they discovered was that some materials (e.g., cellulose) would not undergo this graphitisation process, and instead would form a new type of carbon material that was amorphous, with no long-range order. They also found that the XRD patterns of these non-graphitising carbons were different from those of graphite, and that the c lattice parameter would change for the graphitising carbons, whereas it would remain constant for the non-graphitising carbons.

Here, the c lattice parameter, to which the 002 peak corresponds, is defined as the distance between the graphene layers (or in the authors' words, "*average height of the minute carbon crystallites*"), and the a lattice parameter (10) can be thought of as the in-plane correlations or as defined by the authors as "*the average diameter*".[36] The results of their work are shown in Figure 1.7, and they measured the lattice parameters of various carbon materials, including cellulose, "pitch extract", "Irish peat" and hexachlorobenzene carbon. The results show that the graphitising carbons have a larger c lattice parameter than the non-graphitising carbons, indicating that they have a larger interlayer spacing, which is consistent with the turbostratic structure proposed by Biscoe and Warren. The authors were surprised when cellulose- and hexachlorobenzene-derived carbon samples exhibited a constant c value across temperatures up to 1200 K and remarked:

> "*The constancy of the c dimension is, however, somewhat surprising and suggests that the layer-planes of the individual crystallites are rigidly cross-linked in some way or other.*"[36]

In Figure 1.7, the pitch extract and Irish peat samples are shown to have an increasing c lattice parameter with increasing annealing temperature, indicating that they are graphitising carbons, whereas the cellulose and hexachlorobenzene carbon samples have a constant c lattice parameter, indicating that they are non-graphitising carbons. The increase in the a lattice parameter was attributed to "*the progressive lamellar growth of the hexagon layer-planes of carbon atoms, brought about by a pro-*

cess of aromatic condensation."[36] The authors were unable to fully explain some of these anomalous results, but they concluded:

> "*It has been found possible to prepare carbons which are completely amorphous, i.e., that give no coherent diffraction of X-rays, and indications have been obtained of three-dimensional, covalent cross-linking in such carbons.*"[36]

1.2.2 To graphitise or not to graphitise?

The work of Gibson, Holohan and Riley was a significant step forward in the understanding of amorphous carbon, as it demonstrated that there are two distinct types of amorphous carbon: graphitising and non-graphitising. The former can be

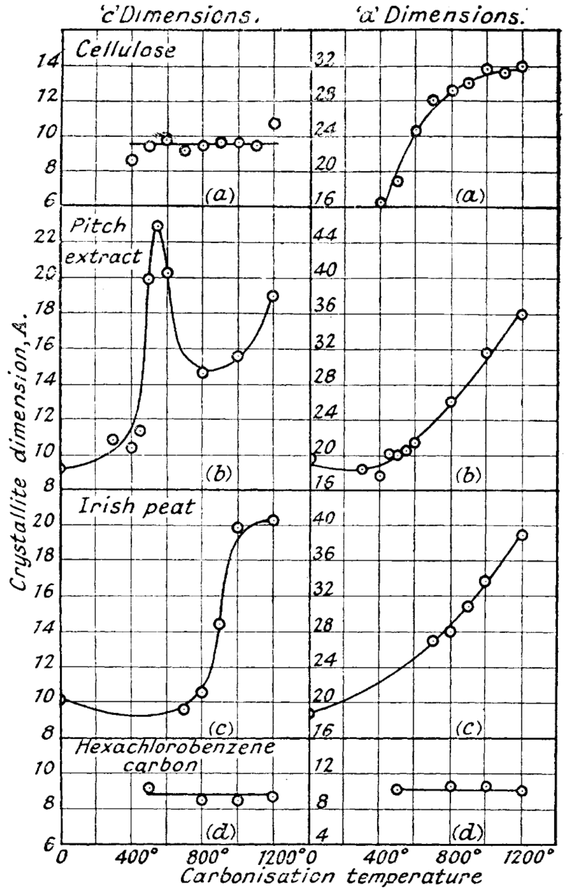

Figure 1.7: Lattice parameters as measured through XRD. **a** Cellulose, **b** Pitch extract, **c** Irish permutation, **d** Hexachlorobenzene carbon. Figure reproduced from Ref. 36 with permission from the Royal Society of Chemistry.

transformed into graphite upon heating, whereas the latter cannot. At the crux of this lies a fundamental question: what determines whether a given carbon material will graphitise or not?

These fundamental questions piqued the interests of a young (and soon to be pioneering) scientist by the name of Rosalind E. Franklin. Her work in unravelling the structure of DNA is well-known, but she also made seminal contributions to the understanding of disordered carbon materials.[38–42] In 1950, Franklin described XRD studies of amorphous carbon prepared by pyrolysis of the polymer polyvinylidene chloride (PVDC) at 1000°C and found that this material was non-graphitsing. 65% of the carbon atoms were found to be in small "graphitic" environments around 1–6 nm in diameter, with the remaining 35% in an amorphous environment.[40]

The graphitising carbon atoms were found to be arranged in a turbostratic structure, similar to that proposed by Biscoe and Warren,[37] whereas the non-graphitising carbon atoms were found to be arranged in a more disordered structure. Prior to this finding, the exact proportion of disorder in amorphous carbon was not known, and Franklin's work provided significant insight into the overall composition of amorphous carbon materials.

Franklin went on to establish the clear distinction between two categories of carbons produced by pyrolysing organic precursors: those capable of transforming

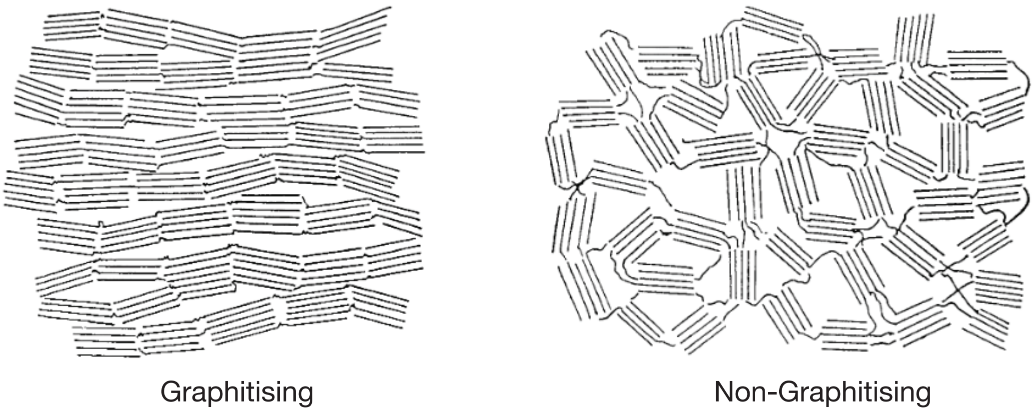

Figure 1.8: Franklin's schematic of graphitising and non-graphitising carbons. Reproduced from Ref. 42 with permission from Harrison and Son.

into graphite upon heat treatment (graphitising carbons) and those that remain disordered even at high temperatures (non-graphitising carbons), terms Franklin introduced. Franklin was the first to use temperatures as high as 3000°C to anneal various precursors where it was expected that the disordered carbons would indeed crystallise into graphite, but instead found that some carbons remained disordered, even at these high temperatures. She concluded that the degree of disorder in amorphous carbon was determined by the precursor material and the conditions under which it was pyrolysed, and that the non-graphitising carbons were likely to be bonded to other elements, such as oxygen or nitrogen, which prevented them from forming the graphitising structure.[42]

As such, Franklin proposed structural models, in which the fundamental units are small graphitic domains containing a few layer planes which are cross-linked. In a graphitising carbon, these units are approximately parallel and the links between them are weak, whereas in non-graphitising carbons, these layers are oriented randomly, with stronger cross-links, however at the time, the precise nature of the cross-links were not known. This can be seen in Figure 1.8, where Franklin's schematic

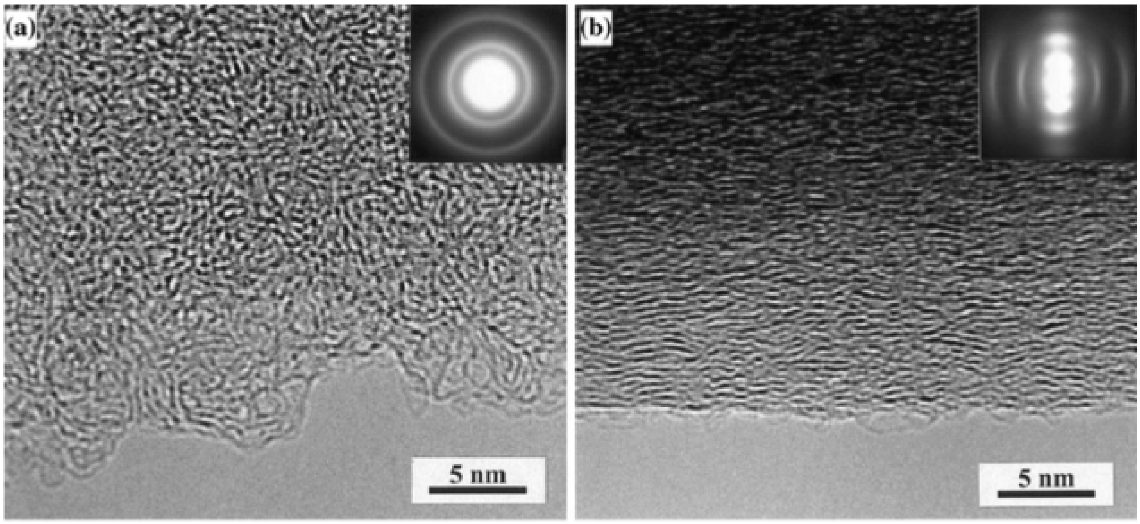

Figure 1.9: Regular TEM images of **(a)** non-graphitising carbon prepared by the pyrolysis of sucrose at 1000°C and **(b)** graphitising carbon prepared by the pyrolysis of anthracene at 1000°C. Insets represent electron diffraction patterns of the respective samples. Reproduced from Ref. 43 with permission from Taylor & Francis.

of graphitising and non-graphitising carbons is shown. The graphitising carbons are shown to have a more ordered structure, with the graphene layers aligned, whereas the non-graphitising carbons are shown to have a more disordered structure, with the graphene layers randomly oriented, resulting in pores and voids in the structure. TEM images of graphitising and non-graphitising carbons are shown in Figure 1.9, which appear to confirm Franklin's structural models.

This binary notion of graphitising and non-graphitising carbons was a remarkable finding for the time, but it has been proved since that the reality is much more complex, with a continuum of disorder. Under high-stress conditions, non-graphitising carbons can be transformed into graphitising carbons,[44, 45] and it was further found that a precursor can be assigned a degree of "graphitisability" to it.[46] It wasn't until the 1960s where the first microscope images of non-graphitising carbon were taken and compared to graphite samples, as well as wider range of physical properties being measured for practical uses in industry,[47] followed by the first electrochemical applications as electrodes owing to the aforementioned pores and voids enabling ion mobility.[48, 49]

Today, the nature of disorder in carbon is still one of the great unanswered questions in materials science, with many researchers still trying to understand the fundamental nature of graphitising and non-graphitising carbons, and in particular the mechanisms involved, which still eludes the community.[45] The work of Franklin and her contemporaries over 70 years ago laid the foundation for this research, and their findings continue to be highly relevant today. Non-graphitising carbons are widely used in a variety of applications, including as electrodes in batteries and supercapacitors, as catalysts in fuel cells, and as reinforcing agents in composites. The ability to control the degree of disorder in carbon materials has led to the development of new materials with tailored properties for specific applications ranging from drug delivery to energy storage.

Thus far we have focused on bulk carbon (extended 3D carbon networks), but it is important to note that carbon can also form one atom thick 2D networks, such as graphene and graphene oxide, which are the focus of this thesis. These materials are of great interest due to their constrained nature enabling theoretical studies to have a simpler "sandbox" to work with. We will now turn our attention to the properties and applications of 2D carbonaceous materials, starting with graphene, followed by graphene oxide.

1.3 2D Carbonaceous Materials

The term "*carbonaceous*" is slightly ambiguous in the sense that it can refer to any material containing or composed of carbon, usually high in carbon content, but not exclusively so. In this thesis, we will use it as an umbrella term to include all-carbon materials, such as graphene, nanotubes, fullerenes, etc., and also materials containing carbon as a major component, such as graphene oxide, hydrocarbons, and polymers for example.

Since the experimental discovery of 0D fullerenes in 1985,[50] 1D carbon nanotubes in 1991,[51] and 2D graphene in 2004,[52] the field of carbon materials has expanded rapidly into a diverse, active research area spanning far and wide. This progress has led to the development of a diverse array of "chemically exotic" carbon allotropes, each exhibiting unique properties and promising a range of potential applications. Research into 2D carbonaceous materials remains highly active, with novel structures continuing to be synthesised, such as the 2D monolayer fullerene network reported in 2022.[53] These exotic allotropes have been widely studied for their mechanical, electrical, thermal, and chemical properties, making them attractive for applications in electronics, energy storage, catalysis, and nanotechnology.

The works of this thesis focus primarily on the atomistic modelling of graphene and graphene oxide, the former being the prototypical ordered 2D material and the latter being the functionalised (disordered) analogue. It must be acknowledged that there are a wide range of 2D materials which are not carbon-based, such as transition metal dichalcogenides, hexagonal boron nitride, transition metal oxides, and black phosphorus which have also been the subject of extensive research, each with their own unique properties and applications.[54] Before we introduce the atomistic simulation methods used to study these materials, a brief overview of graphene and graphene oxide properties, applications, and recent investigations into disorder in these materials is necessary.

1.3.1 Graphene

Graphene is one of a few materials which has entered the cultural *zeitgeist*, gaining mainstream salience for its instantly recognisable "honeycomb" lattice and being heralded as a "wonder material"[55] due to its remarkable properties. This was eventually cemented into history when the Nobel Prize in Physics was awarded to Andre Geim and Konstantin Novoselov in 2010 for their groundbreaking experiments with graphene, involving a pragmatic method of isolating monolayer graphene from graphite using adhesive tape, a technique now known as the "Scotch tape method".[52]

Graphene is a single layer of carbon atoms arranged in a two-dimensional honeycomb lattice, with each carbon atom bonded to three others through strong sp^2

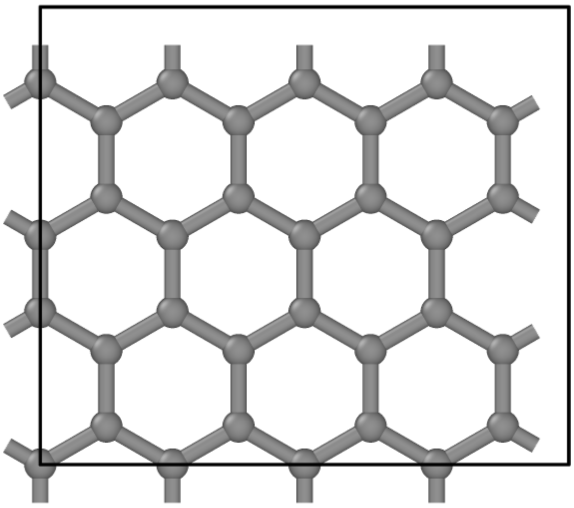

Figure 1.10: Structure of graphene. The term "honeycomb" is used to describe the structure of graphene due to its resemblance to the hexagonal arrangement of honeycomb cells, which is a common motif in nature. This arrangement minimises the amount of used material to reach minimal weight and material cost.

interactions (the C$-$C bond length in graphene is approximately 1.42 Å), resulting in a spare electron from each carbon atom forming an extended π-bonding network, as can be seen in Figure 1.10. This atomic arrangement gives graphene its exceptional in-plane mechanical strength, electrical and thermal conductivity. Graphene is able to form layers through weak π–π interactions, allowing for the formation of graphite, which is composed of many stacked graphene layers, which is the most stable carbon allotrope at ambient conditions.

In 1924, J. D. Bernal[*] proposed the first correct structural model of graphite showing it was composed of stacked graphene layers, where half of the atoms in one layer are directly above those in the layer below and the other half above the holes in the first layer, which became known to be Bernal (ABA) stacking.[56] Another form of stacking, known as rhombohedral (ABC) stacking is also possible, although this is less favourable.

Soon after, attention shifted towards the monolayer, where the electronic properties of graphene were first theoretically explored by Philip R. Wallace in 1947, who applied a tight-binding model to its high-symmetry honeycomb lattice.[57] His work revealed that graphene's band structure features a linear dispersion near the Fermi level, implying that charge carriers (electrons and holes) in graphene behave like massless Dirac fermions, as opposed to the quadratic scaling in most materials. This prediction was experimentally confirmed half a century later following the successful isolation of monolayer graphene and probing of the electric field.[58] The resulting Dirac-like electronic structure underpins many of graphene's remarkable properties, including exceptionally high electron mobility, low effective mass, and the ability to carry extremely high current densities.

---

[*] Bernal was also a prolific writer, authoring the legendary utopian science fiction novel "*The World, the Flesh and the Devil*" in 1929, which was deemed by iconic science fiction writer, Arthur C. Clarke to be "*the most brilliant attempt at scientific prediction ever made*".

1.3.2 Monolayer Amorphous Carbon (MAC)

There is a rich landscape of ongoing research aimed at uncovering new insights into the fundamental nature of graphene when subjected to perturbation from the idealised structure. This includes the study of defects, doping, and functionalisation, which can significantly alter the material's properties and open up new avenues for applications.

As previously discussed, for bulk materials the amorphous nature depends on the precursor material and the conditions under which it was processed, leading to either CRNs or cybotactic structures. But what about 2D materials? With graphene being a single layer of carbon atoms, the question arises: can we have an amorphous monolayer of carbon? The answer is yes, and this material is known as monolayer amorphous carbon (MAC).

Disorder within graphene can manifest in numerous ways, a characteristic that is particularly evident in the material due to its unique ability to host a wide variety of structural defects. One key aspect is topological disorder in graphene, which can arise from various sources such as external irradiation, thermal agitation and harsh synthetic processes. One of the foundational defects in graphene is the Stone–Wales (SW) defect,[59] which involves the rotation of a carbon-carbon bond by 90 degrees, resulting in the formation of a 5- and 7-membered ring pair (Figure 1.11). Such SW defects have been shown to propagate topological disorder in graphene, leading to the formation of larger structural motifs such as pentagons, heptagons, and octagons, which can significantly alter the material's properties. The presence of these defects can lead to localised states in the electronic band structure, which can affect charge transport and optical properties.[61, 62]

The study of MAC is relatively recent, with distinct research into characterising the structural, and for the first time, physical properties emerging in the 2010s, where a graphene sheet was irradiated with an electron beam from a high resolution TEM (HRTEM) to induce disorder and then image the surface.[63, 64] Specialised simulations were also emerging around the same time to study the effects of disorder in graphene through bond switching algorithms[65, 66] which had been popularised in the 1980s to study amorphous silicon.[67][*] These studies revealed that the disorder in graphene could be controlled by varying the irradiation dose, leading to a range of structural configurations from highly disordered to partially ordered structures, however all structures appear to have paracrystallinity resembling the cybotactic theory. One of the variables with both the experiments and simulations is that they can be stopped early to retain crystalline regions and so once again, the process of fabrication and processing conditions play a crucial role in determining the final structure of the material.

In 2017, Joo *et al.* reported the synthesis of a fully CRN-MAC structure using a high temperature process,[69] whilst in 2020, the first free-standing MAC sample was synthesised at lower temperature.[68] Toh *et al.*, used HRTEM to study the structure of MAC and found that depending on the degree of irradiation, it was possible to control the degree of disorder in the material. The TEM images revealed that the more disordered MAC structure was composed of smaller graphitic domains embedded in a matrix of 5-, 7-, and 8-membered rings, whereas the more ordered

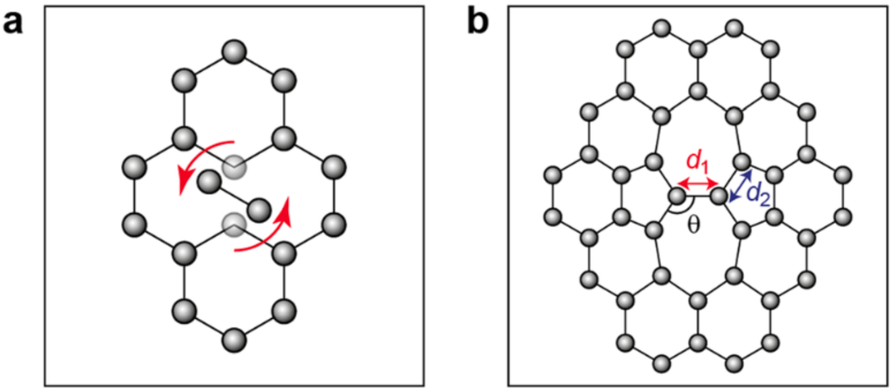

Figure 1.11: Stone–Wales defect in graphene. **a** Schematic of an in-plane single-bond rotation in graphene, leading to an SW defect, **b** A bond rotation by 90° creates two 5-membered and two 7-membered rings. The figure has been adapted from Ref. 60. (CC-BY 3.0)

---

[*] More details can be found in the Monte-Carlo section of Chapter 2.

MAC structure had much larger graphitic domains separated by small regions of 5-, 7-, and 8-membered rings (Figure 1.12). This confirmed that MAC was not a CRN but instead formed "paracrystalline" domains, in agreement with the cybotactic theory.[68] In contrast, Joo *et al.* found that MAC prepared at high temperature closely resembled a CRN,[69] however the resolution of the HRTEM images was not sufficient to resolve the small graphitic domains, and thus it was not possible to determine whether the material was truly amorphous or not. Toh *et al.* make the claim that:

> "*Overall, so far, there exists no solid evidence for the formation of homogeneous, stable, freestanding CRN monolayers, although their existence cannot be ruled out.*"[68]

This statement is still true today, and the question of whether a CRN-MAC can be synthesised remains an open research question. The work of Toh *et al.* was the first to provide solid evidence for the existence of freestanding MAC, and their findings have opened up new avenues for research into the properties and applications of this material. Notably, free-standing MAC is electrically insulating and mechanically robust.[68] Experiments found its resistivity to be comparable to hexagonal boron

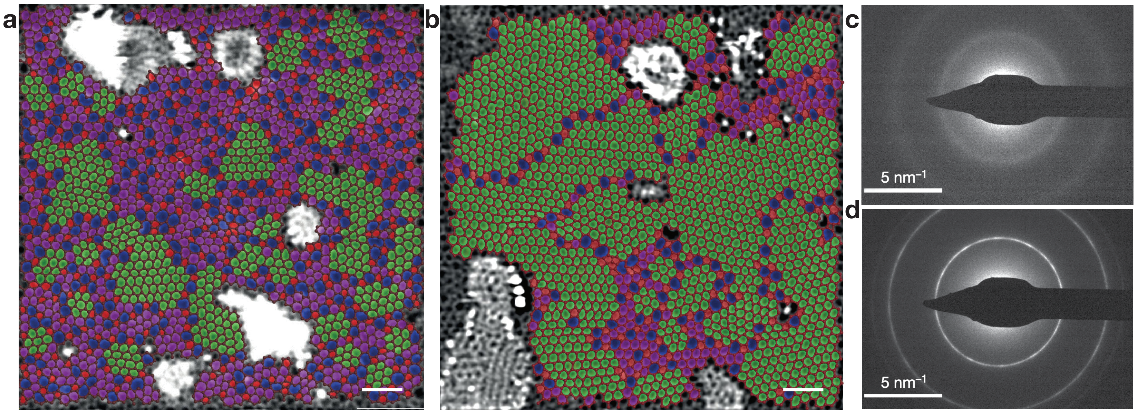

Figure 1.12: Structure of MAC from TEM. Colour overlay shows: pentagons (red), heptagons/octagons (blue), strained hexagons (purple) and crystalline hexagons (green). **a** Scanning TEM image of MAC, **b** Scanning TEM image of nanocrystalline graphene. Scale is 1 nm. **c** Selected area electron diffraction of MAC. **d** Selected area electron diffraction of nanocrystalline graphene. Reproduced from Ref. 68 with permission from Springer Nature.

nitride, reflecting the broken π-network in the amorphous structure.[70] Despite the disorder, MAC is remarkably stable and strong, it can stretch to high strain without brittle fracture. These exceptional properties suggest uses in diffusion barriers or flexible electronics.[68]

Another amorphous 2D carbonaceous material is graphene oxide, a functionalised analogue of graphene, which has been studied for well over 150 years, and is the focus of the next section. Graphene oxide is a fascinating material in its own right, with a rich history and a wide range of applications, but it also serves as a precursor for the synthesis of reduced graphene oxide (rGO), which retains some of the properties of graphene, but with a lower degree of functionalisation. The study of graphene oxide has been instrumental in advancing our understanding of disorder in 2D materials, and we shall explore its structure, properties, and applications in the following section.

1.3.3 Graphene oxide

Graphene oxide (GO) is a functionalised derivative of graphene, characterised by the presence of various oxygen-containing groups situated perpendicular to the basal plane, such as hydroxyl, epoxy, and carboxyl groups, which are introduced during the oxidation process. These functional groups disrupt the sp^2 bonding network of graphene, leading to a more disordered structure and altering its electronic, mechanical, and chemical properties. GO can be produced through several methods, such as Hummers' method, which involves the oxidation of graphite powder using strong oxidising agents such as $NaNO_3$, $KMnO_4$ dissolved in concentrated H_2SO_4.[71]

The origins of GO trace back to the mid-19th century, long before the term graphene was coined. In 1859, British chemist Benjamin C. Brodie reported the first synthesis of what is now recognised as GO (coined 'carbonic acid' in the original paper).[72] Brodie was investigating the action of strong oxidising agents on graphite, seeking to understand its chemical composition and reactivity, and observed the distinct lamellar structure of the newly synthesised material. This reaction functionalised the graphite by introducing oxygen-containing functional groups, namely epoxide and hydroxyl, resulting in a material that was more hydrophilic and dispersible in water compared to pristine graphite.

It was actually research into the graphene oxide which inadvertently led to the discovery of graphene itself. In 1948, Ruess and Vogt used early TEM techniques and observed nm-sized flakes,[73, 74] followed by the work of Boehm and colleagues in 1962, who characterised the structure and properties of GO in detail.[74, 75][*] They found that GO could be reduced to form reduced graphene oxide (rGO), which retained some of the properties of graphene, but with a lower degree of functionalisation and even identified monolayers of the material. Boehm *et al.* went on to coin the name "graphene" in 1986.[76]

Beyond its historical significance in the role it played in the eventual realisation of graphene, GO has garnered considerable attention in recent years due to its unique properties and potential applications. The presence of functional groups in GO enhances its hydrophilicity and dispersibility in aqueous solutions, making it suitable for use in composite materials, as a platform for drug delivery, and as a precursor for the synthesis of reduced graphene oxide (rGO) through chemical or thermal reduction processes. The tuneable electronic properties of GO allow for its integration into various electronic devices, such as field-effect transistors and sensors.

However, there are two questions that remain to be answered: (i) what is the structure of GO and, (ii) what is the relationship between the degree of functionalisation and the resulting properties of GO? The answer to both of these questions are complex and multifaceted, as it depends on a variety of factors, including the method of synthesis, the degree of oxidation, and the specific functional groups present in the material. The structure of GO is still a topic of active research, with various "textbook" models proposed to explain its atomic arrangement.

In 1934, Hofmann *et al.* proposed a model of GO exclusively functionalised with

---

[*] Both papers are originally in German with the historical findings translated in Geim's 2012 graphene retrospective.[74]

epoxy groups in repeating units,[77] followed by Ruess' model in 1947 which introduced hydroxyl groups and a corrugated carbon backbone, typical of sp^3 carbons.[78] In 1969, Scholz and Boehm suggested a structure with alternating, periodic C=C and C−C bonds, notably missing epoxides and replacing them with ketone groups (C=O).[79] It was not until 1998 when the currently most widely accepted model of GO was proposed by Lerf and Klinowski, which introduced a more complex structure with a mixture of hydroxyl, epoxy, and carboxyl groups.[80]

This model has been widely adopted in the literature and is often used as a reference for the structure of GO as it introduced a degree of complexity which had not been explored. The Lerf–Klinowski model suggests that GO consists of a basal plane with hydroxyl and epoxy groups decorated in an amorphous manner perpendicular to the basal plane, while carboxyl groups are located at the edges of the sheets, as shown in Figure 1.13. The structure was characterised through ^{13}C and ^{1}H NMR spectroscopy which unambiguously showed the presence of epoxy and hydroxyl groups which dominated the functionalisation of the structure in conjunction with a significant amount of carboxyl groups at the edges.[80] The conclusions from the model suggest that the carbon backbone of GO consists of two domains: a pristine, crystalline graphene domain and an oxidised sp^3 domain, which is disordered and contains the functional groups, reflecting the cybotactic theory of disorder.[80]

In 2006, Szabo and Dekany revisited the Scholz–Boehm model and made some modifications to the model which suggested a stoichiometric, ordered GO structure.[81] The Szabo–Dekany model represents GO as a periodic ribbon-like structure of aromatic and non-aromatic (cyclohexane-like) carbon atoms. The functionalisation takes the form of hydroxyl and 4-membered ring 1-3 ethers. When C−C bonds are broken, ketones and quinones are suggested to form, which was also a feature of the Scholz–Boehm model. It wasn't until 2010 that the first high-resolution TEM images of GO were taken (Figure 1.14), which revealed the cybotactic structure of GO, confirming the Lerf-Klinowski model.[82, 83] The images showed that the GO sheets were composed of a mixture of crystalline and amorphous regions, with the amorphous regions containing the functional groups, which was consistent with the Lerf-Klinowski model.

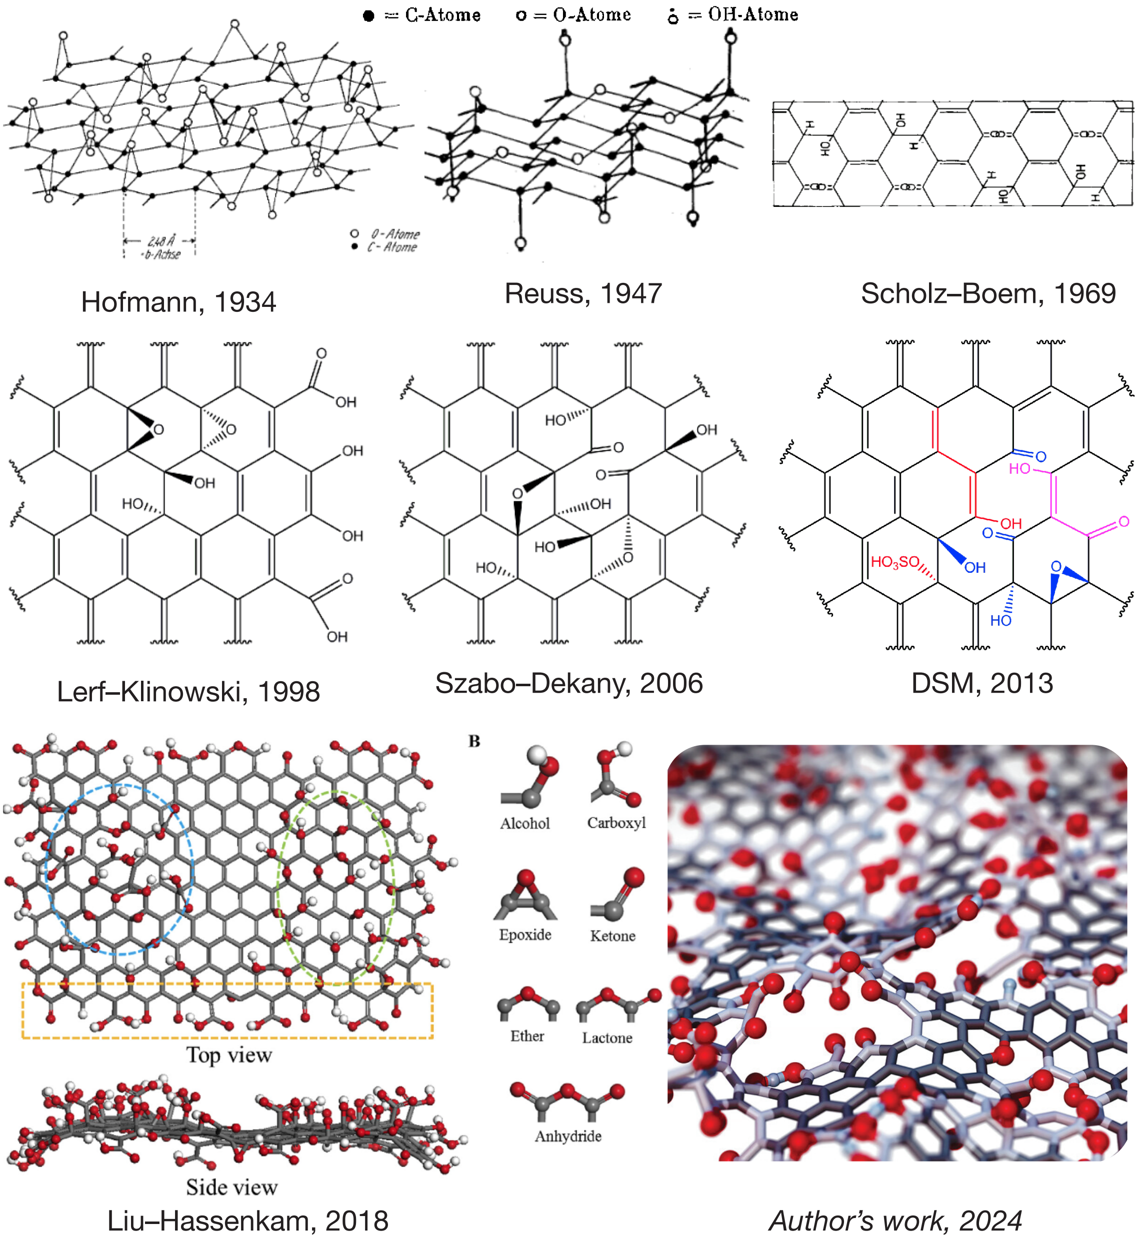

Figure 1.13: Structures of GO. Figures reproduced from Refs. 77–81, 84–87 with permissions from John Wiley and Sons, Springer Nature, and Elsevier.

Despite all of the mounting evidence presented, the structure still cannot be definitively resolved, owing to the fact that HR-TEM cannot determine the functional group type and the fact that HR-TEM images are only able to provide images of smaller areas of the structure and thus cannot provide a full picture of the structure which is statistically significant at the resolution required. There was also the issue of the observed acidity of GO in aqueous solutions which could not be explained by the existing models. In 2013, Dimiev *et al.* proposed a new, "dynamic structural" model (DSM) of GO, which as the name suggests, considers GO as a system, constantly changing its chemical structure due to interaction with water.[84]

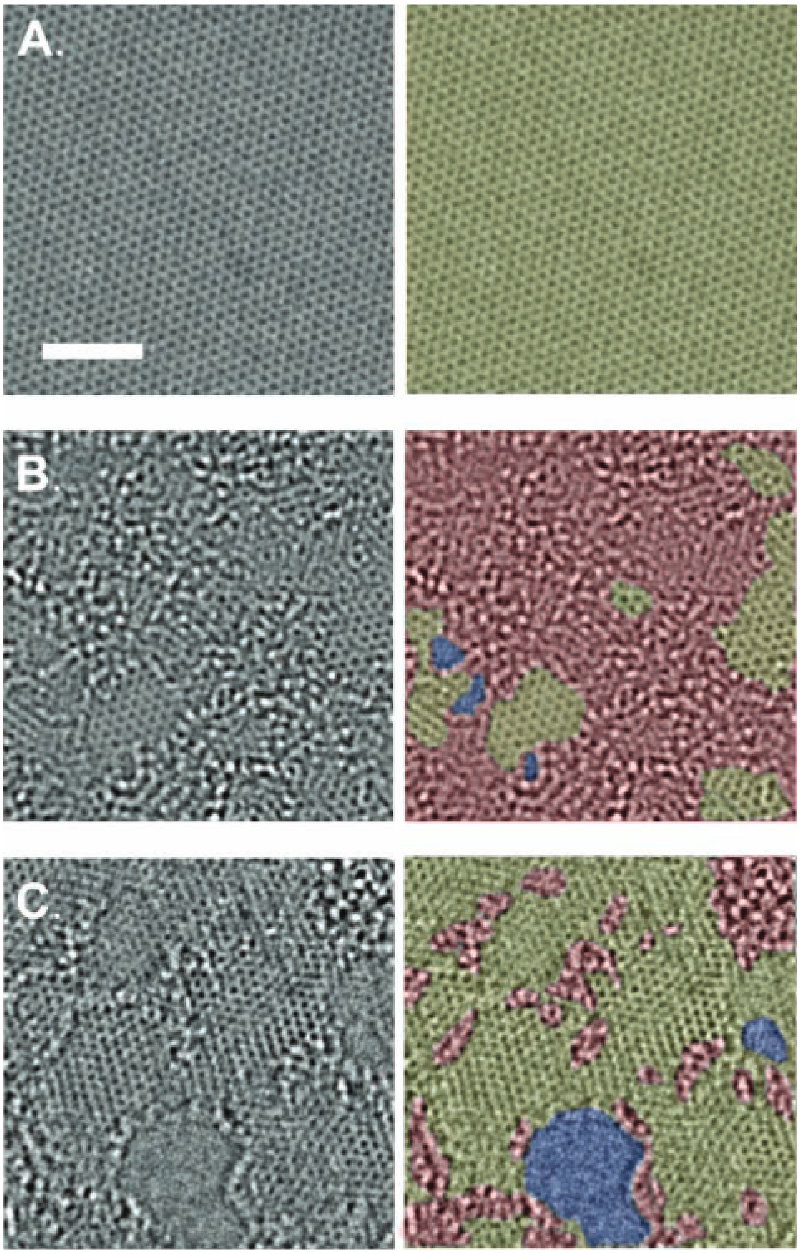

Figure 1.14: TEM images of **a** graphene, **b** GO, and **c** rGO. Graphitic areas are green, amorphous, functionalised regions are red and holes are blue. Scale bar = 2 nm. Figure reproduced from Ref. 83 with permission from John Wiley and Sons.

In 2018, Liu *et al.* suggested there was direct evidence for the observation of carbonyl groups on the edges (aldehydes) and on the basal plane (ketones) of GO, which was a key feature of the Lerf-Klinowski model.[85] The authors used high-resolution atomic force microscopy (AFM) and IR spectroscopy to confirm the presence of epoxy, hydroxyl and also carbonyl groups in GO, which was consistent with the Lerf-Klinowski model, and can be seen in Figure 1.13. And in 2024, colleagues and I published a nanoscale structural model of thermally reduced GO, which also included significant disorder in the carbon backbone, including small pores and holes which has been seen experimentally.[82, 83] This work forms the basis of Chapter 4 and is then applied to the mechanical properties in Chapter 5 of this thesis.

Despite its potential, the presence of disorder in GO can also pose challenges for its practical applications. The disruption of the graphene lattice can lead to a reduction in electrical conductivity and mechanical strength compared to pristine graphene. As such, understanding the relationship between the degree of functionalisation, disorder, and the resulting properties of GO is crucial for optimising its performance in various applications.

1.4 Outline of the Thesis

This thesis is structured as follows. In Chapter 2, I will discuss theoretical and computational methods to study topological disorder in the structure of amorphous materials using atomistic simulation methods, including molecular dynamics and Monte-Carlo. Interatomic potentials will then be explored covering quantum mechanical, empirical and machine learning interatomic potentials. The results Chapters 3, 4 and 5 are self contained, first-author papers written throughout my DPhil, beginning with an exploration of 2D amorphous carbon using Monte-Carlo methods, followed by a study of the structure of graphene oxide using molecular dynamics with machine learning interatomic potentials. The final results chapter turns to an application study of the mechanical properties of graphene oxide and how the local structure and corresponding degree of disorder influences the mechanical properties.

Finally, I will conclude with a summary of the key findings and future directions for research in this area.

Chapter 2

Methodology

This chapter provides a review of the literature and methodologies relevant to the research presented in this thesis. The chapter is divided into several sections, each focusing on a different aspect of the research.

2.1 Atomistic Simulation Methods

While the various experimental methods mentioned in the introduction aim to elucidate the atomic structure through indirect means, atomistic simulations provide a complementary approach to understanding materials at the atomic level *directly*. It is important to understand how representative and accurate structural models of amorphous systems can be made in contrast to highly symmetric crystalline materials which can be expressed through well-defined symmetry operations. Some of the earliest atomistic models of amorphous materials were table-top “ball-and-stick” models, which were meticulously constructed by hand according to Zachariasen’s rules for SiO_2 (Figure 2.1).[88] These models were limited in size and complexity, but they provided a visual representation of the disorder present in these materials.

Beyond hand-built models, computational techniques had advanced significantly during the 1980s, enabling the generation of more realistic and complex models of amorphous materials. This was accomplished through the use of capable atomistic simulations which sufficiently sample the highly complex configurational space of amorphous systems.[90, 91] As mentioned prior, amorphous systems are metastable, meaning that the potential energy surface contains minima which are kinetically stable due to the presence of energy barriers. Thus, traversing the PES of a system of interacting atoms at finite temperature necessitates an approach which samples configurational space thoroughly. These methods can be broadly classified into two categories: molecular dynamics (MD) and Monte-Carlo (MC) simulations (Figure 2.2). MD simulations involve solving the classical equations of motion for a system of interacting atoms. This is achieved by computing the positions and velocities for a set of interacting atoms at the present time step given these quantities one time step prior. In contrast, MC methods attempt to sample configurational space by initialising the atoms in some state followed by proposing random perturbations to the system which are accepted or rejected based on some probability criteria.

Both approaches provide unique ways to sample configurational space. MD is useful to study dynamical properties where temporal averaging over a trajectory can yield information about diffusivity, reactivity, protein folding, etc., but MD suffers from poor sampling of rare events due to low-probability transition states being difficult to traverse. MC methods are highly customisable, where proposed moves need not be physical, enabling a wider coverage of configuration space which is also able to be sampled faster, at the expense of dynamical properties being unavailable.[93]

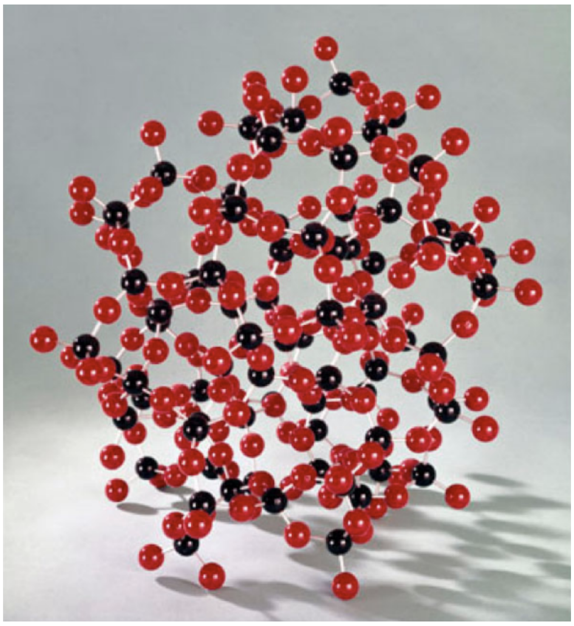

Figure 2.1: A ball-and-stick model of amorphous silica, SiO_2 (red atoms = O, black atoms = Si), constructed according to Zachariasen's rules. Figure reproduced from Ref. [89] with permission from John Wiley and Sons.

2.1.1 Molecular Dynamics

MD approaches result in a trajectory of frames (coordinates of atoms within a periodic cell) which are temporally correlated in configurational space. These frames are connected to each other through an integrator which propagates the system from an initial set of coordinates and velocities according to Newton's equations of motion:

$$F_i = m_i \frac{d^2}{dt^2}\mathbf{R}_i(t), \tag{2.1}$$

where F is the force, i is the atomic index, m is the mass, $\mathbf{R}$ is the Cartesian coordinates and t is time. In order to perform this numerically, time must be discretised to enable an integrator to propagate the system by taking the coordinates and velocities and advancing one time step. This is commonly done using the velocity Verlet algorithm:[94, 95]

$$\mathbf{R}_i(t+\delta t) = \mathbf{R}_i(t) + \mathbf{v}_i(t)\delta t + \frac{1}{2m_i}\mathbf{F}_i(t)\delta t^2, \tag{2.2}$$

$$\mathbf{v}_i(t+\delta t) = \mathbf{v}_i(t) + \frac{1}{2m_i}\left[\mathbf{F}_i(t) + \mathbf{F}_i(t+\delta t)\right]\delta t, \tag{2.3}$$

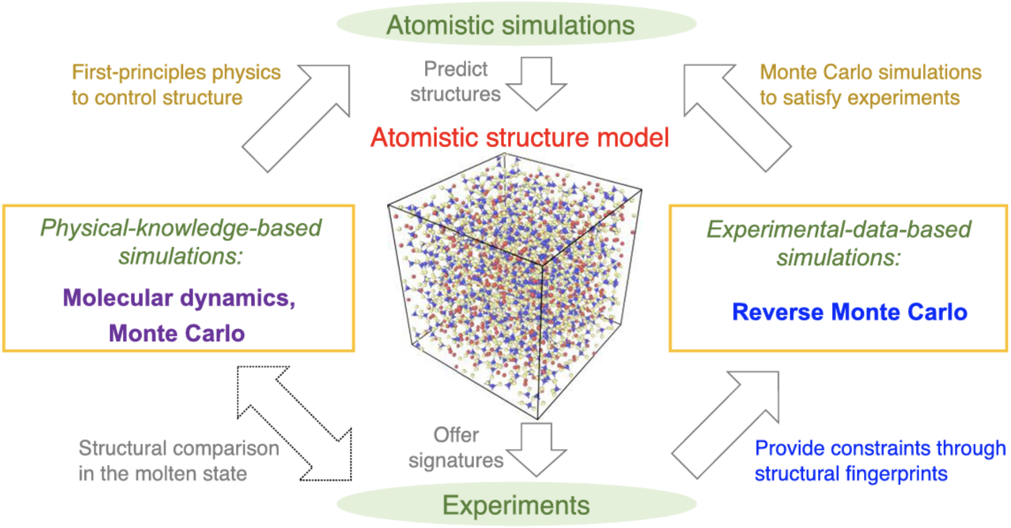

Figure 2.2: Atomistic simulation methods. Figure reproduced from Ref. [92]. (CC-BY-4.0).

where $\mathbf{v}$ is the velocity, δt is the time step, and $\mathbf{F}_i(t+\delta t)$ is the force at the new position, which is computed after updating the positions. The forces are derived from a potential energy function $U(\mathbf{R})$, typically a pairwise potential or a more complex many-body potential, through:

$$\mathbf{F}_i = -\nabla_{\mathbf{R}_i} U(\mathbf{R}), \tag{2.4}$$

where $\nabla_{\mathbf{R}_i}$ denotes the gradient with respect to the coordinates of atom i. This potential energy function drives the MD trajectory by determining the forces acting on each atom based on their positions, and thus is key in defining the dynamics of the system. This will be discussed in more detail in the next section.

The choice of time step δt is crucial: it must be small enough to accurately resolve the system's dynamics – typically on the order of femtoseconds for atomic-scale simulations. Additionally, the integration scheme must ensure both stability and accuracy over long timescales. Time-reversibility is also desirable, as it helps preserve energy and other conserved quantities throughout the simulation.

Since Newton's equations inherently conserve energy rather than temperature, a standard MD simulation naturally samples the microcanonical ensemble (NVE),[93] where the number of particles (N), volume (V), and energy (E) are held constant. An ensemble refers to a collection of microstates that represent the possible configurations of a system under specified macroscopic conditions. Different ensembles correspond to different sets of constraints, such as constant temperature or pressure, which are relevant for various physical scenarios.

To simulate systems at constant temperature requires introducing an external mechanism to regulate thermal fluctuations, typically by coupling the system to a heat reservoir. This thermal control is necessarily an approximation and allows for simulations within the canonical ensemble (NVT), where the number of particles (N), volume (V), and temperature (T) are held constant. Various thermostat schemes exist to achieve this, falling broadly into stochastic and deterministic categories.[96–101] Most influence atomic motion by adjusting velocities, while others, such as the Langevin approach, incorporate both dissipative friction and random noise to emulate the effect of thermal agitation.

In simulations performed within the NVT ensemble, both kinetic and total energies exhibit natural fluctuations.[99, 100] However, not all thermostats capture the correct distribution of these fluctuations, which can lead to inaccuracies, particularly in simulations involving small systems resulting in a phenomena known as "finite system effects".[93]

For systems where volume fluctuations are also significant, such as liquids or materials under pressure, an additional control mechanism is employed. By using a barostat to regulate the pressure, one can perform simulations in the isothermal-isobaric ensemble (NPT), which maintains a constant number of particles (N), pressure (P), and temperature (T).[102] The barostat achieves this by dynamically adjusting the simulation box dimensions in response to the internal pressure, ensuring the system remains at the desired external pressure while allowing for correct thermal and volume fluctuations.

MD has been used to model amorphous materials with great success and can be used to generate structural models of amorphous systems.[9, 32, 103] However, the accuracy of the results is highly dependent on the potential energy function used to describe the interactions between atoms. The choice of potential can significantly affect the resulting structures and properties, especially in complex materials like carbonaceous systems.

Hierarchy of Molecular Dynamics and Enhanced Sampling

MD methods can be categorised into a hierarchy based on the level of accuracy and computational cost associated with the calculation of interatomic forces. This hierarchy spans from highly accurate quantum mechanical methods to more approximate classical and coarse-grained approaches. The main levels in this hierarchy are as follows:

- **Non-Adiabatic Molecular Dynamics (NAMD):** Representing the highest tier of theory, NA-MD calculates the forces on atoms directly from first-principles quantum mechanical (QM) calculations (*see Section 2.2.1*). This approach is used when the Born-Oppenheimer (BO) approximation (the wavefunction of the nuclei and electrons are separable) breaks down, such as in photochemical processes or electron transfer events. These methods allow the system to evolve on and transition between multiple electronic potential energy surfaces (e.g., ground and excited states).[104] However, the computational cost is extremely high, limiting simulations to very small systems (tens of atoms) and short timescales (femtoseconds to picoseconds).

- ***Ab Initio* Molecular Dynamics (AIMD):** At the next level of accuracy, AIMD is also a QM method (*see Section 2.2.1*), except it applies the BO approximation. This approach avoids the need for pre-defined potentials and can accurately model chemical reactions, bond breaking, and electronic effects. However, its immense computational cost limits simulations to small systems (typically a few hundred atoms) and short timescales (picoseconds). The two primary schemes are Born-Oppenheimer MD, which fully converges the electronic ground state at each atomic step to obtain forces,[105] and Car-Parrinello MD, which offers greater efficiency by propagating both the nuclei and electronic orbitals simultaneously within a single dynamical system.[106]

- **Machine Learning Molecular Dynamics (ML-MD):** This approach bridges the gap between AIMD and classical MD. Forces are calculated using a machine learning interatomic potential (MLIP) that has been trained on a large database of QM calculations (*see Section 2.2.3*). ML-MD can achieve near-QM accuracy but at a fraction of the computational cost, enabling simulations of much larger systems (tens of thousands of atoms or more) over longer timescales (nanoseconds). The accuracy of ML-MD is fundamentally dependent on the quality and comprehensiveness of the training data.[107]

- **Empirical Molecular Dynamics:** This is the most common form of MD, where forces are derived from pre-parameterised, analytical functions known as empirical potentials or force fields (as discussed in Section 2.2.2). Empirical MD is computationally efficient, allowing for simulations of millions of atoms over microseconds or longer. Its primary limitation is that the accuracy is fixed by the functional form of the potential, which typically cannot describe chemical reactions unless specifically designed to do so (e.g., reactive force fields) and that these models are often less transferable to different environments.[93]
- **Coarse-Grained Molecular Dynamics (CG-MD):** To access even larger length and time scales, CG-MD simplifies the system by grouping multiple atoms into single interaction sites or "bead". This reduction in degrees of freedom allows for the simulation of massive systems, such as entire proteins or polymers, over milliseconds. The trade-off is a loss of atomistic detail, and the interactions between beads are described by effective potentials that must be carefully parameterised.[108]

In addition to this hierarchy, various enhanced sampling techniques can be combined with these MD methods to overcome the challenge of simulating rare events, such as phase transitions or chemical reactions. Standard MD simulations can become trapped in local energy minima for long periods, making it difficult to sample important but infrequent transitions. Enhanced sampling methods accelerate the exploration of the potential energy surface by introducing a bias.[93] Prominent examples include metadynamics, which adds a history-dependent potential to discourage revisiting old configurations,[109] replica exchange molecular dynamics, which simulates multiple copies of the system at different temperatures to facilitate barrier crossing,[110] and simulated annealing, which is a method for finding low-energy structures by simulating a system at high temperature to overcome energy barriers, followed by a slow, controlled cooling process (also known as quenching) to guide the system into a stable minimum.[93] These techniques are crucial for obtaining statistically meaningful results for processes that occur on timescales far beyond what is accessible with direct simulation.

2.1.2 Monte-Carlo

Monte-Carlo (MC) methods are a class of computational algorithms that rely on repeated random sampling to obtain numerical results. They are particularly useful for systems with a large number of degrees of freedom, where direct analytical solutions are often infeasible. In the context of atomistic simulations, MC methods can be used to sample the configurational space of a system by proposing random moves and accepting or rejecting them based on a probability criterion.[93]

The basic idea behind MC methods is to generate a sequence of configurations (or states) of a system, where each configuration is sampled from a probability distribution that reflects the system's thermodynamic properties. The most common approach is to use the Metropolis algorithm, which involves the following steps:[111]

(i) Start with an initial configuration of the system.

(ii) Propose a random move to a new configuration.

(iii) Calculate the change in energy ΔE between the new and old configurations.

(iv) Accept the new configuration with a probability given by:

$$P = \min\left(1, e^{-\Delta E/k_\mathrm{B}T}\right), \tag{2.5}$$

where k_B is the Boltzmann constant and T is the temperature. If the move is accepted, update the configuration; otherwise, retain the old configuration.

(v) Repeat steps 2–4 for a large number of iterations to sample the configurational space.

In the context of amorphous systems, Monte-Carlo methods have proven especially valuable due to their ability to efficiently explore the vast configurational space characteristic of disordered materials.[112] Unlike crystalline solids, amorphous systems lack long-range order, corresponding to a highly rugged potential energy landscape with numerous local minima. MC techniques, by virtue of their stochastic sampling, can overcome energy barriers that would otherwise trap molecular dynamics trajectories, enabling the generation of representative structural models that capture the inherent disorder of these materials.

Historically, MC simulations have been instrumental in the development and refinement of CRN models, as discussed in Chapter 1. Early work on amorphous silicon and carbon leveraged MC-based bond-switching algorithms to construct large-scale models that reproduce experimental observables such as radial distribution functions and bond angle distributions.[67, 113] These approaches have since been extended and adapted to a variety of amorphous carbon systems, providing insights into their structural motifs, defect populations, and thermodynamic stability.[114–116]

Overall, both MD and MC methods are complementary tools in the modelling of amorphous materials. While MD excels at capturing dynamical processes and relaxation pathways, MC offers flexibility in sampling configurations, making it particularly well-suited for generating and refining models of CRNs and other disordered systems. The interplay between these methods, as well as their integration with advanced interatomic potentials (explored in the following sections), forms the foundation for accurate and predictive simulations of amorphous carbonaceous materials.

2.2 Interatomic Potentials

Whenever phenomena are examined at the atomic level, it becomes necessary to understand the nature of the forces acting between atoms. These interatomic forces play a crucial role in determining many of the behaviours observed in nature, and knowledge of them allows us to modify and utilise natural materials for practical purposes. Describing these forces is inherently complex, as atoms are not fundamental particles but composite systems made up of a nucleus and surrounding electrons that require a quantum mechanical treatment due to low mass. As such, any realistic description of interatomic potentials must account for the intricacies of many-body interactions between many electrons and the nucleus of an atom.

Interatomic potentials are mathematical functions that describe the potential energy of a system of atoms based on their positions (Figure 2.3). They are essential for atomistic simulations, as was alluded to in the previous section, as they provide the total energy of the system and forces acting on each atom, which in turn govern the dynamics of the system. The choice of interatomic potential is crucial for accurately modelling the interactions between atoms, particularly in complex materials like carbonaceous systems. Empirical potentials are parameterised functions that describe the potential energy of a system based on experimental data or high-level quantum mechanical calculations, while quantum mechanical methods provide a first-principles description of the interactions in a system, but are computationally expensive due to the “curse of dimensionality,” and in most cases, impossible to solve analytically for anything other than the hydrogen atom. Thus, it becomes apparent that there is a trade-off between accuracy and computational cost when choosing an interatomic potential, either enabling large-scale simulations at the cost of accuracy, or providing highly accurate descriptions of atomic interactions at the cost of system size and simulation time.

2.2.1 Quantum Mechanics

A quantum mechanical approach provides a true description of the interactions in a system, yet it is prohibitively computationally expensive due to solving the Schrödinger equation (with approximations), limiting the size of systems (or configurations) to a few hundred atoms, which is grossly insufficient to sample the phase space of amorphous systems. Nevertheless, quantum mechanical methods are indispensable for elucidating the fundamental interactions between atoms and electrons, and they establish a computational benchmark (in tandem with experimental data) for the accuracy of interatomic potentials.

In the context of atomistic simulations, quantum mechanical methods are used to calculate the energy and forces of an interacting system of atoms and electrons by solving the Schrödinger equation:[118]

$$i\hbar\frac{\partial}{\partial t}\Psi(\mathbf{r},t) = \hat{H}\Psi(\mathbf{r},t), \tag{2.6}$$

where $\Psi(\mathbf{r},t)$ is the time-dependent wavefunction of the system, $\hbar$ is the reduced Planck's constant, and $\hat{H}$ is the Hamiltonian operator that describes the total energy of the system. The Hamiltonian operator typically includes contributions from the kinetic energy of the particles, potential energy due to interactions between particles, and external potentials (e.g., from nuclei in a molecular system):

$$\hat{H} = -\sum_{i=1}^{N}\frac{\hbar^2}{2m_i}\nabla_i^2 + V(\mathbf{r}), \tag{2.7}$$

where N is the number of particles, m_i is the mass of particle i, ∇_i^2 is the Laplacian operator (and when combined with the prefactors becomes the kinetic energy operator) acting on the coordinates of particle i, and $V(\mathbf{r})$ is the potential energy function

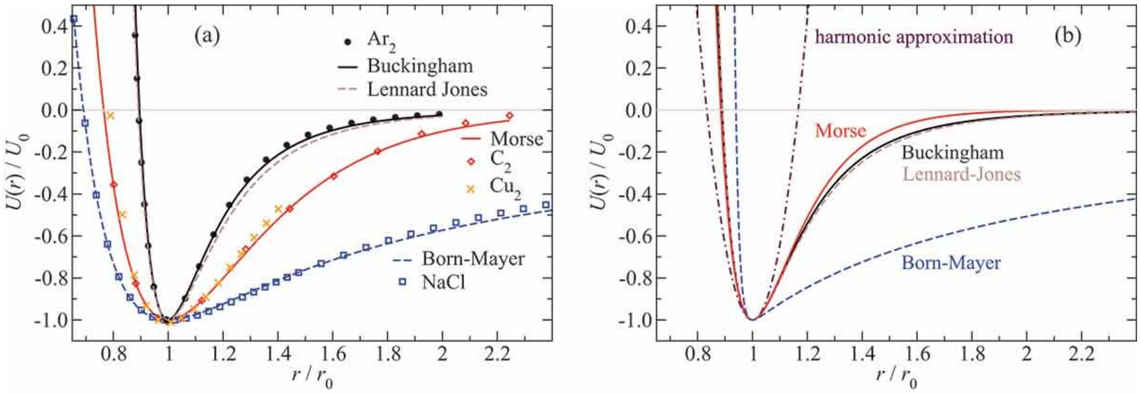

Figure 2.3: **a** Two-body potentials of different diatomic molecules, $U(r)$, normalised to the binding energy U_0 as a function of the bond length r in units of the equilibrium bond length r_0. Symbols reflect high-accuracy reference data, while lines show fits of pair potentials to the energy minimum. Reference data include the Aziz potential for Ar$_2$, and quantum chemical calculations for C$_2$, Cu$_2$, and NaCl. **b** Two-body potentials with a curvature in the minimum corresponding to that of the Lennard–Jones potential. Figure reproduced from Ref. [117]. (CC-BY-4.0).

that describes the interactions between particles. The solution to the Schrödinger equation provides the wavefunction $\Psi(\mathbf{r}, t)$, which contains all the information about the system's quantum state, including its energy and spatial distribution of particles. The time-independent Schrödinger equation is recovered by assuming a stationary state, where the wavefunction can be separated into a spatial part and a time-dependent phase factor:

$$\Psi(\mathbf{r}, t) = \psi(\mathbf{r}) \exp\left(-\frac{iEt}{\hbar}\right), \tag{2.8}$$

where $\psi(\mathbf{r})$ is the spatial wavefunction, E is the total energy of the system and t is time. Substituting this into the time-dependent Schrödinger equation yields the time-independent Schrödinger equation:

$$\hat{H}\psi(\mathbf{r}) = E\psi(\mathbf{r}). \tag{2.9}$$

The interactions of our system are between electrons and nuclei leading to the expanded Hamiltonian:

$$\begin{aligned}\hat{H} = & -\frac{\hbar^2}{2m_{\text{elec}}} \sum_{i=1}^{N} \nabla_i^2 + \sum_{i<j}^{N} \frac{e^2}{4\pi\epsilon_0 |\mathbf{r}_i - \mathbf{r}_j|} - \sum_{i=1}^{N} \sum_{k=1}^{M} \frac{Z_k e^2}{4\pi\epsilon_0 |\mathbf{r}_i - \mathbf{R}_k|} \\ & + \sum_{k<l}^{M} \frac{Z_k Z_l e^2}{4\pi\epsilon_0 |\mathbf{R}_k - \mathbf{R}_l|} - \frac{\hbar^2}{2m_{\text{Nuc}}} \sum_{k=1}^{M} \nabla_k^2,\end{aligned} \tag{2.10}$$

where m_{elec} is the mass of the electron and m_{Nuc} is the mass of the nucleus, N is the number of electrons, M is the number of nuclei, $\mathbf{r}_i$ and $\mathbf{R}_k$ are the positions of the electrons and nuclei, respectively, Z_k is the atomic number of nucleus k, and ϵ_0 is the vacuum permittivity. The first term represents the kinetic energy of the electrons, the second term accounts for electron–electron interactions, the third term describes the interaction between electrons and nuclei, the fourth term accounts for nucleus-nucleus interactions, and the last term represents the kinetic energy of the nuclei.

We can leverage the Born-Oppenheimer approximation[119] to separate the electronic and nuclear degrees of freedom, allowing us to treat the nuclei as fixed points in space while solving for the electronic wavefunction. This approximation is valid when the nuclei are much heavier than the electrons, which is typically the case in systems of interest. Under this approximation, we can rewrite the Hamiltonian as:

$$\hat{H} = \hat{H}_{\text{elec}} + \hat{H}_{\text{Nuc}}, \tag{2.11}$$

where $\hat{H}_{\text{elec}}$ is the electronic Hamiltonian and $\hat{H}_{\text{Nuc}}$ is the nuclear Hamiltonian:

$$\hat{H}_{\text{elec}} = -\frac{\hbar^2}{2m_{\text{elec}}}\sum_{i=1}^{N}\nabla_i^2 + \sum_{i<j}^{N}\frac{e^2}{4\pi\epsilon_0|\mathbf{r}_i - \mathbf{r}_j|} - \sum_{i=1}^{N}\sum_{k=1}^{M}\frac{Z_k e^2}{4\pi\epsilon_0|\mathbf{r}_i - \mathbf{R}_k|}, \tag{2.12}$$

$$\hat{H}_{\text{Nuc}} = -\frac{\hbar^2}{2m_{\text{Nuc}}}\sum_{k=1}^{M}\nabla_k^2 + \sum_{k<l}^{M}\frac{Z_k Z_l e^2}{4\pi\epsilon_0|\mathbf{R}_k - \mathbf{R}_l|}. \tag{2.13}$$

The electronic Schrödinger equation can then be solved for the electronic wavefunction $\psi(\mathbf{r}, \mathbf{R})$ and energy E_{elec}, while treating the nuclei as fixed points in space:

$$\hat{H}_{\text{elec}}\psi(\mathbf{r}, \mathbf{R}) = E_{\text{elec}}\psi(\mathbf{r}, \mathbf{R}), \tag{2.14}$$

where $\psi(\mathbf{r}, \mathbf{R})$ is the electronic wavefunction and E_{elec} is the electronic energy. The solution to this equation yields the PES of the system which can be used to calculate the forces acting on the nuclei:

$$\mathbf{F}_k = -\nabla_{\mathbf{R}_k} E_{\text{elec}}(\mathbf{R}) \tag{2.15}$$

According to the Hellmann–Feynman theorem,[120] these forces can be computed directly from the expectation value of the derivative of the Hamiltonian with respect to the nuclear coordinates, provided the wavefunction is an exact eigenstate of the electronic Hamiltonian. Formally, the force on nucleus k can be written as:

$$\mathbf{F}_k = -\left\langle \psi(\mathbf{r},\mathbf{R}) \middle| \nabla_{\mathbf{R}_k} \hat{H}_{\text{elec}} \middle| \psi(\mathbf{r},\mathbf{R}) \right\rangle . \tag{2.16}$$

This approach allows for the calculation of forces acting on the nuclei based on the electronic structure of the system, enabling the simulation of molecular dynamics or other time-dependent processes. However, solving the electronic Schrödinger equation analytically is impossible for all but the simplest systems like the hydrogen atom, necessitating the use of numerical methods or approximations to obtain the electronic wavefunction and energy.

Hartree–Fock (HF) Theory

The reason why solving the Schrödinger equation is so challenging is due to the many-body nature of the problem. Although the electronic Hamiltonian contains only one- and two-body terms, the full Schrödinger equation remains a many-body problem because the electronic wavefunction depends simultaneously on all electron coordinates, which are coupled through the two-body electron–electron interaction. For a system of N electrons, the wavefunction $\Psi(x_1, x_2, \ldots, x_N)$ depends on the coordinates of all electrons, making it a function in $3N$-dimensional space. Additionally, the wavefunction must satisfy the Pauli exclusion principle,[121] which states that no two identical fermions (such as electrons) can occupy the same quantum state simultaneously. This leads to the requirement that the wavefunction must be antisymmetric under the exchange of any two electrons:

$$\Psi(\ldots, x_i, \ldots, x_j, \ldots) = -\Psi(\ldots, x_j, \ldots, x_i, \ldots), \tag{2.17}$$

where x_i and x_j are the coordinates of electrons i and j.

To replace the many-body interaction between one electron and the remaining electrons, a mean field approximation can be used which approximates these interactions through a spatially-averaged charge distribution summed over those electrons.

It then becomes a simpler task to solve Schrödinger's equation for the motion of the electron within this mean potential. This is otherwise known as the Hartree approximation.[122] There are three interrelated variables in this process: the wavefunction of an N-electron atom, the potential and the charge density. The problem to solve now becomes finding a charge distribution and average potential which are self-consistent, so that this potential and the atomic wavefunction satisfy the Schrödinger equation.

This is accomplished by a method involving successive approximations where an initial wavefunction is chosen and the Hamiltonian of the system setup using this wavefunction. The Schrödinger equation is then solved to find a new wavefunction starting from this Hamiltonian, and if the two functions do not agree to within the specified tolerance the process is recycled starting from the new wavefunction until self-consistency is attained.

It is important for rapid convergence of the self-consistent field method, as for all successive approximation procedures, to choose a reasonable first approximation to the electron wavefunction. One way to form the wavefunction would be simply to consider it as made up of N single-electron wavefunctions (orbitals):

$$\Psi(x_1, x_2, \ldots, x_N) \approx \prod_{i=1}^{N} \psi_i(x_i), \tag{2.18}$$

where ψ_i is the single-electron wavefunction (orbital) for electron i. The orbital is comprised of a spatial part and a spin part, $x_i = (\mathbf{r}_i, s_i)$, where $\mathbf{r}_i$ is the position vector and s_i is the spin coordinate. The spatial part of the atomic orbital is comprised of a set of basis functions:

$$\psi_i(\mathbf{r}) = \sum_{\mu} c_{\mu i} \phi_\mu(\mathbf{r}), \tag{2.19}$$

where $\phi_\mu(\mathbf{r})$ are the basis functions and $c_{\mu i}$ are the expansion coefficients that determine the contribution of each basis function to the molecular orbital ψ_i. The total wavefunction is then constructed as a product of these single-electron wavefunctions. This is known as the Hartree product.[122] This approach effectively describes the motion of each electron in the average field created by all other electrons, leading to a set of coupled equations for the single-electron wavefunctions. Each orbital is solved in the presence of the mean field generated by the other electrons, and the total wavefunction is constructed as a product of these orbitals. The mean field is determined by the charge density (at point $\mathbf{r}$):

$$\rho(\mathbf{r}) = -e \sum_{i=1}^{\text{occ}} |\psi_i(\mathbf{r})|^2, \tag{2.20}$$

where $-e$ is the charge of an electron and $\psi_i(\mathbf{r})$ is the spatial part of the single-electron wavefunction, where the sum runs over occupied orbitals. The potential felt by each electron is then derived from this charge density using Poisson's equation:

$$\nabla^2 V(\mathbf{r}) = -\frac{\rho(\mathbf{r})}{\epsilon_0}, \tag{2.21}$$

which can be written as the Coulomb integral:

$$V(\mathbf{r}) = \frac{1}{4\pi\epsilon_0} \int d^3\mathbf{r}' \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}. \tag{2.22}$$

Substituting this potential into the single-electron Schrödinger equation gives the explicit Hartree equation:

$$\left(-\frac{\hbar^2}{2m}\nabla^2 - \frac{Ze^2}{4\pi\epsilon_0|\mathbf{r}|} - eV(\mathbf{r}) \right) \psi_i(\mathbf{r}) = \epsilon_i \psi_i(\mathbf{r}), \tag{2.23}$$

$$\left(-\frac{\hbar^2}{2m}\nabla^2 - \frac{Ze^2}{4\pi\epsilon_0|\mathbf{r}|} + \frac{e^2}{4\pi\epsilon_0} \sum_j \int d^3\mathbf{r}' \frac{|\psi_j(\mathbf{r}')|^2}{|\mathbf{r} - \mathbf{r}'|} \right) \psi_i(\mathbf{r}) = \epsilon_i \psi_i(\mathbf{r}), \tag{2.24}$$

where ϵ_i is the energy eigenvalue associated with the orbital ψ_i. Here, the first term represents the kinetic energy of the electron, the second term accounts for the attraction between the electron and the nucleus with charge $+Ze$, and the third term represents the repulsive interaction between the electron in orbital ψ_i and the average charge distribution of all other electrons. In the original Hartree formulation the Coulomb sum runs over all orbitals and therefore each orbital contributes to the potential felt by every electron, including itself, leading to unphysical self-interactions.[122] If one wishes to remove that unphysical contribution by hand the Coulomb sum can be written as $\sum_{j \neq i}$.

This equation must be solved self-consistently, as the potential $V(\mathbf{r})$ depends on the charge density, which in turn depends on the wavefunctions $\psi_i(\mathbf{r})$. The process involves iterating between solving for the wavefunctions and updating the potential until convergence is achieved. The variational principle ensures that the energy calculated using an approximate wavefunction will always be greater than or equal to the exact ground-state energy of the system.[123] Consequently, the coefficients associated with the basis functions that form the molecular orbitals can be adjusted iteratively to minimise the total energy. This process is repeated until the input and output electron densities are consistent with one another, a condition known as self-consistency, which gives rise to the term Self-Consistent Field (SCF) method.[122]

Whilst this appears to be a reasonable approximation to the many-body problem, the Hartree SCF method is the result of simple, one-electron wavefunctions which do not fulfill the quantum mechanical requirement of antisymmetry in all electrons, and thus electron exchange is neglected. This means that the ground-state prediction of the wavefunction is fundamentally non-physical for fermions. To satisfy this requirement the wavefunction must be antisymmetric. Hartree-Fock (HF) theory accomplishes this by constructing the many-electron wavefunction not as a simple product, but as a Slater determinant:[124, 125]

$$\Psi(x_1, x_2, \ldots, x_N) \approx \frac{1}{\sqrt{N!}} \begin{vmatrix} \psi_1(x_1) & \psi_1(x_2) & \cdots & \psi_1(x_N) \\ \psi_2(x_1) & \psi_2(x_2) & \cdots & \psi_2(x_N) \\ \vdots & \vdots & \ddots & \vdots \\ \psi_N(x_1) & \psi_N(x_2) & \cdots & \psi_N(x_N) \end{vmatrix}, \tag{2.25}$$

The determinant structure ensures that swapping any two electrons (columns) results in a change of sign of the determinant, thus satisfying the Pauli exclusion principle. This introduces a new term into the energy calculation, the exchange term, which accounts for the antisymmetry requirement. The exchange interaction is a quantum mechanical effect that only occurs between identical particles. Despite sometimes being called an exchange force in analogy to classical force, it is not a true force, as it lacks a force carrier. The effect is due to the wavefunction of indistinguishable particles being subject to exchange symmetry, that is, either remaining unchanged (symmetric) or changing its sign (antisymmetric) when two particles are exchanged.

As with the Hartree method, the goal is to find a set of spin-orbitals $\{\psi_i\}$ that minimise the total energy of the system calculated as the expectation value $E_{\mathrm{HF}} = \langle \varPsi | \hat{H} | \varPsi \rangle$.[*] When the Slater determinant is used with the full electronic Hamiltonian $\hat{H}$, the orbital energies ϵ_i can be expressed as:

$$\epsilon_i = \langle \psi_i | \hat{h}_e | \psi_i \rangle \tag{2.26}$$

$$= \langle \psi_i | \hat{T} + \hat{V} | \psi_i \rangle + \sum_{j(\mathrm{occ})} \langle \psi_i | \hat{J}_j - \hat{K}_j | \psi_i \rangle \tag{2.27}$$

$$= \langle \psi_i | \hat{T} + \hat{V} | \psi_i \rangle + \sum_{j(\mathrm{occ})} (J_{i,j} - K_{i,j}), \tag{2.28}$$

where $\hat{h}_e$ is the one-electron Hamiltonian operator containing $\hat{T}$, the kinetic energy operator ($-\frac{\hbar^2}{2m}\nabla^2$), and $\hat{V}$, the potential energy operator (electron-nucleus attraction). The second term is a sum which runs over all occupied orbitals, and contains two operators: the Coulomb operator $\hat{J}$ and the exchange operator $\hat{K}$, where the indicies reflect the matrix elements. These operators represent the average repulsive

---

[*] The expectation value of an operator returns the average value of a physical quantity represented by that operator when measured on a system described by a given wavefunction. In this case the energy E is the expectation value of the Hamiltonian operator $\hat{H}$ with respect to the wavefunction $\varPsi$. Bra-ket notation is used here for brevity, where $\langle \varPsi |$ is the complex conjugate transpose of the wavefunction $| \varPsi \rangle$. The explicit integral form is given by $\langle \varPsi | \hat{H} | \varPsi \rangle = \int d^3\mathbf{r}_1 \int d^3\mathbf{r}_2 \ldots \int d^3\mathbf{r}_N \varPsi^*(\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N) \hat{H} \varPsi(\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N)$.

interaction between electron i and the charge distribution of electron j (Coulomb), and the quantum mechanical exchange interaction due to the antisymmetry of the wavefunction (exchange). Their explicit forms are:

$$\hat{J}_{i,j} = \langle \psi_i | \hat{J}_j | \psi_i \rangle = \frac{e^2}{4\pi\epsilon_0} \int d^3\mathbf{r} \int d^3\mathbf{r}' \frac{|\psi_i(\mathbf{r})|^2 |\psi_j(\mathbf{r}')|^2}{|\mathbf{r} - \mathbf{r}'|}, \tag{2.29}$$

$$\hat{K}_{i,j} = \langle \psi_i | \hat{K}_j | \psi_i \rangle = \frac{e^2}{4\pi\epsilon_0} \int d^3\mathbf{r} \int d^3\mathbf{r}' \frac{\psi_i^*(\mathbf{r})\psi_j^*(\mathbf{r}')\psi_i(\mathbf{r}')\psi_j(\mathbf{r})}{|\mathbf{r} - \mathbf{r}'|}. \tag{2.30}$$

If ψ_i is an occupied spin-orbital, the $j = i$ term leads to $J_{i,i} = K_{i,i}$, and so it disappears in the sum and the remaining terms represent the Coulomb minus exchange interaction between electron i and $N-1$ other occupied spin-orbitals. The Coulomb operator $\hat{J}_i$ still contains an unphysical self-interaction, but the exchange operator $\hat{K}_i$ has the identical form ($J_{ii} = K_{ii}$). Therefore, the self-interaction is implicitly cancelled ($\hat{J}_i - \hat{K}_i = 0$).

It is convenient to introduce the Fock operator $\hat{F}$, which collects the one-electron Hamiltonian and the mean two-electron contributions:

$$\hat{F} = \hat{h}_e + \sum_{j(\mathrm{occ})} (\hat{J}_j - \hat{K}_j), \tag{2.31}$$

allowing the Hartree-Fock equations to be written in a compact form:

$$\hat{F}\psi_i = \epsilon_i \psi_i. \tag{2.32}$$

These equations must be solved self-consistently, as the Fock operator $\hat{F}$ depends on the occupied orbitals $\{\psi_j\}$ through the Coulomb and exchange operators. The SCF procedure involves iterating between solving for the orbitals and updating the Fock operator until convergence is achieved. The total HF-SCF energy can be written as:

$$E_{\mathrm{HF}} = \sum_{i(\mathrm{occ})} \langle \psi_i | \hat{h}_e | \psi_i \rangle + \sum_{i>j(\mathrm{occ})} (J_{ij} - K_{ij}), \tag{2.33}$$

and the sum of the orbital energies of the occupied spin-orbitals is given by:

$$\sum_{i(\text{occ})} \epsilon_i = \sum_{i(\text{occ})} \langle \psi_i | \hat{h}_e | \psi_i \rangle + \sum_{i,j(\text{occ})} (J_{ij} - K_{ij}). \tag{2.34}$$

A comparison of these two expressions reveals a crucial distinction in Hartree-Fock theory. The total energy, E_{HF}, is not simply the sum of the occupied orbital energies. This is because summing the orbital energies, $\sum_{i(\text{occ})} \epsilon_i$, double-counts the electron-electron interaction terms (J_{ij} and K_{ij}). To obtain the correct total energy, this double-counted interaction energy must be subtracted from the sum of the orbital energies. This highlights that the total energy cannot be naively constructed from orbital energies alone; the inter-electronic interactions must be handled carefully to avoid double-counting.

Post Hartree–Fock Methods

While HF theory is a foundational method that captures the essential physics of electron-nucleus attraction and average electron-electron repulsion, its mean-field approximation introduces significant limitations. The primary source of error is the neglect of electron correlation. In reality, the motion of electrons is correlated; they instantaneously repel each other and actively avoid being close together. HF theory, however, only treats this repulsion in an averaged, or mean-field, sense. An electron in the HF model does not feel the instantaneous position of other electrons, only their time-averaged charge cloud. The difference between the exact non-relativistic ground-state energy, E_{exact}, and the energy obtained from the HF approximation in the limit of a complete basis set, E_{HF}, is defined as the correlation energy:[126]

$$E_{\text{corr}} = E_{\text{exact}} - E_{\text{HF}}. \tag{2.35}$$

By convention, E_{corr} is negative, as the HF energy is an upper bound to the true ground-state energy.

This neglected correlation can be broadly categorised into two types: dynamic correlation and static correlation.[123] Dynamic correlation refers to the short-range interactions where electrons adjust their motion to avoid one another due to Coulombic repulsion.[127] A critical failure stemming from the lack of dynamic correlation is the inability of HF theory to describe dispersion forces (e.g., van der Waals interactions).[128] Dispersion arises from the correlated fluctuations of electron clouds in interacting molecules. Since HF treats these clouds as static, average distributions, it cannot capture this purely correlation-driven attraction, rendering it unsuitable for studying non-covalently bonded systems.

The other type is static (or strong) correlation. This type of correlation becomes important when a single Slater determinant is a fundamentally poor description of the electronic state.[127] This often occurs in systems with nearly degenerate orbitals, such as during bond breaking, in transition metal complexes, or for certain excited states.[129] In these cases, the ground state is better described as a linear combination of several determinants, and the single-determinant HF approximation fails qualitatively.

Relativity is also not accounted for in standard HF theory, which can lead to significant errors in systems containing heavy elements where relativistic effects become pronounced.[130] These effects include spin-orbit coupling and scalar relativistic corrections, which can alter the electronic structure and properties of molecules. While relativistic extensions of HF theory exist, they add complexity and computational cost.

The inherent limitations of the HF approximation necessitate the development of more sophisticated *ab initio* methods that can systematically recover the missing electron correlation energy. These methods are collectively known as post-Hartree–Fock methods. They typically use the HF solution, that is the single Slater determinant and the set of molecular orbitals, as a starting point or a zeroth-order reference. From this reference, they build a more accurate description of the many-electron wavefunction, often by considering it as a linear combination of multiple

Slater determinants. By including configurations that represent the instantaneous positions of electrons relative to one another, these methods can account for dynamic and/or static correlation. Several prominent post-HF approaches exist, such as Møller–Plesset (MP) perturbation theory,[131] Configuration Interaction (CI),[132] and Coupled Cluster (CC) theory,[133] each offering a different strategy for capturing electron correlation and achieving higher accuracy.

CI improves the wavefunction by expressing it as a linear combination of the HF ground-state determinant and determinants representing single, double, and higher electronic excitations. The method is variational, meaning the calculated energy is always an upper bound to the true ground-state energy. However, for practical reasons, the expansion must be truncated, typically at single and double excitations (CISD). This truncation causes a critical failure: the method is not size-consistent. This means the energy of two non-interacting molecules calculated together is not equal to the sum of their energies calculated separately, making it unreliable for many chemical applications.[123]

MP takes a different route. Instead of directly improving the wavefunction, it treats electron correlation as a perturbation to the HF Hamiltonian and calculates the correlation energy as a series of successive corrections. The most common level, MP2, is computationally efficient (relative to other post-HF methods scaling as $\mathcal{O}(N^5)$) and provides a good first estimate of dynamic correlation. Unlike truncated CI, MP theory is size-consistent. However, it is non-variational, meaning the calculated energy is not guaranteed to be an upper bound and can sometimes overcorrect.[123]

CC offers a powerful and robust alternative that combines the strengths of the other methods. It constructs the correlated wavefunction using an exponential operator acting on the HF determinant. This sophisticated mathematical form ensures that CC is size-consistent at any level of truncation, correcting the main deficiency of CI. While it is non-variational like MP theory, its systematic inclusion of electron correlation effects makes it exceptionally accurate. The “gold standard” method, CCSD(T), which includes single and double excitations and a perturbative correction for triples, is widely regarded as one of the most reliable methods for obtaining near-exact results for systems that are well-described by a single HF determinant.[123]

Relativistic effects can be incorporated into post-Hartree–Fock methods through several approaches, such as the Douglas–Kroll–Hess transformation[134, 135] or the use of relativistic effective core potentials.[136, 137] These techniques modify the electronic Hamiltonian to account for relativistic corrections, enabling more accurate descriptions of systems containing heavy elements where such effects are significant. Incorporating relativity is particularly important for accurately predicting properties such as bond lengths, spin–orbit coupling, and electronic spectra in transition-metal and heavy-element compounds.

However, despite their remarkable accuracy, wavefunction-based methods face a significant practical barrier: their steep computational scaling with system size, N. The cost of a CCSD(T) calculation, for instance, scales as $\mathcal{O}(N^7)$. This prohibitive cost restricts the application of these high-accuracy methods to relatively small systems, typically comprising no more than a few dozen atoms.

Density Functional Theory (DFT)

Density Functional Theory (DFT) is an approximation to the explicit quantum mechanical equations listed above, making calculations of the electronic structure of many-body systems computationally tractable, particularly in the context of materials science and chemistry. DFT is based on the principle that the ground-state properties of a many-electron system can be determined from its electron density rather than its wave function, significantly reducing the computational complexity compared to wavefunction based methods. This makes DFT particularly suitable for studying large systems, including solids, liquids, and complex molecules.

DFT begins with the N-electron Hamiltonian, which describes the total energy of a system of N electrons interacting with each other and with an external potential (e.g., the nuclei of atoms), assuming the Born-Oppenheimer approximation holds. The Hamiltonian can be expressed as:

$$\hat{H}_{\text{elec}} = -\frac{\hbar^2}{2m_{\text{elec}}}\sum_{i=1}^{N}\nabla_i^2 + \sum_{i<j}^{N}\frac{e^2}{4\pi\epsilon_0|\mathbf{r}_i - \mathbf{r}_j|} - \sum_{i=1}^{N}\sum_{k=1}^{M}\frac{Z_k e^2}{4\pi\epsilon_0|\mathbf{r}_i - \mathbf{R}_k|} \tag{2.36}$$

We can rewrite the Hamiltonian in a more compact form:

$$\hat{H} = \hat{T} + \hat{U} + \hat{V}_{\text{N-e}}, \tag{2.37}$$

where $\hat{T}$ is the kinetic energy operator, $\hat{U}$ is the electron–electron interaction potential, and $\hat{V}_{\text{N-e}}$ is the contribution from the interaction between the electrons and the nuclei. One could, if desired, incorporate additional contributions into the potential term $\hat{V}_{\text{N-e}}$, such as those arising from nucleus-nucleus interactions or external fields to yield a generic potential energy from the external field, $\hat{V}$. With this Hamiltonian defined, the primary objective becomes solving the Schrödinger equation, typically to determine the ground state of the system. Due to the Pauli exclusion principle,[121] the total wavefunction must be antisymmetric under the exchange of any two identical fermions, meaning it must change sign upon such a swap.

DFT, reformulates the problem in terms of the electron density $\rho(\mathbf{r})$, which is a function of position rather than the many-body wavefunction:

$$n_0(\mathbf{r}) = \sum_{i=1}^{N}\int d\mathbf{r}_\mathbf{1}^\mathbf{3}\int d\mathbf{r}_\mathbf{2}^\mathbf{3}\cdots\int d\mathbf{r}_\mathbf{N}^\mathbf{3}\Psi_0^\dagger(x_1, x_2, \ldots, x_N)\delta(\mathbf{r} - \mathbf{r}_i), \Psi_0(x_1, x_2, \ldots, x_N), \tag{2.38}$$

where $n_0(\mathbf{r})$ is the ground-state electron density, and $\Psi_0(x_1, x_2, \ldots, x_N)$ is the many-body wavefunction of the ground state. The electron density is a scalar field that describes the distribution of electrons in space, and it contains all the information needed to determine the ground-state properties of the system. The central tenet of DFT is that the ground-state energy E_0 can be expressed as a functional of the electron density:

$$E[\rho] = T[\rho] + U[\rho] + V[\rho], \tag{2.39}$$

where $T[\rho]$ is the kinetic energy functional, $U[\rho]$ is the electron–electron interaction energy functional, and $V[\rho]$ is the external potential contribution. The Hohenberg-Kohn theorems establish that there exists a one-to-one mapping between the ground-state electron density and the external potential, allowing for the determination of the ground-state properties from the density alone.[138] The first theorem states that the ground-state energy is a unique functional of the electron density:

$$E_0 = E[n_0] = \langle \varPsi[n_0] | \hat{T} + \hat{U} + \hat{V} | \varPsi[n_0] \rangle, \tag{2.40}$$

where the contribution from the external potential can be written explicitly in terms of the ground-state density $n_0(\mathbf{r})$ as:

$$V[n_0] = \int d^3\mathbf{r}\, n_0(\mathbf{r}) V(\mathbf{r}) \tag{2.41}$$

The second theorem provides a variational principle for finding the ground-state density. The Kohn–Sham (KS) approach is a practical implementation of DFT, where the complex many-body problem is reduced to a set of non-interacting particles moving in an effective potential.[139] The KS equations are given by:

$$\left(-\frac{\hbar^2}{2m} \nabla^2 + V_{\text{eff}}[\rho] \right) \psi_i(\mathbf{r}) = \epsilon_i \psi_i(\mathbf{r}), \tag{2.42}$$

where $\psi_i(\mathbf{r})$ are the KS orbitals, ϵ_i are the corresponding eigenvalues which reconstruct the density of the many-body system:

$$n(\mathbf{r}) = \sum_{i=1}^{N} |\psi_i(\mathbf{r})|^2, \tag{2.43}$$

and $V_{\text{eff}}[\rho]$ is the effective single-particle potential that includes contributions from the electron–electron interaction and the external potential. The effective potential is given by:

$$V_{\text{eff}}(\mathbf{r}) = V(\mathbf{r}) + \int \frac{n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d^3\mathbf{r}' + V_{xc}[n(\mathbf{r})], \tag{2.44}$$

where $V(\mathbf{r})$ is the external potential, the second term represents the classical Coulomb interaction (Hartree term), and $V_{xc}[n(\mathbf{r})]$ is the exchange-correlation potential, which accounts for the many-body effects of electron correlation and exchange interactions.

The key challenge is contained within the effective potential, $V_{\text{eff}}(\mathbf{r})$, which includes the external potential, the classical electrostatic repulsion (Hartree term), and the crucial exchange-correlation (XC) potential, $V_{xc}(\mathbf{r})$. This term, derived from the exchange-correlation energy functional $E_{xc}[\rho]$, encapsulates all the complex many-body quantum mechanical effects. Since the exact form of $E_{xc}[\rho]$ is unknown, it must be approximated. The accuracy of any DFT calculation rests almost entirely on the quality of this approximation.

The development of progressively more accurate and complex XC functionals is often visualised as "Jacob's Ladder", a concept introduced by Perdew (Figure 2.4).[140] Each rung on the ladder represents a new level of physical sophistication, generally offering higher accuracy at an increased computational cost. The first rung on the ladder represents the local density approximation (LDA),[141] which treats the electron density at any point $\mathbf{r}$ as if it were part of a uniform electron gas with the same density $\rho(\mathbf{r})$. While a foundational idea, LDA tends to overbind molecules and severely underestimate band gaps.[142] The next rung involves one of the most widely used functionals, the Perdew-Burke-Ernzerhof (PBE) functional,[143] which is a generalised gradient approximation (GGA) functional. It improves upon LDA by considering not only the local density but also its gradient, $|\nabla\rho(\mathbf{r})|$. This added information allows for a better description of the inhomogeneous electron densities found in atoms and molecules, making GGAs the standard for many solid-state calculations, including those on carbonaceous materials,[144] although they still tend to underestimate band gaps.[145] Meta-GGA functionals are next on the "ladder" and add another ingredient: the kinetic energy density, $\tau(\mathbf{r})$. This allows meta-GGAs to better describe the band gap of materials due to the introduction of non-locality by including the Laplacian.[145] Beyond this, hybrid functionals represent a significant leap in accuracy by mixing a fraction of exact exchange from Hartree-Fock theory with a GGA or meta-GGA functional.[147] Hybrid functionals (e.g., PBE0, B3LYP[148–150]) can significantly correct for one of DFT's most persistent errors, the self-interaction error (discussed later), leading to more accurate predictions of band gaps and reaction barriers.[151] The highest rung incorporates information from unoccupied KS orbitals to calculate a portion of the correlation energy using principles from MP perturbation theory (typically MP2). These offer very high accuracy but at a computational cost that begins to approach that of traditional wavefunction methods.[152]

In practice, DFT calculations are performed using various computational codes that implement the KS equations and solve for the ground-state electron density. These codes typically employ plane-wave basis sets to perform periodic calculations and pseudopotentials which are used to replace the core electrons of atoms with an

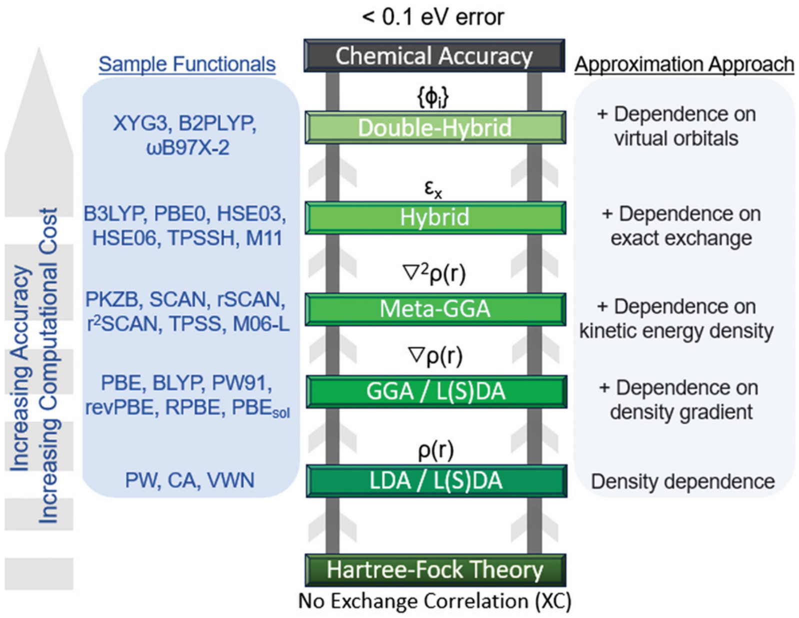

Figure 2.4: Jacob's ladder. A hierarchical level of accuracy and cost for different DFT functionals. Reproduced from Ref. [146] with permission from John Wiley and Sons.

effective potential that captures the essential features of the electron–electron interactions while simplifying the calculations. Pseudopotentials are particularly useful in DFT calculations for materials with many atoms, as they allow for a reduction in the number of basis functions needed to describe the electronic wavefunctions. Plane-wave basis sets are often used in conjunction with pseudopotentials because they provide a systematic way to represent the electronic wavefunctions in a periodic system, enabling efficient calculations of the electronic structure and properties of materials. The choice of pseudopotential and basis set can significantly affect the accuracy of the DFT calculations, particularly for systems with complex bonding environments, such as carbonaceous materials.[153] Among the various DFT codes available, this work employs `CASTEP` for electronic structure calculations.[154] `CASTEP` is a plane-wave pseudopotential code that offers several advantages, particularly its support for on-the-fly (OTF) pseudopotential generation. OTF pseudopotentials allow for the automatic creation of norm-conserving or ultrasoft pseudopotentials tailored to the specific elements and exchange-correlation functional used in a calculation. This flexibility ensures consistency between the pseudopotential and the chosen functional, reduces the risk of transferability errors, and streamlines the setup of calculations for a wide range of elements.

While DFT's balance of cost and accuracy has made it the workhorse of computational materials science, its reliance on approximate XC functionals introduces several systematic errors.[155] It is crucial to be aware of these limitations, especially when modelling extended materials. A fundamental error of DFT is the self-interaction error (SIE). In the exact theory, an electron should not interact with itself, however, in approximate DFT, the Hartree term incorrectly includes an interaction of an electron's charge density with itself, which is not perfectly cancelled by the approximate XC functional. This error is most severe for localised electrons (e.g., in d- or f-orbitals or defect states) and is a primary cause of the chronic underestimation of band gaps in semiconductors and insulators predicted by LDA and

GGA functionals.[155] In addition, there is an absence of Van der Waals interactions in standard DFT functionals. Dispersion forces arise from correlated fluctuations in the electron density between distant regions of space, leading to an attractive interaction that is inherently non-local.[155] Standard local and semi-local functionals (LDA, GGA) cannot capture these long-range correlations, resulting in poor predictions for systems where dispersion is significant, such as molecular crystals, layered materials, and physisorption on surfaces. To address this limitation, various dispersion correction schemes have been developed, such as DFT-D methods (e.g., Grimme's D4[156]), which add an empirical dispersion term to the DFT energy, or non-local correlation functionals (e.g., vdW-DF) that incorporate dispersion effects directly into the XC functional.[157]

Another drawback of DFT is its computational cost, which scales with the cube of the number of electrons in the system, $\mathcal{O}(N^3)$, due to diagonalisation of the Kohn–Sham Hamiltonian. This severely limits the size of systems that can be studied using DFT, particularly for large-scale simulations of materials. In addition, DFT calculations are many orders of magnitude more expensive than empirical potentials, making them prohibitively expensive for large-scale simulations or long time scales. To address these limitations, various approximations and techniques have been developed which reduce the computational overheard at the cost of accuracy, such as the use of pseudopotentials, finite truncation of plane-wave basis sets, and linear scaling methods (such as `ONETEP`)[158], however *ab initio* methods are impractical, if not impossible, for robust and accurate modelling of large-scale amorphous systems, in particular carbonaceous materials which require long melt-quench simulations to generate representative amorphous structural models.

In this thesis, all DFT calculations are performed using `CASTEP` and unless otherwise stated, the PBE functional is used.

2.2.2 Empirical Potentials

One of the key advantages of empirical potentials is that it is possible to bypass the computationally expensive quantum mechanical calculations by using a functional form of the potential energy surface, which can be parameterised to reproduce the results of quantum mechanical calculations or experimental data. This approach is known as empirical potentials, and it has been widely used in atomistic simulations to model the interactions between atoms in materials.

Empirical potentials are simplified, fixed, functional forms which are based on theoretical foundations and chemical intuition, and can be considered as somewhat of an artisanal approach to combining "art and science".[159] These potentials are typically parameterised using a combination of experimental data and high-level quantum mechanical calculations, allowing them to capture the essential features of atomic interactions while remaining computationally efficient.

One of the core principles of empirical potentials is that of *locality*.[159] What this means in practice is that the potential energy of a system can be expressed as a sum of contributions from local interactions between atoms, rather than considering the

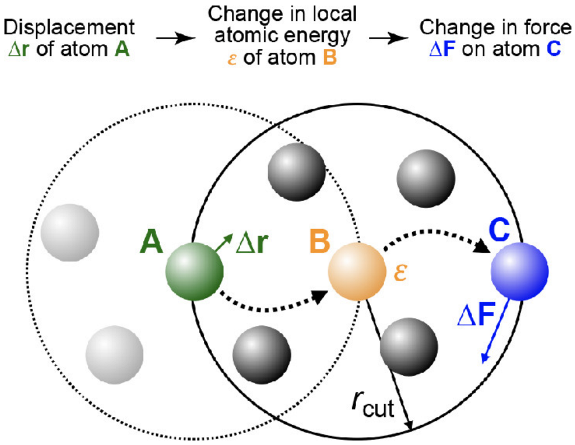

Figure 2.5: Locality of forces in case of atom-centred three- or higher-body order potential terms. The displacement of atom **A** affects the local atomic energy of atom **B**, which in turn affects the force on atom **C**. Both atoms **A** and **C** are just within the cutoff radius of atom **B**, and thus the locality of the forces in the model is twice the cutoff radius. Figure reproduced from Ref. [160]. (CC-BY-4.0).

entire system as a whole. This locality assumption allows for the efficient calculation of forces and energies in large systems, as only a small subset of atoms need to be considered at any given time as is visualised in Figure. 2.5. The potential energy function can be expressed as:

$$E = \sum_{i}^{\text{atoms}} \epsilon_i \tag{2.45}$$

where $\epsilon_i = \epsilon(\mathbf{r}_{ij})$ is the energy contribution from atom i, which is typically a function of the positions of its nearest neighbours j. This locality assumption works surprisingly well in the case of modelling amorphous systems which lack long-range order. Of course, this assumption immediately breaks down in the case of long-range interactions, such as those present in ionic systems, surfaces, hydrogen bonding, etc. In these cases, additional terms must be added to the potential energy function to account for these long-range interactions, such as Coulombic interactions or van der Waals forces at a cost.

Lennard-Jones

Perhaps the most well-known empirical potential is the Lennard–Jones (LJ) potential,[161] which describes the interaction between a pair of neutral atoms or molecules. The LJ potential is given by:

$$U(r) = 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6} \right], \tag{2.46}$$

where r is the distance between the two atoms, ϵ is the depth of the potential well (which corresponds to the strength of the interaction), and σ is the distance at which the potential is zero (which corresponds to the equilibrium distance). The first term represents the repulsive interaction between the atoms at short distances due to the Pauli exclusion principle, while the second term represents the attractive interaction at longer distances due to induced dipoles.

The LJ potential was originally derived to describe the interaction between Ar atoms and despite the simplicity of the functional form, it is surprisingly effective at capturing the essential features of Ar. Beyond this simple system, the LJ potential breaks down in the presence of more complex systems where quantum effects are significant. For the other noble gases, relativistic effects are pronounced, leading to erroneous predictions. Despite this, a system of LJ interacting particles is a useful model for understanding the behaviour of simple fluids and gases, displaying complex phase behaviour, such as phase transitions and critical phenomena. The LJ potential is often used as a starting point for more complex potentials, and it serves as a benchmark for testing the accuracy of new interatomic potentials.

Empirical Potentials for Carbon

Table 2.1: Summary of empirical carbon potentials detailing their main characteristics and capability to describe other chemical species. The potentials are grouped by families and arranged in chronological order. Where relevant the historical basis of the potential is specified. Acronyms: REBO: Reactive Empirical Bond Order; AIREBO: Adaptive Intermolecular REBO; SED-AIREBO: Screened Environment Dependent AIREBO; ABOP: Analytic bond-order potential; EDIP: Environment-dependent interaction potential; ReaxFF: Reactive Force Field; LCBOP: Long-range Carbon Bond Order Potential; MEAM: Modified Embedded Atom Method; RSS: Reactive State Summation; COMB: Charge Optimized Many Body; SW: Stillinger-Weber. Adapted from Ref. 162

Family	Name of potential	Year	Main characteristic	Builds on	Other species
Tersoff	Tersoff	1988	Reparametrisation for carbon	Si Tersoff	–
	Tersoff	1989	Generalisation to multicomponent systems	Si Tersoff	Si
	Tersoff	1990	Reparametrisation to treat C defects in Si	Tersoff 1989	Si
	Tersoff	1994	Minor change in parameters	Tersoff 1990	Si
	Tersoff	1995	Variable cutoff rescaled with system volume	Tersoff 1989	Si
	Nordlund	1996	Adds van der Waals and uses larger cutoff	Tersoff 1988	H
	Tersoff	1998	Adds nitrogen and boron	Tersoff 1989	Si/N/B
	Tersoff	2005	Extended cutoff	Tersoff 1988	–
REBO	REBO	1990	Adds hydrogen. Treats carbon radicals.	–	H
	REBO-II	2002	Adds dihedral. New short-range treatment.	REBO	H
	AIREBO	2000	Adds van der Waals (Lennard-Jones)	REBO-II	H
	AIREBO-II	2002	Minor change in two parameters	AIREBO	H
	REBO-CHO	2004	Adds oxygen	REBO-II	H/O
	AIREBO	2008	Environment-dependent van der Waals	AIREBO	H
	Pastewka	2008	Environment-dependent screening function	REBO-II	H
	qAIREBO	2012	Adds oxygen and charge equilibration	AIREBO	H/O
	SED-REBO	2013	Environment-dependent screening function	REBO-II	–
	AIREBO-M	2015	Replaces Lennard-Jones with Morse	AIREBO	H
Heggie	Heggie	1991	Wigner-Seitz treatment of local bonding	Tersoff 1988	–
	Andribet	1996	Simpler functional form. More efficient.	Heggie	–
	Broughton	1999	Valence bond treatment. No angular terms.	Tersoff 1988	–
ABOP	Pettifor	1999	Analytic approximation of tight-binding	–	H
	Mrovec	2007	Improved a bond term	Pettifor	H
	Zhou	2015	Reparametrisation for pure carbon	Pettifor	–
	EDIP	2000	Long-range π-repulsion. Dihedral terms.	Si EDIP	–
ReaxFF	ReaxFF	2001	Flexible and general functional form	–	H
	RDX	2003	Developed for nitramine explosions	ReaxFF	H/O/N
	ReaxFF$_{\text{CHO}}$	2008	Adds oxygen	ReaxFF	H/O
	ReaxFF-lg	2011	Adds long range London dispersion	ReaxFF	H/O/N
	ReaxFF$_{\text{C-2013}}$	2015	Reparametrisation for solid carbon	ReaxFF$_{\text{CHO}}$	H/O
	ReaxFF$_{\text{CHO-2016}}$	2019	Further reparametrisation for solid carbon	ReaxFF$_{\text{CHO \& C-2013}}$	H/O
	GR-RDX	2022	Reparametrisation for sp^2 carbon	RDX & ReaxFF$_{\text{C-2013}}$	H/O/N
LCBOP	LCBOP-I	2003	Simultaneous bond-order & long-range fit	REBO	–
	LCBOP-I+	2004	Adds torsional term	LCBOP-I	–
	LCBOP-II	2005	Improved bond dissociation energetics	LCBOP-I+	–
	Ehrhart/Albe	2005	Hybrid Tersoff/REBO. Reparametrisation.	Tersoff/REBO	Si
MEAM	Lee/Lee	2005	Reparametrised for C. Long-range term.	MEAM	–
	Uddin	2010	Developed for Ni catalysis of nanotubes.	MEAM	Ni
	Vashishta	2007	Covalent & ionic. Dipoles. Van der Waals.	Si SW	Si
	RSS	2008	Short-range interactions. No sp^3 bonds.	RSS	–
	Kumagai	2009	Screening function and dihedral term.	REBO	–
COMB	COMB3	2013	Charge-dependent short-range interactions	REBO-II	H/O/N

The development of empirical potentials for carbon has been driven by the need to accurately model the various allotropes of carbon, such as diamond, graphite, graphene, and amorphous carbon, each of which exhibits distinct bonding characteristics and structural motifs. Since the late 1980s, several empirical potentials have been developed, as summarised in Table 2.1, to capture the unique bonding properties of carbon, with varying degrees of success. These potentials expand upon the simple pairwise potentials of an earlier generation by incorporating many-body interactions, bond order effects, and angular biases to better represent the covalent bonding nature of carbon, as well as incorporating long-range interactions in some instances. To discuss all empirical potentials for carbon in detail is beyond the scope of this thesis, and so we will focus on the most widely used potentials, namely the Tersoff,[90] EDIP,[163] and ReaxFF[164] potentials which are commonly employed in simulations of carbon-based materials. We note in passing that the REBO[165] (and the later updated version, AIREBO[166, 167]) potential has also been widely used in the literature, but it is not discussed here as it is a variant of the Tersoff potential which is discussed in detail below.

Tersoff

In contrast to the simple pairwise Lennard-Jones potential, the Tersoff interatomic potential, introduced by Jerry Tersoff in a series of papers in the late 1980s,[90] marked a pivotal development in the empirical modelling of carbon. Its fundamental innovation is the explicit incorporation of the bond-order concept, a departure from simpler potential models that often assume fixed bond characteristics. In covalent materials, the nature and strength of a chemical bond are intrinsically linked to its local atomic environment. The Tersoff potential captures this by dynamically adjusting the strength of a bond based on factors such as the number of neighbouring atoms (coordination number) and the angles formed with these neighbours.

This environment-dependent bond order allows the potential to implicitly differentiate between various bonding configurations, such as single, double, or triple bonds, and even different hybridization states (e.g., sp, sp^2, sp^3). This adaptive capability is crucial for simulating processes that involve the breaking and formation of chemical bonds, a common occurrence in chemical reactions, phase transitions, defect dynamics, and surface phenomena in covalent materials. The original formulation aimed to provide a more realistic and transferable description of the energetics and structural properties of Si, C, and Ge than was previously attainable with simpler potentials.

Carbon exhibits a rich diversity of bonding types that depend critically on the local environment. Tersoff addressed this by introducing a bond order term, typically denoted as b_{ij}, which explicitly depends on the local atomic environment, including coordination numbers and bond angles.[90] This b_{ij} term modulates the attractive part of the interatomic potential, effectively changing the bond strength in response to environmental changes. For instance, as the coordination number of an atom increases, or as bond angles deviate from their ideal values, the b_{ij} term decreases, weakening the associated bonds. Consequently, as the local environment changes, for example, during phase transition, or graphitisation, the bond order adapts, enabling the potential to describe bond breaking, bond formation, and transitions between different hybridisation states.

As has been alluded to, the Tersoff potential is a many-body potential, meaning that it accounts for interactions between more than two atoms at a time, in contrast to pair potentials such as the LJ potential, which only consider interactions between pairs of atoms. The total energy of a system described by a many-body potential can be expressed as a sum over all atoms and their interactions with other atoms in the system:

$$E = \sum_{i} V_i(\mathbf{r}_i) + \sum_{i<j} V_2(\mathbf{r}_i, \mathbf{r}_j) + \sum_{i<j<k} V_3(\mathbf{r}_i, \mathbf{r}_j, \mathbf{r}_k) + \cdots , \tag{2.47}$$

where $\mathbf{r}_n$ is the position of the nth atom and V_m is the mth-body potential term. In general, pair potentials favour close packing and are less effective at describing covalent systems which are highly directional, and thus it is evident that going beyond pair potentials (such as the LJ potential) is necessary. The intuition within the process of designing empirical potentials becomes clear when three-body terms are constructed to bias the potential *a priori* into forming geometries which resemble known motifs, such as diamond with a bond angle of 109.5°.

For C, Si and Ge, the Tersoff potential has demonstrated the ability to reproduce bulk properties such as lattice constants, cohesive energies, elastic moduli, and phonon dispersion relations with reasonable accuracy.[90] Its capacity to handle changing coordination numbers also makes it suitable for simulating amorphous and liquid phases of these elements, and has been used in studies of diamond growth[168] (often through its extension, the REBO potential) and in the characterisation of diamond-like carbon films.[165] Simulations of nanoindentation using diamond indenters frequently employ Tersoff potentials for the indenter material due to its good description of diamond's mechanical properties.[169]

The original Tersoff potential, while successful in many areas, exhibited certain limitations that prompted the development of several variants and modifications aimed at improving its accuracy for specific applications or overcoming known deficiencies, which can be found in Table 2.1. Such issues were fundamentally incorrect, for example, the density of graphite is incorrect as the Tersoff potential predicts a density of $\approx$ 3.5 gcm^{-3}, the same as diamond, along with the simulated melting point being $\approx$ 6000 K compared to the experimental value of 4300 K.[162] The Tersoff potential also produces disordered structures at all densities under annealing, as opposed to graphitising which is observed experimentally.[170] This occurs due to the truncation at 3-body (angular) interactions, resulting in the rotation of π-bonds not being penalised, leading to non-planar structures forming entropically and thus preventing graphitisation.

As such, modifications to the Tersoff potential have been made to attempt to overcome such shortcomings, including the introduction of additional parameters to better capture the bond order effects, the incorporation of long-range interactions, and the extension to multicomponent systems.[167] With that being said, the Tersoff potential framework is remarkably cheap yielding somewhat accurate amorphous carbon structures and has been a computational workhorse for simulating carbon-based materials over the past several decades and still finds use today.[171]

The REBO[165] and subsequent AIREBO[167] potentials made significant modifications beyond the original Tersoff framework. REBO was designed to simulate hydrocarbons by introducing a more sophisticated bond-order formulation, making it well-suited for studying reactions involving carbon and hydrogen, such as chain growth and radical chemistry. Building on this, AIREBO incorporates a Lennard-Jones term to capture long-range van der Waals interactions, which are essential for modelling layered structures like graphite and interactions between carbon nanotubes.

However, both potentials come with limitations. REBO, while capable of modelling reactivity, lacks long-range dispersion forces and is therefore unsuitable for systems where such interactions play a dominant role. When applied to amorphous carbon, REBO struggles to capture the graphitisation process, grossly underestimating the sp^2 character of the resulting structures.[162, 170] For AIREBO, the low density structures are highly disordered after annealing, showing a relatively flat ring distribution compared with other potentials. Once again, this is due to the truncation of the potential at 3-body interactions, which does not penalise the rotation of π-bonds.[162, 170]

Environment-Dependent Interatomic Potential (EDIP)

It is apparent that one requires higher body-order terms to describe the intricacies of bonding in amorphous carbon. The Environment-Dependent Interatomic Potential (EDIP) was originally developed to model amorphous Si[172] and was modified for C[163] to simulate the growth of amorphous carbon thin films, accounting for the complex relationship between sp^2 and sp^3 bonding in carbon. EDIP approaches locality in a different manner to the Tersoff potential, using a general coordination term to capture long range π-bonding interactions, and including 4-body interactions in the form of dihedral terms to bias the potential towards planar geometries.

EDIP was parameterised to reproduce the correct energetics of bond breaking and forming in carbon, and was fitted to *ab initio* data from HF calculations for the graphite-diamond transition.[163] As a result, the EDIP potential has a relatively long cutoff radius of 3.2 Å, beyond which there are no interactions. The potential is given by:

$$E = \sum_{i<j} V_{ij}(r_{ij}, Z_i) + \sum_{i<j<k} V_{ijk}(r_{ij}, r_{ik}, \theta_{ijk}, Z_i), \tag{2.48}$$

where θ_{ijk} is the angle formed by atoms i, j, and k, and Z_i is the coordination number of atom i. Briefly, the EDIP potential penalises the rotation of π-bonds through the dihedral term, which is a 4-body interaction. It is this term which yields the correct graphitisation process, as it biases the potential towards planar structures, biasing for the formation of π-bonds. EDIP also includes π-repulsion effects.

EDIP has found successful application in: modelling the graphitisation process of amorphous carbon, capturing the transition from disordered structures to graphitic domains; the growth of amorphous carbon thin films, where it accurately describes the interplay between sp^2 and sp^3 bonding, and provides reasonable defect energies.[162, 170] In particular, the graphitisation process is well captured, with the potential allowing for the formation of planar structures and π-bonding, which is essential for the stability of graphitic materials. This is reflected in the correct position of the peaks in the RDF for low- and high-density annealed structures, as well as a high sp^2 and 6-membered ring content in the final low-density structures, that also show stacking. For the high-density structures, EDIP yields structures which are amorphous with a high degree of sp^3 bonding. Despite this, EDIP performs poorly in the case of predicting lattice constants and mechanical properties of nanotubes, further highlighting the limitations of the flexibility of empirical potentials in capturing the full complexity of carbon bonding.[162, 170]

Reactive Force Field (ReaxFF)

The Reactive Force Field (ReaxFF)[164] was originally developed for hydrocarbons but has since been extended to model a wide range of elements and compounds.[173] Unlike potentials such as Tersoff and EDIP, which embed chemical intuition into their functional forms, ReaxFF takes a more generalised, data-driven approach. It uses a complex set of parameters to capture various interactions including bond order, charge transfer, van der Waals forces, and torsional effects.

A core idea in ReaxFF is that bond order depends directly on interatomic distances and is corrected for over- or under-coordination to reflect chemical valency. The total energy includes contributions from bonds, angles, torsions, long-range electrostatics, and van der Waals interactions, making the model versatile but computationally intensive. ReaxFF has shown success across a broad spectrum of systems, from organic molecules to fuel cells.[173] However, its performance in modelling elemental carbon has been inaccurate; early versions greatly overestimated properties such as the bulk modulus of diamond.[174] Improved carbon-specific versions have been developed, such as ReaxFF$_{\text{C-2013}}$[175], but challenges remain, particularly in accurately reproducing density-dependent graphitisation.

Despite its computational cost and complexity, ReaxFF's flexible and general philosophy has parallels with modern machine learning interatomic potentials (MLIPs). These newer methods replace hand-crafted functions with data-driven models, offering greater accuracy and transferability, albeit often at the cost of physical interpretability and speed. This shift reflects a broader trend in materials modelling: from physically-informed potentials to flexible, data-driven frameworks.

2.2.3 Machine Learning Interatomic Potentials

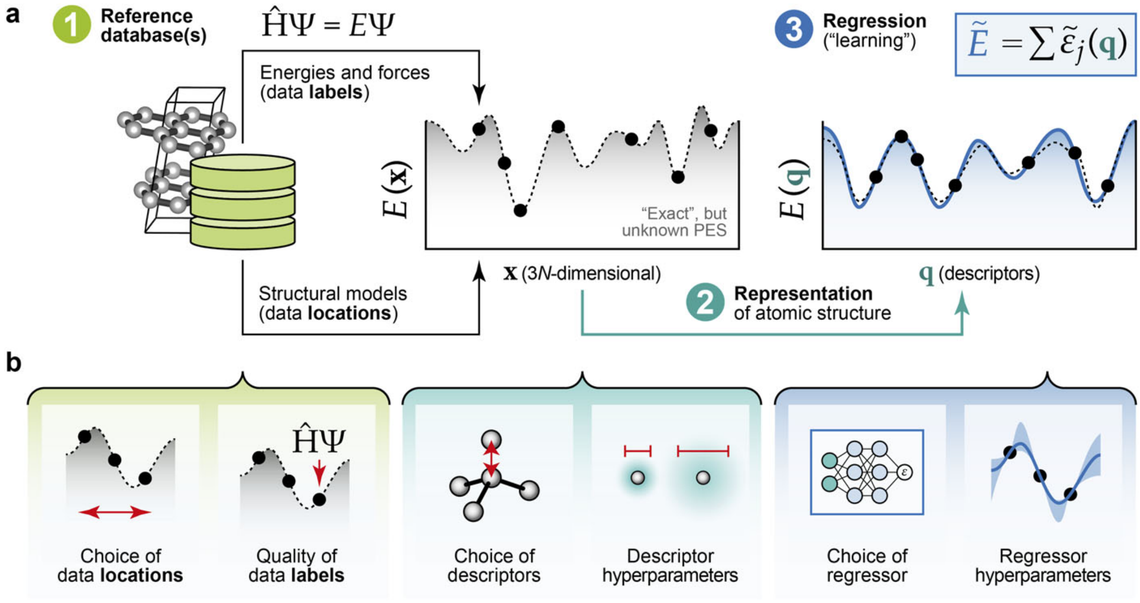

Figure 2.6: Overview of machine learning interatomic potentials. (a) Schematic illustration of the workflow: (1) A database of representative atomic structures is generated and labelled with reference energies and forces, typically from density functional theory (DFT). (2) Atomic environments are encoded as descriptor vectors **q**. (3) A regression model is trained to predict the total energy, $\tilde{E}$, as a sum of machine-learned local atomic energies, $\tilde{\epsilon_j}$. (b) Key factors influencing the quality of an ML potential are shown schematically, with colour coding matching panel (a): green for the reference database, turquoise for the descriptor, and blue for the regression method. Panel (a) adapted from V. L. Deringer, M. A. Caro, and G. Csányi, Adv. Mater. 31, 1902765 (2019), Copyright 2019 Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, and V. L. Deringer, J. Phys. Energy 2, 041003 (2020), Copyright 2020 Author(s), CC-BY 4.0.

Classification vs. Regression

Before delving into the specifics of machine learning interatomic potentials (MLIPs), it is important to clarify the distinction between classification and regression tasks within the field of machine learning (ML). ML encompasses a broad spectrum of algorithms and techniques designed to enable computers to learn from data and make predictions or decisions without being explicitly programmed for specific tasks. In general, ML tasks can be broadly categorised into two main types: classification and regression.

Classification tasks involve predicting discrete labels or categories based on input features. For example, in image recognition, a classification algorithm might be trained to identify whether an image contains a cat or a dog. The output of a classification model is typically a class label (e.g., ‘cat’ or ‘dog’) or a probability distribution over multiple classes (e.g., 70% cat, 30% dog). Algorithms such as decision trees,[176] random forests (RF),[177] support vector machines,[178] and logistic regression[179] are widely used for such tasks. Taking one example algorithm, RF is an ensemble learning method that constructs multiple decision trees during training and combines their outputs to make a final prediction. Each tree is trained on a random subset of the data and features, ensuring diversity among the trees.[177] For classification tasks, the final output is determined by majority voting among all trees. This ensemble approach reduces overfitting and generally leads to more robust and accurate models compared to using a single decision tree.[177]

Regression tasks, in contrast, involve predicting continuous numerical values based on input features. For instance, in house price prediction, a regression algorithm might be trained to estimate the price of a house based on features such as its size, location, and number of bedrooms. The output of a regression model is a continuous value (e.g., the predicted price of a house). Common regression algorithms include linear regression, polynomial regression,[180] and also RF. Here, RF can also be adapted for regression tasks, where the final output is the average of the predictions from all individual trees. Each tree in the ensemble is trained on a random subset of the data and features, similar to the classification case, but instead of voting, the predictions are averaged to produce a continuous output. This approach helps to reduce variance and improve prediction accuracy for regression problems.[177] More complex regression techniques such as Gaussian process regression (GPR)[181] and neural networks (NNs)[182] have also been developed, which can capture non-linear relationships between input features and output values.

GPR is a non-parametric, Bayesian approach to regression that provides not only predictions but also uncertainty estimates for those predictions.[181] It models the underlying function as a Gaussian process, defined by a mean function and a covariance function (kernel). The choice of kernel is crucial, as it encodes assumptions about the smoothness and structure of the function being modeled. GPR is particularly useful for small to medium-sized datasets where uncertainty quantification is important. NNs, on the other hand, are inspired by the structure and function of biological neural networks. They consist of layers of interconnected nodes (neurons) that process input data through weighted connections and activation functions. NNs can model complex, non-linear relationships and are highly flexible, making them suitable for a wide range of regression tasks.[182] They can be trained using backpropagation and gradient descent to minimise a loss function, such as mean squared error, which measures the difference between predicted and actual values.

Chemical Machine Learning

Machine learning is not a new concept in the chemical sciences; it forms the foundation of the entire field of cheminformatics.[183] For decades, this field has focused on predicting molecular properties, often for applications like drug discovery. This is typically achieved by converting a molecular structure, represented by a line notation (e.g., SMILES)[184] or a fixed-size "fingerprint" into a prediction for a single property.[185]

This task, often known as quantitative structure-property relationship modelling,[183] is a direct application of the regression or classification tasks discussed in the previous section. This hierarchy ranges from classical statistical techniques, such as linear regression, to more complex, non-linear models like support vector machines and, increasingly, neural networks.[186] These models are trained to predict a single scalar value, such as a molecule's toxicity (a classification task) or boiling point (a regression task).

The methods central to this thesis, machine learning interatomic Potentials (MLIPs), represent a fundamentally different and more complex application of machine learning.[187] Instead of predicting a single property from a static representation, an MLIP is trained to learn and subsequently approximate the PES of a system, as illustrated in Figure 2.6. The goal is not to predict a single scalar property but rather to model the entire energy landscape of a system as a function of atomic positions.

So the question naturally arises, how can we achieve the accuracy of DFT calculations while maintaining the cost of empirical potentials, yet still ensure flexibility with a system-agnostic framework? The answer lies in the development of MLIPs, which combine the strengths of both approaches. MLIPs leverage the power of ML to learn complex, many-body interactions from high-fidelity reference data, such as DFT calculations, while attempting to retain the computational efficiency of empirical potentials.

MLIPs have transformed atomistic modelling by deviating from empirical, physically inspired functional forms with highly flexible, data-driven representations that can capture complex many-body interactions with minimal prior assumptions.[160] While this flexibility allows MLIPs to approximate arbitrarily intricate PES', it also introduces a key practical difficulty: ensuring reliable performance outside the domain of the training data. In theory, a sufficiently expressive model can reproduce the PES exactly, provided that reference data is available for all relevant atomic configurations.[188] In practice, however, high-fidelity data (e.g., from DFT) is expensive to generate, making it infeasible to densely sample the entire configurational space. To mitigate this limitation, MLIP frameworks must exploit intrinsic properties of the PES, such as its smoothness, symmetry, and locality, which constrain the functional variation and reduce the volume of data needed to achieve accurate interpolation.[189] Leveraging these physical features is especially critical for constructing general-purpose MLIPs that are expected to perform robustly across a wide range of atomic environments and geometries, without requiring prohibitive quantities of training data.

MLIPs are typically trained on energies and forces computed from *ab initio* methods, however it is important to note that spin-aware MLIPs have also been developed for magnetic materials, since the corresponding PES is dependent on the spin configuration as well as the atomic positions.[190–192] This results in the dimensionality problem increasing significantly. For a nonmagnetic potential, there are $3N$ degrees of freedom as input of $E = f(\mathbf{r})$, where $\mathbf{r}$ is the vector of atomic positions, while for a magentic potential, this increases to $6N$ degrees of freedom as input of $E = f(\mathbf{r}, \mathbf{S})$, where $\mathbf{S}$ is the vector of magnetic moments. In addition, the cost of generating reference data increases significantly, as spin-polarised DFT calculations are required which are more expensive than non-spin-polarised calculations. Various approaches have been developed to address this challenge, ranging from highly specific models for particular systems such as iron,[190] to more state-of-the-art general frameworks which can be applied to a wide range of magnetic materials.[191, 192]

Neural Network Potentials (NNPs)

Early attempts to approximate the PES using ML approaches was made in 1995 by Blank *et al.*, where a neural network potential (NNP) was used to compute global properties of CO adsorbed on a Ni(111) surface.[193] Whilst the idea of neural networks had emerged much earlier in the 1940s and 50s, having been inspired by our own human brains interconnectivity with neurons and nodes,[194, 195] and had eventually been applied to various problems in solid-state physics and chemistry,[196–199] it had yet to be used to model a PES. Admittedly, the work of Blank *et al.* was not an orthodox MLIP, as it did not model interatomic interactions directly, but rather used a neural network to compute the energy of a surface and adsorbate based on two degrees of freedom, the lateral position of the center of mass x and the angle of the molecular axis relative perpendicular to the basal plane of the surface

θ.*

These early NNPs represented the first generation as such, and were almost exclusively based on global descriptions of the system and were low-dimensional in nature. The second-generation of NNPs, which emerged in the late 2000s, focused on the assumption of locality, where the total energy of a system can be approximately decomposed into sums of local contributions. This was demonstrated in a groundbreaking paper by Behler in 2011,[200] where he introduced the concept of atom-centred symmetry functions (ACSFs) to encode descriptions of local atomic environments in a way that respects the symmetries of the PES, such as translational, rotational, and permutational invariance which could then be used to predict a local, atomic energy of an atom and once summed, the total energy in a high-dimensional manner:

$$E = \sum_i \epsilon_i \tag{2.49}$$

where $\epsilon_i = \epsilon(\mathbf{r}_{ij})$ is the energy contributions of atom i, which is a function of the positions of its neighbours within some spherical cutoff $f_{\mathrm{cut}(r)}$. Whilst this is clearly non-physical, as "atomic energies" are not observable, it is a very powerful approach, and works surprisingly well for many condensed systems where long-range electrostatic interactions beyond the cutoff are effectively screened out by the local environment, and thus contribute a few meV/atom.[201] Despite being non-physical, these local energies underpin the work done in Chapter 3, where these atomic energies are used to drive a Monte-Carlo simulation of 2D carbon, and provide an intuitive fingerprint of disorder in structures.[60]

It is these atomic decompositions, along with the many-body ACSFs, which result in fitting an atomic energy function in a high-dimensional space, leading to poor extrapolation outside the domain of the training data. This is a key limitation of second-generation NNPs, as they are not guaranteed to be transferable to

---

[*] Training this network of 42 points, each with two degrees of freedom, took "about 10 minutes on a Sun Sparcstation 2/Weitek μP" (80 MHz), on MATLAB version 4.2 in 1993![193]

unseen configurations and can yield non-physical predictions when applied to new systems. This issue is particularly relevant for long simulations at high temperatures (e.g., graphitisation), where the system may explore configurations not well represented in the training set. More recent generations of NNPs have addressed some of these limitations by incorporating local and global charge information.[201] At the time of the initial work for this thesis, however, second generation NNPs were the most widely used. An alternative framework that also relies on the locality assumption, but employs a different approach to representing atomic environments, is the Gaussian Approximation Potential (GAP), discussed below.

Gaussian Approximation Potential (GAP)

The Gaussian Approximation Potential (GAP) is a kernel-based MLIP framework that uses Gaussian process regression (GPR) with a number of different descriptors of the atomic environments. Introduced by Bartók, Csányi, and collaborators,[189, 202] GAP models the total energy of a system as a sum of local atomic energies, each predicted by a non-parametric regression over a database of reference quantum mechanical calculations in a similar manner to the Behler–Parrinello NNPs.

Formally, the total energy is written as a sum over atomic contributions, as is seen in Equation 2.49, where each local energy is expressed as a function of the local environment of atom i. In GAP, this local energy is expressed as a weighted sum over kernel functions comparing the environment of atom i to a set of representative environments (the "sparse points" as determined through CUR decomposition):

$$\epsilon_i = \sum_{s=1}^{N_{\text{sparse}}} \alpha_s K(\mathcal{R}_i, \mathcal{R}_s), \tag{2.50}$$

where $K(\mathcal{R}_i, \mathcal{R}_s)$ is the kernel function measuring the similarity between environments $\mathcal{R}_i$ and $\mathcal{R}_s$, and α_s are regression weights determined by fitting to reference data.

A key innovation in GAP is the use of the Smooth Overlap of Atomic Positions

(SOAP) descriptor, which encodes the local environment, $\mathcal{R}_i$, of each atom i in a way that is invariant to rotations, translations, and permutations of identical atoms. The SOAP descriptor represents the neighbour density around an atom as a sum of Gaussians:

$$\rho_i(\mathbf{r}) = \sum_j \exp\left(-\frac{|\mathbf{r} - \mathbf{r}_{ij}|^2}{2\sigma^2}\right) f_{\text{cut}}(r_{ij}), \tag{2.51}$$

where $\mathbf{r}_{ij}$ is the vector from atom i to neighbour j, $f_{\text{cut}}(r_{ij})$ is a cutoff function that smoothly goes to zero beyond a certain distance, and σ controls the width of the Gaussians, in essence controlling the smoothness of the potential. This density is then expanded in a basis of orthonormal radial functions $g_n(r)$ and spherical harmonics $Y_{lm}(\hat{\mathbf{r}})$:

$$c_{nlm}^{(i)} = \int d\mathbf{r}\, g_n(r)^* Y_{lm}^*(\hat{\mathbf{r}}) \rho_i(\mathbf{r}), \tag{2.52}$$

where $c_{nlm}^{(i)}$ are the coefficients of the expansion for atom i, with n indexing the radial functions and l, m indexing the spherical harmonics and are truncated at some maximum values. The radial functions are typically chosen to be Gaussian-type orbitals or polynomial functions, and the spherical harmonics capture the angular dependence of the atomic environment. The power spectrum of the atomic density, which is rotationally invariant, is then constructed as:

$$p_{nn'l}^{(i)} = \sqrt{\frac{8\pi^2}{2l+1}} \sum_m c_{nlm}^{(i)} \left(c_{n'lm}^{(i)}\right)^*. \tag{2.53}$$

The vectors given by $\mathbf{p}^{(i)} = \{p_{nn'l}^{(i)}\}$ define, after normalisation, the SOAP many-body descriptor for atom i:

$$\mathbf{q}^{(i)} = \frac{\mathbf{p}^{(i)}}{\sqrt{\mathbf{p}^{(i)} \cdot \mathbf{p}^{(i)}}}. \tag{2.54}$$

The similarity between two environments is quantified by the dot product of the SOAP vectors, leading to the SOAP kernel:

$$K(\mathcal{R}_i, \mathcal{R}_j) = \left(\mathbf{q}^{(i)} \cdot \mathbf{q}^{(j)}\right)^\zeta, \tag{2.55}$$

where ζ is a positive integer (typically $\zeta > 1$) that adjusts the sensitivity of the similarity measure, determining how distinctly the kernel differentiates between atomic environments. This differs from the ACSFs used in Behler-Parrinello NNPs, where each atom is assigned its own neural network to predict local energies, and similarity is learned implicitly. While it was initially assumed that SOAP and related descriptors (such as those based on three-body correlations, including ACSFs) were "complete", meaning that any two atomic environments not related by symmetry would yield different descriptors, subsequent work demonstrated that this is not strictly the case. Counterexamples have shown that distinct structures can produce identical descriptors.[203, 204] In practice, however, the resulting errors are comparable to other systematic sources, such as those arising from locality assumptions.[205]

A key advantage of the GAP framework is its flexibility and hierarchical approach to descriptors. As discussed earlier for empirical potentials, capturing the complexity of atomic interactions often requires combining multiple levels of description. GAP achieves this by allowing the simultaneous use of several descriptors such as: two-body, three-body, and SOAP within a single model.[206] The regression is performed using a set of representative environments (sparse points as determined through CUR decomposition), enabling efficient evaluation even for large training datasets. The total energy is then expressed as a sum over the contributions from each descriptor:

$$E = \sum_i \epsilon_i^{(d)} = \sum_{t=1}^{N_t^{(d)}} \alpha_t^{(d)} K^d(\mathcal{R}_i^{(d)}, \mathcal{R}_t^{(d)}), \tag{2.56}$$

where $d \in \{\text{2b, 3b, SOAP}\}$ indicates the type of descriptor, $N_t^{(d)}$ is the number of sparse points for descriptor d, and $\alpha_t^{(d)}$ are the regression weights. The 2b and 3b descriptors correspond to pairwise and three-body interactions, respectively, while the SOAP descriptor captures the many-body interactions of each environment. The pairwise interaction is trivially a scalar value between two atoms, i, j:

$$q^{(\text{2b})} = |\mathbf{r}_i - \mathbf{r}_j|. \tag{2.57}$$

For the three-body interaction, the natural description is not used, but instead a different form is required to enforce permutation invariance of neighbouring atoms, j, k to i:[206]

$$q^{(3\mathrm{b})} = \begin{pmatrix} |\mathbf{r}_i - \mathbf{r}_j| + |\mathbf{r}_i - \mathbf{r}_k| \\ (|\mathbf{r}_i - \mathbf{r}_j| - |\mathbf{r}_i - \mathbf{r}_k|)^2 \\ |\mathbf{r}_j - \mathbf{r}_k| \end{pmatrix} \tag{2.58}$$

These descriptors are explicitly required for most systems (including carbon), as whilst the SOAP descriptor alone yields a very good description for regions where there are data, it provides erroneous descriptions for close-range interactions (sub Å), resulting in non-physical behaviour, predicting a lowering of the energy as the atoms approach each other. In practice, this means that the 2b and 3b descriptors must be included in the model, along with SOAP, to ensure accurate close range, repulsive interactions dominated by 2b+3b interactions in conjunction with the longer-range, many-body interactions captured by the SOAP descriptor. This is illustrated in

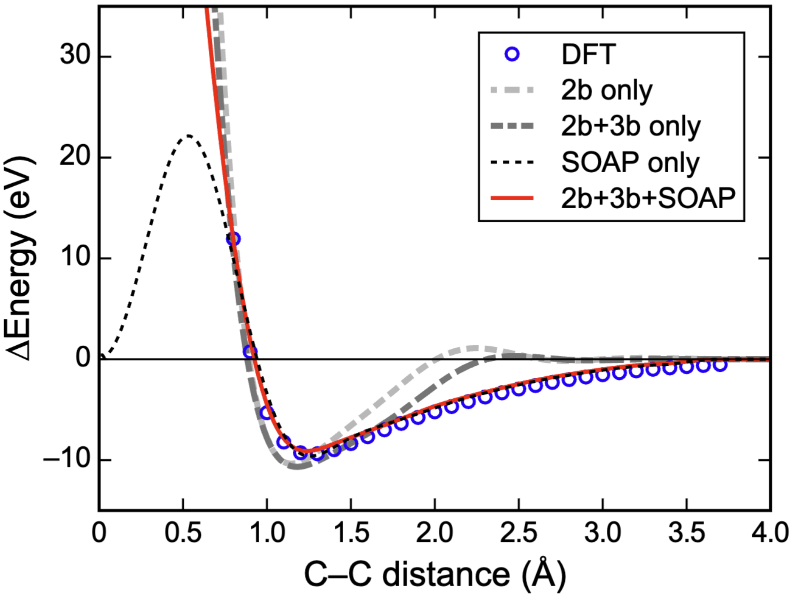

Figure 2.7: Potential-energy scans for an isolated carbon dimer. This plot, with DFT data as reference (blue), allows us to assess the use of different structural descriptors: all three combined are needed for a high-quality fit. Figure reproduced from Ref. 206 with permission from the American Physical Society.

Figure 2.7, which shows the potential energy surface of an isolated carbon dimer, with the DFT reference data shown in blue. The three different descriptors are shown in red (2b+3b+SOAP), black (SOAP only), dark grey (2b+3b) and light grey (2b). It is clear that all three descriptors are required to achieve a high-quality fit to the DFT data, particularly at short distances where the 2b and 3b descriptors dominate.

The resulting potential is able to interpolate smoothly between reference data points, providing DFT-level accuracy for energies and forces within the domain of the training data. It is also this smoothness in high-dimensional space which allows for the potential to be used in a general-purpose manner, as it can be applied (ideally) to configurations which lay outside the domain of the training data, albeit with a loss of numerical accuracy in some instances, but crucially, remaining physically meaningful.

The technical details regarding GPR regression and derivations of forces and energies are beyond the scope of this thesis, but the interested reader is referred to the original papers by Bartók *et al.*,[189, 202] and also the review by Deringer *et al.*[205]

C-GAP-17

The earliest MLIPs for carbon, based on both GAPs and NNPs, were primarily developed to describe crystalline phases such as graphite and diamond. These initial models were evaluated by their ability to reproduce the properties of these crystals and the graphite–diamond phase boundary.[202, 207, 208] Notably, one of the first neural network potentials was used in large-scale simulations to investigate diamond nucleation mechanisms.[209] However, these early MLIPs were trained on a limited range of configurations, resulting in poor transferability and limited performance for disordered carbon systems. By the early 2010s, the only MLIP available for amorphous materials was a NNP for the phase-change material GeTe,[210] which enabled large-scale simulations of thermal transport[211] and atomistic processes during crystallization.[212]

A significant milestone was reached in 2017 with the development of the C-GAP-17 machine learning interatomic potential by Deringer and Csányi.[206] Unlike earlier models, C-GAP-17 was explicitly constructed to capture both crystalline and disordered forms of carbon. The potential was trained and validated on a comprehensive dataset of 4,530 configurations (283,965 atoms) labelled with the local-density approximation (LDA), encompassing a broad spectrum of carbon environments including diamond, graphite, graphene, amorphous phases, surfaces, dimers, and liquid carbon. This diversity enabled C-GAP-17 to achieve accuracy and stability across a wide range of atomic configurations. The potential has been shown to reliably reproduce structural properties of various carbon allotropes and to accurately model the graphitisation process in amorphous carbon;[170] simulate amorphous carbon thin film growth;[103] describe surface properties and chemical reactivity of tetrahedral amorphous carbon;[213, 214] and investigate the structure of porous carbonaceous materials for energy storage applications;[215, 216] as well as support crystal structure prediction.[217]

In 2020, a new GAP potential was developed with a new dataset combining the original C-GAP-17 dataset augmented with newer structures, including nanotubes and fullerenes, and was re-labelled using the optB88-vdW, bringing the total of structures to 6,088 (400,275 atoms).[218] This was done with the aim of improving the accuracy required to capture subtle differences in formation energies for defects and nanostructures, as well as improving phonon predictions explicitly. In addition, dispersion corrections were included within the dataset to improve the description of van der Waals interactions, particularly for layered structures such as graphite and nanotubes. The resulting potential, known as C-GAP-20, has been shown to outperform the original C-GAP-17 in terms of numerical accuracy, particularly for defects and phonons, however the model displayed a fictitious minimum for body-centred cubic carbon, which was not present in the original C-GAP-17 potential.[218]

This minimum was later shown to be an artefact of the training data, and was removed in the C-GAP-20-U potential. As a result, the C-GAP-17 potential was used for carbon only simulations.

In addition, GAP potentials can be fitted and used within `CASTEP` for "on-the-fly" acceleration of AIMD simulations.[219] In brief, GAP potentials can be fit as the AIMD calculation is running, generating a very small initial dataset of highly correlated structures, at which point the GAP potential that is fit is used to compute the forces and energies of the system, which are then used to update the positions of the atoms. This process is repeated until the simulation reaches a desired length or the checks implemented within `CASTEP` trigger a refit based on numerical errors above some threshold and more data is added to the set. This enabled three principle modes of operation: (1) using GAP to accelerate AIMD simulations, (2) generating a single-use GAP potential to be used in other MD engines, and (3) generating a dataset of structures to be used for further training or validation of MLIPs. Such on-the-fly GAP-potentials for carbon have demonstrated remarkable success in modelling the graphitisation process of amorphous carbon despite much fewer structures and cheaper models being used.[219] Further details can be found in the Appendix.

Whilst GAP has demonstrated remarkable success in modelling carbon,[103, 170, 206, 214] it is important to note that the potential is not without limitations. The potential is notably slower than other MLIPs, such as Behler-Parrinello NNPs, due to the computational cost of evaluating the SOAP descriptor and the GPR regression. This can be a significant drawback for large-scale simulations or when high throughput is required. Additionally, while GAP has been shown to be transferable across a wide range of carbon allotropes, it may still struggle with certain configurations that are not well represented in the training data, particularly nanotubes, fullerenes and defects.[218] With the turn of the decade into the 2020s, the field of MLIPs has seen a rapid evolution, with new frameworks and models emerging that aim to address these limitations and further improve the accuracy and efficiency of interatomic potentials. These developments are discussed in the next section.

Atomic Cluster Expansion (ACE)

The Atomic Cluster Expansion (ACE) is a systematic framework for building many-body basis functions cheaply.[220] Put simply, ACE is based on the idea of building higher-body order terms from one- and two-body terms using a body-order expansion along with a mathematical trick which enables the cheap generation of these higher-body order terms. The crux of the problem lies in the expansion of the total energy into sums of local, site energies:

$$\begin{aligned}
\epsilon_i = V_0(z_i) & \\
& + \sum_{j_1} V_1(\boldsymbol{r}_{ij_1}, z_{j_1}; z_i) \\
& + \sum_{j_1 < j_2} V_2(\boldsymbol{r}_{ij_1}, \boldsymbol{r}_{ij_2}, z_{j_1}, z_{j_2}; z_i) \\
& + \ldots \\
& + \sum_{j_1 < \ldots < j_n} V_n(\boldsymbol{r}_{ij_1}, \ldots, \boldsymbol{r}_{ij_n}, z_{j_1}, \ldots, z_{j_n}; z_i),
\end{aligned} \tag{2.59}$$

where V_n is the n-body potential term, $\boldsymbol{r}_{ij}$ is the distance between atoms i and j within some cutoff, and z_i is the chemical element of atom i. As can be seen quite clearly, the number of terms grows exponentially with the body order n, leading to a combinatorial explosion in the number of terms that need to be computed. Counterintuitively, the key idea is to compute all possible interactions, including self-interactions:

$$\begin{aligned}\epsilon_i = {} & U_0(z_i) \\ & + \sum_{j_1} U_1(\boldsymbol{r}_{ij_1}, z_{j_1}; z_i) \\ & + \sum_{j_1,j_2} U_2(\boldsymbol{r}_{ij_1}, \boldsymbol{r}_{ij_2}, z_{j_1}, z_{j_2}; z_i) \\ & + \ldots \\ & + \sum_{j_1,\ldots,j_n} U_n(\boldsymbol{r}_{ij_1}, \ldots, \boldsymbol{r}_{ij_n}, z_{j_1}, \ldots, z_{j_n}; z_i) + U_n(\boldsymbol{r}_{ii}, z_i), \end{aligned} \tag{2.60}$$

so now, the summation is not over unique pairs of atoms, but rather over all possible pairs, including self-interactions. Simplifying our expression $U(\boldsymbol{r}_{ij}, z_j; z_i)$ to $U(x_j)$, the potentials U_n are then expanded into a tensor product basis:

$$\sum_{j_1,\ldots,j_n} U_n(x_{j_1}, \ldots, x_{j_n}) \approx \sum_{j_1,\ldots,j_N} \sum_{k_1,\ldots,k_N} c_{\boldsymbol{k}} \prod_{t=1}^{N} \phi_{k_t}(x_{j_t}), \tag{2.61}$$

where $c_{\boldsymbol{k}}$ are the coefficients of the tensor product basis, ϕ_{k_t} are the one-particle basis functions, which are given as:

$$\phi_{nlm}(x_j) = g_n(r_{ij}, z_j) Y_{lm}(\hat{\mathbf{r}}_{ij}), \tag{2.62}$$

where $g_n(r_{ij}, z_j)$ is a radial function that depends on the distance between atoms i and j, and $Y_{lm}(\hat{\mathbf{r}}_{ij})$ are spherical harmonics. Performing the tensor product with itself gives an N-particle basis function. The next stage is to perform the so-called "density trick", introduced originally to construct SOAP and bispectrum descriptors,[189] where the tensor product and outer summation are interchanged:

$$\sum_{j_1,\ldots,j_N} \sum_{k_1,\ldots,k_N} c_{\boldsymbol{k}} \prod_{t=1}^{N} \phi_{k_t}(x_{j_t}) = \sum_{k_1,\ldots,k_N} c_{\boldsymbol{k}} \prod_{t=1}^{N} \left(\sum_{j} \phi_{k_t}(x_j) \right), \tag{2.63}$$

where the N dimensional sum over the indices $j_1, \dots, j_N$ is now replaced by a single sum over the indices j, which is computationally "well-behaved". The first step is to compute the inner sum over the indices j:

$$A_k = \sum_j \phi_k(x_j), \tag{2.64}$$

followed by forming the product basis:

$$\boldsymbol{A}_{k_1,\dots,k_N} = \prod_{t=1}^{N} A_{k_t}. \tag{2.65}$$

Finally, to enforce invariance, the product basis is symmetrised by averaging over all rotations by integrating over 3D space:

$$B_{k_1,\dots,k_N} = \int_{\mathcal{O}(3)} \boldsymbol{A}_{k_1,\dots,k_N}(R)dR, \tag{2.66}$$

where the integral is over all space. The final step is to compute the coefficients $c_{\boldsymbol{k}}$ by fitting to the reference data, which can be done using a variety of methods, such as least-squares fitting or maximum likelihood estimation, to then obtain the site energies:

$$\epsilon_i = \sum_{k_1,\dots,k_N} c_{\boldsymbol{k}} B_{k_1,\dots,k_N}(x_i). \tag{2.67}$$

The total energy of the system is then obtained by summing over all atomic contributions:

$$E = \sum_i \epsilon_i. \tag{2.68}$$

The resulting ACE descriptors are described as being "complete", meaning that ACE can be expanded to an arbitrary body order and thus in principle can distinguish between any two atomic environments that are not related by symmetry.[203] This is in contrast to the SOAP descriptor used in GAP, which is not strictly complete, as discussed earlier.[203, 204]

MACE

MACE is a state-of-the-art MLIP framework that combines the systematic body-order expansion of ACE with the flexibility and expressivity of message-passing neural networks.[221] We first introduce two concepts that are important to understanding the design choices of MACE: graph neural networks (GNNs) and message passing.

A natural way to represent atomic systems is as graphs, where atoms are nodes connected by edges. Geometric graphs embed additional geometric information, where in our case, the nodes represent atoms in 3D Euclidian space with scalar attributes (such as atom type) and geometric attributes (such as position, velocity, forces, etc.). At the core of GNNs is the message-passing paradigm, which transforms a GNN to a message-passing neural network (MPNN). In this process, each node (atom) in the graph sends "messages" containing information about its own state to its neighbours. These messages are then aggregated by the receiving nodes, which update their own state based on the information they have received. This iterative process allows information to propagate throughout the graph, enabling each node to develop a representation that is informed by its local atomic environment. The number of message-passing iterations, or layers, along with a radial cutoff at each layer determines the size of the receptive field from which each atom gathers information.

A node feature can contain information about the atom's type, position, and other properties. In MACE, the node features of node i at iteration t are:[221]

$$s_i^{(t)} = \left(\mathbf{r}_i, z_i, \mathbf{h}_i^{(t)}\right), \tag{2.69}$$

where $\mathbf{r}_i$ is the position of atom i, z_i is the chemical element, and $\mathbf{h}_i^{(t)}$ are the learnable features. As such, the node features s_i are updated from iteration t to $t+1$ in three steps: (i) message construction, (ii) aggregation, and (iii) update.

The message construction step involves computing messages from each node i and each of its neighbours $\mathcal{N}_i$ in the graph through a message function. The aggregation step combines the messages received by each node from its neighbours through a permutation-invariant function (e.g., sum or mean). Finally, the update step applies a function to the aggregated message to produce the new node feature $s_i^{(t+1)}$ for the next iteration. The form of the message function and the aggregation function can vary from framework to framework, but they are typically neural networks. Formally, the message-passing process can be expressed as:

$$\boldsymbol{m}_i^{(t)} = \bigoplus_{j \in \mathcal{N}_t} M_t(s_i^{(t)}, s_j^{(t)}), \tag{2.70}$$

where $\bigoplus$ is a learnable permutation-invariant aggregation function (e.g., sum or mean) over neighbours $\mathcal{N}_t$, M_t is the learnable message function at iteration t, and $\boldsymbol{m}_i^{(t)}$ is the message sent by node i to its neighbours. The update step then combines the messages received by each node and transforms the node features:

$$s_i^{(t+1)} = U_t\left(s_i^{(t)}, \boldsymbol{m}_i^{(t)}\right), \tag{2.71}$$

where U_t is a learnable update function at iteration t. This process is repeated for T iterations, allowing information to propagate through the graph and enabling each node to develop a representation that is informed by its local atomic environment. After T iterations, the final node features $s_i^{(T)}$ can be used to predict properties of the system, such as site energies:

$$\epsilon_i = \sum_{t=1}^{T} \mathcal{R}_t(s_i^{(t)}), \tag{2.72}$$

where $\mathcal{R}_t$ is a learnable readout function that maps the final node features to the predicted site energy ϵ_i for atom i. The total energy of the system can then be obtained by summing over all atomic contributions:

$$E = \sum_i \epsilon_i. \tag{2.73}$$

One of the outcomes of using MPNNs can be qualitatively thought of as a topological sparsification of relevant basis functions which are "chemically important" environments through iterative message passing. For example, a one-layer linear ACE model with a cutoff of 8 Å, is equivalent to an MPNN with a radial cutoff of 4 Å with two layers (Figure 2.8), and can be generally expressed as:

$$f_{\text{receptive}}(r) = T \times f_{\text{cut}}(r), \tag{2.74}$$

where $f_{\text{cut}}(r)$ is the radial cutoff of the MPNN, and T is the number of message-passing iterations. As such, if an atom is not within the radial cutoff at layer t then it is not considered, whereas for an ACE potential the radial cutoff is fixed and all atoms within the cutoff are considered, potentially leading to a large number of irrelevant basis functions being included in the model. Indeed, the MPNN formalism

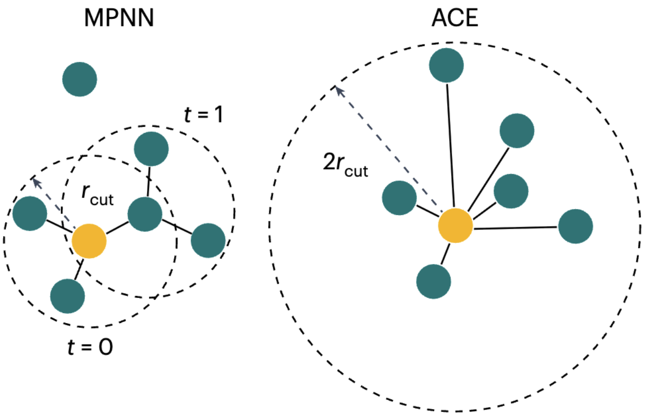

Figure 2.8: Receptive field of MPNNs. Comparison of the clusters formed by two iterations of message passing with cutoff r_{cut} at each iteration on the left for an MPNN and the clusters formed by ACE with cutoff $2r_{cut}$ on the right. In principle, both methods incorporate information from a distance of up to $2r_{cut}$, but in the case of the MPNN, only atoms that can be reached through a chain of closer intermediates contribute. Figure reproduced from Ref. 222. (CC-BY 4.0)

sparsifies according to walks along the graph of atoms, which is chemically intuitive as most physical interactions are short-ranged, and medium-range interactions occur through chain-like mechanisms, where atoms are connected through a series of bonds.[222–224] This is in contrast to ACE, which considers all atoms within the cutoff regardless of their connectivity, leading to a denser representation.

So far, the description of MACE has been similar to that of a standard MPNN. However, MACE distinguishes itself by its ability to construct messages that capture higher-order correlations between atoms, which is a key feature inherited from the ACE framework. MACE leverages this powerful framework by constructing messages that are not limited to pairwise interactions. Instead, it employs a many-body expansion to formulate messages that explicitly include information about triplets, quadruplets, and even higher-order correlations of atoms. These "higher-order" messages provide a *formally* complete description of the local atomic environment. The messages can be expressed as:

$$\boldsymbol{m}_i^{(t)} = \sum_j \boldsymbol{u}_1\left(s_i^{(t)}, s_j^{(t)}\right) + \sum_{j,k} \boldsymbol{u}_2\left(s_i^{(t)}, s_j^{(t)}, s_k^{(t)}\right) + \sum_{j,k,l} \boldsymbol{u}_3\left(s_i^{(t)}, s_j^{(t)}, s_k^{(t)}, s_l^{(t)}\right) + \ldots, \tag{2.75}$$

where $\boldsymbol{u}_1$, $\boldsymbol{u}_2$, $\boldsymbol{u}_3$, etc., are learnable functions that compute messages based on the features of the atoms involved and the sum runs up to a user-defined hyperparameter, the maximum correlation order $\mathcal{V}$, which is the body order - 1, i.e., $\mathcal{V} = 1$ corresponds to pairwise messages, $\mathcal{V} = 2$ to triplet messages, and so on.

The composition of the message is built from a set of one-particle basis functions $\phi_i^{(t)}$ (which involve a two-body interaction between atom i and its neighbour j) comprised of a learnable radial radial functions, $g_n(r)$, and spherical harmonics, $Y_{lm}(\hat{\mathbf{r}})$ and a learnable embedding of the previous node features $W^{(t)} \cdot s_i^{(t)}$, where $W^{(t)}$ is a learnable weight matrix at iteration t. These one-particle basis functions (which can be thought of as our edge features) are then combined to form the atomic basis, which can be expressed as[*]:

$$A_i^{(t)} = \sum_{j \in \mathcal{N}_i} \phi_i(s_i^{(t)}, s_j^{(t)}). \tag{2.76}$$

Then, to build higher body-order messages, $\mathcal{V}$ products of the atomic basis are constructed, which can be expressed as:

$$\boldsymbol{A}_{\boldsymbol{i}}^{(t)} = \prod_{\nu=1}^{\mathcal{V}} A_{i,\nu}^{(t)}. \tag{2.77}$$

Rotational invariance is then achieved by symmetrising the basis functions by integrating over all possible rotations:

$$B_i^{(t)} = \int_{\mathcal{O}(3)} \boldsymbol{A}_i^{(t)} \left([R \cdot (s_i^{(t)}, s_j^{(t)})]_{j \in \mathcal{N}_i} \right) dR, \tag{2.78}$$

where $\mathcal{O}(3)$ is the orthogonal group of rotations in three dimensions, and R is the rotation on a pair of atoms (s_i, s_j). In practice, tensor contractions are used as the angular dependance can be expressed in terms as products of spherical harmonics.

Another important aspect of MACE is its equivariance to rotations, translations, and permutations of identical atoms. Equivariance means that under some transformation of the input data (e.g., a rotation or translation of the atomic positions), the output of the model transforms in a predictable and well-defined manner. In the context of MACE, this means that if the system is rotated then any scalar properties (such as energies) should remain unchanged (invariant), whereas vector properties (such as forces[†]) should transform according to the rotation. As such, the representation of the symmetrised basis through the use of spherical harmonics

---

[*] Note that many indicies are left out for brevity, namely the radial n and angular l indices of the basis functions, as well as the correlation order $\mathcal{V}$, channels k and Clebsch–Gordan indicies used for equivariance. The interested reader is referred to the original MACE paper by Batatia *et al.* for the full details.

[†] The forces which we define are based on the Hellmann–Feynman forces which are defined as the derivatives of the total energy with respect to atomic positions ($F = \frac{\partial E}{\partial \boldsymbol{r}_i}$) and thus also sums of derivatives of the atomic energy function. Thus the forces implicitly rotate with coordinates, and energy is always conserved. These are known as constrained models.

readily extends to enforce equivariance by filtering out basis functions that do not obey the required symmetries. The technical details of how this is done are left out for brevity. Finally, the messages are constructed from the equivariant symmetrised basis functions*:

$$m_i^{(t)} \propto \sum W_B^{(t)} \cdot B_i^{(t)}, \tag{2.79}$$

where the B subscript indicates these are the weights for the basis functions. The next step is the update phase which applies a learnable linear combination of the messages to the node features, which can be expressed as*:

$$s_i^{(t+1)} \propto \sum W_m^{(t)} \cdot m_i^{(t)}, \tag{2.80}$$

where the m subscript indicates these are the weights for the messages. The final step is to readout the atomic energies from the invariant part of the node features, which can be expressed as:

$$\epsilon_i = \sum_{t=1}^{T} \epsilon_i^{(t)}, \tag{2.81}$$

where $\epsilon_i^{(t)}$ is the readout of the node features at layer t. The final readout function is a a non-linear, one layer MLP.

2.2.4 MACE foundation models

The MACE framework has demonstrated high accuracy and transferability across a wide range of systems, including carbon, silicon, and other materials.[225] It has been shown to outperform existing MLIPs in terms of accuracy and transferability, particularly for disordered systems.[221, 222] Recently, the MACE framework has been used to develop a "foundation model" for atomistic modelling of materials spanning 89 elements of the periodic table trained on a public dataset provided by the Materials Project.[225]† This dataset contains 1.58 million structures from

---

*Proportionality is used here to emphasise that there are missing indicies.

†The results shown below were completed as part of a wider collaboration for Ref. 225. The author of this thesis contributed the graphitisation simulations seen in Figure 2.9 (b–h) as a

146,000 materials labelled at the PBE level of theory, and is referred to as "MACE-MP-0b3". An updated model was trained on an additional 10.4 million structures from 3.23 million materials from a subset of the Alexandria dataset[226] and is referred to as "MACE-MPA-0".[225]

Melt–quench simulations

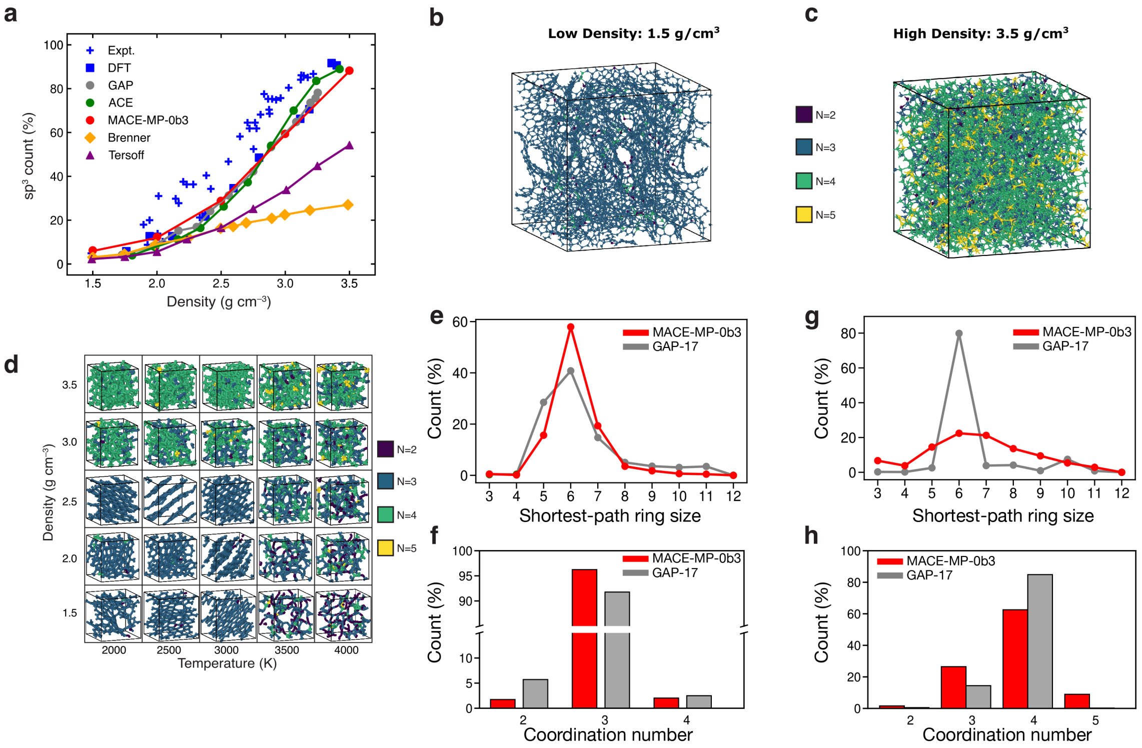

Figure 2.9: (a) Count of sp^3 (fourfold coordinated) carbon atoms in melt-quenched carbon structures as a function of density. The results obtained with the MACE-MP-0b3 model are compared to computational and experimental data compiled in Ref. [206] and references therein, as well as Refs. [227] and [228]. (b–c) 4096-atom structures generated using the MACE-MP-0b3 potential. (d) Results from 25 × 200 atom graphitisation simulations spanning relevant temperature and density ranges, similar to Ref. [219]. The structures in panels (b–d) are colour-coded according to coordination numbers as indicated in the legends. (e) Shortest-path ring size count for 4096-atom low-density structures as determined using `matscipy`.[229] (f) Coordination number count for 4096-atom low-density structures. (g–h) Same as panels (e) and (f) but for high-density structures. The results in panels (e–h) are shown for a C-GAP-17-driven simulation (grey) and for a comparable simulation driven by the MACE-MP-0b3 potential (red).

---

benchmark to assess the performance of the MACE foundation models for amorphous carbon, and the author also contributed the analysis of the results in Figure 2.10. The author also contributed to some of the writing of the manuscript.

The accuracy of the MACE-MP-0b3 model in reproducing the structural complexity of amorphous phases is demonstrated by plotting the concentration of four-fold coordinated atoms (sp^3) as a function of density in Figure 2.9a. To generate amorphous structures with a given density using the MACE-MP-0b3 model, we perform melt-quench simulations. We start by melting diamond structures at a given density by running NVT simulations at 8000 K for 3 ps. We then perform a fast quench, reducing the temperature from 8000 K to 300 K at a cooling rate of 1000 K s^{-1}. Finally, we optimise the geometry using the limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm (LBFGS)[230] to obtain the final structure and determine the count of sp^3 atoms using a bond-length cutoff of 1.85 Å. We observe in Figure 2.9 that the MACE-MP-0b3 model predictions reproduce the trend observed in both the DFT [227] and the experimental data extracted from [206]. We also see good agreement with results of quenches using the carbon ACE potential reported by Qamar *et al.* in Ref. [228], and using C-GAP-17 reported in Ref. [227], noting that both potentials had been specifically trained on large carbon datasets.

Graphitisation

MACE-MP-0b3 was used for two annealing runs for simulation cells containing 4096 atoms at low density (1.5 gcm^{-3}) and high density (3.5 gcm^{-3}), respectively to sample graphitisation behaviour at densities resembling both graphite (low density) and diamond (high density). The low-density simulation was run at 2000 K and the high-density simulation was run at 4000 K. Additionally, 25 × 200 atom annealing runs spanning from 2000 to 4000 K and 1.5 to 3.5 gcm^{-3} were conducted to explore a finer grid of parameters.

Both 4096-atom structures were compared to structures generated using C-GAP-17 using the same protocol, which was also used recently for an "on-the-fly" generated GAP potential. [219] The protocol has two stages. The goal of stage I is to prepare the starting configuration for the annealing in stage II, and it begins with a random structure with a hard-sphere constraint of $r_{\mathrm{min}} \geq 1$ Å and equilibrating it at

9000 K for 40 ps, followed by cooling to 5500 K over 40 ps and subsequent quenching to 300 K over 10 ps. The structures are then held at 300 K over 50 ps before being rapidly heated up to the annealing temperature over 10 ps. This concludes stage I. In stage II, the structure is annealed at 2000 K or 4000 K for 350 ps using a time step of 1 fs. We used C-GAP-17 to perform stage I because MACE-MP-0b3 was found to be unstable at 9000 K.

Figures 2.9b–c show the final structures generated using MACE-MP-0b3. Qualitatively, the low-density structure contains graphitic regions along with large pores of a few nm and sp chains, and the high-density structure contains highly ordered diamond-like regions. This is similar to structures generated with C-GAP-17 where similar features are also observed – see, for example, a comprehensive survey by de Tomas *et al.* (Ref. [170]). The same can be observed for a more comprehensive set of smaller structures (Figure 2.9d). For the low-density case, the structures generated by MACE-MP-0b3 in annealing simulations agree with the predictions of the established C-GAP-17 model in terms of overall trends (Figures 2.9b and 2.9d). For more detailed insight, Figure 2.9e shows the shortest-path ring distribution for the low-density structure, indicating that MACE-MP-0b3 predicts a greater number of 6-membered rings and fewer large rings for the low-density structure compared to C-GAP-17 (thus suggesting a higher degree of crystallinity in the MACE-MP-0b3 prediction) – this might be correlated with the higher relative count of graphite versus amorphous-like structures in the training dataset, although further analysis is required.

Figure 2.9g shows poor agreement between MACE-MP-0b3 and C-GAP-17 in terms of ring count for the high-density structure. Whereas annealing with C-GAP-17 led to a partly crystallised structure under the conditions chosen (indicated by the large abundance of 6-membered rings, as are found in the diamond structure), the MACE-MP-0b3 model gave rise to a highly disordered structure – this is evident from a visual inspection of Figure 2.9c, and reflected in the ring-count plot in Figure 2.9g showing a notable number of (presumably strained) 3-membered rings as well as larger ring sizes of > 7. In other words, the MACE-MP-0b3 appears to fail to describe the crystallisation in this case (we note that a partially ordered structure was obtained in a 200-atom simulation at 3000 K but not at 4000 K at the same density; Figure 2.9d).

Finally, Figures 2.9f and 2.9h show the coordination number for both 4096 atom structures. MACE-MP-0b3 leads to more sp^2 environments and fewer sp and sp^3 environments compared to C-GAP-17 (again likely consistent with a higher degree of crystallinity, as suggested by the ring count in Figure 2.9e). In the high-density 4096-atom structure obtained with MACE-MP-0b3, a notable amount of 5-fold coordinated atoms are observed, on the order of about 10% – a behaviour that would not be expected in a simulation of carbon at diamond-like density.

We note that these issues are not expected to be a fundamental shortcoming of the MACE architecture: we show that training on a wider-ranging dataset with more crystal structures like the Alexandria dataset (which notably has no amorphous structures) results in a model, MACE-MPA-0, that describes the structure of amorphous carbon in better agreement with C-GAP-17 at high and low density, as shown in the next section.

MACE-MPA-0 model

To investigate the effect of training set size, composition, and method of generation on model performance, we developed a MACE model, MPA-0, using identical hyperparameters to the MP-0b3 medium model but trained on an expanded dataset combining MPtraj and sAlex [226]. The sAlex dataset comprises 10.5M structures extracted from the original Alexandria dataset, curated for compatibility with the Matbench Discovery benchmark. The Alexandria dataset consists only of randomly generated and relaxed crystal structures. Here, we demonstrate that increased data volume on just crystal structures significantly enhances model robustness across multiple applications. Notably, we observe strong enhancement in the accuracy of the MPA-0 model compared to MP-0b3 on applications for which coverage was not increased, namely amorphous phases and small molecules on surfaces.

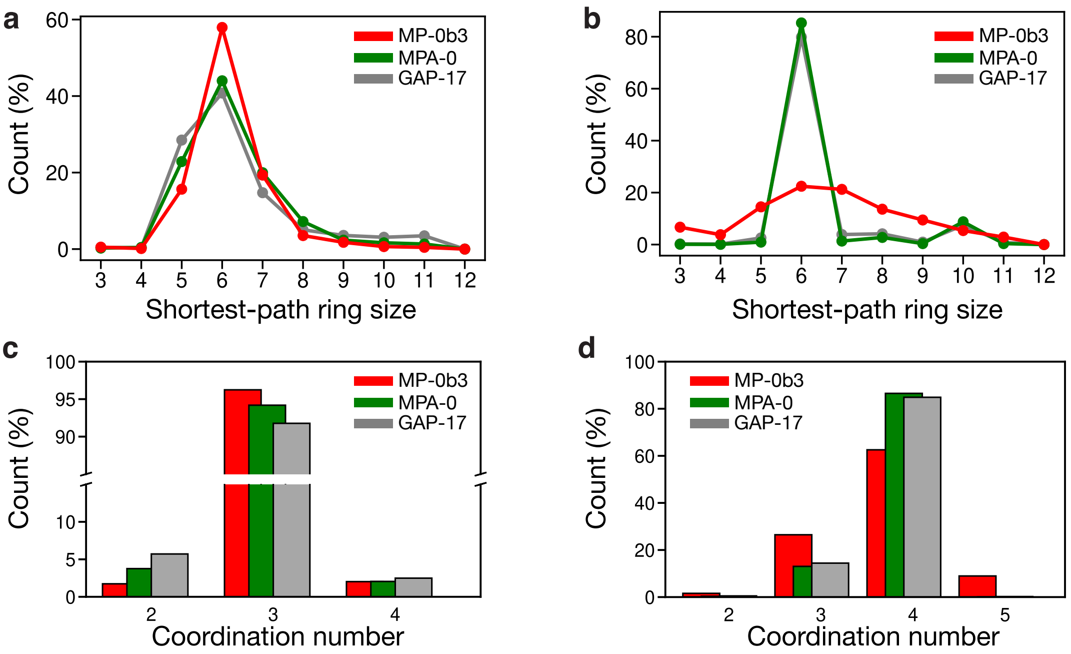

Figure 2.10: Structural analysis of 4096-atom amorphous carbon configurations. (a-b) Shortest-path ring size distributions and (c-d) coordination number distributions for MP-0b3 (red), MPA-0 (green), and GAP-17 reference model (grey). MPA-0 better captures the correct bonding environments and medium-range order compared to MP-0b3.

Accurate modelling of amorphous materials requires precise representation of local atomic environments and bonding. Figure 2.10 compares the structural properties of amorphous carbon predicted by MP-0b3, MPA-0, and the reference GAP-17 model. While MP-0b3 performs qualitatively well at low density (panel a-c), MPA-0 substantially improves the distribution of shortest-path ring sizes for the high density phases, showing a sharper peak at 6-membered rings (panel b) that closely aligns with the result for GAP-17, whereas MP-0b3 displays a broader, less defined distribution across ring sizes 6-8. For coordination numbers at high density (panel d), MPA-0 reproduces the correct sp^3:sp^2 ratio with approximately 85% 4-coordinated and 15% 3-coordinated carbon atoms, closely matching GAP-17. In contrast, MP-0b3 shows excess 3-coordinated (sp^2) and 5-coordinated carbon atoms, indicating structural deficiencies in the predicted amorphous networks. The Alexandria dataset does not contain any amorphous structures, demonstrating that an increase in broad chemical coverage of crystals helps in predicting amorphous phases.

Overall, both foundation models show reasonable accuracy and transferability across a wide range of materials, with MACE-MPA-0 outperforming MACE-MP-0b3 in terms of accuracy and transferability. However, caution must be exercised as these models do not reach *chemically* accuracy ($\leq$ 1 kcal mol^{-}1 $\approx$ 0.043 eV At.$^{-1}$) for some systems (e.g., atomisation energy of carbon has a RMSE of 0.459 eV At.$^{-1}$ when using MACE-MP-0b3).[225] The accuracy can be improved by including additional data (as is the case with MACE-MPA-0) or the foundation model can be "fine-tuned", which is when a pre-trained model is further trained on a smaller, more relevant dataset to improve its performance on a particular task.

2.3 Remaining structure of the thesis

The remainder of this thesis is comprised of three results chapters which are self-contained, first-author published papers. The chapters are arranged in chronological order of publication and as a result reflect the technology available at the time. The third chapter of the thesis uses C-GAP-17 to investigate the structure of amorphous carbon driven by atomic energies via a Monte-Carlo simulation. The fourth chapter applies recent developments in on-the-fly acceleration within `CASTEP`, using a GAP potential to efficiently "seed" an initial dataset. This dataset is then used to train a preliminary MACE model, which is iteratively improved by generating new graphene oxide structures—featuring varying degrees of functionalisation and at progressively higher temperatures—via MD and retraining on the resulting data until the model has converged on numerical energy and force errors. This MACE model (which we coin "`GO-MACE-23`") was then used to thermally reduce a nanoscale graphene oxide structural model and investigate the dynamic transformation along with simulated X-ray photoelectron spectroscopy (XPS). The fifth chapter of the thesis applies `GO-MACE-23` to investigate the relationship between mechanical properties and functionalisation of graphene oxide, and the effect of thermal reduction on these properties. The final chapter of the thesis provides a summary of the work presented in this thesis, along with an outlook on future work and potential applications of the MACE framework to other materials systems.

Chapter 3

Exploring the configurational space of amorphous graphene with machine-learned atomic energies

This chapter is reproduced in its entirety from the following publication:

> Zakariya El-Machachi, Mark Wilson, and Volker L. Deringer, "Exploring the configurational space of amorphous graphene with machine-learned atomic energies", *Chemical Science*, 2022, **13**, 13720.

Authors' contributions: Z. E.-M. developed the Monte-Carlo protocols and carried out the computational work. V. L. D. initiated and coordinated the study. All authors made substantial contributions to data analysis and discussions. Z. E.-M. and V. L. D. wrote the paper with input from M. W., and all authors approved the final version.

Volume 13
Number 46
14 December 2022
Pages 13635–13968

Chemical Science

rsc.li/chemical-science

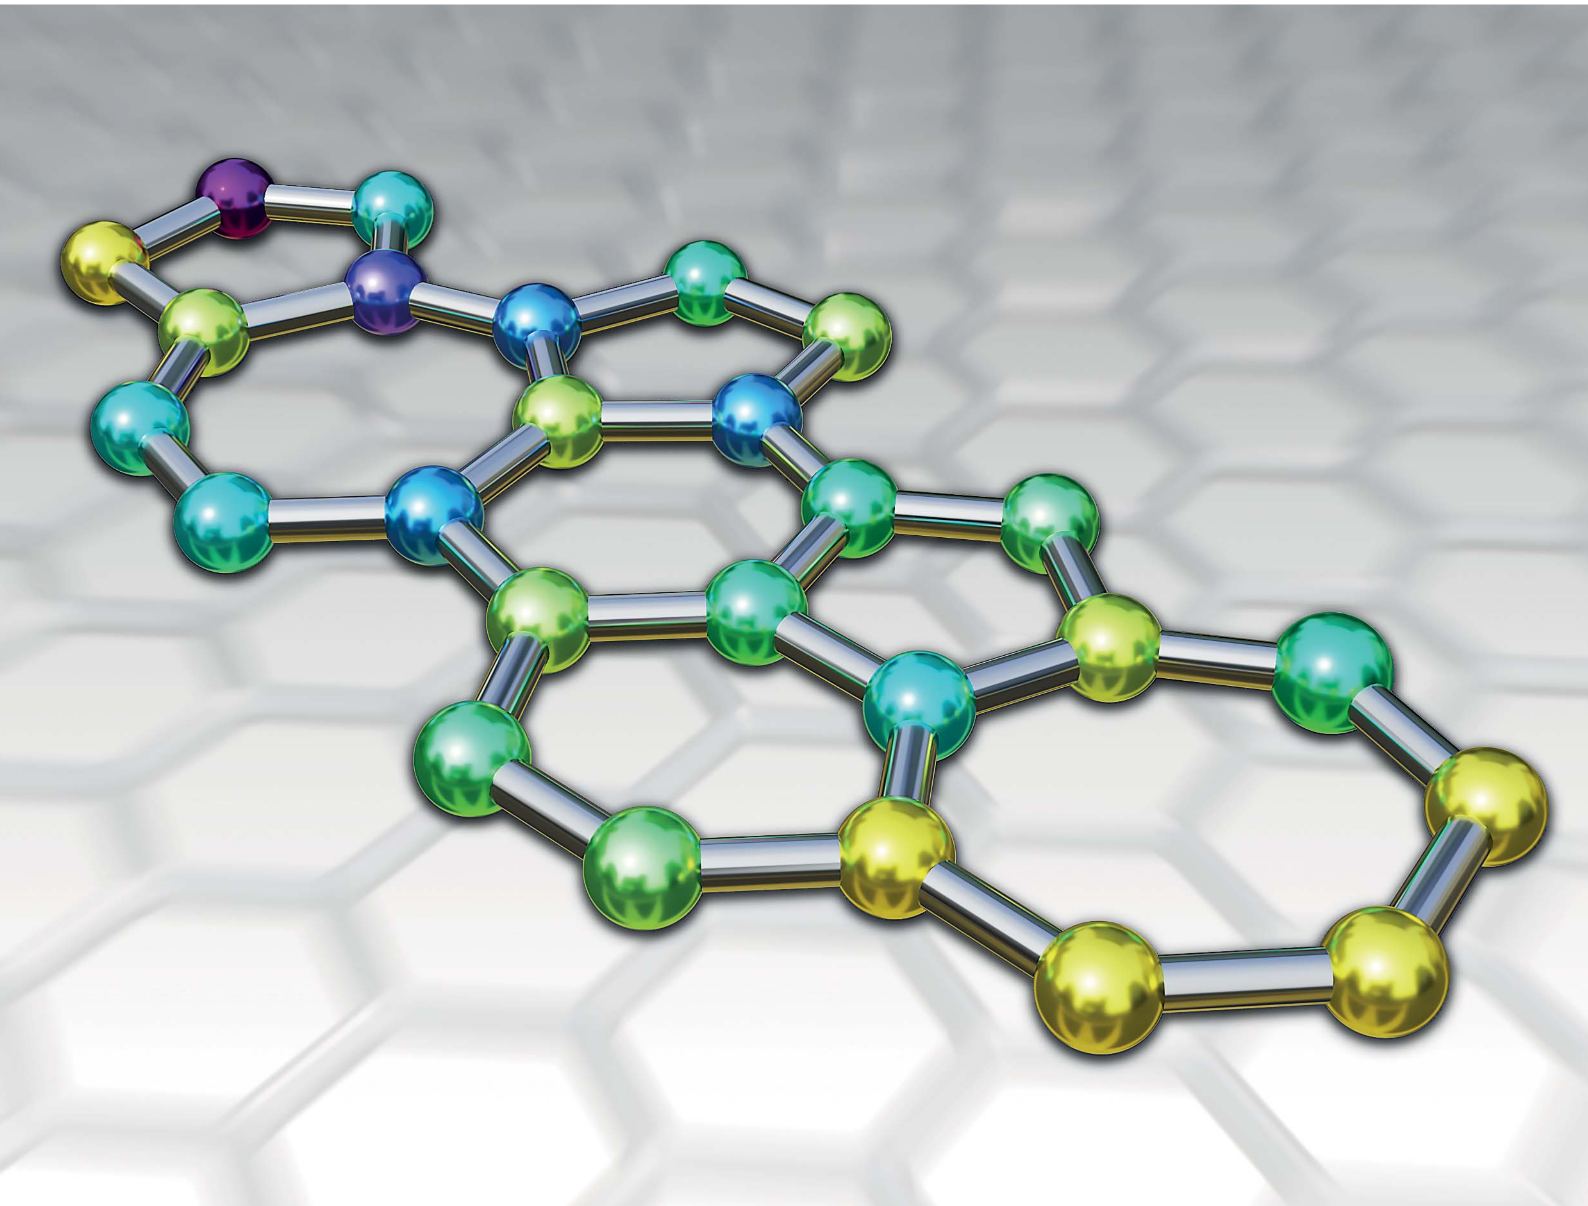

ISSN 2041-6539

ROYAL SOCIETY
OF **CHEMISTRY**

EDGE ARTICLE
Volker L. Deringer *et al.*
Exploring the configurational space of amorphous graphene with machine-learned atomic energies

Chemical Science

ROYAL SOCIETY OF CHEMISTRY

EDGE ARTICLE

View Article Online
View Journal | View Issue

Check for updates

Cite this: *Chem. Sci.*, 2022, **13**, 13720

All publication charges for this article have been paid for by the Royal Society of Chemistry

Received 3rd August 2022
Accepted 14th October 2022

DOI: 10.1039/d2sc04326b

rsc.li/chemical-science

Exploring the configurational space of amorphous graphene with machine-learned atomic energies

Zakariya El-Machachi,[a] Mark Wilson[b] and Volker L. Deringer*[a]

Two-dimensionally extended amorphous carbon ("amorphous graphene") is a prototype system for disorder in 2D, showing a rich and complex configurational space that is yet to be fully understood. Here we explore the nature of amorphous graphene with an atomistic machine-learning (ML) model. We create structural models by introducing defects into ordered graphene through Monte-Carlo bond switching, defining acceptance criteria using the machine-learned local, atomic energies associated with a defect, as well as the nearest-neighbor (NN) environments. We find that physically meaningful structural models arise from ML atomic energies in this way, ranging from continuous random networks to paracrystalline structures. Our results show that ML atomic energies can be used to guide Monte-Carlo structural searches in principle, and that their predictions of local stability can be linked to short- and medium-range order in amorphous graphene. We expect that the former point will be relevant more generally to the study of amorphous materials, and that the latter has wider implications for the interpretation of ML potential models.

Introduction

The study of the amorphous state has long been of fundamental research interest.[1] It is also increasingly important to understand structure–property correlations in amorphous materials, owing to ubiquitous applications in solar cells,[2] transparent electronic devices,[3] or phase-change memories.[4] Whilst bulk amorphous phases are challenging to study structurally, two-dimensional (2D) amorphous materials can be directly visualized by atomistic imaging techniques such as high-resolution transmission electron microscopy (HRTEM).[5–11] And just like graphene is the prototypical ordered 2D material, there is ongoing research interest in its disordered analogue(s). Indeed, Toh *et al.* recently synthesized a centimeter-scale sample of free-standing monolayer amorphous carbon[12] and characterized the structure based on the interpretation of HRTEM images.

The amorphous forms of carbon have been widely studied using computer simulations. Most commonly, these studies have been carried out in the framework of molecular dynamics (MD), from early work on melt-quenching[13–19] to direct simulations of thin-film growth by ion deposition.[20–23] Very recently, the properties of "amorphous graphite", as a 3D extended arrangement of individual amorphous graphene sheets, were investigated with density-functional theory (DFT) based MD and electronic-structure computations.[24] Earlier simulation studies had emphasized the connection between low-density amorphous carbon and the idealized case of 2D amorphous graphene ("aG" in the following).[25,26]

Graphene itself, as a 2D system, is rather well-confined and allows for Monte Carlo (MC) simulations to model topological defects. The Wooten–Winer–Weaire (WWW) algorithm, initially proposed for silicon,[27] remains a simple yet robust approach to generating continuous random network (CRN) models. In this, disorder is gradually introduced through local bond transpositions and structural relaxation. To simulate aG, bond transpositions are introduced as Stone–Wales (SW) defects,[28] with moves being accepted or rejected based on a suitable Metropolis criterion, typically the total-energy difference between the new and old configuration. SW defects are created by a formal in-plane 90° rotation of two atoms around the mid-point of the bond[28] and are the foundational example of (intrinsic) topological disorder in 2D carbon.[29–32] This approach has long been used to study aG.[33] Various implementations of the WWW algorithm exist: D'Ambrosio *et al.* showed that the majority of bond transpositions are rejected during annealing, and that an early decision scheme can enhance computation speed by rejecting unfavorable transpositions;[34] Ormrod Morley *et al.* constructed Metropolis criteria from topological metrics such as ring distributions (that is, energy changes were not considered in that case).[35]

Machine learning (ML) based interatomic potential models are increasingly being used to accelerate materials simulations.[36–38] ML potentials are typically "trained" with DFT data and can achieve similar accuracy for a small fraction of the cost. One key assumption in many of these methods is that the total energy can be separated into sums of machine-learned

---

[a] Department of Chemistry, Inorganic Chemistry Laboratory, University of Oxford, Oxford OX1 3QR, UK. E-mail: volker.deringer@chem.ox.ac.uk

[b] Department of Chemistry, Physical and Theoretical Chemistry Laboratory, University of Oxford, Oxford OX1 3QZ, UK

atomic energies.[39,40] Whilst being an approximation in the first place, it was argued that these atomic energies may in fact be amenable to interpretation: a connection between ML atomic energies and local chemical structure was made for partly occupied crystallographic sites in β-rhombohedral boron,[41] and for atomic environments with different coordination numbers in amorphous silicon.[42] ML models for other properties of local atomic environments have recently been investigated, ranging from the local electronic density of states[43,44] to local distortion factors in grain boundaries.[45] The nature of these local ML properties (including atomic energies), and their usefulness in predicting physical properties, remains an interesting research question. (See, *e.g.*, ref. 46 for a discussion of ML atomic energies in a chemically complex system.)

In the present work, we explore the configurational space of amorphous graphene based on an ML potential model that gives access to total and local energies. We use an MC bond-switching algorithm where ML atomic energies are used in the acceptance criterion and show that doing so leads to physically sound structural models. Depending on the details of the algorithm, we obtain CRN-like or paracrystalline structures. Our work shows that ML atomic energies can be used in different ways to "drive" MC simulations based on local and nearest-neighbor (NN) energy contributions, with implications for research on amorphous graphene and likely on other disordered structures and materials.

Methods

Potential-energy models

A common ansatz in developing potential-energy models is that the total energy, E, can be separated into a sum of atomic contributions:

$$E = \sum_{i}^{\text{atoms}} \varepsilon_i, \tag{1}$$

where $\varepsilon_i = \varepsilon(\{\mathbf{r}_{ij}\})$, with i, j being atomic indices, and $r_{ij} \leq r_{\text{cut}}$. This ansatz can be applied to many systems as short- and medium-range interactions predominantly determine the total energy. However, the question of how to formulate $\varepsilon(\{\mathbf{r}_{ij}\})$ is not trivial: in carbon, there is a vast configurational space with a subtle interplay between structure and energetics, such as in dihedral and torsional forces, weak interlayer interactions, and so on. Moreover, locality can depend on the structure: numerical experiments showed a large difference in the locality of atomic forces for diamond *versus* graphite.[47] (We note in passing that extracting local energies directly from DFT is not trivial, although attempts have previously been made.[48])

Empirical potentials are typically parameterized for a specific composition and phase of a material. For example, the original Tersoff potential for carbon[13] was parameterized by fitting parameters of repulsive and attractive pair potentials to cohesive energies of carbon polymorphs along with the lattice parameter and bulk modulus of diamond. This potential[13] used a bond-order approach, where bond strengths are modified according to the number of neighbors. The original reactive empirical bond-order (REBO) potential was an update of the Tersoff potential incorporating hydrogen,[49] with REBO-II adding further improvements.[50] A long-range Lennard-Jones term was added to REBO-II, creating the adaptive intermolecular reactive bond order (AIREBO) potential.[51] The long-range bond order potential (LCBOP)[52] is similar to AIREBO, having a long-range term in conjunction with a bond-order description—albeit here it is built in from the beginning. The environment-dependent interaction potential (EDIP) for carbon[53,54] was initially developed from an earlier silicon EDIP model.[55,56] EDIP has been shown to be successful in describing various properties of amorphous carbon, including the graphitization at low and high densities.[57–59]

Gaussian approximation potential (GAP)

The Gaussian approximation potential (GAP) framework[40] is used to "machine-learn" interatomic potential models from quantum-mechanical data, often based on DFT. Unlike empirical potentials which are constructed based on physical knowledge, GAP makes a non-parametric fit. This means that the model can adjust to complex input data—however, it also means that the selection and quantity of reference ("training") data is critically important.[60]

In brief, the local energy for a GAP model is

$$\varepsilon(\mathbf{q}) = \sum_{t=1}^{N_t} \alpha_t K(\mathbf{q}, \mathbf{q}_t), \tag{2}$$

where the sum runs over N_t training configurations represented by the local-environment descriptor $\mathbf{q}_t$, with a corresponding weighting coefficient α_t attributed during fitting. K is a covariance kernel which measures the similarity between the input and training configurations, represented by $\mathbf{q}$ and $\mathbf{q}_t$, respectively. A commonly used approach for the latter task is the Smooth Overlap of Atomic Positions (SOAP) descriptor and kernel.[61]

For the present work we use the amorphous carbon potential, GAP-17.[47] This model has been shown to predict energies within tens of meV per atom compared to DFT, as well as providing a good description of structural and elastic properties of amorphous carbon.[47]

Initialization

We generate a pristine, 200-atom layer of crystalline graphene using the experimental bond length of 1.42 Å (ref. 62 and 63). A spacing of 20 Å between layers ensures that there is no interaction between periodic replicas. Where noted, NN-energy-based runs start from structural models that have been "thermalized", *i.e.*, disordered, using local energies at $\beta = 2.0\ \text{eV}^{-1}$ (see below). The cell parameters are allowed to relax prior to the first bond switch, and then kept fixed for the duration of the MC simulation. The structures are kept planar, simplifying the present study to the idealized 2D case, and noting that puckering can lower the energy further.[64]

Monte-Carlo protocol

A kinetic MC algorithm is used to generate aG structural models. The simulations begin either from cG or from a thermalized structure. At each step, a random atom is chosen along with a random neighbor, determined using a cutoff of 1.85 Å. The atoms undergo an SW rotation (90° about the bond center) and the new structure is then relaxed using the conjugate-gradient algorithm. For initial testing of structural relaxations, energy evaluations, and MC runs, we used the Atomic Simulation Environment (ASE)[65] interfaced to quippy[66,67] (**https://github.com/libAtoms/QUIP**). The production MC runs shown in this work used LAMMPS.[68] The force tolerance for structural relaxations was 1 meV $Å^{-1}$ with a maximum of 150 relaxation steps. A topological constraint is imposed where the newly relaxed structure must be 3-coordinate, otherwise the move is rejected. (This constraint is included because there will be moves which create stubborn coordination defects, and we found that these defects hindered the progress of MC annealing where the simulation would get "stuck" in local minima.) If the constraints are met, the new structure is accepted with the following probability:

$$w = \begin{cases} 1, & \text{if } \Delta E \leq 0, \\ \exp(-\beta \Delta E), & \text{if } \Delta E > 0, \end{cases} \tag{3}$$

where $\beta = (k_{\mathrm{B}}T)^{-1}$ and $\Delta E = E_{\text{new}} - E_{\text{old}}$. If $\Delta E > 0$, we generate a random number z in [0,1), and if $w > z$, then the move is accepted, else it is rejected and the previous configuration is kept. In this work, β has no physical relation to temperature since ΔE is not measured for all atoms about their equilibrium position, as the atoms are being held fixed after the bond transposition and thus thermal fluctuations are ignored (unlike with MD). Fixing particle positions at every move results in ergodicity being broken and thus samples are not taken from a Boltzmann distribution.[69] Furthermore, the local-energy and NN-energy framework explicitly do not include all atomic energies which is another source of ergodicity being broken. This is not a concern for the present work as temperature-driven dynamics are not relevant (see ref. 69 for an approach to introducing ergodicity to a similar problem) and results in β becoming a tunable parameter. β was initially chosen to be 2.0 eV^{-1} to correspond to the study by Toh *et al.*,[12] and then tuned heuristically for the local- and total-energy-based MC runs.

We use the atomic energies for the defect pair in the Metropolis criterion, defining the local defect energy as

$$\varepsilon_{\text{local}} = \varepsilon_1 + \varepsilon_2 \tag{4}$$

and the NN defect energy as

$$\varepsilon_{\text{NN}} = \varepsilon_{\text{local}} + \sum_{i}^{\text{NN}} \varepsilon_i, \tag{5}$$

where ε_i is the local atomic energy of the i-th atom, with $i = 1$ and $i = 2$ denoting the two atoms in the SW defect pair, and $i = 3 \ldots 6$ the topological NNs of the SW defect pair (Fig. 1c and d). MC runs based on total energies were also carried out for reference. For the 200-atom systems, 25 independent and parallel simulations were conducted over 10 000 MC steps each for different β values and for the 612-atom system, 20 000 MC steps were taken.

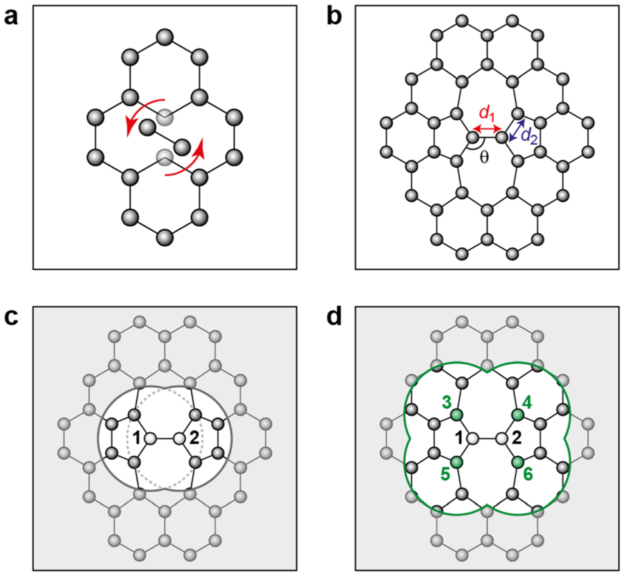

Fig. 1 Stone–Wales defect in graphene and definition of nearest-neighbor (NN) atoms. (a) Schematic of an in-plane single-bond rotation in graphene, leading to an SW defect. (b) A bond rotation by 90° creates two 5-membered and two 7-membered rings. Definitions are given for the defect bond length, d_1, the defect–NN bond length, d_2, and the defect bond angle, θ (*cf.* Table 1). (c) Definition of local energy for the Metropolis criterion. The gray atoms labeled **1** and **2** are the SW defect pair, with shading indicating the overlap of the two cutoff spheres up to which atoms contribute to the local energy (this value is 3.7 Å for the GAP-17 model, larger than sketched here). (d) The green atoms (**3**–**6**) are the topological NNs of the defect pair. Green lines indicate the overlap of the cutoff spheres for each NN atom.

Structural analysis

Ring statistics were determined using a shortest-path algorithm[70] as implemented in Matscipy.[71] Topological metrics typically used in network analysis were applied to the 612-atom structures, namely, Lemaître's law[72] and the assortativity[73] (as discussed below). For SOAP structural analysis, we compared each individual atom in a given aG structural model to an atom in cG. The SOAP parameters were: radial cutoff, 5.5 Å; cutoff transition width, 0.5 Å; neighbor-density smoothness, $\sigma_{\text{atom}} = 0.5$ Å; basis-set convergence parameters, $n_{\max} = l_{\max} = 16$; dot-product kernel raised to a power of $\zeta = 4$. All structures were visualized using OVITO.[74]

Phonon calculations

A geometry optimization was performed on the local, NN and Toh *et al.* structures keeping the cell fixed ($F_{\max} < 1$ meV $Å^{-1}$). We applied a finite-difference approach with a displacement of 0.01 Å. The forces in the displaced structures were computed using GAP-17 and second-order force constants were computed using phonopy.[75] A 50 × 50 × 1 k-point mesh was used to determine the vibrational density of states (VDOS). We observed minor imaginary modes in the aG models, presumably due to the constraint of planarity; these are omitted from VDOS plots and will be investigated in future work.

Results and discussion

Local energies for a single defect

Fig. 2 characterizes the structure and energetics of a single SW defect in graphene. GAP-17 predicts that the energies of the SW defect pair of atoms are considerably lower in energy than those of the NN atoms. This behavior is in marked contrast with that for the empirical potentials (Fig. 2b and Table 1): the latter suggest that the SW-pair atoms are relatively higher in energy (with the exception of LCBOP), with atoms in the 5-membered rings being higher in energy overall and those in 7-membered rings being lower. LCBOP and GAP-17 predictions are similar when averaged over the pentagon-pair environments, only differing slightly in structural parameters. REBO and AIREBO show slightly larger SW-pair energies and more strained structures (Table 1).

The total defect formation energy ΔE_{total} for GAP-17 is similar overall to that predicted by most empirical potentials—and to DFT values, which we obtained using CASTEP.[79] In an earlier study, hybrid-DFT and quantum Monte Carlo calculations predicted values of 5.69 eV and 5.92 ± 0.03 eV, respectively.[30] EDIP, LCBOP, and REBO were parameterized for sp^2 environments, and their ΔE_{total} is in close agreement with that for GAP-17. AIREBO was also parameterized for sp^2 environments, however the added long-range term may not be optimized for 2D planar carbon. In contrast, $\Delta\varepsilon_{\text{local}}$ and $\Delta\varepsilon_{\text{NN}}$ are not in close agreement, where $\Delta\varepsilon_{\text{local}}$ differs the most across all the potentials (except for REBO and AIREBO which are parameterized identically for short-range terms). Overall, GAP-17 predicts the lowest energy for the NN defect formation energy among the potential models investigated.

Table 1 Structure and formation energy of a single SW defect in graphene from ML and empirical potentials as well as DFT. The table shows the computed bond length in pristine graphene (d_0), the defect-pair (d_1) and defect-NN bond lengths (d_2), the defect bond angle (θ; *cf.* Fig. 1b), and three energies associated with defect formation: the local energy change in the rotated atoms only ($\Delta\varepsilon_{\text{local}}$), the NN energy change ($\Delta\varepsilon_{\text{NN}}$), and the total defect formation energy (ΔE_{total})

	d_0 (Å)	d_1 (Å)	d_2 (Å)	θ (deg)	$\Delta\varepsilon_{\text{local}}$ (eV)	$\Delta\varepsilon_{\text{NN}}$ (eV)	ΔE_{total} (eV)
GAP-17	1.41	1.32	1.44	121.89	0.14	2.25	5.58
EDIP	1.50	1.39	1.58	118.26	1.69	3.32	5.58
LCBOP	1.42	1.34	1.45	121.76	0.73	2.39	5.33
REBO	1.42	1.34	1.46	120.64	0.99	2.85	5.61
AIREBO	1.40	1.32	1.44	120.88	1.09	3.15	6.27
LDA[76]	1.41	1.31	1.45	122.34	—	—	5.14
PBE[77]	1.43	1.32	1.46	122.49	—	—	5.03
PBEsol[78]	1.42	1.32	1.46	122.40	—	—	5.02

Structural parameters as predicted by the respective potentials are also shown in Table 1. The defect-pair bond length (d_1), defect–NN bond length (d_2), and defect bond angle, θ (Fig. 1b) are in close agreement across all potentials—with the exception of EDIP, which notably increases the cell parameters upon relaxation. This is shown by the difference in cG bond length (d_0) between EDIP and the other potentials when the cell parameters are allowed to relax, with EDIP predicting an elongated d_0. GAP-17 and LCBOP are in close agreement with regard to θ, predicting a slightly wider bond angle for the 5|7 pair, whereas EDIP predicts it slightly lower, notwithstanding the fact that the EDIP defect formation energy is in close agreement with that for GAP-17. Furthermore, GAP-17 is in excellent agreement with the bond lengths and angles from our DFT calculations.

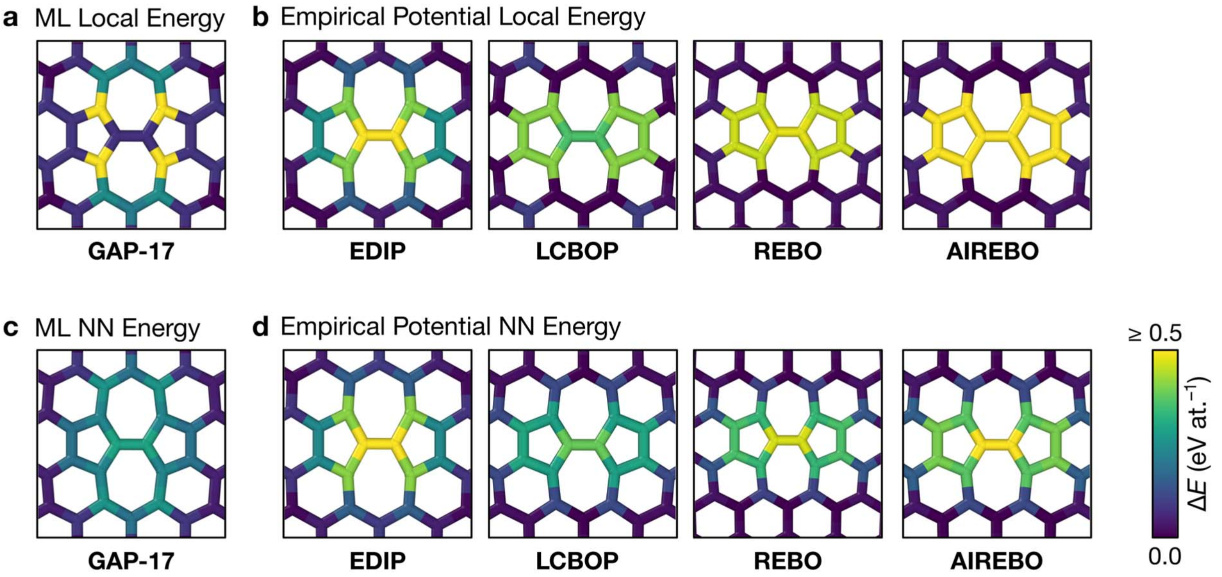

Fig. 2 Structure and local energy of a Stone–Wales defect in graphene. The figure compares local energies (top row) and nearest-neighbor averaged energies ("NN", bottom row) for a single Stone–Wales (SW) defect as relaxed using the respective interatomic potential. (a, b) Local atomic energies from GAP-17 and empirical potentials, respectively. (c, d) Same but for NN energies (eqn (5)), from GAP-17 and empirical potentials, respectively.

Using local energies to drive Monte-Carlo annealing

Fig. 3 shows the energy profiles of independent parallel MC simulations *via* successive SW transformations, which we use to create structural models of aG. Panel (a) shows the evolution of the average energy of the ensemble when using (only) the atomic energies of the SW defect pair in the Metropolis criterion. Clearly, using a β value of 2.0 eV^{-1} leads to a highly disordered 3-coordinate structure (as shown in the inset), with energy convergence occurring after about 1000 MC steps. Such structures contain high-energy 3- and 4-membered rings with significant strain. We call the resulting structures "thermalized" and use them as starting points for the runs characterized in blue in Fig. 3b, as indicated by an arrow.

Nearest-neighbor (NN) atoms are defined topologically in the present work, and so we include the sum over the SW pair and its bonded neighbors, including six atoms in total (eqn (5)). The energy difference before and after the SW transformation is used in the Metropolis criterion. Results from two protocols, starting either from cG or from the previously thermalized structures, converge to similar energy values (green and blue curves in Fig. 3b). It seems that runs from both starting points tend toward a mutual limiting distribution. This is reflected in Fig. 3b, as both curves have means and standard deviations in very close agreement after approximately 4000 MC steps. However, it is apparent that this protocol is not ergodic, as discussed in the Methods section. The principle of ergodicity states that all possible configurations of the system should be attainable[80] and thus it is clear that given the topological constraints imposed, it is impossible for certain proposed moves to be accepted even though they may be energetically favorable.

Fig. 3c illustrates the effect of varying β on the resulting energy profile when using the local-energy criterion. It is clear

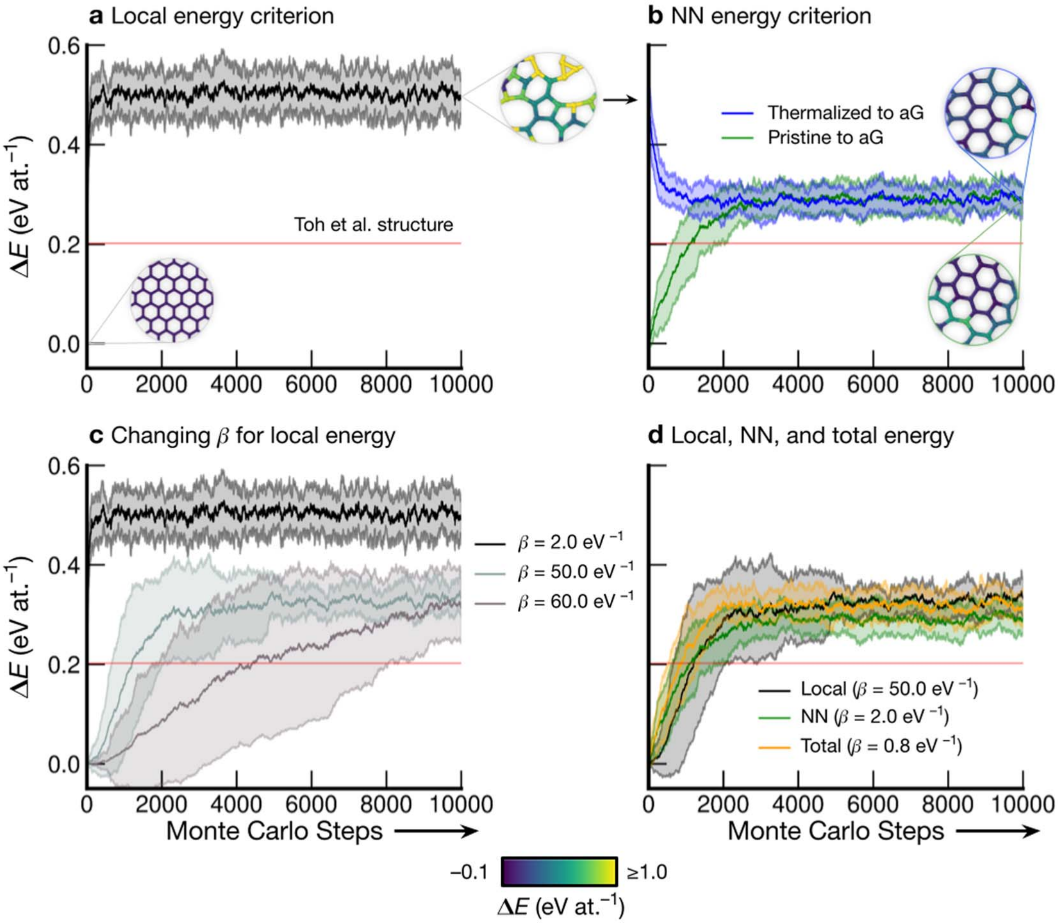

Fig. 3 Evolution of disordered graphene structures during Monte-Carlo simulations. Lines indicate mean energies for an ensemble of 25 separate runs, with shaded regions indicating a standard deviation of 1σ. Representative structural snapshots are shown, with atoms color-coded by local energy. (a) The atomic energies of the two atoms involved in the SW transformation are used for the Metropolis criterion, creating a "thermalized" structure rapidly at $\beta = 2.0$ eV^{-1}. These structures are used as initial configurations in (b). (b) MC annealing *via* NN energies ($\beta = 2.0$ eV^{-1}) for the Metropolis criterion. The curves converge at slightly above the energy of an MC-annealed paracrystalline structure from Toh *et al.* (ref. [12]), which is included for reference (red line). (c) The effect of β on the MC energy profile. It is seen that a smaller β (higher simulation "temperature") leads to a higher-energy structure and faster convergence. We emphasize that despite the high standard deviations, the lowest occurring energy value is $\Delta E = 0$ eV at.$^{-1}$ corresponding to ideal crystalline graphene. (d) Heuristically tuning β to find a mutual limiting distribution between MC approaches (within statistical fluctuations).

that using a lower β value corresponds to a higher-energy final structure as well as rapid convergence: for $\beta = 2.0$ eV^{-1}, convergence occurs in fewer than 1000 MC steps (black line). Increasing β results in lower-energy structures and slower convergence, as shown by the $\beta = 50.0$ eV^{-1} and $\beta = 60.0$ eV^{-1} results. For $\beta = 60.0$ eV^{-1}, convergence has not occurred after 10 000 steps, however the mean appears to be tending towards that for $\beta = 50.0$ eV^{-1}, implying that there might be a minimum level of disorder that is "stable" for the local-energy criterion. Additionally, as β was increased beyond 60.0 eV^{-1}, proposed moves were always rejected, further suggesting a minimum convergence energy using this framework.

Fig. 3d shows the results from three different runs with heuristically tuned β values for local-, NN-, and total-energy criteria, respectively, converging towards a mutual distribution (within statistical fluctuations). This result allows the direct comparison of structures generated by the different frameworks side-by-side, as the effect of β has been removed. Such a comparison is particularly instructive for lower-energy structures, as we will show below: they show a richer configurational space, including medium-range order (as compared to the highly disordered high-energy structures), and this space may be traversed differently by different MC runs.

The configurational space of amorphous graphene

Having explored different protocols for generating aG structural models, we next created larger-scale structures: the system size was increased to 612 atoms and structures from local- and NN-energy-based searches were studied. Additionally, the 610-atom structural model published by Toh *et al.* (ref. 12) was included for comparison. Fig. 4 shows these three configurations. As we keep the cells fixed throughout the Monte-Carlo searches (Methods), we computed the tensile stress for the final structures, using GAP-17, and found them to correlate with the degree of disorder. The structure generated using local energies had an in-plane tensile stress of 3.53 GPa; the NN structure had 2.85 GPa, and for the Toh *et al.* structure we obtained 0.87 GPa.

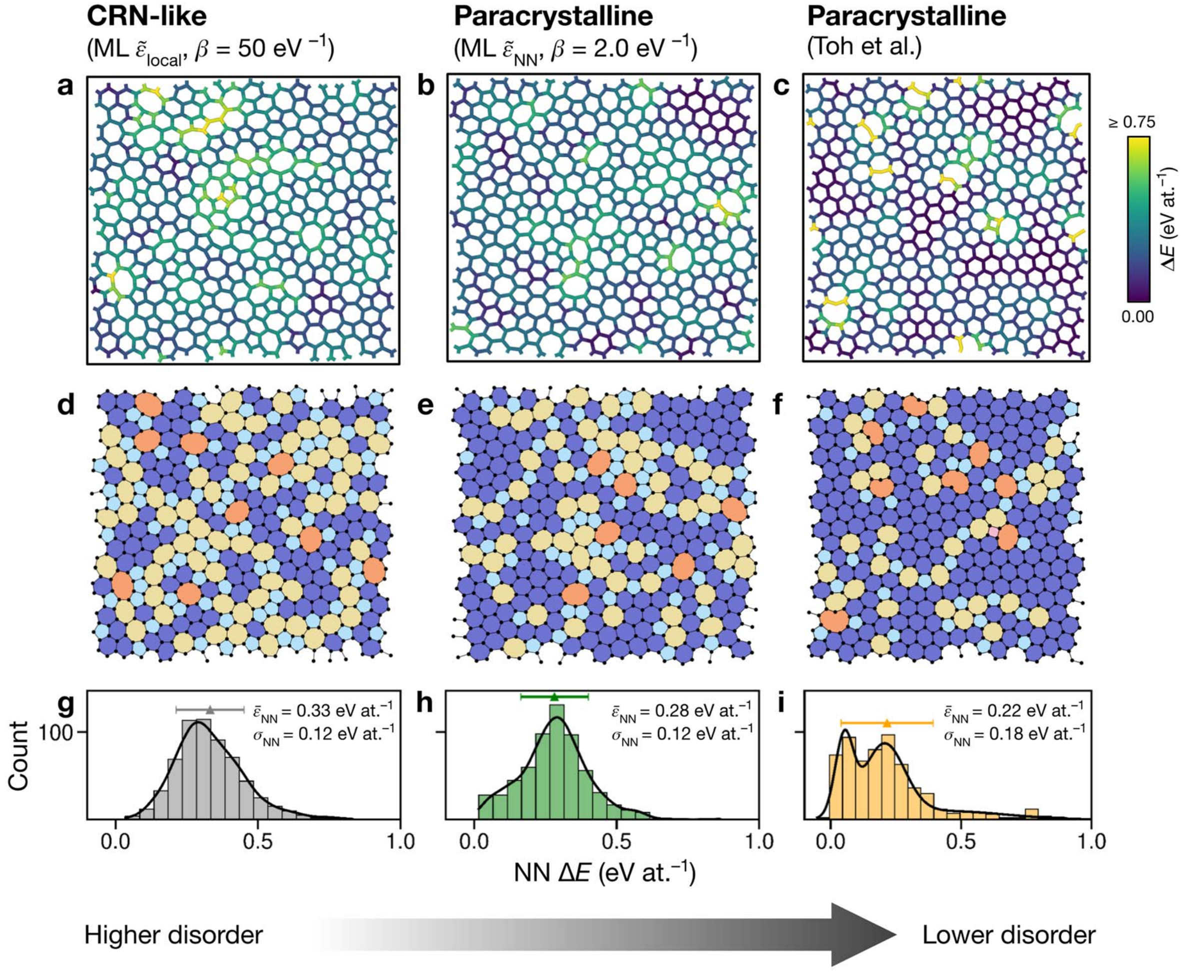

Fig. 4 Structural models of amorphous graphene. Structures in panels (a–c) are color-coded according to the average atomic energy over a chosen atom and the corresponding NNs. (a) Structure after the final MC step for the local-energy criterion and $\beta = 50$ eV^{-1}. (b) As for (a), but with NN energies used in the Metropolis criterion. (c) Structure taken from Toh *et al.*[12] and optimized with GAP-17. (d–f) Structures as in (a–c), now color-coded by ring size (pink, ≤4; light blue, 5; dark blue, 6; yellow, 7; orange, ≥8). (g–i) Distributions of NN energies in the respective structures, with mean values and standard deviations given. Curves were obtained as kernel density estimates.

For the local-energy framework, the absence of crystal-like pockets is clear as seen in panels (a) and (d). The structure resembles a CRN, with chains of 5- and 7-membered rings running across the cell, and there are more large rings compared with the other structures. The pronounced disorder is likely a result of the low ML energy for the SW defect itself (Fig. 2a), and thus of the low energy cost for these transpositions (given that medium-range order is not captured in the local energies alone).

With NN energies used in the Metropolis criterion, panels (b) and (e) suggest that using an NN criterion encourages small pockets of crystal-like regions forming, indicating that direct contributions from NNs maintain medium-range order, whilst retaining an amorphous structure at $\beta = 2.0$ eV^{-1}. Visually, the locally averaged energies (up to NNs) in panels (a–c) suggest that using just the local defect-pair energy gives more regions of higher energy (shown in yellow in Fig. 4a) compared with the NN framework (Fig. 4b). This appears consistent with the nature of CRNs *versus* paracrystalline structures, and color-coding by NN energies shows a clear difference between the two structures.

In the NN-energy-based structure (Fig. 4b), there are pockets of crystallinity, indicated by regions of ordered 6-membered rings. Interestingly, there is an aggregation of more disordered regions, with 5-membered rings surrounding larger 7- and 8-membered rings. For the structure shown in Fig. 4c, the authors started from a randomized, hard-sphere constrained structure and worked towards a paracrystalline sample using the AIREBO potential. In the resulting structure, we note the presence of coordination defects, since the randomized initial structure was not topologically constrained. There is also a 4-membered ring (pink in Fig. 4f). It is evident that this sample is paracrystalline with regions of locally crystal-like order separated by 5|7 grain boundaries, and that it is more ordered than the NN-based structure (Fig. 4e). Larger defects appear to congregate as seen, for example, in the top of the figure. These defective environments are high in energy, as shown by the color-coding.

Fig. 4g–i shows the distribution of NN energies for the respective structures. The average energy for the atoms in the local-energy-based structure is 0.33 eV at.$^{-1}$ with a standard deviation of $\sigma = 0.12$ eV at.$^{-1}$, and the energy of the NN-based structure is 0.28 ± 0.12 eV at.$^{-1}$ above pristine graphene. Both values are close to the mean energy resulting from the independent runs for the 200-atom system (Fig. 3b and d): hence, an individual 200-atom run will not likely be sufficient to describe aG, but an ensemble of multiple independent runs will be—just like many small-scale structural models of 3D amorphous carbon have been used in the fitting of GAP-17.[47] The distribution in the locally averaged energies is similar overall between Fig. 4g and h.

For Fig. 4i, we find a bimodal distribution in the energy, indicating that this structure shows higher paracrystalline order compared to that generated by the NN framework. There exist pockets of more ordered, more energetically favorable regions in between more disordered ones, as seen in the structural model shown in panels (c) and (f). Very recently, a paracrystalline sample of diamond has been synthesized.[82] This discovery, along with the synthesis and characterization of paracrystalline graphene[12,83] suggests that the landscape of disorder in carbon may have a link between "fully" amorphous and crystalline. This is further reflected in the NN energy distributions in Fig. 4, where panel (g) clearly shows a single distribution, panel (i) a bimodal distribution, and panel (h) a small hint of one. These ML energy distributions may provide a quantitative distinction between CRN and paracrystalline graphene, where a prominent bimodal distribution indicates the latter and the lack thereof indicates the former.

Quantifying disorder

We show further, quantitative structural indicators in Fig. 5. The bond-angle distribution (panel a) characterizes medium-range order. The local-energy-based structures (gray) display a wide range of bond angles, reflecting the relatively high level of disorder. The absence of a clear peak at 120° also emphasizes the reduction of medium-range order as is characteristic of CRNs. The NN curves (green) are narrower and centered around 120°, consistent with locally crystal-like environments. A shoulder peak at $\approx 109°$ shows the strong presence of 5-membered rings. The Toh *et al.* structure has a pronounced peak at 120° with a smaller shoulder at $\approx 109°$, indicative of the larger degree of locally crystal-like order. The bond angles in the 200-atom models (solid lines) agree well with those in the 612-atom ones (dashed).

The count of shortest-path rings is another metric for medium-range order. Ring counts for the local-energy MC runs reflect the large disorder, showing more 5- and 7-membered rings than in the other structures. The structures from the NN runs have ring counts centered around 6, as for the Toh *et al.* structure, with a larger count suggesting increased crystallinity in the latter case. As with the bond-angle distributions (Fig. 5a), the ring statistics for the 200-atom *versus* 612-atom structures are in close agreement within each other, within the standard deviation of the values for the former (Fig. 5a).

SOAP is a structural similarity metric for atomic environments.[61] 2D plots can reveal correlations between SOAP similarity on the one hand, and locally averaged energies on the other hand.[42] Fig. 5d shows the distribution of locally averaged energies (up to NN) and SOAP similarity to cG. The wider the spread in both axes, the greater the structural disorder. Comparing local-energy- with NN-energy-based data, both KDE curves are in close agreement, being skewed slightly toward the left (SOAP) and to higher energies for the runs using local energies only. This is expected, since the ML model predicts relatively low energies for the SW pair (Fig. 2), allowing structures to become more disordered. When averaging over NN environments, the distribution shifts and narrows as we observe fewer highly disordered environments. Data for the structure from Toh *et al.*[12] are provided to show how ordered this structure is. Defective environments are clearly identified with a series of data points at higher energies. As with Fig. 4, there is a clear bimodal distribution in the KDE curve. As seen for the structural indicators characterized in Fig. 5a and b, the dashed lines representing the 612-atom structures agree well with the set of separate 200-atom structures.

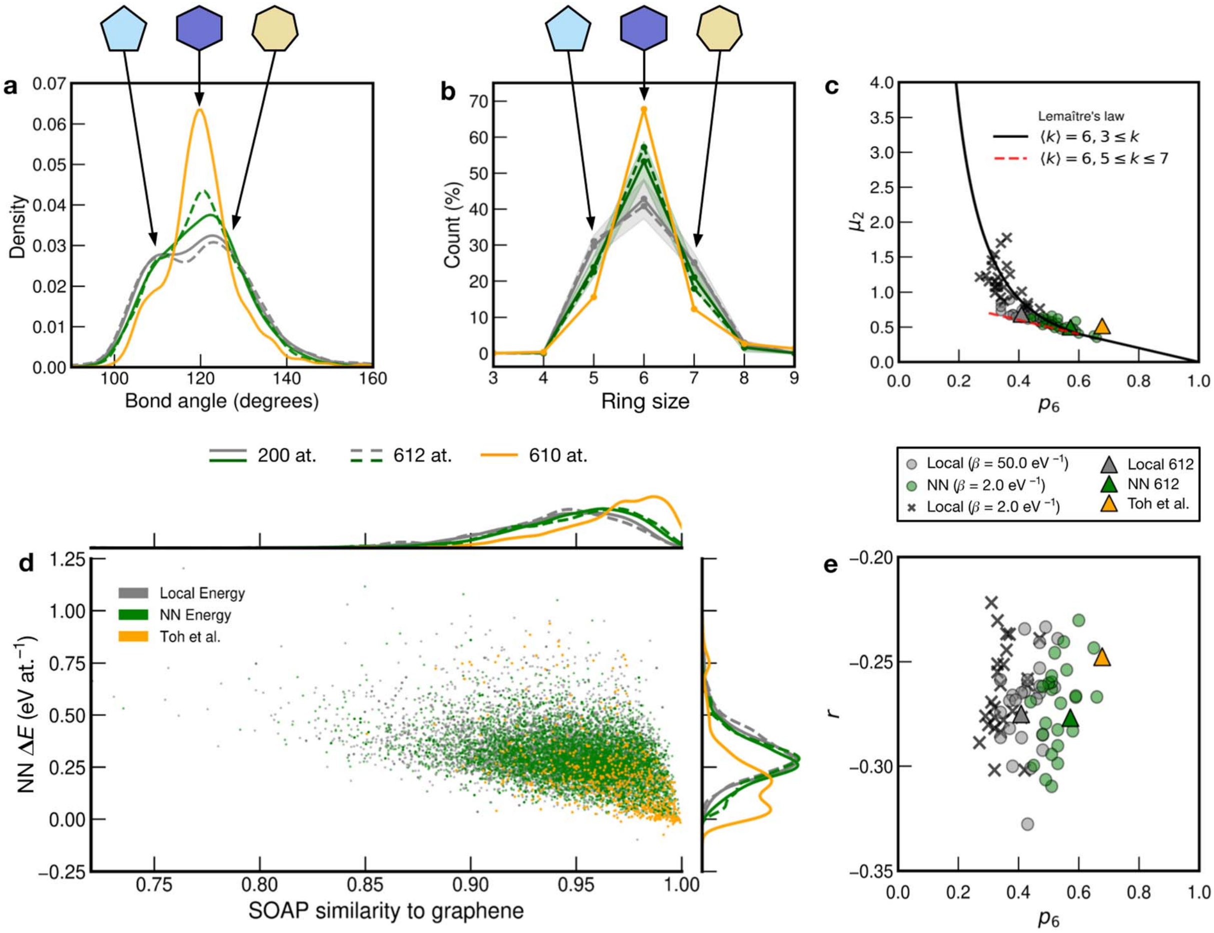

Fig. 5 Local structure and stability in models of amorphous graphene. Results for 25 local-energy-based ($\beta = 50.0$ eV^{-1}, grey) and NN-energy-based ($\beta = 2.0$ eV^{-1}, green) runs, *i.e.*, for 25 × 200 atoms each, and for the 612-atom structures based on local ($\beta = 50.0$ eV^{-1}, gray dashed) and NN ($\beta = 2.0$ eV^{-1}, green dashed) energies are given. The structural model from Toh *et al.* is also analyzed (orange, 610 atoms). Data from the 612-atom and Toh *et al.* structures are scaled arbitrarily to fit. (a) Bond-angle distributions for the different structures. (b) Ring statistics for the different structures. (d) An NN energy *vs.* SOAP plot, following ref. 42. Kernel density estimates (KDEs) are used to show distributions of properties. The bandwidth is determined following Scott's rule,[81] with a grid size of 200. Additionally, panels (c) and (e) show plots of network-topology metrics, namely, Lemaître's law and assortativity, respectively.

In addition to the well-established ring statistics and the SOAP similarity, we analyzed the structural models with two topological metrics typically used in network theory, *viz.* Lemaître's law[72] and the assortativity metric.[73] Lemaître's law connects the second moment of the ring distribution ($\mu_2 = \langle k^2 \rangle - \langle k \rangle^2$, where k is the ring size) with the fraction of six-membered rings, p_6. Since we impose the constraint that all atoms must be 3-fold connected, it follows that the distribution can be explained well by a single maximum-entropy distribution leading to a characteristic curve.[72] As $\lim_{p_6 \to 0.6}$ from $p_6 = 1$, μ_2 increases linearly as $1 - p_6$ in the region of $p_6 \gtrapprox 0.6$. This line is extended (red line) to show that most of the points from both local ($\beta = 50.0$ eV^{-1}) and NN frameworks are located on this curve. This curve corresponds to the maximum-entropy solution if only 5-, 6-, and 7-membered rings are present, which is the case in both these frameworks. μ_2 then increases exponentially beyond this where structures from the local framework at $\beta = 2.0$ eV^{-1} are located. These structures show greater disorder and a wider spread of μ_2 values, as expected due to the presence of 3- and 4-membered rings, where the maximum-entropy solution follows an exponential profile. For the 612-atom structures, the local framework ($\beta = 50.0$ eV^{-1}) based structure is more disordered than the NN-based one and the Toh *et al.* structure (orange). The latter does not yield a datapoint on the curve as it has a (small) number of 2-coordinate sites and hence would be located on a different Lemaître curve.[35]

The assortativity, r, measures how likely a large ring is to be next to a smaller ring (disassortative, $r < 0$) or to other large rings (assortative, $r > 0$). In many physical systems, one finds a preference for disassortative configurations. This is, here, reflected in the range of r values (Fig. 5e). For the 200-atom systems, we find no discernible correlation between r and p_6. For the 612-atom ones, both the local and NN frameworks yield structures with similar r values. The Toh *et al.* structure has a slightly less negative r, implying a slight preference for a more random arrangement. This may be due to coordination defects in the structure skewing r toward zero.

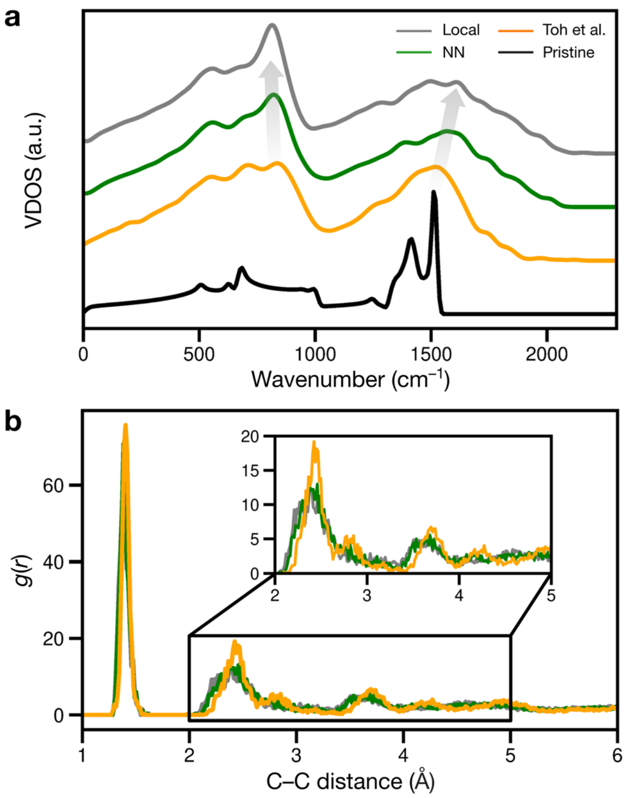

Fig. 6 Signatures of structural disorder. (a) Vibrational density of states (VDOS) for the different structures and pristine graphene (black line). The gray arrows indicate trends in peak shifts. (b) Radial distribution functions, $g(r)$, for the different structures.

We finally considered properties that can be compared to previous or future experiments. Vibrational properties of graphene are readily characterized by Raman spectroscopy[84,85] and computations,[26,86] and we show predicted VDOS in Fig. 6a. For pristine graphene (as reference), we obtain a characteristic peak at 1510 cm^{-1} and a smaller one at 1410 cm^{-1}, in close agreement with previous computations[26,86] and experimental spectra.[84,85] As disorder increases, the 1510 cm^{-1} peak shifts to higher frequencies, as shown in Fig. 6a: to about 1520, 1570, and finally 1610 cm^{-1}.

The peaks at 1510–1610 cm^{-1} seen in Fig. 6a for the disordered structures are close to Raman data by Toh *et al.*, with the corresponding peak at 1558 cm^{-1}.[12] The blue shift of the 1510 cm^{-1} peak with increasing disorder is a well-known effect and has been observed in VDOS and Raman spectra.[26,86] Furthermore, we observe an "amorphous" peak in the acoustic (low-frequency) region, at 814–836 cm^{-1}, red-shifting with increasing disorder.

Another way of validating against experiment is to evaluate the radial distribution function (RDF) and X-ray diffraction (XRD) patterns: in disordered structures, characteristic peaks will become wider and less intense.[12,87] However, when comparing amorphous structures, differences in RDF and XRD are subtle:[18,87] in Fig. 6b, the second RDF peak at 2.41 Å broadens with increasing disorder, but this effect is subdued for the remaining peaks at 2.82 and 3.62 Å for the local and NN structures. The Toh *et al.* structure has longer-range order, reflected in more prominent RDF peaks at greater distances (4.23 and 4.93 Å). The local and NN structures have peak locations comparable to that of the pair correlation function recreated from TEM in ref. 12.

Conclusions

Atomic energies predicted by the GAP ML framework can be used to drive Monte-Carlo structural exploration in principle. We have shown this by creating structural models of amorphous graphene, one of the prototypical disordered systems in physics and chemistry. We found that using (only) ML atomic energies leads to structures resembling continuous random networks, whereas including nearest-neighbor energies tends to drive the simulations toward some degree of paracrystalline order. We suggest that histograms of local energies, as shown in Fig. 4, can give insight into the degree of "randomness" in different amorphous networks: deviations from random order are visible as peaks at either low (paracrystalline) or high (coordination defects) energy.

The demonstrated ability to use ML local energies in MC annealing indicates potential for future research on amorphous materials. Describing local environments using ML methods can provide insight into the relation between atomic structure and energetics, and therefore structural stability, in amorphous materials—including amorphous carbon, which has emerging applications in biosensing[88] or batteries.[89] The fact that NN-averaged energies yield reasonable, partly paracrystalline structural models may be attributed to the fact that they provide "smoothing" over the variance in local atomic energies. This finding is consistent with earlier findings for the electronic DOS[43,44] and might have wider consequences for ML predictions of local properties, which are yet to be fully explored.

Data availability

Data supporting this work are available at **https://doi.org/10.5281/zenodo.7221166**.

Author contributions

Z. E.-M. developed the Monte-Carlo protocols and carried out the computational work. V. L. D. initiated and coordinated the study. All authors made substantial contributions to data analysis and discussions. Z. E.-M. and V. L. D. wrote the paper with input from M. W., and all authors approved the final version.

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

We thank Prof. G. Csányi for helpful comments on the Monte-Carlo simulations. We are grateful for support from the

EPSRC Centre for Doctoral Training in Theory and Modelling in Chemical Sciences (TMCS), under grant EP/L015722/1. V. L. D. acknowledges support from the Engineering and Physical Sciences Research Council through a New Investigator Award [grant number EP/V049178/1]. This paper conforms to the RCUK data management requirements. The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work (**http://dx.doi.org/10.5281/zenodo.22558**).

References

1 A. C. Wright, The great crystallite *versus* random network controversy: A personal perspective, *Int. J. Appl. Glass Sci.*, 2014, **5**, 31–56.

2 D. Shi, Z. Guo and N. Bedford, 10 - Nanoenergy Materials, in *Nanomaterials and Devices*, William Andrew Publishing, Oxford, 2015, pp. 255–291.

3 J. Zhao, G. Zhu, W. Huang, Z. He, X. Feng, Y. Ma, X. Dong, Q. Fan, L. Wang, Z. Hu, Y. Lü and W. Huang, Synthesis of large-scale undoped and nitrogen-doped amorphous graphene on mgo substrate by chemical vapor deposition, *J. Mater. Chem.*, 2012, **22**, 19679–19683.

4 M. L. Gallo and A. Sebastian, An overview of phase-change memory device physics, *J. Phys. D: Appl. Phys.*, 2020, **53**, 213002.

5 J. Kotakoski, A. V. Krasheninnikov, U. Kaiser and J. C. Meyer, From point defects in graphene to two-dimensional amorphous carbon, *Phys. Rev. Lett.*, 2011, **106**, 105505.

6 P. Y. Huang, S. Kurasch, A. Srivastava, V. Skakalova, J. Kotakoski, A. V. Krasheninnikov, R. Hovden, Q. Mao, J. C. Meyer, J. Smet, D. A. Muller and U. Kaiser, Direct imaging of a two-dimensional silica glass on graphene, *Nano Lett.*, 2012, **12**, 1081–1086.

7 P. Y. Huang, S. Kurasch, J. S. Alden, A. Shekhawat, A. A. Alemi, P. L. McEuen, J. P. Sethna, U. Kaiser and D. A. Muller, Imaging atomic rearrangements in two-dimensional silica glass: Watching silica's dance, *Science*, 2013, **342**, 224–227.

8 Z. Yang, J. Hao, S. Yuan, S. Lin, H. M. Yau, J. Dai and S. P. Lau, Field-effect transistors based on amorphous black phosphorus ultrathin films by pulsed laser deposition, *Adv. Mater.*, 2015, **27**, 3748–3754.

9 W.-J. Joo, J.-H. Lee, Y. Jang, S.-G. Kang, Y.-N. Kwon, J. Chung, S. Lee, C. Kim, T.-H. Kim, C.-W. Yang, U. J. Kim, B. L. Choi, D. Whang and S.-W. Hwang, Realization of continuous Zachariasen carbon monolayer, *Sci. Adv.*, 2017, **3**, e1601821.

10 Z. Yang, J. Hao and S. P. Lau, Synthesis, properties, and applications of 2D amorphous inorganic materials, *J. Appl. Phys.*, 2020, **127**, 220901.

11 S. Hong, C.-S. Lee, M.-H. Lee, Y. Lee, K. Y. Ma, G. Kim, S. I. Yoon, K. Ihm, K.-J. Kim, T. J. Shin, S. W. Kim, E.-c. Jeon, H. Jeon, J.-Y. Kim, H.-I. Lee, Z. Lee, A. Antidormi, S. Roche, M. Chhowalla, H.-J. Shin and H. S. Shin, Ultralow-dielectric-constant amorphous boron nitride, *Nature*, 2020, **582**, 511–514.

12 C.-T. Toh, H. Zhang, J. Lin, A. S. Mayorov, Y.-P. Wang, C. M. Orofeo, D. B. Ferry, H. Andersen, N. Kakenov, Z. Guo, I. H. Abidi, H. Sims, K. Suenaga, S. T. Pantelides and B. Özyilmaz, Synthesis and properties of free-standing monolayer amorphous carbon, *Nature*, 2020, **577**, 199–203.

13 J. Tersoff, Empirical interatomic potential for carbon, with applications to amorphous carbon, *Phys. Rev. Lett.*, 1988, **61**, 2879–2882.

14 G. Galli, R. M. Martin, R. Car and M. Parrinello, Structural and electronic properties of amorphous carbon, *Phys. Rev. Lett.*, 1989, **62**, 555–558.

15 D. A. Drabold, P. A. Fedders and P. Stumm, Theory of diamond like amorphous carbon, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1994, **49**, 16415–16422.

16 N. A. Marks, D. R. McKenzie, B. A. Pailthorpe, M. Bernasconi and M. Parrinello, Microscopic structure of tetrahedral amorphous carbon, *Phys. Rev. Lett.*, 1996, **76**, 768–771.

17 D. G. McCulloch, D. R. McKenzie and C. M. Goringe, Ab initio simulations of the structure of amorphous carbon, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2000, **61**, 2349–2355.

18 A. Kumar, M. Wilson and M. F. Thorpe, Amorphous graphene: a realization of Zachariasen's glass, *J. Phys.: Condens. Matter*, 2012, **24**, 485003.

19 D. R. Robinson and M. Wilson, The liquid↔amorphous transition and the high pressure phase diagram of carbon, *J. Phys.: Condens. Matter*, 2013, **25**, 155101.

20 H.-P. Kaukonen and R. M. Nieminen, Molecular-dynamics simulation of the growth of diamond like films by energetic carbon-atom beams, *Phys. Rev. Lett.*, 1992, **68**, 620–623.

21 N. A. Marks, Thin film deposition of tetrahedral amorphous carbon: a molecular dynamics study, *Diamond Relat. Mater.*, 2005, **14**, 1223–1231.

22 M. A. Caro, V. L. Deringer, J. Koskinen, T. Laurila and G. Csányi, Growth mechanism and origin of high sp^3 content in tetrahedral amorphous carbon, *Phys. Rev. Lett.*, 2018, **120**, 166101.

23 M. A. Caro, G. Csányi, T. Laurila and V. L. Deringer, Machine learning driven simulated deposition of carbon films: From low-density to diamond like amorphous carbon, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2020, **102**, 174201.

24 R. Thapa, C. Ugwumadu, K. Nepal, J. Trembly and D. A. Drabold, Ab initio simulation of amorphous graphite, *Phys. Rev. Lett.*, 2022, **128**, 236402.

25 B. Bhattarai, A. Pandey and D. A. Drabold, Evolution of amorphous carbon across densities: An inferential study, *Carbon*, 2018, **131**, 168–174.

26 B. Bhattarai, P. Biswas, R. Atta-Fynn and D. A. Drabold, Amorphous graphene: a constituent part of low density amorphous carbon, *Phys. Chem. Chem. Phys.*, 2018, **20**, 19546–19551.

27 F. Wooten, K. Winer and D. Weaire, Computer generation of structural models of amorphous Si and Ge, *Phys. Rev. Lett.*, 1985, **54**, 1392–1395.

28 A. J. Stone and D. J. Wales, Theoretical studies of icosahedral C_{60} and some related species, *Chem. Phys. Lett.*, 1986, **128**, 501–503.

29 J. C. Meyer, C. Kisielowski, R. Erni, M. D. Rossell, M. F. Crommie and A. Zettl, Direct imaging of lattice atoms and topological defects in graphene membranes, *Nano Lett.*, 2008, **8**, 3582–3586.

30 J. Ma, D. Alfè, A. Michaelides and E. Wang, Stone-Wales defects in graphene and other planar sp^2-bonded materials, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2009, **80**, 033407.

31 F. Banhart, J. Kotakoski and A. V. Krasheninnikov, Structural defects in graphene, *ACS Nano*, 2011, **5**, 26–41.

32 F. L. Thiemann, P. Rowe, A. Zen, E. A. Müller and A. Michaelides, Defect-dependent corrugation in graphene, *Nano Lett.*, 2021, **21**, 8143–8150.

33 V. Kapko, D. A. Drabold and M. F. Thorpe, Electronic structure of a realistic model of amorphous graphene, *Phys. Status Solidi B*, 2010, **247**, 1197–1200.

34 F. D'Ambrosio, J. Barkema and G. T. Barkema, Efficient structural relaxation of polycrystalline graphene models, *Nanomaterials*, 2021, **11**, 1242.

35 D. Ormrod Morley and M. Wilson, Controlling disorder in two-dimensional networks, *J. Phys.: Condens. Matter*, 2018, **30**, 50LT02.

36 J. Behler, First principles neural network potentials for reactive simulations of large molecular and condensed systems, *Angew. Chem., Int. Ed.*, 2017, **56**, 12828–12840.

37 V. L. Deringer, M. A. Caro and G. Csányi, Machine learning interatomic potentials as emerging tools for materials science, *Adv. Mater.*, 2019, **31**, 1902765.

38 P. Friederich, F. Häse, J. Proppe and A. Aspuru-Guzik, Machine-learned potentials for next-generation matter simulations, *Nat. Mater.*, 2021, **20**, 750–761.

39 J. Behler and M. Parrinello, Generalized neural-network representation of high-dimensional potential-energy surfaces, *Phys. Rev. Lett.*, 2007, **98**, 146401.

40 A. P. Bartók, M. C. Payne, R. Kondor and G. Csányi, Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons, *Phys. Rev. Lett.*, 2010, **104**, 136403.

41 V. L. Deringer, C. J. Pickard and G. Csányi, Data-driven learning of total and local energies in elemental boron, *Phys. Rev. Lett.*, 2018, **120**, 156001.

42 N. Bernstein, B. Bhattarai, G. Csányi, D. A. Drabold, S. R. Elliott and V. L. Deringer, Quantifying chemical structure and machine-learned atomic energies in amorphous and liquid silicon, *Angew. Chem., Int. Ed.*, 2019, **58**, 7057–7061.

43 C. Ben Mahmoud, A. Anelli, G. Csányi and M. Ceriotti, Learning the electronic density of states in condensed matter, *Phys. Rev. B*, 2020, **102**, 235130.

44 V. L. Deringer, N. Bernstein, G. Csányi, C. Ben Mahmoud, M. Ceriotti, M. Wilson, D. A. Drabold and S. R. Elliott, Origins of structural and electronic transitions in disordered silicon, *Nature*, 2021, **589**, 59–64.

45 X. Song and C. Deng, Atomic energy in grain boundaries studied by machine learning, *Phys. Rev. B*, 2022, **6**, 043601.

46 M. Eckhoff and J. Behler, From molecular fragments to the bulk: Development of a neural network potential for MOF-5, *J. Chem. Theory Comput.*, 2019, **15**, 3793–3809.

47 V. L. Deringer and G. Csányi, Machine learning based interatomic potential for amorphous carbon, *Phys. Rev. B*, 2017, **95**, 094203.

48 N. Chetty and R. M. Martin, First-principles energy density and its applications to selected polar surfaces, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1992, **45**, 6074–6088.

49 D. W. Brenner, Empirical potential for hydrocarbons for use in simulating the chemical vapor deposition of diamond films, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1990, **42**, 9458–9471.

50 D. W. Brenner, O. A. Shenderova, J. A. Harrison, S. J. Stuart, B. Ni and S. B. Sinnott, A second-generation reactive empirical bond order (REBO) potential energy expression for hydrocarbons, *J. Phys.: Condens. Matter*, 2002, **14**, 783–802.

51 S. J. Stuart, A. B. Tutein and J. A. Harrison, A reactive potential for hydrocarbons with intermolecular interactions, *J. Chem. Phys.*, 2000, **112**, 6472–6486.

52 J. H. Los and A. Fasolino, Intrinsic long-range bond-order potential for carbon: Performance in Monte Carlo simulations of graphitization, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2003, **68**, 024107.

53 N. A. Marks, Generalizing the environment-dependent interaction potential for carbon, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2000, **63**, 035401.

54 N. A. Marks, Modelling diamond-like carbon with the environment-dependent interaction potential, *J. Phys.: Condens. Matter*, 2002, **14**, 2901–2927.

55 J. F. Justo, M. Z. Bazant, E. Kaxiras, V. V. Bulatov and S. Yip, Interatomic potential for silicon defects and disordered phases, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1998, **58**, 2539–2550.

56 M. Z. Bazant, E. Kaxiras and J. F. Justo, Environment-dependent interatomic potential for bulk silicon, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1997, **56**, 8542–8552.

57 C. de Tomas, I. Suarez-Martinez and N. A. Marks, Graphitization of amorphous carbons: A comparative study of interatomic potentials, *Carbon*, 2016, **109**, 681–693.

58 C. de Tomas, I. Suarez-Martinez, F. Vallejos-Burgos, M. J. López, K. Kaneko and N. A. Marks, Structural prediction of graphitization and porosity in carbide-derived carbons, *Carbon*, 2017, **119**, 1–9.

59 T. B. Shiell, D. G. McCulloch, D. R. McKenzie, M. R. Field, B. Haberl, R. Boehler, B. A. Cook, C. de Tomas, I. Suarez-Martinez, N. A. Marks and J. E. Bradby, Graphitization of glassy carbon after compression at room temperature, *Phys. Rev. Lett.*, 2018, **120**, 215701.

60 V. L. Deringer, A. P. Bartók, N. Bernstein, D. M. Wilkins, M. Ceriotti and G. Csányi, Gaussian process regression for materials and molecules, *Chem. Rev.*, 2021, **121**, 10073–10141.

61 A. P. Bartók, R. Kondor and G. Csányi, On representing chemical environments, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2013, **87**, 184115.

62 G. Yang, L. Li, W. B. Lee and M. C. Ng, Structure of graphene and its disorders: a review, *Sci. Technol. Adv. Mater.*, 2018, **19**, 613–648.

63 D. R. Cooper, B. D'Anjou, N. Ghattamaneni, B. Harack, M. Hilke, A. Horth, N. Majlis, M. Massicotte, L. Vandsburger, E. Whiteway and V. Yu, Experimental review of graphene, *ISRN Condens. Matter Phys.*, 2012, 501686.

64 Y. Li, F. Inam, A. Kumar, M. F. Thorpe and D. A. Drabold, Pentagonal puckering in a sheet of amorphous graphene, *Phys. Status Solidi B*, 2011, **248**, 2082–2086.

65 A. Hjorth 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, *J. Phys.: Condens. Matter*, 2017, **29**, 273002.

66 G. Csányi, S. Winfield, J. R. Kermode, A. De Vita, A. Comisso, N. Bernstein and M. C. Payne, Expressive programming for computational physics in Fortran 95+, *IoP Comput. Phys. Newsletter*, Spring, 2007.

67 J. R. Kermode, f90wrap: an automated tool for constructing deep python interfaces to modern Fortran codes, *J. Phys.: Condens. Matter*, 2020, **32**, 305901.

68 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, *Comput. Phys. Commun.*, 2022, **271**, 108171.

69 R. L. C. Vink, A finite-temperature Monte Carlo algorithm for network forming materials, *J. Chem. Phys.*, 2014, **140**, 104509.

70 D. S. Franzblau, Computation of ring statistics for network models of solids, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1991, **44**, 4925–4930.

71 Matscipy, version 0.7.0 **https://github.com/libAtoms/matscipy**, 2022.

72 A. Gervois, J. P. Troadec and J. Lemaître, Universal properties of Voronoi tessellations of hard discs, *J. Phys. A: Math. Gen.*, 1992, **25**, 6169–6177.

73 M. E. J. Newman, Assortative mixing in networks, *Phys. Rev. Lett.*, 2002, **89**, 208701.

74 A. Stukowski, Visualization and analysis of atomistic simulation data with OVITO–the Open Visualization Tool, *Modell. Simul. Mater. Sci. Eng.*, 2010, **18**, 015012.

75 A. Togo and I. Tanaka, First principles phonon calculations in materials science, *Scr. Mater.*, 2015, **108**, 1–5.

76 J. P. Perdew and A. Zunger, Self-interaction correction to density-functional approximations for many-electron systems, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1981, **23**, 5048–5079.

77 J. P. Perdew, K. Burke and M. Ernzerhof, Generalized gradient approximation made simple, *Phys. Rev. Lett.*, 1996, **77**, 3865–3868.

78 J. P. Perdew, A. Ruzsinszky, G. I. Csonka, O. A. Vydrov, G. E. Scuseria, L. A. Constantin, X. Zhou and K. Burke, Restoring the density-gradient expansion for exchange in solids and surfaces, *Phys. Rev. Lett.*, 2008, **100**, 136406.

79 S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. J. Probert, K. Refson and M. Payne, First principles methods using CASTEP, *Z. Kristall.*, 2005, **220**, 567–570.

80 D. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics*, Cambridge University Press, 2005.

81 D. W. Scott, *Multivariate density estimation: theory, practice, and visualization*, John Wiley & Sons, 2015.

82 H. Tang, X. Yuan, Y. Cheng, H. Fei, F. Liu, T. Liang, Z. Zeng, T. Ishii, M.-S. Wang, T. Katsura, H. Sheng and H. Gou, Synthesis of paracrystalline diamond, *Nature*, 2021, **599**, 605–610.

83 Z. Liu, D. Panja and G. T. Barkema, Structural dynamics of polycrystalline graphene, *Phys. Rev. E*, 2022, **105**, 044116.

84 A. C. Ferrari and J. Robertson, Interpretation of Raman spectra of disordered and amorphous carbon, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2000, **61**, 14095–14107.

85 A. C. Ferrari, J. C. Meyer, V. Scardaci, C. Casiraghi, M. Lazzeri, F. Mauri, S. Piscanec, D. Jiang, K. S. Novoselov, S. Roth and A. K. Geim, Raman spectrum of graphene and graphene layers, *Phys. Rev. Lett.*, 2006, **97**, 187401.

86 S. N. Shirodkar and U. V. Waghmare, Electronic and vibrational signatures of Stone-Wales defects in graphene: First-principles analysis, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2012, **86**, 165401.

87 Y. Wang, Z. Fan, P. Qian, T. Ala-Nissila and M. A. Caro, Structure and pore size distribution in nanoporous carbon, *Chem. Mater.*, 2022, **34**, 617–628.

88 T. Laurila, S. Sainio and M. A. Caro, Hybrid carbon based nanomaterials for electrochemical detection of biomolecules, *Prog. Mater. Sci.*, 2017, **88**, 499–594.

89 E. Olsson, J. Yu, H. Zhang, H.-M. Cheng and Q. Cai, Atomic-scale design of anode materials for alkali metal (Li/Na/K)-ion batteries: Progress and perspectives, *Adv. Energy Mater.*, 2022, **12**, 2200662.

Chapter 4

Accelerated First-Principles Exploration of Structure and Reactivity in Graphene Oxide

This chapter is reproduced in its entirety from the following publication:

> Zakariya El-Machachi, Damyan Frantzov, A. Nijamudheen, Tigany Zarrouk, Miguel A. Caro and Volker L. Deringer, "Accelerated First-Principles Exploration of Structure and Reactivity in Graphene Oxide", *Angewandte Chemie International Edition*, 2024, **63**, e202410088.

Authors' contributions: Z. E.-M. developed the structure generation protocols, trained the machine learning models and performed and analysed production simulations. D. F. contributed to the structure generation code. A. N. performed pilot studies. T. Z. and M. A. C. performed XPS simulations and analysis. V. L. D. initiated and coordinated the study. All authors made substantial contributions to data analysis and discussions. Z. E.-M. and V. L. D. wrote the paper with input from T. Z. and M. A. C., and all authors approved the final version.

How to cite: *Angew. Chem. Int. Ed.* **2024**, *63*, e202410088
doi.org/10.1002/anie.202410088

Machine Learning

Accelerated First-Principles Exploration of Structure and Reactivity in Graphene Oxide

*Zakariya El-Machachi, Damyan Frantzov, A. Nijamudheen, Tigany Zarrouk, Miguel A. Caro, and Volker L. Deringer**

Abstract: Graphene oxide (GO) materials are widely studied, and yet their atomic-scale structures remain to be fully understood. Here we show that the chemical and configurational space of GO can be rapidly explored by advanced machine-learning methods, combining on-the-fly acceleration for first-principles molecular dynamics with message-passing neural-network potentials. The first step allows for the rapid sampling of chemical structures with very little prior knowledge required; the second step affords state-of-the-art accuracy and predictive power. We apply the method to the thermal reduction of GO, which we describe in a realistic (ten-nanometre scale) structural model. Our simulations are consistent with recent experimental findings, including X-ray photoelectron spectroscopy (XPS), and help to rationalise them in atomistic and mechanistic detail. More generally, our work provides a platform for routine, accurate, and predictive simulations of diverse carbonaceous materials.

Graphene oxide (GO) is a summary term for a range of layered materials created by reacting graphite with aggressive agents, such as $KMnO_4$, typically followed by partial reduction and sometimes functionalisation.[1–4] Today, GO materials can be controllably prepared[5] and find emerging applications in catalysis,[6] membranes,[7] electronics,[8] and photonics.[9] Despite decades of work, however, the precise chemical structure of these materials has remained elusive. The ordered regions of GO sheets can be directly visualised using high-resolution electron microscopy,[10,11] but the nature of the more disordered regions can only be inferred from indirect observations, such as vibrational and NMR spectroscopy. The properties of GO materials cannot be unambiguously linked to chemical structure if this structure itself is not precisely known (which functional groups are present; in what amounts?).

To complement experimental techniques, GO has been widely studied by computational chemistry methods. For example, Kumar et al. combined reactive-force-field simulations with density-functional theory (DFT) to show how varying functional groups affect the stability and electronic structure of thermally reduced graphene oxide (rGO), and how rGO forms graphitic and oxidised domains during thermal annealing.[16,17] Atomistic modelling of rGO further revealed that defects formed during thermal reduction can lead to pores for applications in water desalination and natural gas purification.[18] Explicit water molecules have been incorporated into computational models of GO membranes to simulate interlayer separation and water diffusivity, providing insights for applications.[19–21] Theoretical studies delved into aspects such as the excess surface charge in hydrated GO and the dynamic evolution of functional groups, employing *ab initio* molecular dynamics (AIMD) for a comprehensive understanding of GO in water.[22]

Despite these advances, there remains an inherent limit to the length and time scales accessible to AIMD. Machine learning (ML) based interatomic potentials provide an emerging alternative approach that promises much faster simulations while retaining quantum-mechanical accuracy.[23–26] In the context of carbon materials, ML-driven simulations have been used to describe defective[27] and fully amorphous graphene,[28] the growth of carbon thin films,[29] and the formation of voids in low-density porous forms.[30–32]

In the present work, we show how one can rapidly explore a wide range of functional groups and disorder in GO materials by combining two recent innovations in atomistic ML (Figure 1). First, we use on-the-fly-accelerated AIMD[12,33,34] to efficiently sample configurations for an initial training dataset. Second, we show that this approach can be used to kick-start a much wider-ranging exploration using state-of-the-art neural-network potentials. For the first task, we use CASTEP+ML[12,35] coupled to the Gaussian approximation potential (GAP) ML framework;[15,36,37] for the latter, we use an equivariant neural-network architecture based on the message-passing atomic cluster expansion (MACE),[38–40] which enables highly accurate predictions beyond the system-size limits of AIMD. A key point of our study is that those two principally different methodologies can be synergistically combined. The predictions of the final ML model agree remarkably well with experimental obser-

---

[*] Z. El-Machachi, D. Frantzov, Dr. A. Nijamudheen, Prof. V. L. Deringer
Inorganic Chemistry Laboratory, Department of Chemistry
University of Oxford, Oxford OX1 3QR, United Kingdom
E-mail: volker.deringer@chem.ox.ac.uk

Dr. T. Zarrouk, Dr. M. A. Caro
Department of Chemistry and Materials Science
Aalto University, 02150 Espoo, Finland

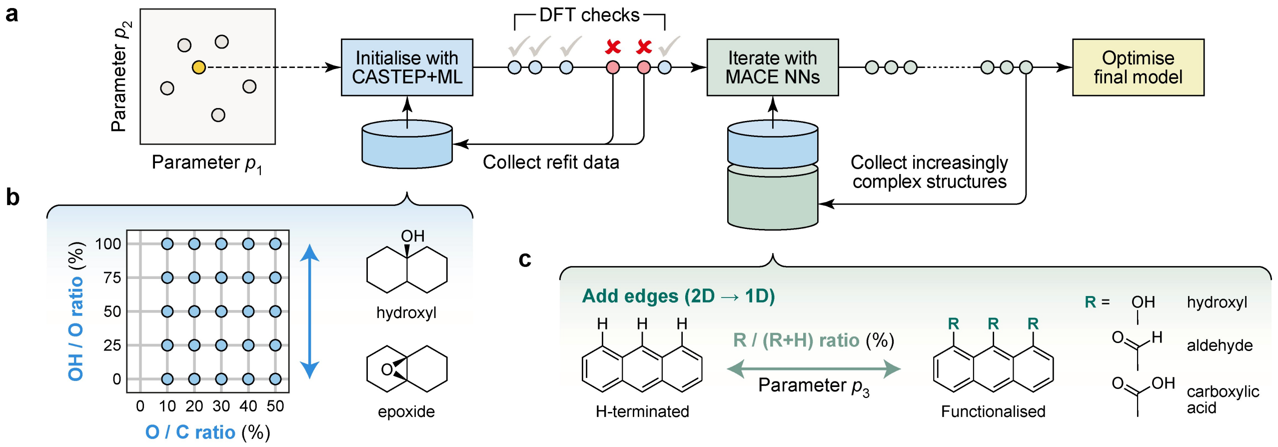

Figure 1. Accelerated exploration of functional groups in GO with machine-learning-driven simulations. (a) Schematic overview of the overall approach. We initialise the search for structures with CASTEP+ML trajectories, which combine first-principles MD with on-the-fly fitting of ML potentials[12] (blue). Once this initial generation is complete, iterative training kicks in, exploring increasingly complex structural spaces (green). The data are then used to train and optimise the final model (yellow). (b) Parameter space of 2D functionalised GO, with a schematic sketch of how the OH/O ratio controls the ratio of hydroxyl and epoxy groups in the initial structures. (c) Extension of the parameter space to include 1D structures (edges and ribbons), which can be hydrogen-terminated or functionalised with different groups, R.

vations, showing promise for future applications to the chemistry of carbon-based materials.

To sample the wide variety of possible GO structures, we define a space of N parameters, $\mathcal{P} = [p_1, \ldots, p_N]$, that determines the composition of an initial candidate structure. Here, in line with existing knowledge in the field,[1–4] we choose these parameters to be: (1) the ratio of O to C atoms in the initial sheet, determining the degree of oxidation; (2) the OH/O ratio, *i.e.*, the concentration of hydroxyl groups relative to all O atoms; (3) the ratio of functionalised edges (–OH, –CHO, or $–CO_2H$) to hydrogen-terminated ones. We explored up to p_2 in CASTEP+ML runs (that is, we functionalised only 2D graphene sheets), and up to p_2 and then p_3 in MACE iterations.

We started the process with 25 CASTEP+ML runs, at 300 K for 10 ps each, corresponding to the grid shown in Figure 1b. Most of those simulations (20 of 25) ran to completion; five terminated early due to erroneously lost atoms. The latter results are still valuable, as they contain high-energy and -force structures which can be used to guide early models away from unphysical configurations during iterations. In all, 820 CASTEP+ML structures were used for the initial training dataset. Next, equivariant MACE potentials were trained in an iterative fashion that gradually extended the scope of the model: (i) exploring higher temperatures in MD runs, gradually increasing from 600 to 1,500 K; (ii) repeating the protocol at 1,500 K for the next four iterations; (iii) finally, exploring 1D structures at 1,500 K. After a total of 12 iterations, training structures with any force component >50 eV/Å were filtered out to further enhance the model. The final training dataset contains 3,016 simulation snapshots (605,204 atoms) and is described in detail in the Supporting Information.

We test the performance of our ML models on external data not seen in training: two AIMD trajectories of functionalised 2D sheets and 1D nanoribbons, respectively, as well as single-point calculations for molecular (0D) fragments taken from Ref. [13]. Figure 2a shows how the force error—our main performance metric—evolves during iterations. As more data are added, the errors decrease for the 2D test set, as expected. For 1D and 0D structures, the errors are initially high since early models have not "seen" edges, specifically C–H bonds which are explicitly included only from model 10 onwards. Adding edge structures rapidly reduces the corresponding errors (dashed line in Figure 2a). The final force accuracy is similar across all benchmarks, just over 0.1 eV/Å.

To illustrate the gradual exploration of chemical and configurational space, we show four SOAP similarity maps[15,43] in Figure 2b: the CASTEP+ML seed at 300 K, the dataset after gradually ramping to 1,500 K, the inclusion of the first edge structures, and the final dataset. The initial structures form two distinct clusters on the map; at higher T, one cluster grows and a third, smaller one appears. Finally, including edges adds a distinct set of structures (green).

We now describe an application of the final MACE model to a challenging problem in materials chemistry—namely, to large-scale MD simulations of the thermal reduction of GO to rGO. This process involves a vast number of functional groups which transform and eventually disappear, accompanied by the evolution of gaseous species such as CO_2. Experimentally, reduction temperatures of 1,100 C yielded resistivity values of ~$10^{-5}\,\Omega$ m,[44] on par with that of graphite.[45] Understanding how functional groups evolve during thermal reduction could help to correlate the structure of the sheet with its properties. We show in the following that our ML-accelerated approach can provide such an atomic-scale understanding.

Our starting structure is a partially disordered, fully sp^2-bonded graphene sheet with 10,368 atoms (17.7×15.3 nm^2 in a single layer). The sheet was generated using Monte-Carlo bond switching driven by ML local-environment

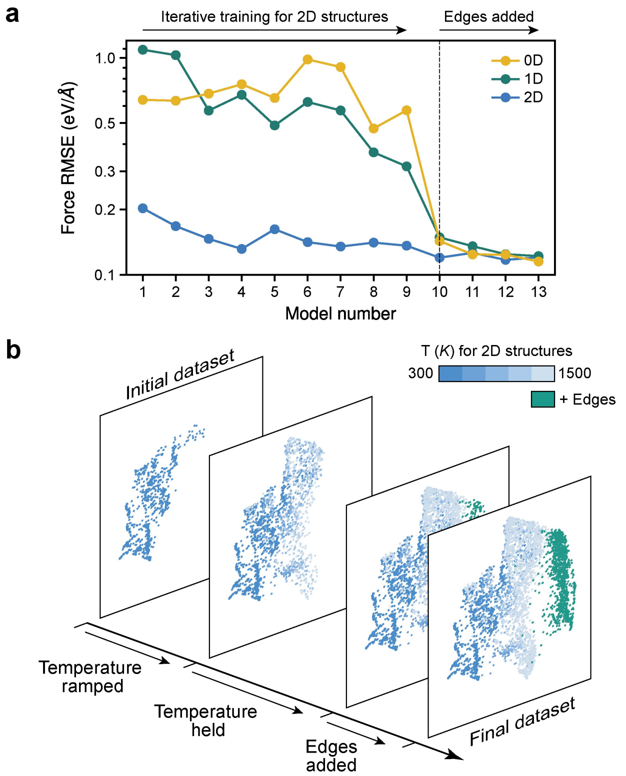

Figure 2. Training ML models for GO. (a) Root-mean-square error (RMSE) of forces predicted by iteratively trained MACE models. Errors are evaluated on an external set of DFT data not included in the training, comprising 100 snapshots each of a 2D sheet (blue) and a 1D nanoribbon (green) sampled in separate AIMD trajectories at 300 K, as well as 0D nanoflake structures (yellow) taken from Ref. [13] and re-evaluated at the relevant level of DFT. (b) Structural diversity during iterative training. The different iterations are visualised by UMAP embedding[14] of a kernel-based structural similarity metric.[15] Points are colour-coded according to the temperature set in the MD simulation for 2D structures (blue), whereas they are shown in a single different colour for 1D edge structures (green).

energies, following Ref. [28], and then functionalised with $\mathcal{P} = [0.4, 0.5, 0]$, raising the atom count to 16,645. This structural model represents features of GO including the topological disordering of the carbon backbone (presence of non-6-membered rings), although it is constrained to threefold coordination for all carbon atoms, and therefore does not initially contain large pores.

The thermal reduction was studied in three independent MD simulations at temperatures of 900, 1,200, and 1,500 K, respectively. The structures were rapidly heated over 100 ps and then held at the respective annealing temperature for 1.9 ns. We note that experimental protocols for thermal reduction of GO span a wide range of parameters: temperatures from 80 to 1,100 C[44] and times from 10 minutes[46] to 5 days.[17] Computationally, we are limited by the timescales accessible to MD (on the order of nanoseconds); thus, more aggressive heating is used to overcome local energy barriers.[47] We found that annealing at 1,500 K yields a structure in good agreement with experiment, which we discuss below; results for the other MD runs are given in the Supporting Information.

Figure 3a shows the rGO structural model during annealing at 1,500 K. "Graphene-like" regions, shown in dark grey, form small islands embedded within disordered and porous regions. (We quantify "graphene-like" content through polyhedral template matching, a method to identify crystal-like local environments.[41]) This result agrees qualitatively with electron microscopy images clearly showing amorphous regions together with holes and pores in the structure.[48]

The formation of this structure is accompanied by a mass loss of $>20\,\%$ as gaseous species leave the surface (Figure 3b), which can be qualitatively correlated with thermogravimetric experiments.[49] Initially, as the temperature ramps up over the first 100 ps, nearly all mass loss is due to H_2O (light blue shading in the stacked plot of Figure 3b). Having reached 1,500 K (dashed line), the first CO_2 molecules detach—and this species, indicated by red shading, quickly begins to dominate the mass loss as the sheet is reduced. CO was also evolved in notable amounts (magenta); other gaseous species such as OH, C_3O_2, H_2O_2, etc., were occasionally observed but were mostly rare and short-lived.

In addition to the mass loss, our simulations allow us to address the changes in the structure of the GO sheet itself. Figure 3c shows the fraction of "graphene-like" atoms in ordered local environments: initially, their percentage decreases, correlating with the loss of H_2O; then, as CO_2 is released, there is a clear and concomitant increase in graphenic content. These observations agree with experimental findings where CO_2 and, to a lesser extent, CO loss leads to defects in the GO sheet,[49–51] inducing structural rearrangements and increasingly more graphene-like environments. Interestingly, the count of 6-membered rings containing only carbon decreases slightly (purple in Figure 3c), whereas if we consider all atoms in the sheet in the ring analysis, the 6-membered rings increase from just above 65 % to $\approx 85\,\%$ (orange). This analysis suggests that oxygen atoms take part in substantial rearrangements during annealing, from sp^3 environments perpendicular to the basal plane such as epoxide and alcohol groups, to sp^2 environments parallel to the plane—for example, ethers and esters (cf. Figure 3a).

Beyond the overall "graphene-like-ness", we can also trace the evolution of individual functional groups, enabling qualitative comparison with ^{13}C NMR results from Ref. [49]. Figure 3d shows how the C_{sp^2}–C_{sp^2} count increases exponentially during the initial 100 ps—in contrast to the graphenic content, which decreases during this period (Figure 3c). This observation suggests that as water molecules leave the sheet, the carbon backbone will initially form defective rather than ordered sp^2 environments. Then, during high-T annealing, the sp^2 count grows more gradually, indicating a transformation to a more graphene-like carbon backbone, consistent with Figure 3c. The loss of oxygen-based molecules (cf. Figure 3b) is clearly mirrored in a declining C_{sp^3}–O–C_{sp^3} (epoxide) and C_{sp^3}–OH (hydroxyl) count during the first 300 ps of the simulation. Concomitantly, C_{sp^2}–O

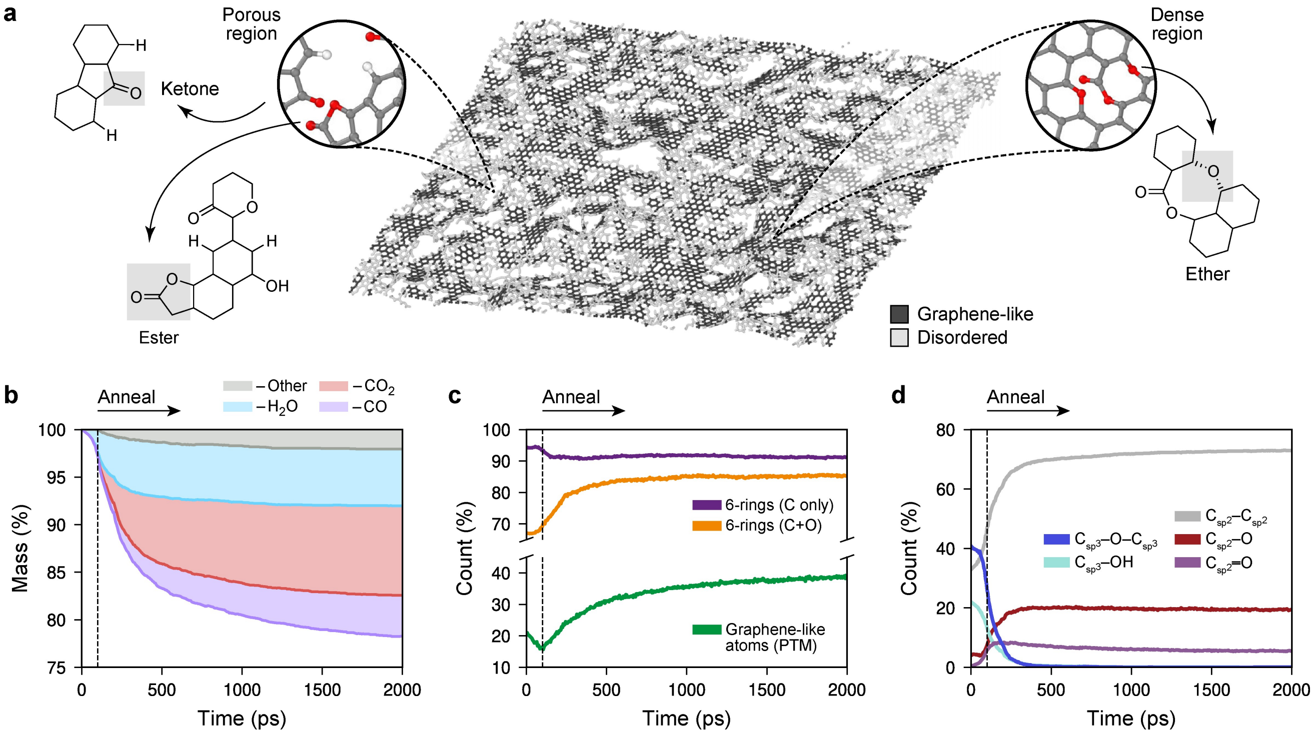

Figure 3. A large-scale structural model of reduced graphene oxide (rGO) generated through simulated thermal reduction. (a) The rGO structure after 1.5 ns of simulation time (3 million timesteps). The dark grey regions highlight graphene-like regions identified using polyhedral template matching (PTM).[41,42] Insets show close-ups with C atoms in grey, O in red, and H in white. (b) The change in mass of the GO sheet as it is thermally reduced in the 1,500 K simulation. The three most common leaving molecules (H_2O, blue; CO_2, red; CO, purple) are tracked in the stacked plot, along with other species (grey). (c) Evolution of structural indicators: the count of 6-membered rings (shown separately for only C-based rings, purple, and for all rings, orange), and the count of graphene-like atoms identified by PTM. (d) Evolution of functional groups bonded to sp^2 and sp^3 carbon atoms, respectively, obtained using a topological bond-counting algorithm.

groups form, which has been proposed as a mechanism by which the material drops in resistivity without a significant change in mass.[49] Once the sp^3-bonded epoxide and hydroxyl groups are removed, the C_{sp^2}–O count remains at an almost steady state. Finally, the C_{sp^2}=O (carbonyl) count peaks at ≈ 250 ps before gently decreasing for the remainder of the annealing simulation.

With a quantum-mechanically accurate description of the chemical structure in hand, the newly created structural models can now be analysed with advanced X-ray spectroscopy predictions,[52,53] as we have previously exemplified for small-scale GO models (with DFT-level predictions at the time).[54] Here, applying the ML spectroscopy model of Ref. [52] reveals two peaks in the predicted X-ray photoelectron spectroscopy (XPS) data for the initial structure (Figure 4a). The first peak corresponds to unmodified "sp^2" carbon atoms, whilst the second relates to oxygen/hydrogen-based functional groups, with the largest contribution arising from C_{sp^3}–O–C_{sp^3} (epoxide, dark blue) and C_{sp^3}–OH (alcohol, light blue) groups. The core electron binding energies (CEBEs) of all functional groups are shifted upwards from experimental reference energies due to the electronegativity of oxygen. The eventual removal of these oxygen-based groups during annealing reduces the magnitude of the aforementioned CEBE shifts: all motifs decrease in CEBE.

Comparison with experimental data from Ref. [49], shown in Figure 4b, reveals good agreement between theory and experiment for rGO at 300 K (in air). Experimentally, XPS spectra use a fixed reference value for deconvolution, which does not take into account the effect of local interactions from electronegative species. As a result, the spectra are shifted accordingly to align well with experimental data. We refer the reader to Refs. [52] and [53] for a detailed discussion on this matter.

In conclusion, we have reported a computational approach to modelling and understanding the highly diverse chemical structure of GO, and we have shown an initial application to the thermal reduction of this material. Our work combines two recent developments in atomistic ML. For CASTEP+ML, we view its main advantage in this context to be in saving "human time": it allows the researcher to create, from scratch, a chemically diverse training dataset to seed a new ML potential with minimal manual input.[12] For MACE, our work builds on recent capability demonstrations,[55,56] showing that this architecture can be combined with efficient dataset-building workflows to readily deploy to new, challenging modelling problems in chemistry.

Looking forward, we expect this combined methodology to provide a powerful platform for further studies of GO materials. For example, the present proof-of-concept for

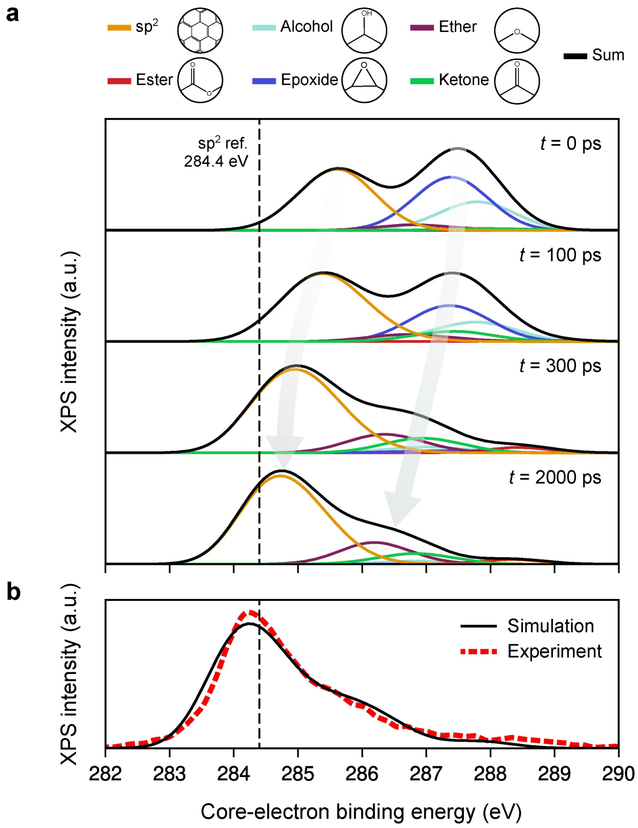

Figure 4. X-ray photoelectron spectroscopy (XPS) predictions for the rGO structure annealed at 1,500 K. (a) ML-predicted XPS spectra at different points of the simulation. The grey arrows indicate a clear shift to lower core-electron binding energy values and also a reduction in the second peak during the simulation as oxygen-based functional groups are removed. The vertical, dashed black line indicates the reference sp^2 value at 284.4 eV. (b) XPS prediction for the final structural model obtained after full structural optimisation. The simulation data in panel (b) are shifted horizontally to align with the experimental data from Ref. [49].

XPS prediction during (simulated) structural transformations could motivate future *in situ* experiments—as an example of advanced, experimentally-compatible modelling in which both the simulation *and* the XPS model are based on machine-learned quantum-mechanical data. Beyond the simulations of GO in vacuum reported here, the interaction of the material with water has been studied using empirical[19,20] and DFT methods,[22] and it would now be interesting to use ML-accelerated modelling to more fully explore the nature of water between GO sheets—building on combined experimental and simulation studies in this area,[57] and also on recent ML-driven work on unconventional phases of water "sandwiched" between sheets of pristine graphene.[58] A long-term vision could be to use predictive ML-driven simulations to find ways to optimise the nanoscale structure and thus the properties of GO materials—for example, porosity or catalytic activity—directly informing the preparation of samples in the laboratory.

Acknowledgements

We thank I. Batatia for technical discussions about MACE model fitting and code implementation, and Dr C. Ben Mahmoud for valuable feedback on the manuscript. We are grateful for support from the EPSRC Centre for Doctoral Training in Theory and Modelling in Chemical Sciences (TMCS), under grant EP/L015722/1. This paper conforms to the RCUK data management requirements. This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/V049178/1]. We are grateful for computational support from the UK national high performance computing service, ARCHER2, for which access was obtained via the UKCP consortium and funded by EPSRC grant ref EP/X035891/1. T.Z. and M.A.C. acknowledge financial support from the Research Council of Finland under grants no. 330488, 347252 and 355301, as well as computational resources from CSC (the Finnish IT Center for Science) and Aalto University's Science IT project.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability Statement

ata supporting this publication are available at https://doi.org/10.5281/zenodo.14066557.

Keywords: carbon materials · computational chemistry · graphene · machine learning · neural-network potentials

[1] A. M. Dimiev, S. Eigler, *Graphene Oxide: Fundamentals and Applications*, John Wiley & Sons, Chichester, 1st edition **2017**.
[2] D. R. Dreyer, S. Park, C. W. Bielawski, R. S. Ruoff, *Chem. Soc. Rev.* **2010**, *39*, 228.
[3] S. Guo, S. Garaj, A. Bianco, C. Ménard-Moyon, *Nat. Rev. Phys.* **2022**, *4*, 247.
[4] J. Wu, H. Lin, D. J. Moss, K. P. Loh, B. Jia, *Nat. Chem. Rev.* **2023**, *7*, 162.
[5] D. A. Dikin, S. Stankovich, E. J. Zimney, R. D. Piner, G. H. B. Dommett, G. Evmenenko, S. T. Nguyen, R. S. Ruoff, *Nature* **2007**, *448*, 457.
[6] C. Su, M. Acik, K. Takai, J. Lu, S.-j. Hao, Y. Zheng, P. Wu, Q. Bao, T. Enoki, Y. J. Chabal, K. Ping Loh, *Nat. Commun.* **2012**, *3*, 1298.
[7] R. K. Joshi, P. Carbone, F. C. Wang, V. G. Kravets, Y. Su, I. V. Grigorieva, H. A. Wu, A. K. Geim, R. R. Nair, *Science* **2014**, *343*, 752.
[8] G. Eda, M. Chhowalla, *Adv. Mater.* **2010**, *22*, 2392.
[9] J. Wu, L. Jia, Y. Zhang, Y. Qu, B. Jia, D. J. Moss, *Adv. Mater.* **2021**, *33*, 2006415.
[10] K. Erickson, R. Erni, Z. Lee, N. Alem, W. Gannett, A. Zettl, *Adv. Mater.* **2010**, *22*, 4467.
[11] S. H. Dave, C. Gong, A. W. Robertson, J. H. Warner, J. C. Grossman, *ACS Nano* **2016**, *10*, 7515.

[12] T. K. Stenczel, Z. El-Machachi, G. Liepuoniute, J. D. Morrow, A. P. Bartók, M. I. J. Probert, G. Csányi, V. L. Deringer, *J. Chem. Phys.* **2023**, *159*, 044803.

[13] B. Motevalli, A. J. Parker, B. Sun, A. S. Barnard, *Nano Futur.* **2019**, *3*, 045001.

[14] L. McInnes, J. Healy, N. Saul, L. Grossberger, *J. Open Source Softw.* **2018**, *3*, 861.

[15] A. P. Bartók, R. Kondor, G. Csányi, *Phys. Rev. B* **2013**, *87*, 184115.

[16] P. V. Kumar, M. Bernardi, J. C. Grossman, *ACS Nano* **2013**, *7*, 1638.

[17] P. V. Kumar, N. M. Bardhan, S. Tongay, J. Wu, A. M. Belcher, J. C. Grossman, *Nat. Chem.* **2014**, *6*, 151.

[18] L.-C. Lin, J. C. Grossman, *Nat. Commun.* **2015**, *6*, 8335.

[19] C. D. Williams, P. Carbone, F. R. Siperstein, *Nanoscale* **2018**, *10*, 1946.

[20] C. D. Williams, P. Carbone, F. R. Siperstein, *ACS Nano* **2019**, *13*, 2995.

[21] R. Futamura, T. Iiyama, T. Ueda, P. A. Bonnaud, F.-X. Coudert, A. Furuse, H. Tanaka, R. J. M. Pellenq, K. Kaneko, *Nat. Commun.* **2024**, *15*, 3585.

[22] F. Mouhat, F.-X. Coudert, M.-L. Bocquet, *Nat. Commun.* **2020**, *11*, 1566.

[23] J. Behler, *Angew. Chem. Int. Ed.* **2017**, *56*, 12828.

[24] V. L. Deringer, M. A. Caro, G. Csányi, *Adv. Mater.* **2019**, *31*, 1902765.

[25] P. Friederich, F. Häse, J. Proppe, A. Aspuru-Guzik, *Nat. Mater.* **2021**, *20*, 750.

[26] O. T. Unke, S. Chmiela, H. E. Sauceda, M. Gastegger, I. Poltavsky, K. T. Schütt, A. Tkatchenko, K.-R. Müller, *Chem. Rev.* **2021**, *121*, 10142.

[27] F. L. Thiemann, P. Rowe, A. Zen, E. A. Müller, A. Michaelides, *Nano Lett.* **2021**, *21*, 8143.

[28] Z. El-Machachi, M. Wilson, V. L. Deringer, *Chem. Sci.* **2022**, *13*, 13720.

[29] M. A. Caro, V. L. Deringer, J. Koskinen, T. Laurila, G. Csányi, *Phys. Rev. Lett.* **2018**, *120*, 166101.

[30] V. L. Deringer, C. Merlet, Y. Hu, T. H. Lee, J. A. Kattirtzi, O. Pecher, G. Csányi, S. R. Elliott, C. P. Grey, *Chem. Commun.* **2018**, *54*, 5988.

[31] Y. Wang, Z. Fan, P. Qian, T. Ala-Nissila, M. A. Caro, *Chem. Mater.* **2022**, *34*, 617.

[32] C. Ugwumadu, R. Thapa, K. Nepal, A. Gautam, Y. Al-Majali, J. Trembly, D. A. Drabold, *J. Chem. Theory Comput.* **2024**, *20*, 1753.

[33] Z. Li, J. R. Kermode, A. De Vita, *Phys. Rev. Lett.* **2015**, *114*, 096405.

[34] R. Jinnouchi, K. Miwa, F. Karsai, G. Kresse, R. Asahi, *J. Phys. Chem. Lett.* **2020**, *11*, 6946.

[35] S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. I. J. Probert, K. Refson, M. C. Payne, *Z. Kristallogr.* **2005**, *220*, 567.

[36] A. P. Bartók, M. C. Payne, R. Kondor, G. Csányi, *Phys. Rev. Lett.* **2010**, *104*, 136403.

[37] V. L. Deringer, A. P. Bartók, N. Bernstein, D. M. Wilkins, M. Ceriotti, G. Csányi, *Chem. Rev.* **2021**, *121*, 10073.

[38] I. Batatia, D. P. Kovacs, G. Simm, C. Ortner, G. Csanyi, MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields, in S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Editors), *Advances in Neural Information Processing Systems*, volume 35, Curran Associates, Inc. **2022** pages 11423–11436.

[39] I. Batatia, S. Batzner, D. P. Kovács, A. Musaelian, G. N. C. Simm, R. Drautz, C. Ortner, B. Kozinsky, G. Csányi, The Design Space of E(3)- Equivariant Atom-Centered Interatomic Potentials, arXiv:2205.06643 [stat.ML].

[40] D. P. Kovács, I. Batatia, E. S. Arany, G. Csányi, *J. Chem. Phys.* **2023**, *159*, 044118.

[41] P. M. Larsen, S. Schmidt, J. Schiøtz, *Model. Simul. Mater. Sci. Eng.* **2016**, *24*, 055007.

[42] A. Stukowski, *Model. Simul. Mater. Sci. Eng.* **2009**, *18*, 015012.

[43] B. Cheng, R.-R. Griffiths, S. Wengert, C. Kunkel, T. Stenczel, B. Zhu, V. L. Deringer, N. Bernstein, J. T. Margraf, K. Reuter, G. Csanyi, *Acc. Chem. Res.* **2020**, *53*, 1981.

[44] C. Mattevi, G. Eda, S. Agnoli, S. Miller, K. A. Mkhoyan, O. Celik, D. Mastrogiovanni, G. Granozzi, E. Garfunkel, M. Chhowalla, *Adv. Funct. Mater.* **2009**, *19*, 2577.

[45] W. W. Tyler, A. C. Wilson, *Phys. Rev.* **1953**, *89*, 870.

[46] V. M. Maphiri, G. Rutavi, N. F. Sylla, S. A. Adewinbi, O. Fasakin, N. Manyala, *Nanomaterials* **2021**, *11*, 1909.

[47] C. de Tomas, I. Suarez-Martinez, F. Vallejos-Burgos, M. J. López, K. Kaneko, N. A. Marks, *Carbon* **2017**, *119*, 1.

[48] M. Ceniceros-Reyes, K. Marín-Hernández, U. Sierra, E. Saucedo-Salazar, R. Mendoza-Resendez, C. Luna, P. Hernández-Belmares, O. Rodríguez-Fernández, S. Fernández-Tavizón, E. Hernández-Hernández, E. D. Barriga-Castro, *Surf. Interfaces* **2022**, *35*, 102448.

[49] C. Valentini, V. Montes-García, P. A. Livio, T. Chudziak, J. Raya, A. Ciesielski, P. Samorì, *Nanoscale* **2023**, *15*, 5743.

[50] N. D. K. Tu, J. Choi, C. R. Park, H. Kim, *Chem. Mater.* **2015**, *27*, 7362.

[51] M. Pelaez-Fernandez, A. Bermejo, A. Benito, W. Maser, R. Arenal, *Carbon* **2021**, *178*, 477.

[52] D. Golze, M. Hirvensalo, P. Hernández-León, A. Aarva, J. Etula, T. Susi, P. Rinke, T. Laurila, M. A. Caro, *Chem. Mater.* **2022**, *34*, 6240.

[53] T. Zarrouk, R. Ibragimova, A. P. Bartók, M. A. Caro, *J. Am. Chem. Soc.* **2024**, *146*, 14645.

[54] A. Aarva, S. Sainio, V. L. Deringer, M. A. Caro, T. Laurila, *J. Phys. Chem. C* **2021**, *125*, 18234.

[55] D. P. Kovács, J. H. Moore, N. J. Browning, I. Batatia, J. T. Horton, V. Kapil, W. C. Witt, I.-B. Magdău, D. J. Cole, G. Csányi, MACE-OFF23: Transferable Machine Learning Force Fields for Organic Molecules, arXiv:2312.15211 [physics.chem-ph].

[56] 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. Della 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, G. Csányi, A foundation model for atomistic materials chemistry, arXiv:2401.00096 [physics.chem-ph].

[57] S. Zheng, Q. Tu, J. J. Urban, S. Li, B. Mi, *ACS Nano* **2017**, *11*, 6440.

[58] V. Kapil, C. Schran, A. Zen, J. Chen, C. J. Pickard, A. Michaelides, *Nature* **2022**, *609*, 512.

Manuscript received: May 28, 2024
Accepted manuscript online: August 12, 2024
Version of record online: November 13, 2024

Angewandte Chemie
Eine Zeitschrift der Gesellschaft Deutscher Chemiker

Supporting Information

Accelerated First-Principles Exploration of Structure and Reactivity in Graphene Oxide

*Z. El-Machachi, D. Frantzov, A. Nijamudheen, T. Zarrouk, M. A. Caro, V. L. Deringer**

Supporting Information for

"Accelerated First-Principles Exploration of Structure and Reactivity in Graphene Oxide"

Zakariya El-Machachi[1], Damyan Frantzov[1], A. Nijamudheen[1], Tigany Zarrouk[2], Miguel A. Caro[2], and Volker L. Deringer*[1]

[1]*Inorganic Chemistry Laboratory, Department of Chemistry, University of Oxford, Oxford OX1 3QR, United Kingdom*
[2]*Department of Chemistry and Materials Science, Aalto University, 02150 Espoo, Finland*

---

*volker.deringer@chem.ox.ac.uk

Computational methods

ML acceleration for CASTEP. The CASTEP+ML scheme described in Ref. S1 was used to accelerate the sampling of the relevant configurational space (as compared to full *ab initio* molecular dynamics) and thus to construct an initial dataset for ML potential fitting. All DFT computations were performed at the Γ point and used the PBE functional with a plane-wave cutoff of 550 eV and a SCF halting criterion of $\Delta E < 10^{-5}$ eV at.$^{-1}$. A Gaussian smearing width of 0.1 eV was applied and all computations performed are non-spin-polarized.

For CASTEP+ML, the adaptive fitting method (increasing and decreasing the number of steps between DFT checks, n, depending on model performance; Ref. S1) was used, where $n_{\mathrm{min}} = 1$ and $n_{\mathrm{max}} = 10000$ with an adaptive scaling factor of 2. The tolerances used for refitting checks were an energy difference of < 0.01 eV at.$^{-1}$, a maximum force difference < 3.00 eV Å^{-1}, and a force RMSE of < 0.50 eV Å^{-1}.

The GAP models fitted by CASTEP+ML are based on a combination of 2-body, 3-body, and SOAP terms, similar to the C-GAP-17 model [S2], however with modified hyperparameters and much fewer sparse (representative) points particularly for the SOAP term. The GAP fitting string used for on-the-fly potential fitting is given in Listing 1 below.

Listing 1: GAP fitting string for CASTEP+ML runs.

```
default_sigma: "0.008 0.04 0 0"
descriptor_str: "distance_Nb order=2 cutoff=4.5 covariance_type=ard_se delta
    =2.0 theta_uniform=1.0 sparse_method=uniform add_species=T n_sparse=15 :
    distance_Nb order=3 cutoff=2.8 covariance_type=ard_se delta=0.5
    theta_uniform=1.0 add_species=T n_sparse=50 sparse_method=uniform : soap
    cutoff=4.5 covariance_type=dot_product zeta=4.0 delta=0.05 atom_sigma=0.5
     l_max=8 n_max=8 n_sparse=200 sparse_method=cur_points"
extra_gap_opts: "sparse_jitter = 1.0e-8"
```

MACE fitting. Equivariant MACE models [S3] were fitted on a single NVIDIA RTX A6000 graphics card in a Linux workstation. The final model required about 53 hours to train. We used the MACE code version 0.2.0, development branch, commit `55f7411`. The input used for MACE fitting is given in Listing 2 below.

Listing 2: MACE fitting input. “x” denotes the iteration number in the training and testing file.

```
CUDA_VISIBLE_DEVICES="$gpu_id" python mace/scripts/run_train.py \
--name="MACE_model" \
--train_file="structures/iter-x-train.xyz" \
--valid_fraction=0.10 \
--test_file="structures/iter-x-test.xyz" \
--config_type_weights='{"Default":1.0}' \
--model="ScaleShiftMACE" \
--hidden_irreps='128x0e+128x1o' \
--loss='huber' \
--r_max=3.7 \
--batch_size=30 \
--max_num_epochs=2000 \
--swa \
--default_dtype='float32' \
--energy_key='QM_energy' \
--forces_key='QM_forces' \
--stress_key=None \
--start_swa=1000 \
--swa_energy_weight=1000 \
--swa_forces_weight=100 \
--lr=0.001 \
--ema \
--ema_decay=0.99 \
--amsgrad \
--restart_latest \
--device=cuda \
--seed=123 \
```

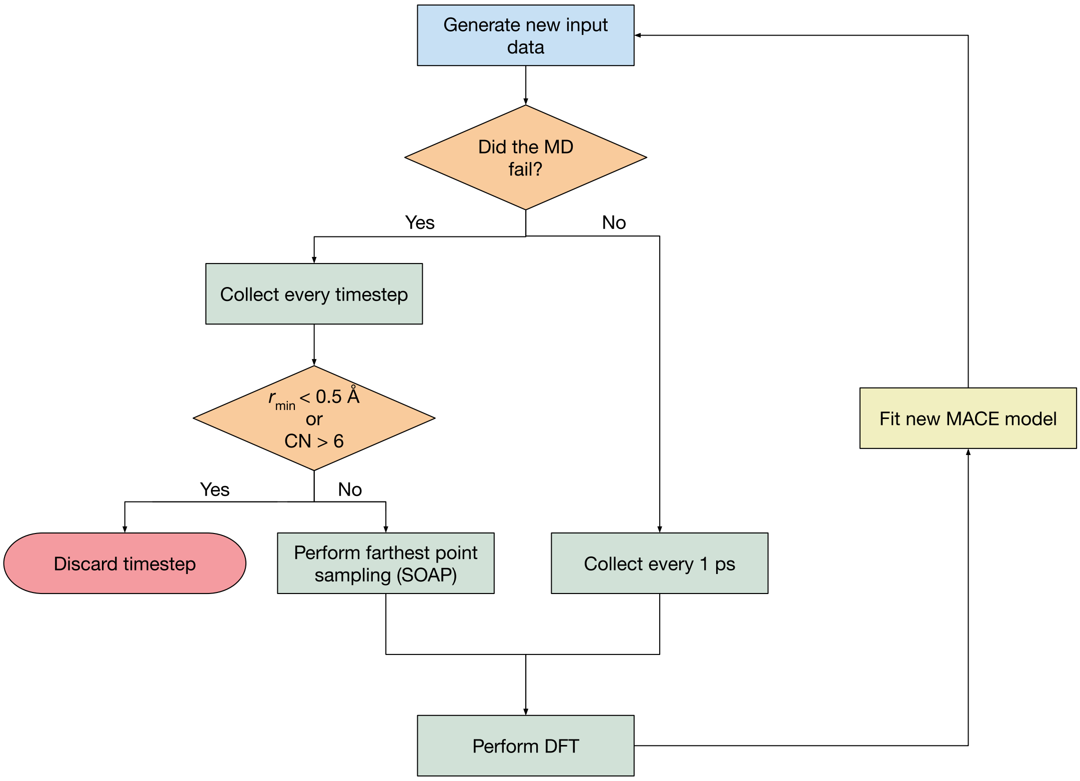

Figure S1: Flowchart outlining the process for selecting structures during iterative training.

Data filtering. The iterative process occasionally generated structures with very (even unreasonably) high energies and forces. We found that a process for filtering high-energy and -force structures throughout the iterative training process was important for the stability of early MACE models. Including highly unfavourable structures resulted in a less smooth regression—however, without inclusion of these unfavourable structures, some early models displayed instabilities and led to catastrophic failure during molecular dynamics (MD), such as lost atoms or nonphysical clustering. The approach to filtering structures during iterative training is illustrated in Figure S1.

Initially, we check whether a given MD run has completed without failing, in which case we simply collect structures at every 1 ps from the MD trajectory. However, if a run did fail, we investigate the trajectory in more detail. We heuristically define two criteria to determine whether a structure is fully nonphysical from the iterative MD (and should therefore not be included in the training dataset). We deem a structure to be fully nonphysical if it meets one of the following criteria:

- Any two atoms come closer than 0.5 Å (denoted “r_{min}” in Figure S1)
- The coordination number (CN) of any atom is greater than 6

The cutoff values were experimented with and proved to be sensitive to the types of structures filtered. Values of > 0.5 Å did not include enough unfavourable structures, whereas values of < 0.5 Å would include structures that then failed to converge in DFT.

Farthest point sampling (FPS) was used on the structures from the trajectory that remained after the filtering step. FPS was performed on the per-structure SOAP vectors constructed from the average atomic SOAP vectors of that structure. This enabled an automated and rigorous sampling of configurational space related to these failed MD runs. The final dataset was filtered of structures containing force components above 50 eV/Å.

Iterative training. MD simulations using the Atomic Simulation Environment (ASE) were used to generate configurations iteratively. The Nosé–Hoover thermostat in the NVT ensemble was used with an MD timestep of 0.2 fs. For the first four iterations, the temperature of the MD simulation was increased in increments of 300 K, starting from 600 K where model-1, which was trained on the CASTEP+ML data generated at 300 K, was used to drive MD across the parameter space and ending at 1,500 K for model-4, which had been trained iteratively from data generated from subsequent models. This iterative approach was chosen to enable a wider-ranging sampling of configuration space, without jeopardising stability in early models. We found that up to 1,200 K, the resulting MACE potentials were stable for at least 10 ps. At 1,200 K, the MACE potential at this iteration (model-3) was not stable for two runs at high O content and low OH/O ratio ($p_1 = 0.40$, $p_2 = 0.00$ and $p_1 = 0.50$, $p_2 = 0.25$), where model-3 incorrectly predicts a minimum in the potential energy surface, leading to unphysical clustering of atoms. This behaviour is somewhat expected during early stages of the iterative training of ML potentials, where the configurational space is still actively being explored. Furthermore, explicit models of isolated dimers are not included in the training data and thus extremely short-range interactions ($r \leq 0.5$ Å) are poorly described by early models. This was also observed at 1,500 K in the first instance (iteration 4) where two configurations failed ($p_1 = 0.50, p_2 = 0.00$ and $p_1 = 0.50, p_2 = 0.25$). 7 structures were removed from the final training set and 0 were removed from the test set.

Training and test sets. The total number of structures in the training set is 3,013 (+ 3 cells containing isolated atoms), for a total of 605,204 atoms. The testing set contains 800 structures (163,052 atoms). We also report the number of gaseous species from the training and testing set in Table S1. These gaseous molecules were identified using the average ionic radius to build

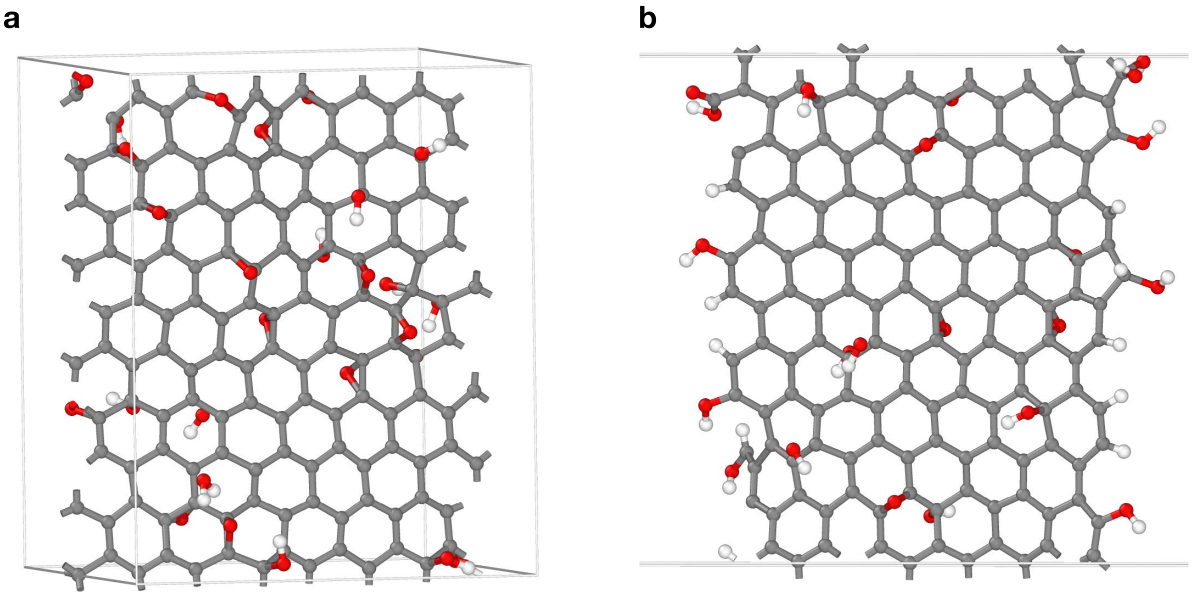

Figure S2: Representative structures from the training set. Panel (a) shows an example of a 2D structure whereas panel (b) shows an example of a 1D structure. Grey atoms are C, red atoms are O, and white atoms are H.

Table S1: Total gaseous species count from training and test set as determined through topological bond counting using pairwise cutoff radii and cluster analysis within OVITO [S4].

Gaseous Species	OH	H_2O	Other	CO	H	CO_2	O	O_2
Training Set Count	853	782	46	32	27	18	16	13
Test Set Count	213	213	13	9	8	7	4	2

pairwise cutoffs (with the exception of C–O, which required an increased cutoff of 1.85 Å, see Figure S13 below) to define bonds. These bonds were passed to the cluster analysis tool within OVITO [S4] to identify unique cluster fragments which could then be separately counted.

We note a relatively high count of OH and H_2O species in comparison to CO_2 and CO is observed due to the short MD time and high MD temperature (1,500 K for iter-3 and beyond) during iterative training. As seen in the main text in Figure 3, H_2O groups are lost early on during the simulation, whereas CO_2 dominates later on. Species labelled 'Other' are generated due to the high temperature and are short-lived.

Production MD. The same MD parameters as in the iterative training process were used for the production runs characterised in Figure 3, with the exception of the timestep, which was increased to 0.5 fs. Three production runs were conducted at 900 K, 1,200 K, and 1,500 K for the same initial structure.

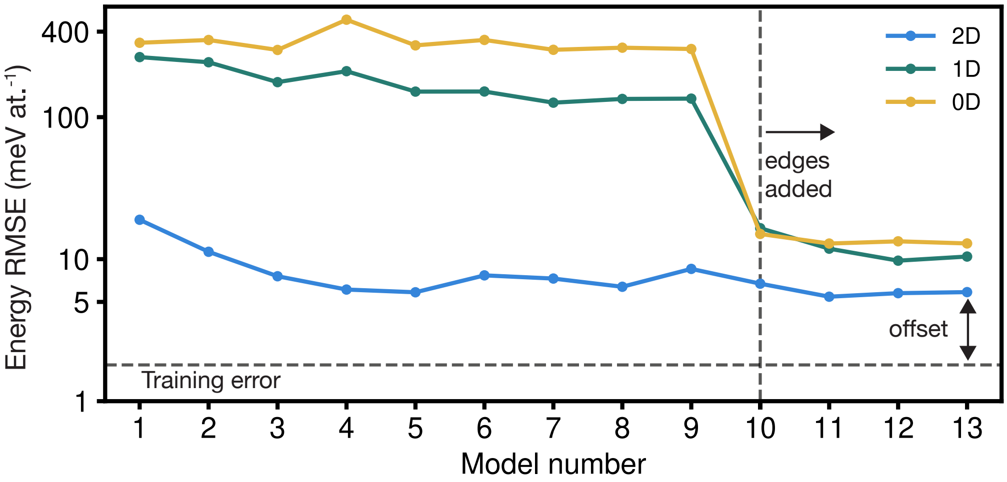

Figure S3: Energy RMSE for the external test sets discussed in the main text. The training error is shown as a dashed line at 1.8 meV at.$^{-1}$.

Energy RMSE. The energy RMSE values of the MACE models are characterised in Figures S3 to S6. Briefly, the energy predictions are offset by a notable amount, shifting the energy predictions up by approximately 6 meV at.$^{-1}$ for the external AIMD 2D test set and approximately 10 meV at.$^{-1}$ for the external AIMD 1D test set. The external 0D case (structures taken from Ref. S5) is not shown in the parity plots, as the structures have varying numbers of atoms, however the highest error is approximately 30 meV at.$^{-1}$ with the average being approximately 8 meV/atom as seen in Figure S6.

We believe that these offsets do not meaningfully impact the potentials' performance in dynamics since only the *relative* positions of minima and maxima are significant, not their absolute values. The variance of the data is perhaps another useful metric here, however this is not typically reported and thus we will refrain from doing so. However, it is important to note when reporting the energy RMSE, these potentials appear to perform worse than they actually do. We do not think that this offset is an issue specific to the MACE framework, and we will continue to investigate its origin.

At the end of the process, our final MACE model yielded an accuracy of 0.174 kJ mol^{-1} (1.8 meV at.$^{-1}$) for energies on the train set and 98.1 meV Å^{-1} for forces on the train set, the latter being in close agreement with the results for the external test set.

We also report the energy difference for the "internal" test set in Figure S6, which are defined as structures generated by the MACE model during iterative MD but not included in the training. In this internal test set, we have 800 structures and observe that there is no offset present as the

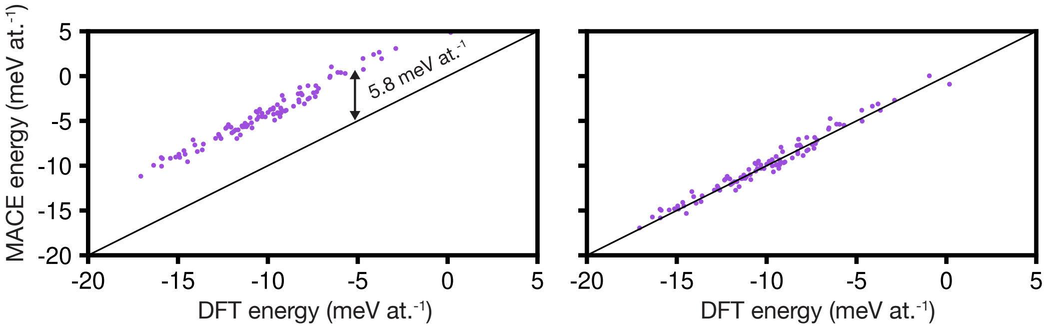

Figure S4: Energy parity plot for the 2D external test set. The left plot is the raw data, the right plot subtracts the shift to centre the plot.

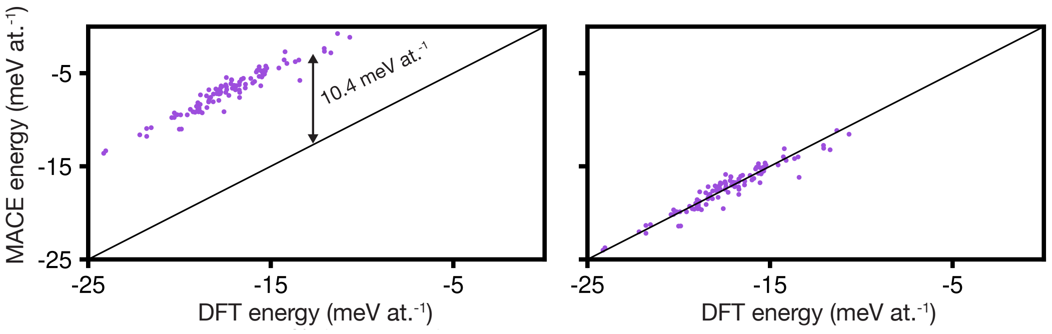

Figure S5: Energy parity plot for the 1D external test set. The left plot is the raw data, the right plot subtracts the shift to centre the plot.

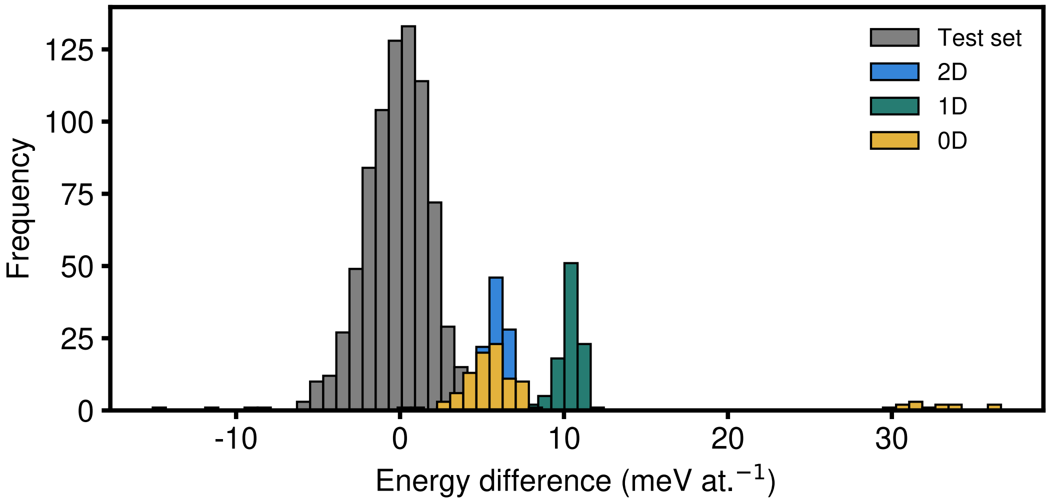

Figure S6: Histograms showing energy difference between final MACE model and DFT on the internal test set (grey); external AIMD 2D (blue) and 1D (green) test set and external 0D (yellow) test set using structures from Ref. S5.

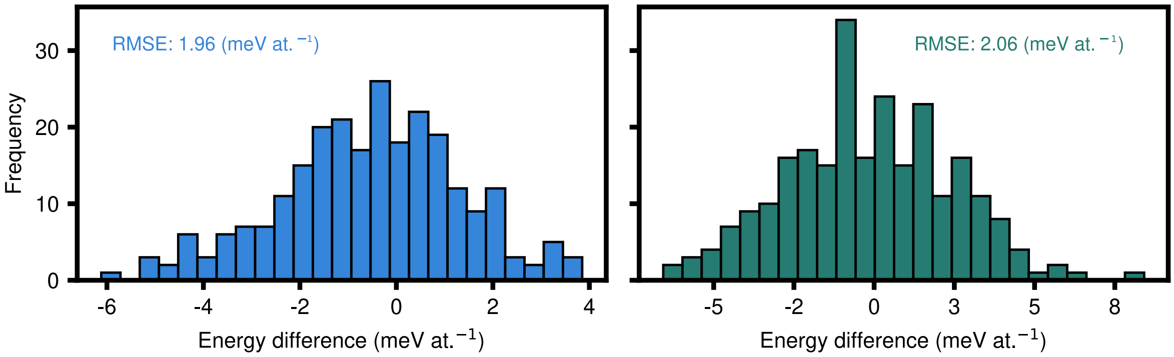

Figure S7: Distribution of energy differences with respect to DFT single-point computations, for separate test sets of 250 structures each, generated by MACE-driven MD at 1,500 K across $p_1 \in \{0.10, 0.20, 0.30, 0.40, 0.50\}$, $p_2 \in \{0.00, 0.25, 0.50, 0.75, 1.00\}$ (2D, blue), and $p_1 \in \{0.10, 0.20, 0.30, 0.40, 0.50\}$, $p_2 = 0.50$, $p_3 \in \{0.10, 0.20, 0.30, 0.40, 0.50\}$ (1D, green), respectively.

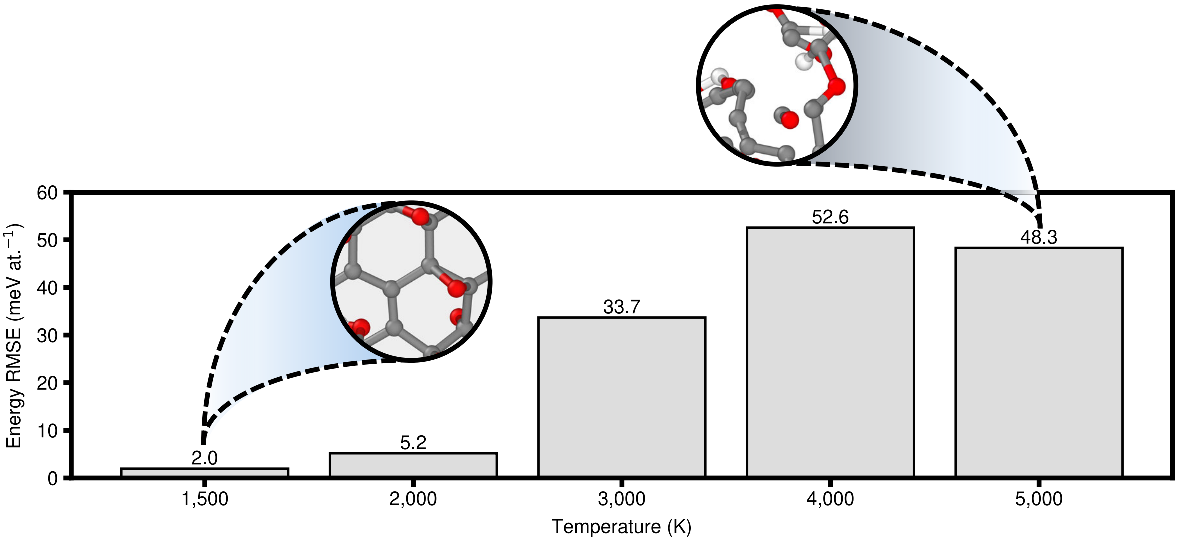

Figure S8: Energy RMSE relative to DFT single-point data for structures generated using MACE-driven MD in the high-temperature regime.

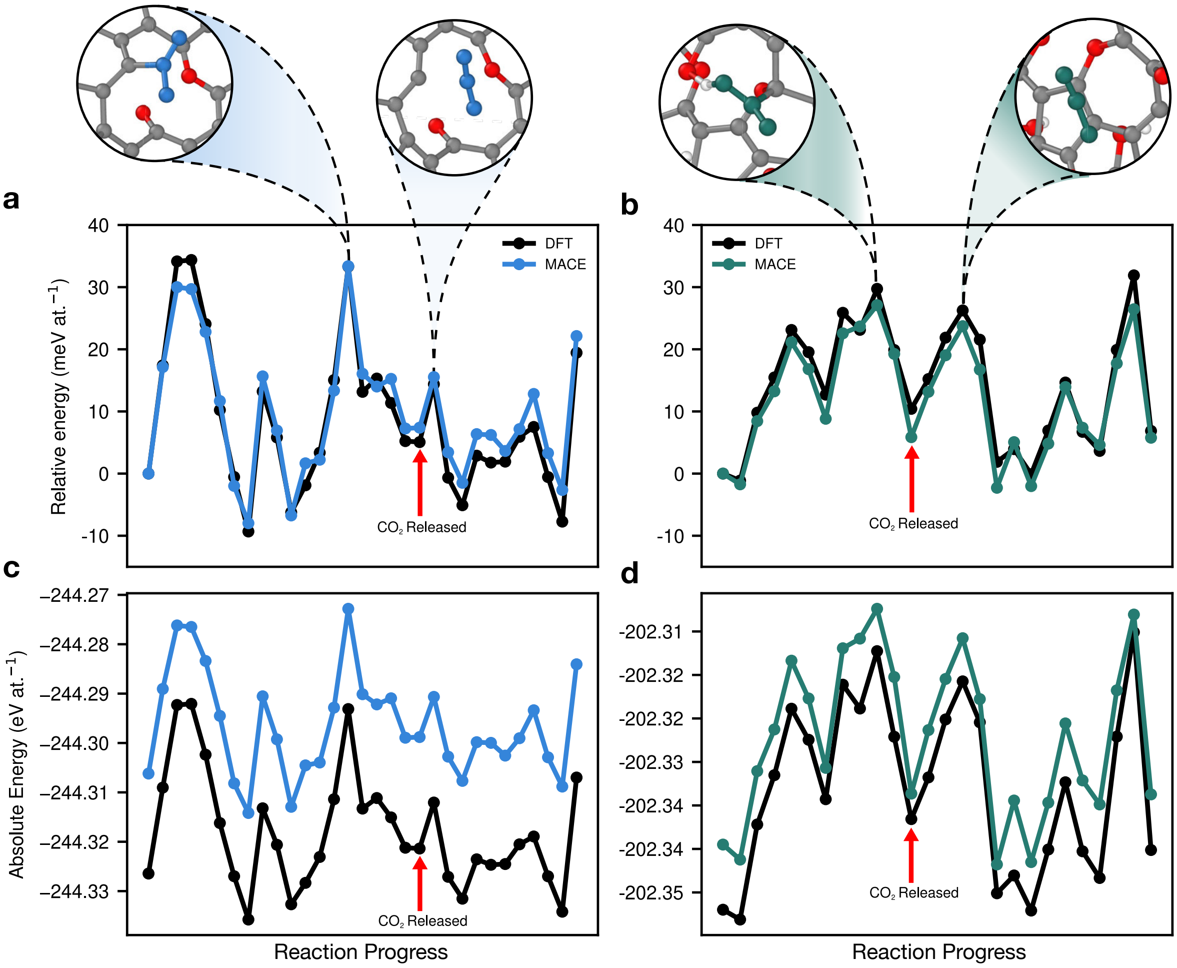

Figure S9: Energy profiles for two selected decarbonylation events from two separate 10-ps simulations. Panels (a–b) show the relative energy profiles of the decarbonylation events for a 2D (a, blue) structure and a 1D (b, green) structure compared with DFT single-point data (black), each shifted with respect to the first frame characterised here. Panels (c–d) show the absolute energy profiles.

data are centred around 0 meV at.$^{-1}$. As mentioned earlier, the 0D data show two regions, one centred at $\approx$ 8 meV at.$^{-1}$, which corresponds to H-terminated edges, and another region centred at $\approx$ 33 meV at.$^{-1}$, which corresponds to edges with no H atoms.

When MD is run using the MACE model and the results compared with DFT, we see that there is no offset present as seen in Figure S7. A set of structures spanning the p_1–p_2 (2D) and p_1–p_3 (1D) parameter space were used as an external test (data not explicitly included in training) and show the accuracy of the model on both 2D and 1D structures.

Additionally, we generate a set of high temperature configurations using the MACE model as seen in Figure S8. We ran MD for 10 ps using the same settings as the iterative training protocol and find that the model is stable in MD for 10 ps at these high temperatures. At 3,000 K and above, the sheet begins to disintegrate into chains and eventually gaseous species at higher temperatures. These configurations are very high in energy relative to the systems trained on and thus are outside of the scope of the training data. Whilst these energies are relatively high compared to 1,500 K and 2,000 K (also not trained on), we find that these errors are acceptable for these exotic configurations.

Finally, we track two decarbonylation events in a 2D and 1D structure as shown in Figure S9. The relative energy is in close agreement with DFT, as seen in Figure S9(a,b), yet the offset is apparent when using the absolute energies (c,d). This example is used to highlight that the relative energies are in good agreement with DFT. In order to thoroughly explore the reaction coordinate and barrier for a decarbonylation event, one would perform a nudged elastic band calculation which could constitute possible future work.

Polyhedral template matching. Polyhedral Template Matching (PTM) is used to identify atoms resembling simple crystalline structures at a local level [S6]. Within PTM, each input particle establishes a correlation between its local environment and the template of choice. If a correlation is identified, an RMSD (Root Mean Square Deviation) value is computed to measure spatial deviation from the ideal structural template, determining the match quality. In this work we used an RMSD cutoff value of 0.15 to determine whether a given local environment was "graphenic" or not. This method was used as implemented in OVITO 3.6.0 [S4].

Ring statistics. Ring statistics were determined using a shortest-path algorithm [S7] as implemented in `matscipy` [S8, S9].

Cell optimisation. The structures from the production runs were cooled down to 300 K over 100 ps and then underwent a full cell relaxation. This was done within ASE using the `FrechetCellFilter` class and the LBFGS optimizer. The force threshold was 10 meV at.$^{-1}$.

Methodology for X-ray photoelectron spectroscopy. The GW-corrected delta Kohn–Sham core-electron binding energy model from Ref. S10 was used to predict the C 1s core-electron binding energies (CEBE). This is a single `soap_turbo` [S11] model which is trained on neutral bulk delta Kohn–Sham CEBEs corrected by the difference of a GW core-electron binding energy and a charged delta Kohn–Sham core-electron binding energy on a configuration "carved" from an extended structure:

$$\Delta \mathrm{KS}^{0}_{\mathrm{ext}} + GW_{\mathrm{carv}} - \mathrm{KS}^{+}_{\mathrm{carv}}. \tag{S1}$$

Details of the methodology are given in Ref. S10.

Deconvolutions were performed by considering each structure at a particular timestep as an undirected graph, where edges are bonds, and nodes are atoms differentiated by species. Bonds cutoffs were defined in ASE using `natural_cutoffs(atoms, mult=1.2)`. A local subgraph was made for each carbon atom by adding edges and nodes initially from the atom's first neighbors, and if oxygen was found, adding the additional edges and nodes from the oxygen neighbors (which is necessary to discern epoxide and ether groups). Motifs were found by seeking subgraph isomorphisms between a reference dataset of motif graphs (e.g. aldehyde, alcohol, ketone etc). These were sought hierarchically, in the order: carboxylic acid, aldehyde, alcohol, carbonate, peroxide, ester, epoxide, ether, ketone, CH, CO_2, CO, sp^3, sp^2, sp.

Supplementary results

Mass-loss profiles (supplement to Figure 3b)

We carried out two further annealing simulations in parallel, starting from the same large-scale structural model as described in the main text, but now annealing at 900 and 1,200 K, respectively. The mass loss profiles for those lower-temperature simulations are shown in Figure S10. Overall, the change in mass at 1,200 K begins to plateau, converging at approximately 17%. This mass loss profile agrees qualitatively with experimental thermogravimetric data from Ref. [S12], where at the start of the thermal reduction, the mass loss follows a rapid exponential curve before reaching a linear phase for the remainder of the reduction process. In this way it is qualitatively similar to the 1,500 K simulation reported in the main text. The 900 K annealing run (left-hand side of Figure S10), however, shows much a much less pronounced mass loss which is largely due to loss of H_2O.

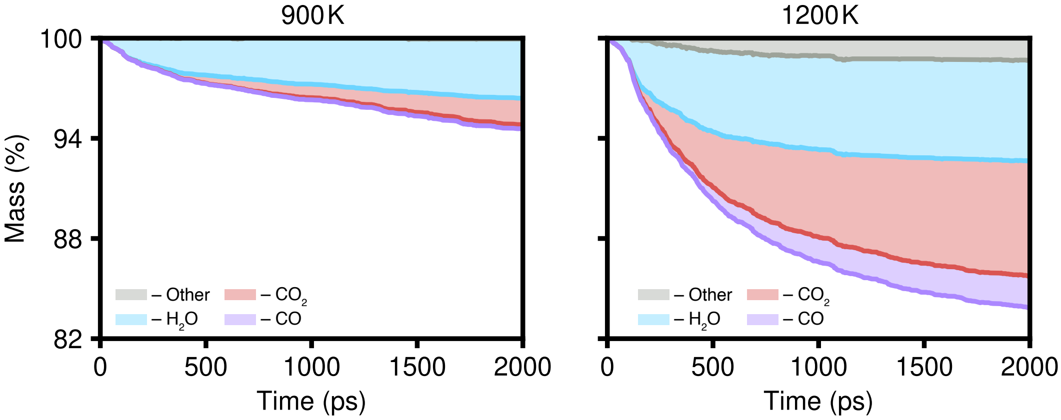

Figure S10: Mass loss profiles for the 900 K and 1,200 K simulations.

Graphene-like atoms and 6-membered rings (supplement to Figure 3c)

Figure S11 shows the evolution of the count of graphene-like motifs and 6-membered rings in the 900 and 1,200 K simulations. In our simulations, the “graphene-like” content eventually reaches ≈ 40 % for the 1,500 K MD run (Figure 3c) whereas it is still increasing for the 1,200 K run at just over 30 % (Figure S11). 900 K appears to be too low of a temperature to effectively reduce GO in nanosecond simulations: the degree of “graphene-like” similarity remains essentially steady, with only a very slight increase over 2 ns. The fraction of “graphene-like” atoms (dark grey) in combination with the percentage of 6-membered rings (blue) provides in-

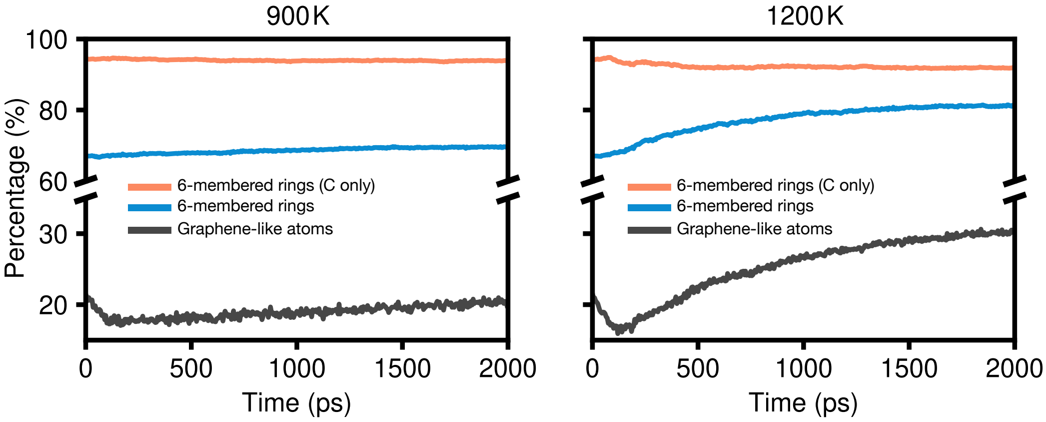

Figure S11: Graphene-like atoms and 6-membered rings profiles for the 900 K and 1,200 K simulations.

sight into the degree of crystallinity attained during the annealing process. We also show the percentage of 6-membered rings with only carbon atoms (orange) which is naturally quite high, but does not indicate how graphitic the structure is.

Annealing temperatures of 900 K appear to have no significant effect on the percentage of 6-membered rings which stays approximately constant. In addition to this, the count of "graphene-like" atoms decreases slightly during the initial heating phase. This is then followed by a very gentle linear increase which is still less than the initial percentage.

At 1,200 K, we observe more dynamic behaviour, where the percentage of 6-membered rings increases whereas that of 6-membered rings containing only carbon decreases, suggesting that more oxygen atoms are incorporated into the graphene network. This is supported by the X-ray photoelectron spectroscopy (XPS) data seen in Figure S14. In addition to this, the percentage of "graphene-like" atoms decreases at a faster rate when compared to the 900 K data over the heating period, before increasing at a faster rate before beginning to plateau after 2 ns. Overall, we see the percentage of "graphene-like" atoms increase with that of the 6-membered rings, as expected.

Functional groups (supplement to Figure 3d)

Tracking the functional groups during the annealing process provides insight into the mechanisms by which the structure becomes more graphene-like. For the 900 K simulation, there is a gradual decrease of sp^3 groups (dark and light blue) and a corresponding gradual increase of sp^2 groups (orange, red, purple). Overall, this profile further emphasises that 900 K is too low a temperature to effectively anneal GO structural models on the ns timescale using our

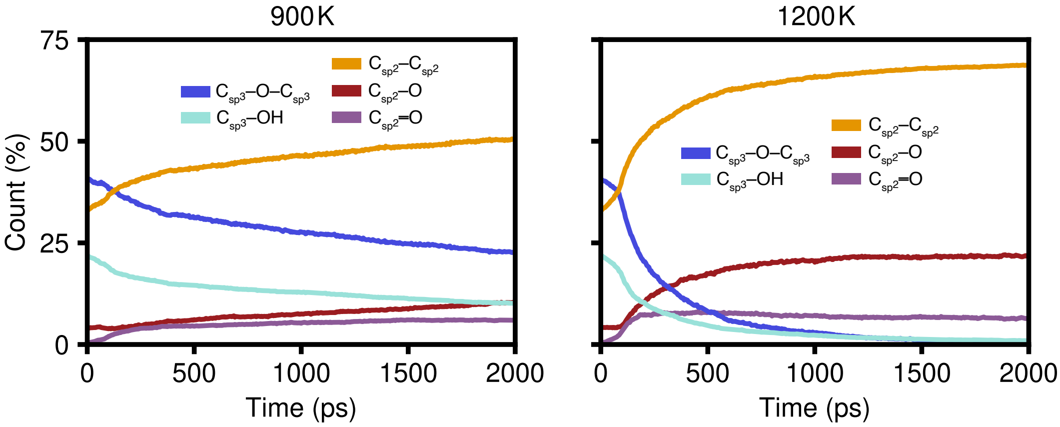

Figure S12: Functional group profiles for the 900 K and 1,200 K simulations.

model.

In contrast, the 1,200 K simulation follows a similar profile to the 1,500 K simulation, albeit at a slower rate. Here it is clear that sp^3 groups are nearly completely removed after 2 ns with a similar growth profile of sp^2 groups, including a characteristic steady-state period for C_{sp^2}–O groups.

We note that using the average ionic radius for the C–O bond ($C_{ionic}/2 + O_{ionic}/2$) gives oscillatory behaviour for the C_{sp^2}-C_{sp^2} count, as indicated by the black line in the top panel of Figure S13. As the defined bond cutoff is increased beyond the ionic radius (from 1.685 to 1.750 to 1.850 Å), we see the magnitude of oscillations decrease. We observe changes in the oscillatory behaviour during the course of the simulation, which may be due to gaseous species being removed and simulations restarted at some points. Much less pronounced oscillations are seen for the 1,500 K simulation, which is discussed in the main text (Figures S13 and 3d). For other species, all structural analysis was performed using the ionic-radius-based cutoff values.

X-ray photoelectron spectroscopy (supplement to Figure 4)

After the structures had been annealed, they were cooled down to 300 K over 100 ps and then underwent a full cell optimisation to relax the structure. XPS calculations were performed on these fully relaxed structures along with selected structures along the trajectory.

We summarize additional results in Figure S14. The top panel shows the shifted simulated spectra, where the shift has been introduced to better fit the experimental data (red, dashed). This is done since deconvolutions of experimental data are fixed at reference values (sp^2 =

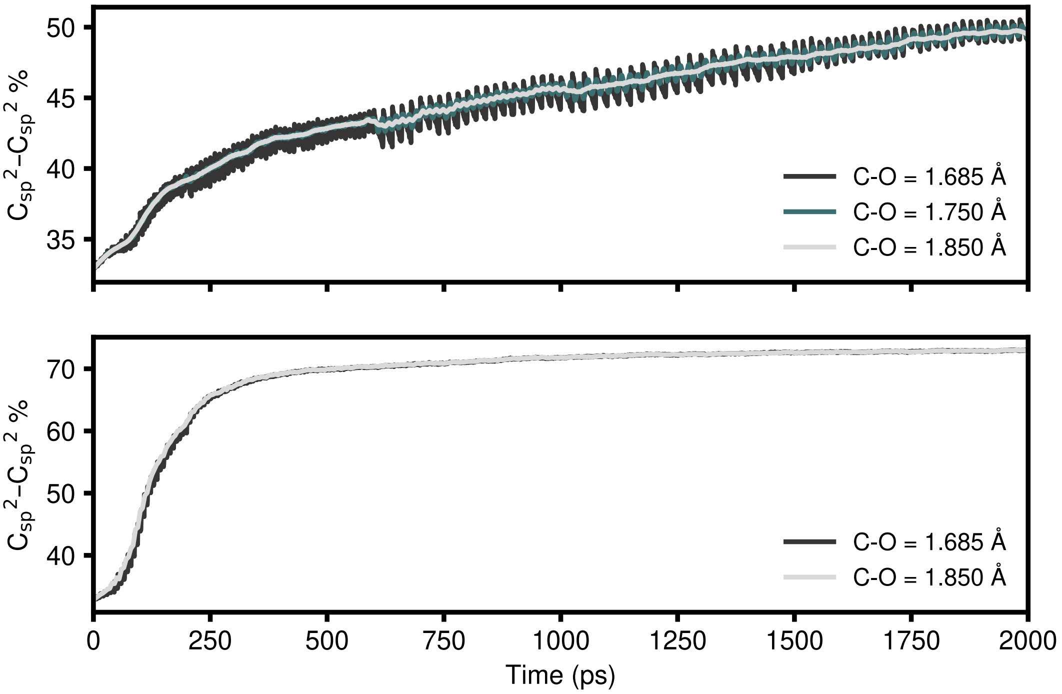

Figure S13: *Top:* Oscillations for the C_{sp^2}-C_{sp^2} count at 900 K, as evaluated using different C–O bond lengths for the cutoff definition. *Bottom:* Oscillations at 1,500 K. The data with a cutoff value of 1.850 Å are shown in Figure 3d of the main text.

284.4 eV, black dashed lines) to perform this task. In contrast, the simulated spectra are not constrained by this, and so we report both types of data in Figure S14; the data in the top panel are shifted towards the reference peak and the bottom panel shows the un-shifted simulation data. The experimental data taken from Ref. S12 have had their background removed and have been normalised in order to be compared to simulated spectra.

Both approaches provide useful information—the former giving comparison to experimental data, the latter providing information about the shifting of the peak towards the reference sp^2 value as oxygen is removed from the system after annealing at different temperatures. We have chosen to display only the six most important groups as other functional groups did not contribute substantially.

At 900 K, we see a spectrum which appears very different to the experimental data from Ref. S12 (red dashed line). This is due to the increased presence of oxygen groups in the structure, in particular, epoxide and alcohol groups which are yet to be removed. There are further contributions from ether, ester, and ketone groups. There is also a shift to higher CEBEs due to the increased electronegativity from the oxygen groups present.

At 1,200 K, the oxygen peak has now been reduced to a shoulder peak with the epoxide peak being nearly removed entirely and the alcohol peak being greatly lowered. In contrast the ether group, ketone group, and ester group peaks have all increased greatly. This agrees with the wider trend of transitioning from sp^3 to sp^2. The shift down to the sp^2 reference peak is also observed in the bottom panel due to fewer oxygen groups being present and a greater presence of sp^2 environments.

Finally, at 1,500 K, we observe the best agreement with the experimental XPS data of Ref. S12. There are no significant contributions from epoxide environments and, in total, fewer oxygen groups present. This can be clearly seen by the reduction in the shoulder peak, giving closer agreement experiment than for the two other simulations at lower temperatures. Additionally, there is a closer shift to the reference sp^2 value, further highlighting the removal of oxygen. The results in Figure S14 therefore provide further justification that 1,500 K is an appropriate choice of annealing temperature for the simulations reported in the main text.

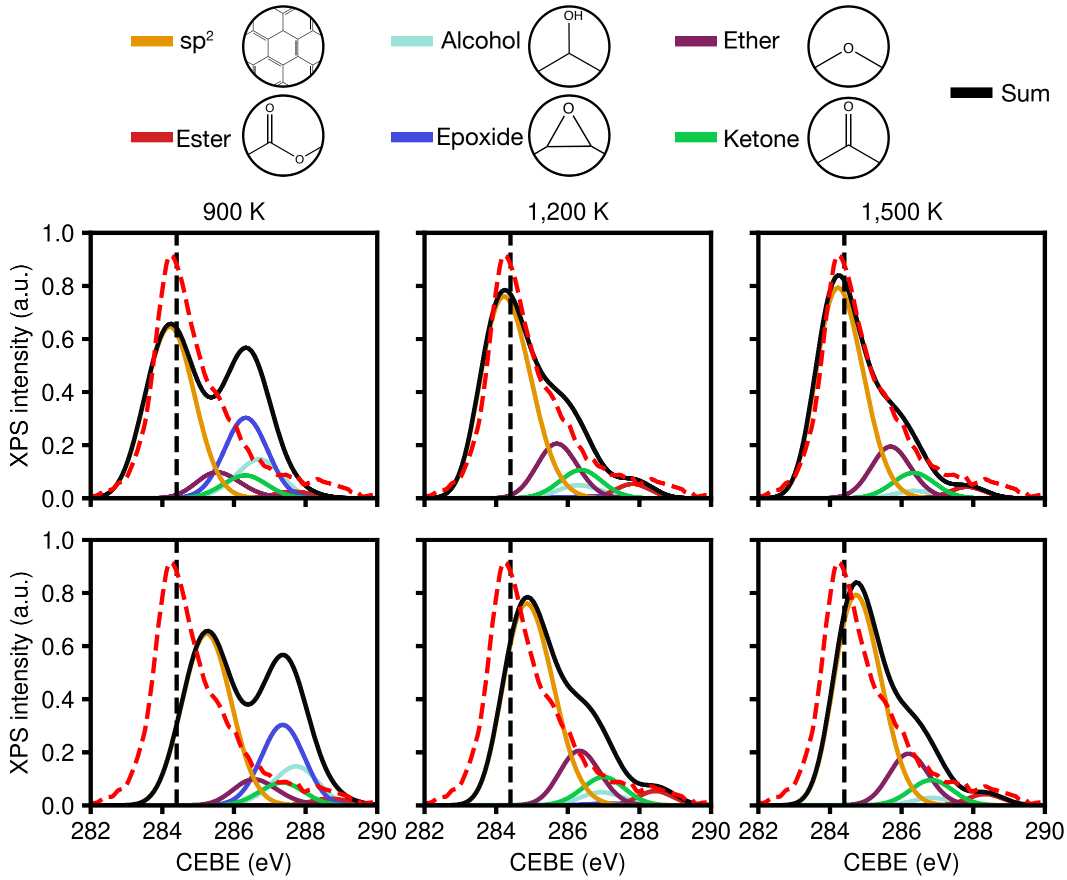

Figure S14: ML-model-predicted XPS spectra for the 900 K, 1,200 K, and 1,500 K simulations. **Top:** Shifted XPS spectra to experimental sp^2 reference value at 284.4 eV. The red dashed line shows experimental data taken from Ref. S12. **Bottom:** Unshifted XPS data showing the effect of electronegativity on CEBE values.

Supplementary references

[S1] T. K. Stenczel, Z. El-Machachi, G. Liepuoniute, J. D. Morrow, A. P. Bartók, M. I. J. Probert, G. Csányi, V. L. Deringer, *J. Chem. Phys.* **2023**, *159*, 044803.

[S2] V. L. Deringer, G. Csányi, *Phys. Rev. B* **2017**, *95*, 094203.

[S3] I. Batatia, D. P. Kovacs, G. Simm, C. Ortner, G. Csanyi, MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields, in S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Editors), *Advances in Neural Information Processing Systems*, volume 35, Curran Associates, Inc. **2022** pages 11423–11436.

[S4] A. Stukowski, *Model. Simul. Mater. Sci. Eng.* **2009**, *18*, 015012.

[S5] B. Motevalli, A. J. Parker, B. Sun, A. S. Barnard, *Nano Futures* **2019**, *3*, 045001.

[S6] P. M. Larsen, S. Schmidt, J. Schiøtz, *Model. Simul. Mater. Sci. Eng.* **2016**, *24*, 055007.

[S7] D. S. Franzblau, *Phys. Rev. B* **1991**, *44*, 4925.

[S8] J. R. Kermode, L. Pastewka, Matscipy (version 0.7.0), `https://github.com/libAtoms/matscipy` **2022**.

[S9] P. Grigorev, L. Frérot, F. Birks, A. Gola, J. Golebiowski, J. Grießer, J. L. Hörmann, A. Klemenz, G. Moras, W. G. Nöhring, J. A. Oldenstaedt, P. Patel, T. Reichenbach, T. Rocke, L. Shenoy, M. Walter, S. Wengert, L. Zhang, J. R. Kermode, L. Pastewka, *J. Open Source Softw.* **2024**, *9*, 5668.

[S10] D. Golze, M. Hirvensalo, P. Hernández-León, A. Aarva, J. Etula, T. Susi, P. Rinke, T. Laurila, M. A. Caro, *Chem. Mater.* **2022**, *34*, 6240.

[S11] M. A. Caro, *Phys. Rev. B* **2019**, *100*, 024112.

[S12] C. Valentini, V. Montes-García, P. A. Livio, T. Chudziak, J. Raya, A. Ciesielski, P. Samorì, *Nanoscale* **2023**, *15*, 5743.

Chapter 5

Mechanical properties of graphene oxide from machine-learning-driven simulations

This chapter is reproduced in its entirety from the following publication:

> Zakariya El-Machachi, Bowen Cheng and Volker L. Deringer, "Mechanical properties of graphene oxide from machine-learning-driven simulations", *Chemical Communications*, 2025, **61**, 11405.

Authors' contributions: Z. E.-M. performed and analysed production simulations. B. C. contributed to pilot studies. V. L. D. initiated and coordinated the study. All authors made substantial contributions to data analysis and discussions. Z. E.-M. and V. L. D. wrote the paper with input from B. C., and all authors approved the final version.

ChemComm

ROYAL SOCIETY OF CHEMISTRY

COMMUNICATION

View Article Online
View Journal

Cite this: DOI: 10.1039/d5cc02753e

Received 15th May 2025,
Accepted 16th June 2025

DOI: 10.1039/d5cc02753e

rsc.li/chemcomm

Mechanical properties of graphene oxide from machine-learning-driven simulations

Zakariya El-Machachi, Bowen Cheng and Volker L. Deringer *

Graphene oxide (GO) materials have complex chemical structures that are linked to their macroscopic properties. Here we show that first-principles simulations with a machine-learned interatomic potential can predict the mechanical properties of GO sheets in agreement with experiment and provide atomistic insights into the mechanisms of strain and fracture. Our work marks a step towards understanding and controlling mechanical properties of carbon-based materials with the help of atomistic machine learning.

The term ‘graphene oxide’ (GO) encompasses a class of carbonaceous 2D materials with applications in many fields.[1] GO materials show various degrees of chemical functionalisation, introduced by oxidative and reductive processes, and it is important to understand how these structural aspects are connected to macroscopic properties. A key point here is the response of GO sheets to mechanical strain, and indeed it has been argued that the mechanical properties depend strongly on what functional groups (epoxide, hydroxyl, *etc.*) are present.[2] GO and reduced GO (‘rGO’) have been used as sensors in biological applications,[3] where small mechanically-induced structural changes from electromechanically linked cells alter the electronic properties.[4]

In 2007, there began renewed interest in multilayer GO and rGO, primarily focusing on mechanical, electrical, and catalytic properties.[5] Atomic force microscopy indentation was used to study the mechanical properties of monolayer rGO.[6] Independent, high-resolution TEM images of monolayer GO and rGO revealed the atomic structure,[7,8] confirming a two-type domain structure of GO, where oxidised regions form a continuous network with smaller graphitic domains interspersed.[9] Since then, various authors have investigated mechanical properties of (r)GO with varying degrees of functionalisation.[10–13]

In tandem, computational studies have been reported for GO including its mechanical properties,[2,12–18] but have typically been limited by the length and time scales accessible to first-principles computations, or by the accuracy limits of empirical potentials.[19] Machine learning (ML)-based interatomic potentials approximate the predictions of quantum-mechanical methods whilst adopting key aspects of empirical approaches, such as locality, to achieve first-principles accuracy with *near*-empirical speed.[20] Such ML-driven simulations have been used to model defective,[21] nanoporous,[19] and amorphous graphene.[22]

Here, we study the links of atomic structure and mechanical properties in GO and rGO with ML-driven simulations, using the recently introduced GO-MACE-23 model.[23] This model was trained using a domain-specific protocol that gradually explores relevant configurations, first through CASTEP + ML,[24] then using the MACE architecture.[25] We simulate mechanically straining a nano-scale rGO sheet and smaller-scale (r)GO structures across a range of parameters, *viz.* oxygen content and functionalisation. The predictions from our GO-MACE-23 model agree with experiments and with other computational simulations for GO. To our knowledge, limited work has been conducted on the mechanical properties of reduced GO, and so computations could provide new insight and lead to design rules based on the degree of functionalisation for next-generation (r)GO-based materials.

Our starting structure is taken from ref. 23 and is a partially disordered, fully sp^2-bonded graphene sheet with 10 368 atoms (17.7×15.3 nm^2) generated using Monte–Carlo bond switching, following ref. 22. The structure was then functionalised with an O/C ratio of 40% and OH/O ratio of 0.5 decorated randomly, raising the atom count to 16 645, then thermally reduced *via* MD over 1.9 ns at 1500 K, reducing the atom count to 11 984 and yielding a simulated XPS spectrum in agreement with experimental spectra.[23] To compute mechanical properties, we apply uniaxial tensile strain in the basal plane (Fig. 1a), then a fixed-volume geometry optimisation is then conducted, the uniaxial stress component recorded at that point, and the strain increased again in steps of 0.05%; the process repeats up to a strain of 0.50 (50%). A correction to the recorded stress value, σ_{cell}, is required to remove the dependence on the out-of-plane cell parameter (see, *e.g.*, ref. 16): we report $\sigma = (c/t) \times \sigma_{\mathrm{cell}}$,

---

Inorganic Chemistry Laboratory, Department of Chemistry, University of Oxford, Oxford OX1 3QR, UK. E-mail: volker.deringer@chem.ox.ac.uk

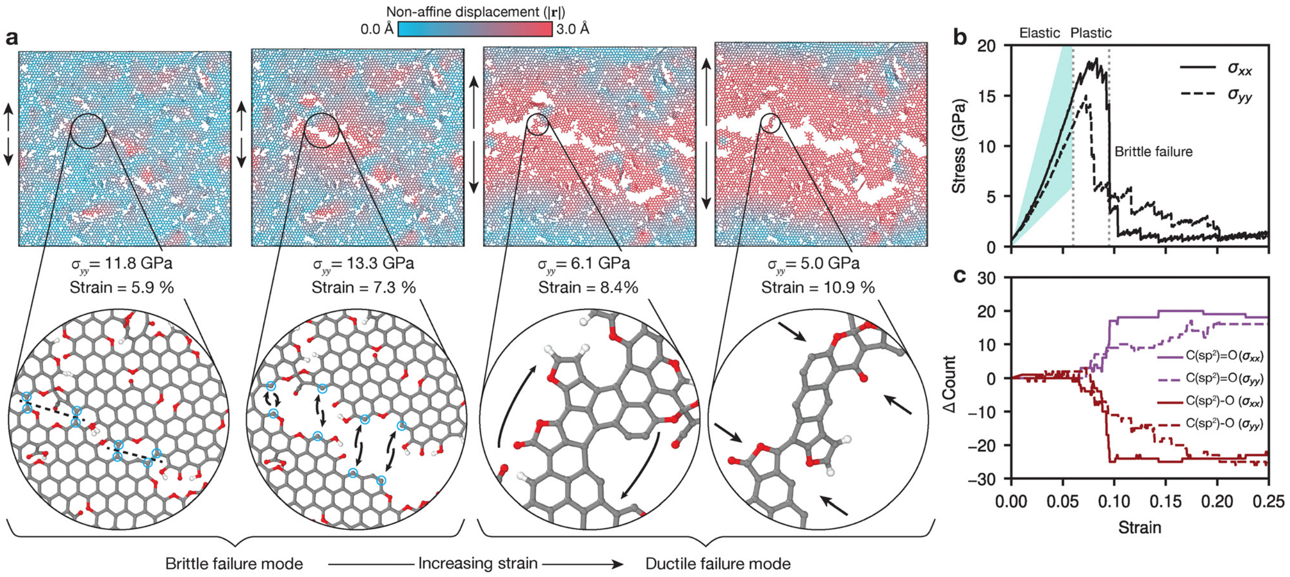

Fig. 1 Mechanical properties of a reduced GO sheet from machine-learning-driven simulations. (a) Visualisation of the rGO sheet taken from ref. 23 now placed under progressive strain (arrows). Atoms and bonds are colour-coded by the non-affine displacement magnitude relative to the initial configuration using OVITO.[26] Insets are shown below (C, grey; H, white; O, red). (b) Stress–strain curves for the rGO sheet axially strained in x or y direction. The fitted Young's moduli are E_{xx} = 261 GPa, E_{yy} = 205 GPa; we include a value for rGO from ref. 6 (shaded area). (c) Change in aldehyde and ketone (purple) *vs.* ether count (red) relative to the unstrained structure. Solid (dashed) lines indicate strain in x (y) direction, respectively.

where c is the cell parameter perpendicular to the basal plane and t is the thickness of the sheet (determined as the full width at half maximum for the z coordinates of all atoms in the sheet). We use 3.5 Å as a minimum thickness, since our small-scale monolayer structural models do not reach the experimentally-derived 7 Å for (r)GO,[6,10] even at high O content: their effective thickness in our computations is therefore closer to graphene, as in ref. 14. (We note that the concept of a monolayer thickness must rely on an assumption,[13] and that previous attempts of measurement for graphene have yielded various results.[27])

We colour-code the rGO nano-structure in Fig. 1a by the non-affine displacement magnitude, which decouples displacement due to external strain and internal, non-affine deformations. This is done by mapping atomic positions onto the reference (initial) frame and computing displacement vectors using OVITO.[26] In Fig. 1a, initially there is a uniform blue colour indicating large areas experiencing elastic deformation, with fewer pockets of red regions, which indicate local deformations due to non-affine displacements. These red regions involve plastic deformation in the sheet, arising due to rearrangements, defects, and bond cleavage. As the external strain is increased, there is an increase in local strain around holes and defects, such as O atoms in the carbon backbone. This can be seen in the first and second snapshot in Fig. 1a, where the sheet is still predominantly under elastic deformation—with the exception of the highlighted region which undergoes brittle fracture (dashed lines in the first inset), cleaving a large hole. In the second inset, arrows shows the motion of the 8 atoms (circled) previously involved in 4 C–C bonds which have been broken in a brittle manner. As the strain is increased further, we observe an increasing area of 'red' regions, indicating a change from elastic to plastic behaviour. In the third panel, we begin to see a few nanoribbons which join parts of the mostly fractured sheet. With increasing strain, the crack grows and the nanoribbon thins to just one ring thick: such structures, including their stress–strain curves, were recently studied with ML potentials.[19] Here, there is a large degree of plastic deformation on the sheet (red regions in Fig. 1a).

The stress–strain curves, recorded separately for both directions, are shown in Fig. 1b. We note that in the ideal large-size limit, both directions would be equivalent. However, as the sheet has been thermally reduced for 2 ns,[23] the introduction of holes and defects in the sheet appears to have introduced a (slight) stochastic anisotropy, which is revealed under uniaxial straining. The curves in Fig. 1b qualitatively have the same profile as other stress–strain curves for GO: (i) an initial elastic regime showing a linear response; (ii) a brief plastic region where the curve plateaus; and (iii) a sharp, brittle failure followed by plastic deformation.[6,10,13] There is a notable difference between the two curves: one (σ_{xx}) indicates higher ultimate tensile strength and a single catastrophic brittle-failure event, where the sheet fractures completely; for the other (σ_{yy}), we find that the sheet has lower ultimate tensile strength in this direction and undergoes staggered brittle-fracture events leading to a wider plastic deformation phase. Our estimation of the Young's modulus for the rGO sheet agrees with ref. 6 (Fig. 1b).

Given that our GO-MACE-23-driven simulations agree well with experiments, we are now able to extract and analyse structural information from the simulation trajectory. One of the key structural features in (r)GO is the oxygen functionalisation. In our previous work, we highlighted the sp^3-to-sp^2

transition at high temperatures resulting in the epoxide and hydroxyl groups transforming into C_{sp^2}–O groups, such as ether.[23] These groups become part of the carbon backbone in the material, and thus, we show here that these groups play a key role in the response to mechanical strain. In Fig. 1c, we show the evolution of the C_{sp^2}–O and C_{sp^2}=O count, indicating a change from O being present within the carbon backbone to cleaving off to form a carbonyl group such as a ketone. This rapid decrease of C_{sp^2}–O groups coinciding with the brittle failure of the material (under straining in x; solid lines in Fig. 1c) suggests that these groups contribute substantially to the failure mode of (r)GO. In the y direction (dashed lines in Fig. 1c), the stress–strain profile differs slightly and shows an extended period of plastic deformation. Here, we see a more gradual decrease in C_{sp^2}–O groups as the latter slowly transform into carbonyl groups.

We then tested whether the properties of GO can be controlled by varying functional groups. Simulations were run for an ensemble of graphene decorated randomly with oxygen groups above and below the basal plane. We varied the O/C ratio from 10% to 50%, and the initial OH/O ratio (*cf.* ref. 23), in line with experimental reports.[1] We observe that generally, with increasing O/C ratio, there is a lower ultimate tensile strength in the material, and the gradient in the elastic region decreases with increasing O content, corresponding to a lower Young's modulus. Another important feature is the type of functional group. We observe that sheets with 100% epoxide and 0% hydroxyl (Fig. 2a) display plastic behaviour, being able to withstand higher strain at higher O content. At equal ratio (Fig. 2b), we find a reduced plastic response, with failure at lower strain. Finally, for fully hydroxyl-decorated sheets (Fig. 2c), the behaviour shifts towards an elastic response and failure at even lower strain. This difference is due to the epoxide groups flattening under strain and incorporating into the sheet, whereas hydroxyl groups cannot be incorporated into the sheet in the same manner, resulting in brittle failure due to the cleavage of C–C bonds.

We note briefly that this dependence of mechanical properties on functionalisation type (Fig. 2) was not observed in earlier simulations using an empirically fitted ReaxFF potential, originally developed for hydrocarbon combustion:[2a] in the latter case, the weakening with higher O content was independent of the functional group type (hydroxyl or epoxide), and the material was predicted to exclusively show plastic behaviour.[2a] We speculate that the (more accurate, but more costly) ML potential can uncover additional nuance here which is difficult or impossible to describe with empirically fitted interatomic potentials.

Beyond the overall oxygen content (which can be tuned by gradual thermal reduction), the properties of the sheets could be controlled more finely if one found ways to modify the chemical groups present. We explore this idea in Fig. 3a: we plot averaged Young's moduli for independent structures which start from a pristine graphene backbone and are subsequently decorated with different O/C and OH/O ratios, with 10 strained in x (filled) and 10 strained in y (Fig. 3a). Fig. 3b shows the same, but for structures which have been thermally reduced for 10 ps at 1500 K using GO-MACE-23, resulting in gas evolution. There is a clear relationship between increasing hydroxyl-to-epoxide ratio and Young's modulus, with a more profound effect at higher O content. Literature values reported for GO are 290–430 GPa with O/C ratio from 0.1–0.7;[10,12,13,16] for rGO, 250 ± 150 GPa has been reported in ref. 6.

Once our GO structures have been heated for a relatively short time, we observe a significant change in the Young's modulus. These heated structures, which we call rGO, have a lower Young's modulus at higher O content compared to that of GO. This is because oxygen leaves the sheet upon heating, mainly as H_2O and CO_2.[23] It is the latter which causes a

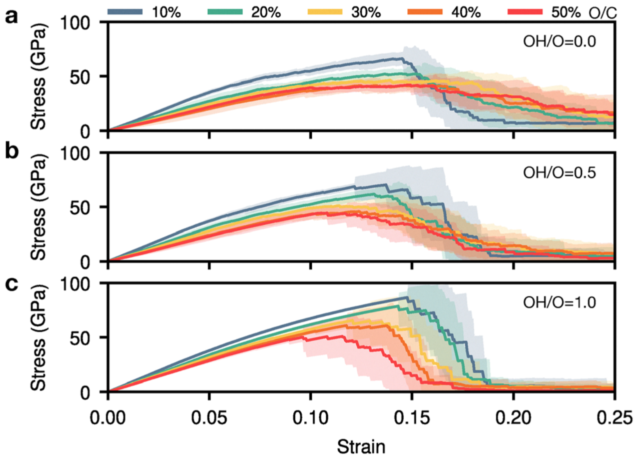

Fig. 2 Mechanical behaviour of GO and its dependence on the oxygen content. We show stress–strain curves for 20 GO sheets of ∼200 atoms each, at varying O/C ratio (indicated by colour-coding). The shaded areas indicate one standard deviation. We characterise sheets that have been functionalised: (a) fully with epoxide groups; (b) with half epoxide, half hydroxyl groups; and (c) fully with hydroxyl groups.

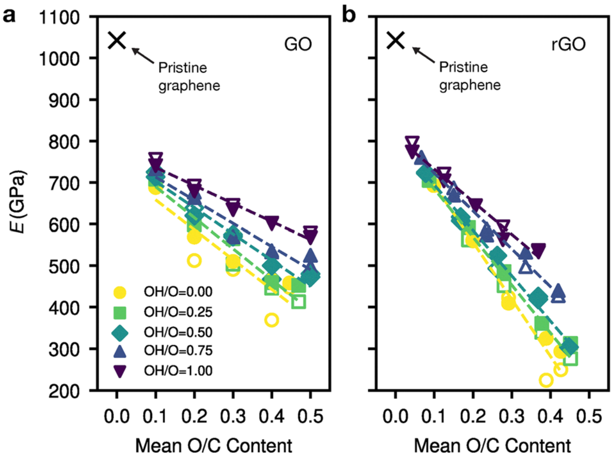

Fig. 3 Controlling the Young's modulus, E, of GO by tuning the hydroxyl content. Results are shown for varying initial OH/O ratios. (a) E values for GO: averaged for 10 structures of ∼200 atoms each, strained uniaxially in x (filled markers), and 10 strained uniaxially in y (unfilled), respectively. The cross represents pristine graphene. (b) As for panel (a), except that the structures have been thermally reduced for 10 ps.

rearrangement of the graphene backbone, introducing defects into the sheet in the form of point defects or topological defects. There is a drastic shift in the Young's modulus for structures with more epoxide groups once heated. These defects reduce the homogeneity of the sheet, resulting in 'weak points' in the sheet, thus lowering the Young's modulus. At the other end, at low initial O/C ratio (0.1) there is minimal disruption to the sheet as gas evolves, and we observe an increase in the Young's modulus compared with GO, due to lower O content as the system is heated.

Looking forward, our results support the idea that the mechanical properties of GO might be amenable to computationally-guided 'design'[2] (see ref. 28 for a general perspective). This goal could be achieved if the chemical structure can be fine-tuned, in this case between hydroxyl (more elastic behaviour) and epoxide groups (more plastic behaviour), or through nitrogen doping which increases the Young's modulus.[29] Applications such as wearable electronics[30] may benefit from a flexible GO structure with more epoxide groups. Conversely, high OH content may be useful for reinforcing composites: *e.g.*, adding 1 vol% of GO to an Al–Mg alloy enhanced the bending strength substantially.[31]

We thank F. C. Mocanu for helpful comments and technical advice. This work was supported by the Engineering and Physical Sciences Research Council [grant numbers EP/L015722/1, EP/V049178/1, EP/W524311/1] and UK Research and Innovation [grant number EP/X016188/1]. We acknowledge the use of the University of Oxford Advanced Research Computing facility (**https://dx.doi.org/10.5281/zenodo.22558**).

Conflicts of interest

There are no conflicts to declare.

Data availability

Data and code supporting this work are openly available at **https://github.com/zakmachachi/GO-mechanical-properties**.

References

1 (*a*) D. R. Dreyer, S. Park, C. W. Bielawski and R. S. Ruoff, *Chem. Soc. Rev.*, 2010, **39**, 228; (*b*) *Graphene Oxide: Fundamentals and Applications*, ed. A. M. Dimiev and S. Eigler, John Wiley & Sons, Chichester, 2017; (*c*) S. Guo, S. Garaj, A. Bianco and C. Ménard-Moyon, *Nat. Rev. Phys.*, 2022, **4**, 247; (*d*) J. Wu, H. Lin, D. J. Moss, K. P. Loh and B. Jia, *Nat. Rev. Chem.*, 2023, **7**, 162.

2 (*a*) Y. Sun, X. Tang, H. Bao, Z. Yang and F. Ma, *RSC Adv.*, 2020, **10**, 29610; (*b*) M. Tavakol, A. Montazeri, S. H. Aboutalebi and R. Asgari, *Appl. Surf. Sci.*, 2020, **525**, 146554.

3 T. Laurila, S. Sainio and M. A. Caro, *Prog. Mater. Sci.*, 2017, **88**, 499.

4 R. Kempaiah, A. Chung and V. Maheshwari, *ACS Nano*, 2011, **5**, 6025.

5 (*a*) D. A. Dikin, S. Stankovich, E. J. Zimney, R. D. Piner, G. H. B. Dommett, G. Evmenenko, S. T. Nguyen and R. S. Ruoff, *Nature*, 2007, **448**, 457; (*b*) S. Stankovich, D. A. Dikin, R. D. Piner, K. A. Kohlhaas, A. Kleinhammes, Y. Jia, Y. Wu, S. T. Nguyen and R. S. Ruoff, *Carbon*, 2007, **45**, 1558.

6 C. Gómez-Navarro, M. Burghard and K. Kern, *Nano Lett.*, 2008, **8**, 2045.

7 C. Gómez-Navarro, J. C. Meyer, R. S. Sundaram, A. Chuvilin, S. Kurasch, M. Burghard, K. Kern and U. Kaiser, *Nano Lett.*, 2010, **10**, 1144.

8 K. Erickson, R. Erni, Z. Lee, N. Alem, W. Gannett and A. Zettl, *Adv. Mater.*, 2010, **22**, 4467.

9 A. Lerf, H. He, M. Forster and J. Klinowski, *J. Phys. Chem. B*, 1998, **102**, 4477.

10 J. W. Suk, R. D. Piner, J. An and R. S. Ruoff, *ACS Nano*, 2010, **4**, 6557.

11 T. Sakuma, R. Sato, A. Yamaguchi, H. Imai, N. Arai and Y. Oaki, *J. Am. Chem. Soc.*, 2025, **147**, 11564.

12 X. Wei, L. Mao, R. A. Soler-Crespo, J. T. Paci, J. Huang, S. T. Nguyen and H. D. Espinosa, *Nat. Commun.*, 2015, **6**, 8029.

13 C. Cao, M. Daly, C. V. Singh, Y. Sun and T. Filleter, *Carbon*, 2015, **81**, 497.

14 J. T. Paci, T. Belytschko and G. C. Schatz, *J. Phys. Chem. C*, 2007, **111**, 18099.

15 A. Bagri, C. Mattevi, M. Acik, Y. J. Chabal, M. Chhowalla and V. B. Shenoy, *Nat. Chem.*, 2010, **2**, 581.

16 L. Liu, J. Zhang, J. Zhao and F. Liu, *Nanoscale*, 2012, **4**, 5910.

17 B. D. Jensen, K. E. Wise and G. M. Odegard, *J. Phys. Chem. A*, 2015, **119**, 9710.

18 C. D. Williams, P. Carbone and F. R. Siperstein, *Nanoscale*, 2018, **10**, 1946.

19 A. Kabylda, B. Mortazavi, X. Zhuang and A. Tkatchenko, *Adv. Funct. Mater.*, 2024, **35**, 2417891.

20 (*a*) J. Behler, *Angew. Chem., Int. Ed.*, 2017, **56**, 12828; (*b*) V. L. Deringer, M. A. Caro and G. Csányi, *Adv. Mater.*, 2019, **31**, 1902765; (*c*) P. Friederich, F. Häse, J. Proppe and A. Aspuru-Guzik, *Nat. Mater.*, 2021, **20**, 750.

21 F. L. Thiemann, P. Rowe, A. Zen, E. A. Müller and A. Michaelides, *Nano Lett.*, 2021, **21**, 8143.

22 Z. El-Machachi, M. Wilson and V. L. Deringer, *Chem. Sci.*, 2022, **13**, 13720.

23 Z. El-Machachi, D. Frantzov, A. Nijamudheen, T. Zarrouk, M. A. Caro and V. L. Deringer, *Angew. Chem., Int. Ed.*, 2024, **63**, e202410088.

24 T. K. Stenczel, Z. El-Machachi, G. Liepuoniute, J. D. Morrow, A. P. Bartók, M. I. J. Probert, G. Csányi and V. L. Deringer, *J. Chem. Phys.*, 2023, **159**, 044803.

25 (*a*) I. Batatia, D. P. Kovács, G. N. C. Simm, C. Ortner and G. Csányi, *Adv. Neural Inf. Process Syst.*, 2022, 11423–11436; (*b*) I. Batatia, S. Batzner, D. P. Kovács, A. Musaelian, G. N. C. Simm, R. Drautz, C. Ortner, B. Kozinsky and G. Csányi, *Nat. Mach. Intell.*, 2025, **7**, 56.

26 A. Stukowski, *Model. Simul. Mater. Sci. Eng.*, 2009, **18**, 015012.

27 C. J. Shearer, A. D. Slattery, A. J. Stapleton, J. G. Shapter and C. T. Gibson, *Nanotechnology*, 2016, **27**, 125704.

28 Y. Liu, A. Madanchi, A. S. Anker, L. Simine and V. L. Deringer, *Nat. Rev. Mater.*, 2025, **10**, 228.

29 K. Tschoppe, F. Beckert, M. Beckert and R. Mülhaupt, *Macromol. Mater. Eng.*, 2015, **300**, 140.

30 H. Jang, Y. J. Park, X. Chen, T. Das, M.-S. Kim and J.-H. Ahn, *Adv. Mater.*, 2016, **28**, 4184.

31 H. Kwon, J. Mondal, K. A. AlOgab, V. Sammelselg, M. Takamichi, A. Kawaski and M. Leparoux, *J. Alloys Compd.*, 2017, **698**, 807.

Chapter 6

Conclusion and Outlook

At the start of this thesis, we set out to explore fundamental questions pertaining to the nature of the amorphous solid. Whilst this question may never definitively be answered, we have made significant progress as a community since Brodie's isolation of graphite oxide in 1859. With the advent of advanced experimental techniques in conjunction with the exponential growth of computational power and simulation capability, we are now able to probe the structure and dynamics of these materials in unprecedented detail.

In Chapter 1, we provided a brief overview of the definitions and historical context of amorphous solids, highlighting the challenges in characterising these materials. We further discussed topological disorder in two closely related but distinct amorphous materials: graphene and GO and outlined that the cybotactic nature of these materials is most likely opposed to the CRN model.

In Chapter 2, we derived the ideas behind using atomistic simulations to study amorphous solids, focusing on the techniques such as MD and MC along with the importance of accurate interatomic potentials and the role of statistical sampling in capturing the inherent disorder of these materials. A chronological description of the methods used in this thesis was provided, starting with first-principles DFT calculations, followed by the use of empirical potentials, and culminating in the use of state-of-the-art MLIPs to achieve DFT-level accuracy at a fraction of the computational cost.

In Chapter 3, my first paper is presented with a central focus on the fundamental understanding of two key concepts: (i) what is the nature of monolayer amorphous carbon (MAC)?, (ii) can we use nonphysical, ML atomic (site) energies to "drive" a bond-switching MC simulation and subsequently use these ML atomic energies to characterise the subtle differences between CRN and cybotactic (referred to as paracrystalline in Ref. 60) models of MAC? On point (i), we found that our structural models of MAC could be "tuned" to a varying degree of disorder, where in the extreme limit, the structure is indistinguishable from a CRN model, with the constraint that all carbon atoms are sp^2 coordinated.

For point (ii), we found that the ML atomic energies can effectively guide a bond-switching Monte Carlo simulation into distinct regions of configurational space, depending on the defect selection range—whether restricted to the defect pair itself (which we call "local" in Ref. 60), topological nearest neighbours (NN) or extended to topological next-nearest neighbours (NNN). When considering the energies of the defect pair, the structural models closely resembled a CRN model; NN interactions were found to favour a 'CRN-like' model with relatively small regions of order embedded in a much larger CRN network; NNN interactions closely resembled a cybotactic model with larger ordered domains and smaller CRN networks connecting these regions. Furthermore, the ML atomic energies could provide a unique and characteristic fingerprint depending on whether or not the structure was cybotactic or 'CRN-like', where conventional methods such as RDF and bond angle distributions were unable to distinguish between the two models easily. These findings highlighted and provided theoretical evidence for the presence of a CRN model in MAC, owing to the dynamic stability of the structure (no imaginary modes in the phonon spectrum), however the kinetics was not explored and thus the question of whether or not the CRN model is kinetically stable and whether MAC can undergo a glass transition remains an open question.

In Chapter 4, my second paper is presented, which focuses on development of a bespoke MLIP to model GO using a synergistic ML approach combining two state-of-the-art ML methods. The first stage was to "seed" an initial dataset of representative GO structures resembling the Lerf–Klinowski model through GAP accelerated AIMD simulations in `CASTEP`. This process used a naïve implementation of active learning, where the AIMD simulations were run in "GAP mode" until the

GAP model exceeded a numerical error threshold, at which point `CASTEP` would perform an electronic structure calculation and this structure was added to the dataset. The second stage was to train a MACE model on this initial dataset and to perform iterative training, increasing the temperature, and then eventually include edge structures.

With a robust, accurate MACE model (which we coined "`GO-MACE-23`"), we thermally reduced a nanoscale, amorphous GO structure to yield rGO, and the dynamical transformation was rigorously studied. The resulting structure was found to be cybotactic in nature, containing regions of graphitic domains separated by regions of amorphous, functionalised domains. The structure contained pores and holes, similar to what has previously been observed experimentally, and although it is not necessarily a 1:1 correspondence with TEM images in the literature,[82, 83] the overall morphology was consistent. The subsequent structure was then analysed using simulated XPS calculations and showed excellent agreement with experimental XPS data from an rGO sample which contained a similar profile indicative of increased sp^2 character along with a shift from epoxide and hydroxyl groups to carbonyl, ether and ester groups, highlighting oxygen groups incorporating into the basal plane of the rGO structure. This work directly demonstrates the power of simulation to explain experimental observations, providing a 1:1 atomic-scale interpretation of the chemical signatures seen in XPS that would be otherwise impossible to obtain.

In Chapter 5, my third paper is presented, which applied the `GO-MACE-23` model to study the mechanical properties of rGO. The structure was uniaxially strained with the corresponding stress–strain curves calculated and were found to be in agreement with experimental values in the literature.[231] The mechanical properties were found to be slightly anisotropic induced by thermal stochasticity in the structure, with the mechanism of failure being attributed to C$-$O bonds in the basal plane being cleaved (O atoms which had formed part of the backbone of the rGO structure). This was then expanded to perform a statistical analysis of the mechanical properties of GO and rGO, where smaller structural models were decorated with varying ratios of epoxide/hydroxyl groups and various O/C content. The results showed that lower O/C ratios resulted in structures with a higher Young's modulus, and vice versa. It was also discovered that structures with more epoxide groups displayed a plastic response to uniaxial strain, whereas structures with more hydroxyl groups displayed a brittle response. This was attributed to the fact that epoxide groups are more likely to form cross-links in the structure, which can dissipate energy and allow for plastic deformation, whereas hydroxyl groups do not provide this same level of connectivity.

This thesis provides a solid foundation for future research and clearly demonstrates the utility of advanced simulation for explaining, replacing, and guiding the experimental effort. The work in Chapter 4 provides a direct, atomistic rationalisation for the complex C 1s peaks observed in experimental XPS of rGO. Similarly, the work in Chapter 5 explains why different GO samples exhibit different mechanical properties (brittle vs. plastic) by linking the macroscopic response directly to the microscopic functional-group chemistry. The findings on mechanical properties (Chapter 5) provide a clear directive for experimentalists. To create rGO-based materials with enhanced flexibility for applications in flexible electronics, synthesis routes should be tuned to favour a higher epoxide-to-hydroxyl ratio. This predictive power allows for the rational design of new materials, saving significant laboratory time and resources. The thermal reduction in Chapter 4 and the uniaxial strain tests in Chapter 5 are examples of simulations replacing experiments that are exceptionally difficult, costly, or impossible to perform. Observing the real-time dynamics of bond-breaking during thermal annealing or precisely quantifying failure mechanisms atom-by-atom on a perfectly characterised nanoscale sample are feats unachievable by current experimental probes.

Furthermore, the outlook for this work is broad, with clear paths for both specific extensions and general transferability. Regarding possible extensions, the development of the `GO-MACE-23` model has provided a powerful tool for studying the structure and dynamics of GO and rGO, enabling us to explore their properties with unprecedented accuracy and efficiency. Beyond the papers presented in this thesis, the transferability of the `GO-MACE-23` has been further investigated for "out-of-domain" applications, as part of a wider work assessing the zero-shot generalisation behaviour of GNN potentials.[232] Located at the interface of materials and molecular modelling, GO offers an opportunity to connect these different domains of atomistic machine learning. In this work, we explored whether our `GO-MACE-23` model trained on a bulk material with organic functionalisation can be used to model molecular systems, fullerenes and reactivity with reasonable accuracy compared to MACE-MP-0[225] and a benchmark model, MACE-OFF-23, a foundation model trained on a wide range of organic molecules.[233] We found good—perhaps *surprisingly* good—zero-shot performance compared to MACE-OFF, a pre-trained model for molecular chemistry. The accuracy of both models decreases when describing reaction pathways, especially when radical species are involved.

Separately, `GO-MACE-23` has already been used to kickstart the new `GO-MACE-24` model which has been trained on roughly 10x the DFT data, extending to more domains including layered structures and bulk water between the layers. This model is expected to be used for simulations requiring a more accurate description due to the higher cost of MACE, and thus model distillation is currently being explored to create a more efficient `GO-ACE-25` model which will be used to model GO structures which are fabricable experimentally and thus, much larger than what was done before. Ongoing work is being done in this area to further improve the ACE model's stability. The aim of this future study will be to return to the textbook question of the nature of the atomistic structure of GO and rGO compared to existing models covered in Chapter 1, such as the Lerf–Klinowski model, and to broaden the simulated spectra to include NMR and Raman spectroscopy, further enhancing the model's applicability to experimental data.

Beyond the specific extensions for GO, the methodologies developed in this thesis are highly transferable to other challenging materials systems. The synergistic workflow from Chapter 4—using accelerated AIMD to "seed" an initial dataset, training a robust MACE model, and then iteratively refining it—is a general-purpose pipeline. It could be directly applied to model other complex, multi-component amorphous systems where empirical potentials fail, such as the degradation of amorphous metal-organic frameworks,[234] the dynamics of ion transport in amorphous solid-state electrolytes,[235] or the reactive aging of polymer composites.[236] Likewise, the ML-driven Monte Carlo strategy from Chapter 3 is a transferable method, offering a new, computationally efficient way to explore energy landscapes and identify hidden structural motifs in other disordered network materials, such as amorphous silicates or chalcogenide glasses.[237]

Finally, the advent of more advanced ML methods, such as diffusion models[238] and normalising flow matching,[239] has the potential to revolutionise the way we study and understand amorphous solids owing to the sampling techniques used in these approaches which can generate configurations orders of magnitude faster than using MD. Of course, these methods are not without their challenges, such as the need for large datasets and the difficulty in interpreting dynamic properties from such trajectories, but they offer a promising avenue for future research.

In conclusion, this thesis has contributed to our understanding of the nature of amorphous solids and in particular has laid the groundwork for future investigations into the fundamental nature and properties of amorphous graphene, GO and rGO. As we look to the future, I am excited about the possibilities that lie ahead and the role that amorphous materials will play in shaping the next generation of technologies. Whether it is MAC being used in future flexible electronics due to its exceptional mechanical properties, or GO and rGO finding applications in fields ranging from energy storage to biomedicine, the journey of discovery continues — and I am grateful to have been a part of it.

Bibliography

[1] W. H. Zachariasen, *J. Am. Chem. Soc* **54**, 3841 (1932).

[2] S. Elliott, *Physics of Amorphous Materials* (John Wiley & Sons, Incorporated, 1986), pp. 1–133.

[3] K. De Corte, A. Korsakov, W. Taylor, P. Cartigny, M. Ader, P. De Paepe, *Island Arc* **9**, 428 (2000).

[4] S. Balibar, *J. Low Temp. Phys* **146**, 441 (2007).

[5] T. Varga, A. Navrotsky, J. L. Moats, R. M. Morcos, F. Poli, K. Müller, A. Saha, R. Raj, *J. Amer. Ceram. Soc.* **90**, 3213 (2007).

[6] L. Yu, R. Raj, *Sci. Rep.* **5**, 14550 (2015).

[7] M. Powell, *IEEE Trans. Electron Devices* **36**, 2753 (1989).

[8] B. Fischer, W. Beyer, A. Lambertz, M. Nuys, W. Duan, K. Ding, U. Rau, *Sol. RRL* **7**, 2300103 (2023).

[9] L. A. Rosset, D. A. Drabold, V. L. Deringer, *Nat. Commun.* **16**, 2360 (2025).

[10] W. Vogel, *Glass Chemistry* (Springer, 1994), pp. 1–21.

[11] M. Faraday, *Phil. Trans. R. Soc.* **120**, 1 (1830).

[12] A. C. Wright, *Int. J. Appl. Glass Sci.* **5**, 31 (2014).

[13] S. Kyropoulos, *Z. Anorg. Allgem. Chem.* **99**, 197 (1917).

[14] A. A. Lebedev, *Trudy Gos. Opt. Inst* **2**, 1 (1921).

[15] J. Horbach, W. Kob, *Phys. Rev. B* **60**, 3169 (1999).

[16] S. Cheng, *Sci. Rep.* **12**, 13151 (2022).

[17] A. A. Lebedev, *Trudy Gos. Opt. Inst* **3**, 1 (1924).

[18] H. Macey, *et al.*, *J. Soc. Glass Technol. Trans* **15**, 54 (1931).

[19] N. Valenkov, E. Poray-Koshitz, *Z. Kristallogr.* **95**, 195 (1936).

[20] B. E. Warren, J. Biscce, *J. Amer. Ceram. Soc.* **21**, 49 (1938).

[21] W. Rosenhain, *J. Soc. Glass Technol. Trans.* **11**, 77 (1927).

[22] B. E. Warren, *Z. Kristallogr.* **86**, 349 (1933).

[23] B. E. Warren, *Phys. Rev.* **45**, 657 (1934).

[24] B. E. Warren, C. F. Hill, *Z. Kristallogr.* **89**, 481 (1934).

[25] B. E. Warren, A. D. Loring, *J. Amer. Ceram. Soc.* **18**, 269 (1935).

[26] B. E. Warren, *J. Amer. Ceram. Soc.* **17**, 249 (1934).

[27] A. Leadbetter, A. Wright, *Philos. Mag* **33**, 105 (1976).

[28] B. E. Warren, H. Krutter, O. Morningstar, *J. Amer. Ceram. Soc.* **19**, 202 (1936).

[29] M. Treacy, J. Gibson, P. Keblinski, *J. Non-Cryst. Solids* **231**, 99 (1998).

[30] J. Gibson, M. Treacy, T. Sun, N. Zaluzec, *Phys. Rev. Lett.* **105**, 125504 (2010).

[31] M. Treacy, K. Borisenko, *Science* **335**, 950 (2012).

[32] V. L. Deringer, N. Bernstein, G. Csányi, C. Ben Mahmoud, M. Ceriotti, M. Wilson, D. A. Drabold, S. R. Elliott, *Nature* **589**, 59 (2021).

[33] Y. Shang, Z. Liu, J. Dong, M. Yao, Z. Yang, Q. Li, C. Zhai, F. Shen, X. Hou, L. Wang, *et al.*, *Nature* **599**, 599 (2021).

[34] H. Tang, X. Yuan, Y. Cheng, H. Fei, F. Liu, T. Liang, Z. Zeng, T. Ishii, M.-S. Wang, T. Katsura, *et al.*, *Nature* **599**, 605 (2021).

[35] P. Debye, P. Scherrer, *Physik. Z.* **18** (1917).

[36] J. Gibson, M. Holohan, H. Riley, *J. Chem. Soc.* pp. 456–461 (1946).

[37] J. Biscoe, B. Warren, *J. Appl. Phys.* **13**, 364 (1942).

[38] R. E. Franklin, *Trans. Faraday Soc.* **45**, 274 (1949).

[39] R. E. Franklin, *Trans. Faraday Soc.* **45**, 668 (1949).

[40] R. E. Franklin, *Acta Crystallogr.* **3**, 107 (1950).

[41] R. E. Franklin, *Acta Crystallogr.* **4**, 253 (1951).

[42] R. E. Franklin, *Proc. R. Soc. Lond. A* **209**, 196 (1951).

[43] P. Harris, A. Burian, S. Duber, *Philos. Mag. Lett.* **80**, 381 (2000).

[44] M. Ghazinejad, S. Holmberg, O. Pilloni, L. Oropeza-Ramos, M. Madou, *Sci. Rep.* **7**, 16551 (2017).

[45] J. L. Fogg, K. J. Putman, T. Zhang, Y. Lei, M. Terrones, P. J. Harris, N. A. Marks, I. Suarez-Martinez, *Commun. Mater.* **1**, 47 (2020).

[46] J. Kipling, J. Sherwood, P. Shooter, N. Thompson, *Carbon* **1**, 315 (1964).

[47] S. Yamada, H. Sato, *Nature* **193**, 261 (1962).

[48] T. Yoshimori, M. Arakawa, T. Takeuchi, *Talanta* **12**, 147 (1965).

[49] H. Zittel, F. Miller, *Anal. Chem* **37**, 200 (1965).

[50] H. W. Kroto, J. R. Heath, S. C. O'Brien, R. F. Curl, R. E. Smalley, *Nature* **318**, 162 (1985).

[51] S. Iijima, *Nature* **354**, 56 (1991).

[52] K. S. Novoselov, A. K. Geim, S. V. Morozov, D.-e. Jiang, Y. Zhang, S. V. Dubonos, I. V. Grigorieva, A. A. Firsov, *Science* **306**, 666 (2004).

[53] L. Hou, X. Cui, B. Guan, S. Wang, R. Li, Y. Liu, D. Zhu, J. Zheng, *Nature* **606**, 507 (2022).

[54] Z. Yang, J. Hao, S. P. Lau, *J. Appl. Phys.* **127**, 220901 (2020).

[55] M. Chakraborty, M. S. J. Hashmi, *Advances in materials and processing technologies* **4**, 573 (2018).

[56] J. D. Bernal, *Proc. R. Soc. Lond. A* **106**, 749 (1924).

[57] P. R. Wallace, *Phys. Rev.* **71**, 622 (1947).

[58] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, M. I. Katsnelson, I. V. Grigorieva, S. V. Dubonos, A. A. Firsov, *Nature* **438**, 197 (2005).

[59] A. J. Stone, D. J. Wales, *Chem. Phys. Lett.* **128**, 501 (1986).

[60] Z. El-Machachi, M. Wilson, V. L. Deringer, *Chem. Sci.* **13**, 13720 (2022).

[61] J. Ma, D. Alfè, A. Michaelides, E. Wang, *Phys. Rev. B* **80**, 033407 (2009).

[62] M. I. Heggie, G. L. Haffenden, C. D. Latham, T. Trevethan, *Philos. Trans. R. Soc. A* **374**, 20150317 (2016).

[63] J. Kotakoski, A. V. Krasheninnikov, U. Kaiser, J. C. Meyer, *Phys. Rev. Lett.* **106**, 105505 (2011).

[64] F. R. Eder, J. Kotakoski, U. Kaiser, J. C. Meyer, *Sci. Rep.* **4**, 4060 (2014).

[65] V. Kapko, D. A. Drabold, M. F. Thorpe, *Phys. Status Solidi B* **247**, 1197 (2010).

[66] A. Kumar, M. Wilson, M. F. Thorpe, *J. Phys.: Condens. Matter* **24**, 485003 (2012).

[67] F. Wooten, K. Winer, D. Weaire, *Phys. Rev. Lett.* **54**, 1392 (1985).

[68] C.-T. Toh, H. Zhang, J. Lin, A. S. Mayorov, Y.-P. Wang, C. M. Orofeo, D. B. Ferry, H. Andersen, N. Kakenov, Z. Guo, I. H. Abidi, H. Sims, K. Suenaga, S. T. Pantelides, B. Özyilmaz, *Nature* **577**, 199 (2020).

[69] W.-J. Joo, J.-H. Lee, Y. Jang, S.-G. Kang, Y.-N. Kwon, J. Chung, S. Lee, C. Kim, T.-H. Kim, C.-W. Yang, *et al.*, *Sci. Adv.* **3**, e1601821 (2017).

[70] X. Zhang, Y.-T. Zhang, Y.-P. Wang, S. Li, S. Du, Y.-Y. Zhang, S. T. Pantelides, *Phys. Rev. B* **109**, 174106 (2024).

[71] W. S. Hummers Jr, R. E. Offeman, *J. Amer. Chem. Soc.* **80**, 1339 (1958).

[72] B. C. Brodie, *Phil. Trans. R. Soc.* pp. 249–259 (1859).

[73] G. Ruess, F. Vogt, *Monatsh. Chem.* **78**, 222 (1948).

[74] A. K. Geim, *Phys. Scr.* **2012**, 014003 (2012).

[75] H.-P. Boehm, A. Clauss, G. Fischer, U. Hofmann, *Z. Anorg. Allg. Chem.* **316**, 119 (1962).

[76] H. Boehm, R. Setton, E. Stumpp, *Carbon* **24**, 241 (1986).

[77] U. Hofmann, A. Frenzel, E. Csalán, *Liebigs Ann* **510**, 1 (1934).

[78] G. Ruess, *Monatsh. Chem.* **76**, 381 (1947).

[79] W. Scholz, H. Boehm, *Z. Anorg. Allg. Chem.* **369**, 327 (1969).

[80] A. Lerf, H. He, M. Forster, J. Klinowski, *J. Phys. Chem. B* **102**, 4477 (1998).

[81] T. Szabó, O. Berkesi, P. Forgó, K. Josepovits, Y. Sanakis, D. Petridis, I. Dékány, *Chem. Mater.* **18**, 2740 (2006).

[82] C. Gómez-Navarro, J. C. Meyer, R. S. Sundaram, A. Chuvilin, S. Kurasch, M. Burghard, K. Kern, U. Kaiser, *Nano Lett.* **10**, 1144 (2010).

[83] K. Erickson, R. Erni, Z. Lee, N. Alem, W. Gannett, A. Zettl, *Adv. Mater.* **22**, 4467 (2010).

[84] A. M. Dimiev, L. B. Alemany, J. M. Tour, *ACS Nano* **7**, 576 (2013).

[85] Z. Liu, K. Nørgaard, M. H. Overgaard, M. Ceccato, D. M. Mackenzie, N. Stenger, S. L. Stipp, T. Hassenkam, *Carbon* **127**, 141 (2018).

[86] Z. El-Machachi, D. Frantzov, A. Nijamudheen, T. Zarrouk, M. A. Caro, V. L. Deringer, *Angew. Chem. Int. Ed.* **63**, e202410088 (2024).

[87] A. M. Dimiev, S. Eigler, eds., *Graphene Oxide: Fundamentals and Applications* (John Wiley & Sons, Chichester, 2017), first edn.

[88] R. Bell, P. Dean, *Phil. Mag.* **25**, 1381 (1972).

[89] A. Wright, M. Thorpe, *Phys. Status Solidi B* **250**, 931 (2013).

[90] J. Tersoff, *Phys. Rev. Lett.* **61**, 2879 (1988).

[91] G. Galli, R. M. Martin, R. Car, M. Parrinello, *Phys. Rev. Lett.* **62**, 555 (1989).

[92] H. Liu, Z. Zhao, Q. Zhou, R. Chen, K. Yang, Z. Wang, L. Tang, M. Bauchy, *Comptes Rendus. Géoscience* **354**, 35 (2022).

[93] D. Frenkel, B. Smit, *Understanding molecular simulation: from algorithms to applications* (Elsevier, 2023).

[94] L. Verlet, *Phys. Rev.* **159**, 98 (1967).

[95] W. C. Swope, H. C. Andersen, P. H. Berens, K. R. Wilson, *J. Chem. Phys.* **76**, 637 (1982).

[96] T. Schneider, E. Stoll, *Phys. Rev. B* **17**, 1302 (1978).

[97] H. C. Andersen, *J. Chem. Phys.* **72**, 2384 (1980).

[98] H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, J. R. Haak, *J. Chem. Phys.* **81**, 3684 (1984).

[99] S. Nosé, *J. Chem. Phys.* **81**, 511 (1984).

[100] W. G. Hoover, *Phys. Rev. A* **31**, 1695 (1985).

[101] G. Bussi, D. Donadio, M. Parrinello, *J. Chem. Phys.* **126** (2007).

[102] D. Heyes, *Chem. Phys.* **82**, 285 (1983).

[103] M. A. Caro, V. L. Deringer, J. Koskinen, T. Laurila, G. Csányi, *Phys. Rev. Lett.* **120**, 166101 (2018).

[104] E. Tapavicza, G. D. Bellchambers, J. C. Vincent, F. Furche, *Phys. Chem. Chem. Phys.* **15**, 18336 (2013).

[105] R. Iftimie, P. Minary, M. E. Tuckerman, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 6654 (2005).

[106] R. Car, M. Parrinello, *Phys. Rev. Lett.* **55**, 2471 (1985).

[107] V. L. Deringer, M. A. Caro, G. Csányi, *Adv. Mater.* **31**, 1902765 (2019).

[108] S. Kmiecik, D. Gront, M. Kolinski, L. Wieteska, A. E. Dawid, A. Kolinski, *Chem. Rev.* **116**, 7898 (2016).

[109] G. Bussi, A. Laio, *Nat. Rev. Phys.* **2**, 200 (2020).

[110] D. J. Earl, M. W. Deem, *Phys. Chem. Chem. Phys.* **7**, 3910 (2005).

[111] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).

[112] T. C. Nicholas, A. E. Stones, A. Patel, F. M. Michel, R. J. Reeder, D. G. Aarts, V. L. Deringer, A. L. Goodwin, *Nat. Chem.* **16**, 36 (2024).

[113] N. Mousseau, G. Barkema, *Curr. Opin. Solid State Mater. Sci.* **5**, 497 (2001).

[114] S. K. Jain, G. T. Barkema, N. Mousseau, C.-M. Fang, M. A. van Huis, *The Journal of Physical Chemistry C* **119**, 9646 (2015).

[115] F. D'Ambrosio, J. Barkema, G. T. Barkema, *Nanomaterials* **11**, 1242 (2021).

[116] Z. Liu, D. Panja, G. T. Barkema, *Phys. Rev. E* **105**, 044116 (2022).

[117] M. H. Müser, S. V. Sukhomlinov, L. Pastewka, *Adv. Phys.: X* **8**, 2093129 (2023).

[118] E. Schrödinger, *Phys. Rev.* **28**, 1049 (1926).

[119] M. Born, R. Oppenheimer, *Ann. Phys.* **389**, 457 (1927).

[120] R. P. Feynman, *Phys. Rev.* **56**, 340 (1939).

[121] W. Pauli, *Z. Phys.* **31**, 765 (1925).

[122] D. R. Hartree, *Math. Proc. Camb. Philos. Soc.* **24**, 89 (1928).

[123] A. Szabo, N. S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover Publications, 1996), second edn.

[124] J. C. Slater, *Phys. Rev.* **34**, 1293 (1929).

[125] V. Fock, *Z. Phys.* **61**, 126 (1930).

[126] P.-O. Löwdin, *Phys. Rev.* **97**, 1509 (1955).

[127] J. McGuire, *Electron Correlation Dynamics in Atomic Collisions*, Cambridge Monographs on Atomic, Molecular and Chemical Physics (Cambridge University Press, 1997).

[128] C. D. Sherrill, *Acc. Chem. Res.* **46**, 1020 (2013).

[129] B. O. Roos, *Advances in Chemical Physics* (John Wiley & Sons, Ltd, 1987), pp. 399–445.

[130] P. Pyykkö, *Ann. Rev. Phys. Chem.* **63**, 45 (2012).

[131] C. Møller, M. S. Plesset, *Phys. Rev.* **46**, 618 (1934).

[132] R. Nesbet, *Proc. R. Soc. Lond. A* **230**, 312 (1955).

[133] J. Čížek, *J. Chem. Phys.* **45**, 4256 (1966).

[134] M. Douglas, N. M. Kroll, *Ann. Phys.* **82**, 89 (1974).

[135] B. A. Hess, *Phys. Rev. A* **33**, 3742 (1986).

[136] H. Hellmann, *J. Chem. Phys.* **3**, 61 (1935).

[137] L. Kleinman, *Phys. Rev. B* **21**, 2630 (1980).

[138] P. Hohenberg, W. Kohn, *Phys. Rev.* **136**, B864 (1964).

[139] W. Kohn, L. J. Sham, *Phys. Rev.* **140**, A1133 (1965).

[140] J. P. Perdew, K. Schmidt, *AIP Conf. Proc.* **577**, 1 (2001).

[141] J. P. Perdew, A. Zunger, *Phys. Rev. B* **23**, 5048 (1981).

[142] R. O. Jones, O. Gunnarsson, *Rev. Mod. Phys.* **61**, 689 (1989).

[143] J. P. Perdew, K. Burke, M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).

[144] P. Borlido, T. Aull, A. W. Huran, F. Tran, M. A. Marques, S. Botti, *J. Chem. Theory Comput.* **15**, 5069 (2019).

[145] V. Sokolovskiy, D. Baigutlin, O. Miroshkina, V. Buchelnikov, *Metals* **13**, 728 (2023).

[146] A. Badreldin, O. Bouhali, A. Abdel-Wahab, *Adv. Funct. Mater.* **34**, 2312425 (2024).

[147] M. A. Marques, J. Vidal, M. J. Oliveira, L. Reining, S. Botti, *Phys. Rev. B* **83**, 035119 (2011).

[148] A. D. Becke, *J. Chem. Phys.* **98**, 5648 (1993).

[149] P. J. Stephens, F. J. Devlin, C. F. Chabalowski, M. J. Frisch, *J. Phys. Chem.* **98**, 11623 (1994).

[150] K. Kim, K. Jordan, *J. Phys. Chem.* **98**, 10089 (1994).

[151] X. Wu, E. J. Walter, A. M. Rappe, R. Car, A. Selloni, *Phys. Rev. B* **80**, 115201 (2009).

[152] I. Y. Zhang, X. Ren, *arXiv preprint arXiv:2301.12119 [cond-mat.mtrl-sci]* (2023).

[153] G. Prandini, A. Marrazzo, I. E. Castelli, N. Mounet, N. Marzari, *npj Comput. Mater.* **4**, 72 (2018).

[154] S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. J. Probert, K. Refson, M. Payne, *Z. Kristall.* **220**, 567 (2005).

[155] A. J. Cohen, P. Mori-Sánchez, W. Yang, *Chem. Rev.* **112**, 289 (2012).

[156] E. Caldeweyher, S. Ehlert, A. Hansen, H. Neugebauer, S. Spicher, C. Bannwarth, S. Grimme, *J. Chem. Phys.* **150** (2019).

[157] M. Dion, H. Rydberg, E. Schröder, D. C. Langreth, B. I. Lundqvist, *Phys. Rev. Lett.* **92**, 246401 (2004).

[158] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, M. C. Payne, *J. Chem. Phys.* **122**, 084119 (2005).

[159] D. W. Brenner, *Phys. Status Solidi B* **217**, 23 (2000).

[160] V. L. Deringer, A. P. Bartók, N. Bernstein, D. M. Wilkins, M. Ceriotti, G. Csányi, *Chem. Rev.* **121**, 10073 (2021).

[161] J. E. Jones, *Proc. R. Soc. Lond. A* **106**, 441 (1924).

[162] C. de Tomas, I. Suarez-Martinez, N. A. Marks, *Carbon* **109**, 681 (2016).

[163] N. A. Marks, *Phys. Rev. B* **63**, 035401 (2000).

[164] A. C. Van Duin, S. Dasgupta, F. Lorant, W. A. Goddard, *J. Phys. Chem. A* **105**, 9396 (2001).

[165] D. W. Brenner, *Phys. Rev. B* **42**, 9458 (1990).

[166] S. J. Stuart, A. B. Tutein, J. A. Harrison, *J. Chem. Phys.* **112**, 6472 (2000).

[167] D. W. Brenner, O. A. Shenderova, J. A. Harrison, S. J. Stuart, B. Ni, S. B. Sinnott, *J. Phys.: Condens. Matter* **14**, 783 (2002).

[168] J. Tersoff, *Phys. Rev. Lett.* **74**, 5080 (1995).

[169] A. Noreyan, J. Amar, I. Marinescu, *Mater. Sci. Eng. A* **117**, 235 (2005).

[170] C. de Tomas, A. Aghajamali, J. L. Jones, D. J. Lim, M. J. López, I. Suarez-Martinez, N. A. Marks, *Carbon* **155**, 624 (2019).

[171] M. Walter, F. Mangolini, J. B. McClimon, R. W. Carpick, M. Moseler, *Phys. Rev. Mater.* **9**, 035601 (2025).

[172] M. Z. Bazant, E. Kaxiras, J. F. Justo, *Phys. Rev. B* **56**, 8542 (1997).

[173] T. P. Senftle, S. Hong, M. M. Islam, S. B. Kylasa, Y. Zheng, Y. K. Shin, C. Junkermeier, R. Engel-Herbert, M. J. Janik, H. M. Aktulga, *et al.*, *npj Comput. Mater.* **2**, 1 (2016).

[174] B. D. Jensen, K. E. Wise, G. M. Odegard, *J. Phys. Chem. A* **119**, 9710 (2015).

[175] S. G. Srinivasan, A. C. Van Duin, P. Ganesh, *J. Phys. Chem. A* **119**, 571 (2015).

[176] J. R. Quinlan, *Mach. Learn.* **1**, 81 (1986).

[177] L. Breiman, *Mach. Learn.* **45**, 5 (2001).

[178] C. Cortes, V. Vapnik, *Mach. Learn.* **20**, 273 (1995).

[179] D. R. Cox, *J. R. Stat. Soc. Ser. B Stat. Methodol.* **20**, 215 (1958).

[180] T. Hastie, R. Tibshirani, *Stat. Sci.* **1**, 297 (1986).

[181] C. Rasmussen, C. Williams, *Gaussian Processes for Machine Learning*, Adaptive Computation and Machine Learning series (MIT Press, 2005).

[182] D. E. Rumelhart, G. E. Hinton, R. J. Williams, *Nature* **323**, 533 (1986).

[183] A. Cherkasov, E. N. Muratov, D. Fourches, A. Varnek, I. I. Baskin, M. Cronin, J. Dearden, P. Gramatica, Y. C. Martin, R. Todeschini, *et al.*, *J. Med. Chem.* **57**, 4977 (2014).

[184] D. Weininger, *J. Chem. Inf. Comput. Sci.* **28**, 31 (1988).

[185] D. Rogers, M. Hahn, *J. Chem. Inf. Model.* **50**, 742 (2010).

[186] C. M. Bishop, N. M. Nasrabadi, *Pattern recognition and machine learning*, vol. 4 (Springer, 2006).

[187] J. Behler, *J. Chem. Phys.* **145** (2016).

[188] C. Zeni, A. Anelli, A. Glielmo, K. Rossi, *Phys. Rev. B* **105**, 165141 (2022).

[189] A. P. Bartók, R. Kondor, G. Csányi, *Phys. Rev. B* **87**, 184115 (2013).

[190] S. Dudarev, P. Derlet, *J. Phys.: Condens. Matter* **17**, 7097 (2005).

[191] M. Eckhoff, J. Behler, *npj Comput. Mater.* **7**, 170 (2021).

[192] W. Xu, R. Y. Sanspeur, A. Kolluru, B. Deng, P. Harrington, S. Farrell, K. Reuter, J. R. Kitchin, *Proc. Natl. Acad. Sci. U.S.A.* **122**, e2422973122 (2025).

[193] T. B. Blank, S. D. Brown, A. W. Calhoun, D. J. Doren, *J. Chem. Phys.* **103**, 4129 (1995).

[194] W. S. McCulloch, W. Pitts, *Bull. Math. Biophys.* **5**, 115 (1943).

[195] F. Rosenblatt, *Psychol. Rev.* **65**, 386 (1958).

[196] W. Morgan, J. Larsen, W. Goldstein, *J. Quant. Spectrosc. Radiat. Transf.* **51**, 247 (1994).

[197] J. Androsiuk, L. Kułak, K. Sienicki, *Chem. Phys.* **173**, 377 (1993).

[198] B. G. Sumpter, C. Getino, D. Noid, *J. Chem. Phys.* **97**, 293 (1992).

[199] M. Tusar, J. Zupan, J. Gasteiger, *J. Chim. Phys.* **89**, 1517 (1992).

[200] J. Behler, *J. Chem. Phys.* **134** (2011).

[201] J. Behler, *Chem. Rev.* **121**, 10037 (2021).

[202] A. P. Bartók, M. C. Payne, R. Kondor, G. Csányi, *Phys. Rev. Lett.* **104**, 136403 (2010).

[203] S. N. Pozdnyakov, M. J. Willatt, A. P. Bartók, C. Ortner, G. Csányi, M. Ceriotti, *Phys. Rev. Lett.* **125**, 166001 (2020).

[204] F. Musil, A. Grisafi, A. P. Bartók, C. Ortner, G. Csányi, M. Ceriotti, *Chem. Rev.* **121**, 9759 (2021).

[205] V. L. Deringer, A. P. Bartók, N. Bernstein, D. M. Wilkins, M. Ceriotti, G. Csányi, *Chem. Rev.* **121**, 10073 (2021).

[206] V. L. Deringer, G. Csányi, *Phys. Rev. B* **95**, 094203 (2017).

[207] J. Behler, M. Parrinello, *Phys. Rev. Lett.* **98**, 146401 (2007).

[208] R. Z. Khaliullin, H. Eshet, T. D. Kühne, J. Behler, M. Parrinello, *Phys. Rev. B.* **81**, 100103 (2010).

[209] R. Z. Khaliullin, H. Eshet, T. D. Kühne, J. Behler, M. Parrinello, *Nat. Mater.* **10**, 693 (2011).

[210] G. C. Sosso, G. Miceli, S. Caravati, J. Behler, M. Bernasconi, *Phys. Rev. B.* **85**, 174103 (2012).

[211] G. C. Sosso, D. Donadio, S. Caravati, J. Behler, M. Bernasconi, *Phys. Rev. B.* **86**, 104301 (2012).

[212] G. C. Sosso, G. Miceli, S. Caravati, F. Giberti, J. Behler, M. Bernasconi, *J. Phys. Chem. Lett.* **4**, 4241 (2013).

[213] V. L. Deringer, M. A. Caro, R. Jana, A. Aarva, S. R. Elliott, T. Laurila, G. Csányi, L. Pastewka, *Chem. Mater.* **30**, 7438 (2018).

[214] M. A. Caro, A. Aarva, V. L. Deringer, G. Csányi, T. Laurila, *Chem. Mater.* **30**, 7446 (2018).

[215] V. L. Deringer, C. Merlet, Y. Hu, T. H. Lee, J. A. Kattirtzi, O. Pecher, G. Csányi, S. R. Elliott, C. P. Grey, *Chem. Commun.* **54**, 5988 (2018).

[216] J.-X. Huang, G. Csányi, J.-B. Zhao, J. Cheng, V. L. Deringer, *J. Mater. Chem. A* **7**, 19070 (2019).

[217] V. L. Deringer, G. Csányi, D. M. Proserpio, *ChemPhysChem* **18**, 873 (2017).

[218] P. Rowe, V. L. Deringer, P. Gasparotto, G. Csányi, A. Michaelides, *J. Chem. Phys.* **153**, 034702 (2020).

[219] T. K. Stenczel, Z. El-Machachi, G. Liepuoniute, J. D. Morrow, A. P. Bartók, M. I. J. Probert, G. Csányi, V. L. Deringer, *J. Chem. Phys.* **159**, 044803 (2023).

[220] R. Drautz, *Phys. Rev. B* **99**, 014104 (2019).

[221] I. Batatia, D. P. Kovacs, G. Simm, C. Ortner, G. Csanyi, *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh, eds. (Curran Associates, Inc., 2022), vol. 35, pp. 11423–11436.

[222] I. Batatia, S. Batzner, D. P. Kovács, A. Musaelian, G. N. C. Simm, R. Drautz, C. Ortner, B. Kozinsky, G. Csányi, *Nat. Mach. Intell.* **7**, 56 (2025).

[223] J. Nigam, S. Pozdnyakov, G. Fraux, M. Ceriotti, *J. Chem. Phys.* **156** (2022).

[224] A. Bochkarev, Y. Lysogorskiy, C. Ortner, G. Csányi, R. Drautz, *Phys. Rev. Res.* **4**, L042019 (2022).

[225] I. Batatia, *et al.*, A foundation model for atomistic materials chemistry (2023). ArXiv:2401.00096 [physics.chem-ph].

[226] L. Barroso-Luque, M. Shuaibi, X. Fu, B. M. Wood, M. Dzamba, M. Gao, A. Rizvi, C. L. Zitnick, Z. W. Ulissi, Open Materials 2024 (OMat24) Inorganic Materials Dataset and Models (2024). ArXiv:2410.12771v1 [cond-mat.mtrl-sci].

[227] R. Jana, D. Savio, V. L. Deringer, L. Pastewka, *Model. Simul. Mater. Sci. Eng.* **27**, 085009 (2019).

[228] M. Qamar, M. Mrovec, Y. Lysogorskiy, A. Bochkarev, R. Drautz, *J. Chem. Theory Comput.* **19**, 5151–5167 (2023).

[229] P. Grigorev, L. Frérot, F. Birks, A. Gola, J. Golebiowski, J. Grießer, J. L. Hörmann, A. Klemenz, G. Moras, W. G. Nöhring, J. A. Oldenstaedt, P. Patel, T. Reichenbach, T. Rocke, L. Shenoy, M. Walter, S. Wengert, L. Zhang, J. R. Kermode, L. Pastewka, *J. Open Source Softw.* **9**, 5668 (2024).

[230] D. C. Liu, J. Nocedal, *Math. Program* **45**, 503 (1989).

[231] C. Gómez-Navarro, M. Burghard, K. Kern, *Nano Lett.* **8**, 2045 (2008).

[232] C. Ben Mahmoud, Z. El-Machachi, K. A. Gierczak, J. L. A. Gardner, V. L. Deringer, *Digit. Discov.* p. Accepted Manuscript (2025).

[233] D. P. Kovács, J. H. Moore, N. J. Browning, I. Batatia, J. T. Horton, Y. Pu, V. Kapil, W. C. Witt, I.-B. Magdau, D. J. Cole, *et al.*, *J. Am. Chem. Soc.* **147**, 17598 (2025).

[234] T. D. Bennett, A. K. Cheetham, *Acc. Chem. Res.* **47**, 1555 (2014).

[235] Q. Yang, J. Xu, X. Fu, J. Lian, L. Wang, X. Gong, R. Xiao, H. Li, *J. Mater. Chem. A* **13**, 2309 (2025).

[236] F. Zhang, R. Yang, D. Lu, *Polymers* **15**, 1928 (2023).

[237] Y. Liu, A. Madanchi, A. S. Anker, L. Simine, V. L. Deringer, *Nat. Rev. Mater.* **10**, 228 (2025).

[238] H. Kwon, T. Hsu, W. Sun, W. Jeong, F. Aydin, J. Chapman, X. Chen, V. Lordi, M. R. Carbone, D. Lu, *et al.*, *Mach. Learn.: Sci. Technol.* **5**, 045037 (2024).

[239] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, M. Le, Flow matching for generative modeling (2023). ArXiv:2210.02747 [cs.LG].

[240] N. Marks, *J. Phys.: Condens. Matter* **14**, 2901 (2002).

[241] M. A. Caro, G. Csányi, T. Laurila, V. L. Deringer, *Phys. Rev. B* **102**, 174201 (2020).

[242] E. H. Lahrar, A. Belhboub, P. Simon, C. Merlet, *ACS Appl. Mater. Interfaces* **12**, 1789 (2020).

[243] X. Wang, *et al.*, *Nat. Energy* **5**, 478 (2020).

[244] Y. Wang, Z. Fan, P. Qian, T. Ala-Nissila, M. A. Caro, *Chem. Mater.* **34**, 617 (2022).

[245] R. C. Powles, N. A. Marks, D. W. M. Lau, *Phys. Rev. B* **79**, 075430 (2009).

[246] Y. Zuo, C. Chen, X. Li, Z. Deng, Y. Chen, J. Behler, G. Csányi, A. V. Shapeev, A. P. Thompson, M. A. Wood, *et al.*, *J. Phys. Chem. A* **124**, 731 (2020).

[247] J. Lahnsteiner, M. Bokdam, *Phys. Rev. B* **105**, 024302 (2022).

[248] R. Jinnouchi, *Phys. Chem. Chem. Phys.* **24**, 15522 (2022).

[249] 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, S. J. Plimpton, *Comput. Phys. Commun.* **271**, 108171 (2022).

[250] A. Aghajamali, C. de Tomas, I. Suarez-Martinez, N. A. Marks, *Mol. Simul.* **44**, 164 (2018).

[251] S. Le Roux, P. Jund, *Comput. Mater. Sci.* **49**, 70 (2010).

[252] D. R. Cooper, B. D'Anjou, N. Ghattamaneni, B. Harack, M. Hilke, A. Horth, N. Majlis, M. Massicotte, L. Vandsburger, E. Whiteway, V. Yu, *ISRN Condens. Matter Phys.* **2012**, 501686 (2012).

[253] G. Yang, L. Li, W. B. Lee, M. C. Ng, *Sci. Technol. Adv. Mater.* **19**, 613 (2018).

Appendix

Appendix A

Machine-learned acceleration for molecular dynamics in CASTEP

This work was conducted as part of a wider collaboration with CASTEP developers to incorporate and benchmark machine-learned interatomic potentials for accelerated AIMD simulations. My contributions were validation of the “on-the-fly” generated GAP potential for an AIMD graphitisation simulation. This validation included numerical error analysis, chemical error analysis, and benchmarking of the on-the-fly GAP potential against C-GAP-17, along with transferability tests to test low density and high density structures. The findings are copied from the original paper Ref. [219] and are presented here as an appendix to this thesis.

A Case Study: Graphitisation of Amorphous Carbon

We present a case study for a complex, disordered material, as an example of a simulation that will ultimately require larger simulation cell sizes than would be amenable to DFT at all. This section, therefore, exemplifies the second “mode of operation”: generating a potential with CASTEP + ML, which is then taken out and used to drive stand-alone simulations. Our application case is amorphous carbon (a-C), whose structures can have pore radii of several nm, and a-C films can be described in deposition simulations.[240, 241] The ML-driven computational modeling of a-C materials is relevant for diverse applications, e.g., in energy storage: porous carbon structures can host metal ions (for battery anodes), molecular species (for supercapacitors), and so on.[215, 216, 242–244]

The thermal annealing of a-C structures is a benchmark challenge for empirical and ML interatomic potentials.[162, 170, 245] Starting from a disordered precursor, these simulations involve the breaking and forming of bonds typically over hundreds of thousands of MD steps, resulting in a gradually more ordered network that tends toward an sp^2-rich, graphitic structure at lower densities (1.5–2.0 g cm^{-3}) or an sp^3-rich, diamond-like structure at higher densities (3.0–3.5 g cm^{-3}). Meeting this benchmark is challenging for carbon potentials because of the diverse atomic environments in the amorphous state.

As mentioned above, on-the-fly fitting needs a different strategy compared to a "typical" GAP fit to an existing database of structures. Here, rather than starting from the existing C-GAP-17 database,[206] we show that it is possible to start with a small amount of fitting data sampled at relatively few sparse points when fitting on-the-fly and yet achieve accuracy to within tens of meV for a specific use-case. While this error would normally be considered relatively high for ML potentials,[246] we note that modeling a-C is challenging due to its complex configurational space and that C-GAP-17 has been shown to be a robust model for a-C.[170] CASTEP-GAP produces potentials of comparable quality, but will be less transferable due to the much smaller training database.

For this test, we took a 200-atom a-C structure generated by melt-quenching with C-GAP-17 and annealed it at 3000 K for 350 ps using two different methods: (i) direct MD with C-GAP-17 and (ii) CASTEP MD accelerated with on-the-fly GAP fitting. We chose a density of 2.5 g cm^{-3}, slightly higher than that of graphite, but not high enough to form a dense diamond-like structure—it might, therefore, be viewed as one of the more challenging cases.

Figure A1 shows that the reference C-GAP-17 simulation and the new CASTEP + ML run exhibit qualitatively similar trends and physically correct behavior with respect to the fraction of sp^2 atoms. The latter increases rapidly during the first 25 ps, followed by a slower increase until converging at ∼180 ps, up to numerical fluctuations. The simulation driven by the original C-GAP-17 model shows a similar profile while reaching convergence faster, at around 60 ps. Figure A1a indicates by arrows when the CASTEP-GAP model was refit, with a total of 255 refits out of 404 DFT checks throughout the simulation. Several refitting events occurred near the beginning of the run; they then gradually spread out according to the adaptive refitting scheme, up until around 170 ps where an increasing number of refitting events occurs. After ∼180 ps, the simulation required restarting, and at that point, refitting was manually disabled due to the perceived stability of the potential, allowing for the simulation to continue with the ML potential without further refitting.

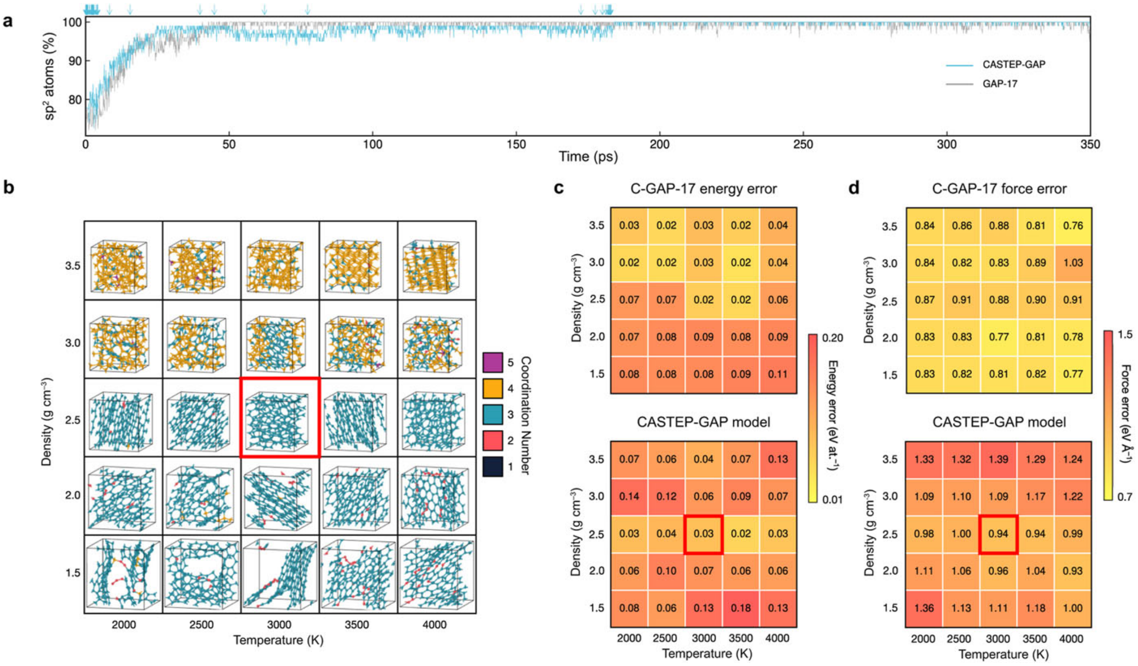

Figure A1: Graphitization of amorphous carbon as a benchmark for ML-accelerated MD. (a) Evolution of the count of sp^2 (threefold-connected) atoms at 3000 K for a 200-atom system at a density of 2.5 g cm^{-3}. Annealing within CASTEP + ML (on-the-fly GAP model fit; blue line) and C-GAP-17 in LAMMPS (gray line) shows similar profiles starting from the same structure. Blue arrows indicate refitting events within CASTEP + ML. (b) A matrix visualizing annealed carbon structures obtained using C-GAP-17, which are here used to create a DFT benchmark set. Structures are color-coded by the atomic coordination numbers. Higher-density structures appear more diamond-like in nature, whereas lower-density structures appear graphitic. The central structure is highlighted in red: this corresponds to the density and temperature settings chosen for the CASTEP + ML simulation. (c) Energy error matrix comparing potentials with DFT. Top: Energy RMSE of C-GAP-17 on 10 unique structures at 500 ps. Bottom: Same for the CASTEP-GAP potential. (d) Force error matrix comparing potentials with DFT. Top: Force RMSE of C-GAP-17 on 10 unique structures at 500 ps. Bottom: Same for the CASTEP-GAP potential.

To analyze the quality of the resulting potential, we introduce here a systematic benchmark for disordered carbon across the range of densities and degrees of ordering [Figure A1b]. These structures consist of 200 atoms and were generated using C-GAP-17 following the same annealing protocol. Ten independent structures were DFT-labeled per density-temperature run at the end of the annealing process (500 ps) for a total of 250 uncorrelated cells (50,000 atoms) in this benchmark set. The DFT parameters were adjusted for the benchmark set to match that used in C-GAP-17[206] for direct comparison; we note that, in future work, we expect to re-label the dataset at higher levels of DFT. The root mean square error (RMSE) against DFT for the energies and forces were computed for both the CASTEP-GAP model [Figs. A1(c) and A1(d), bottom] and for C-GAP-17 (top). At the chosen temperature and density, the energy accuracy of the CASTEP-GAP model is within 30 meV/at., and the force error is below 1 eV/Å, as highlighted by a square in Fig. A1(c) (bottom). C-GAP-17 reaches an energy error of $\approx$ 20 meV/at. and a force error below 0.9 eV/at., being fitted with many more SOAP sparse points and a larger database. As a result of the latter, C-GAP-17 is better suited to capturing a range of amorphous structures as indicated by the lower energy errors across most of the structures with no anomalous results [Fig. A1(c), top]. This is further illustrated by the force errors observed in C-GAP-17 [Fig. A1(d), top] where the majority of errors are below 1 eV/Å, compared with the CASTEP-GAP model [Fig. A1(d), top], which has relatively poor force accuracy since force errors were not accounted for during the refitting criterion. Changes in temperature and density have a significant impact on the numerical accuracy of the CASTEP-GAP potential. Notably, at densities of 1.5 and 3.5 g cm^{-3}, we observe significant errors. The reason behind these errors is the absence of higher- and lower-density structures in the CASTEP-GAP database. Consequently, diamond-like structures are evaluated with poor numerical accuracy in terms of forces when compared to C-GAP-17. Additionally, structures at lower densities display a greater variety of chemical environments (e.g., chains, pores, and low-coordination environments) that are absent from the CASTEP-GAP database. This lack of representation results in higher energy and force errors.

One centrally important aspect of on-the-fly ML potential fitting is the ability to extract the potential for use in an external MD engine to run larger simulations.[247, 248] We exemplify this approach here: we scaled up the system size to 10,000 atoms [Figs. 5(a) and 5(b)] and generated an a-C precursor at 2.5 g cm^{-3} using the same protocol as for the 200-atom one. The CASTEP-GAP potential was then used in LAMMPS[249] to anneal the 10,000-atom structure at 3000 K for 350 ps. The final annealed structure is shown in Fig. A2(b), color-coded by SOAP similarity to diamond,[189] where a value of 1 indicates a local chemical environment identical to that in bulk diamond. Qualitatively, the structure in Fig. 5(b) shows the expected warped graphitic layers at this density, where it is too dense to form flat, ordered graphitic sheets, but not dense enough to form a bulk tetrahedral (diamond-like) structure. Notably, there are locally two regions with that characteristic, highlighted in yellow in Fig. 5(b). This type of behavior has been observed with C-GAP-17 previously (along with other empirical potentials)[170, 250] and ascribed to numerical artifacts arising from interactions between the cutoff function and the second-neighbor peak in the radial distribution function.[250]

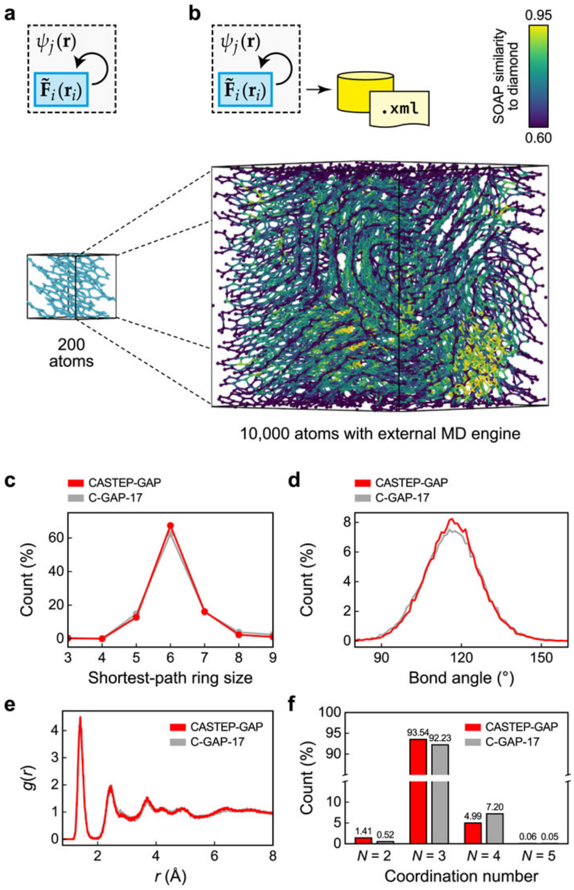

Figure A2: Using the on-the-fly generated potential to simulate a 10,000-atom graphitization. (a) 200-atom structure annealed using CASTEP + ML. (b) 10,000-atom structure generated using the on-the-fly potential within LAMMPS. The structure is color coded according to per-atom SOAP similarity to bulk diamond. (c) Shortest-path ring size count as determined using the R.I.N.G.S software.[251] (d) Bond-angle distribution. (e) Radial distribution function. (f) Coordination number count. The results in (c)–(f) are shown for a C-GAP-17-driven simulation (gray) and for a comparable simulation driven by the on-the-fly-generated CASTEP + ML potential (red).

For comparison, the same 10,000-atom structure was also annealed using C-GAP-17. The shortest-path ring counts for both structures give nearly identical trends, where the structure generated from the on-the-fly potential has slightly more six-membered rings [Fig. A2(c)]. Both bond angle distributions [Fig. A2(d)] are centered around 120°, emphasizing the presence of sp^2 environments; both radial distribution functions [Fig. A2(e)] have a first peak at 1.42 Å, consistent with the experimental bond length in graphene (1.42 Å),[252, 253] along with a second, slightly broader peak centered at 2.44 Å, corresponding to next-nearest neighbors, and a smaller peak at 3.77 Å. Finally, an assessment of coordination numbers [Fig. A2(f)] shows that both annealed structures have a high percentage of sp^2 environments. However, the structure generated with the on-the-fly potential has fewer sp^3 environments compared to the C-GAP-17 one. Note that there is a very small amount of five-coordinated atoms in both structures (0.06% and 0.05%, respectively); for a discussion of five-coordinated atoms in simulations of amorphous carbon, see Ref. 241.

Overall, the on-the-fly ML potential shows good performance in a scaled-up simulation using an external MD engine. As demonstrated in Fig. A2, the potential yields structural metrics that are comparable to the predictions of C-GAP-17 while requiring only a fraction of the fitting cost. Additionally, little human effort was expended for the construction of the training database and the fitting procedure itself. Numerical validation across densities and temperatures [Fig. A1(c)] showed that the generated potential performs similarly to C-GAP-17 for some configurations, but has a higher error for others for which it has not been "trained." We emphasize that to be considered fully validated (for production simulations), the robustness of the potential would now need to be more comprehensively tested—e.g., by repeating the same structural validation across various densities and temperatures.