



Machine Learning Force Fields

for Modelling Reactions at Complex Interfaces

Lars Leon Schaaf

University of Cambridge

Wolfson College

This dissertation is submitted for the degree of

Doctor of Philosophy

April 2025

To my parents, for showing me the joy of asking 'why?' ...

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface or specified in the text. It is not substantially the same as any work that has already been submitted, or, is being concurrently submitted, for any degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the preface or specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Lars Leon Schaaf

April 2025

Abstract

Computational modelling at the atomic scale has traditionally been constrained by a stark trade-off between accuracy and computational cost. Recently, machine learning (ML) architectures have been trained on highly accurate quantum mechanical calculations to sidestep this constraint [1, 2]. So-called machine learning force fields (MLFFs) predict energies and forces on atomic configurations at near quantum mechanical accuracy with orders-of-magnitude reduction in computational cost. Furthermore, under the assumption of locality, MLFFs scale linearly with system size, while even approximate electronic structure methods such as density functional theory (DFT) scale cubically.

In this thesis, we explore how MLFFs can be used to model complex reactive systems. We examine both catalytic reactions at oxide interfaces and carbon capture in porous materials. A key obstacle to the widespread use of MLFFs for complex systems is the need to curate relevant training datasets that lead to accurate and trustworthy results. We present an automated framework for training force fields for reactive systems which uses model uncertainty to iteratively improve and select new configurations for evaluation with the reference method. The protocol is validated on the extensively explored hydrogenation of carbon dioxide to methanol over indium oxide. We demonstrate that our workflow can be used to determine energy barriers with quantum mechanical accuracy, requiring minimal human supervision.

Furthermore, we show that machine learning surrogate modelling not only reduces the computational cost of routine *in silico* catalytic simulation tasks but also allows for an entirely new approach to computational modelling. We capture entropic finite-temperature effects by computing free-energy barriers. For a single barrier calculation of formaldehyde conversion over indium oxide, this requires 10^7 energy evaluations. Although quantum mechanical calculations would take more than a century to run on a modern supercomputer, we can obtain the answers within a single day.

Moreover, the fractional computational cost allows us to explore reactions in greater detail. Our automated reaction path search identifies an alternative reaction

progression with a 40% reduction in activation energy for the previously assumed rate-limiting step in CO₂ hydrogenation to methanol on indium oxide. The ability of MLFFs to enhance our understanding of extensively studied catalysts underscores the need for fast and accurate alternatives to direct *ab-initio* simulations.

Next, we show that the training workflow is also applicable for curating training data for porous metal-organic frameworks to simulate carbon capture. With the help of this workflow, we can decipher previously unexplained NMR spectra, leading to a more accurate understanding of the carbon capture mechanism. Additionally, we explore how recent developments in atomistic foundation models can be used to accelerate the MLFF training workflow through fine-tuning and initial dataset curation.

Finally, we address one of the key limitations of prevalent machine learning architectures: their assumption of locality. One reason for MLFFs' linear scaling with system size and transferability is the assumption that the energy is a local function of the atomic environment. We introduce a new approach to capturing non-local interactions called matrix function neural networks (MFNs). By mimicking the ground truth quantum mechanical methods, MFNs can model highly non-local systems. To date, no other architectures can capture the non-locality of cumulene chains and extrapolate to unseen chain lengths, including global transformer networks. We anticipate that non-local methods will play a critical role in modelling electro-catalytic reactions and charge transfer.

The developments presented in this thesis make MLFFs more accessible for routine computational catalysis applications. Force fields with quantum mechanical accuracy are set to significantly transform computational catalysis, enabling a more accurate representation of realistic surfaces and ultimately leading to improved predictions and stronger agreement with experimental results.

Acknowledgements

First and foremost, I would like to express my deepest gratitude to Prof. Gábor Csányi, my supervisor, for inspiring me with his passion for atomistic modelling. I am sincerely thankful for your guidance, mentorship, and for always making time to answer questions and explore ideas.

My PhD experience would not have been the same without the vibrant and stimulating environment created by the researchers around me. I am truly grateful to all members of the Gábor group, past and present, for fostering a fun atmosphere while producing cutting-edge research. Special thanks to Elena Gelžinytė for helping me to get acquainted with the GAP framework during the early days of my PhD. I owe particular thanks to Ilyes Batatia. Working together on the matrix function neural networks was not only immensely rewarding but also incredibly fun. I am grateful to the many colleagues that I had the pleasure of bouncing ideas back and forth with. Working with Rokas Elijošius was a particular highlight, making me wish we had only started collaborating sooner. I also thank Will Baldwin for the countless coffee breaks, for sharing your passion for electrostatics, and for keeping spirits high through the ups and downs.

I extend my thanks to our visiting collaborators, including Will Lin and Carsten Staacke, whose fresh perspectives were much appreciated.

I am also deeply thankful to the external collaborators I had the privilege of working with. Benjamin J. Rhodes, thank you for teaching me about experimental NMR and for the many long discussions where we taught each other our respective crafts. Seeing how computational methods can shed light on experimental findings was truly eye opening. I am grateful to Edvin Fako and Sandip De from BASF for patiently teaching me about catalysis and to Christoph Ortner for providing a more mathematical viewpoint. I am also thankful to Chris Pickard, my advisor, for offering inspiration on computational–experimental collaboration and Igor Lerner for encouraging me to go into research and giving me certainty in uncertain times.

My journey would not have been possible without the encouragement of my school teachers. I am especially grateful to Julie Krejci for nurturing my curiosity

and encouraging me to ask questions, and to Elena Vojtelová and Eric Neuhold for their support during my secondary school education.

Outside of research, I am thankful to the friends who made the past years so memorable. Thank you Marcus Köttering, Willem Jongert, Mathilde Allard, and Harrison Roberts for making Cambridge feel like home. I am also grateful to Wolfson College for its vibrant, multidisciplinary community, where I had the pleasure of meeting Simon Mathis, Constantin Kogler, and Theodore de Pomereu amongst many more. I am particularly grateful to An Tran, Charlotte Lefèvre, and Javier de Haro Abrona for sharing the PhD journey, for long walks during COVID, and for all the laughter along the way. I would like to thank Daniel Greenhouse, whose opinions and friendship I value greatly. Above all, I would like to thank Mathilde Lavanant. Thank you for all the adventures beyond the keyboard.

I am also grateful to the Wolfson College Football Club, and to captain Mike Yates. During the pandemic especially, the sense of camaraderie was a much-enjoyed refuge.

The people I owe the most gratitude to are my parents. Among countless things, I am especially grateful for the curiosity you have given me. Thank you for always being a phone call away, your perspective, and your unparalleled support. I want to thank my grandmother. Thank you for your humour and reminding me to slow down from time to time. Finally, I want to thank my sister for bringing a bit of home to Cambridge and for all the joy you brought with you.

Table of contents

List of abbreviations	xv
1 Introduction	1
1.1 Atomic processes	1
1.2 Computational modelling and rational design	2
1.3 Outline of the thesis	4
2 Background	7
2.1 Quantum chemistry	7
2.1.1 Schrödinger equation	8
2.1.2 Born-Oppenheimer approximation	9
2.1.3 Quantum many-body problem	11
2.1.4 Density functional theory	11
2.2 Empirical force fields	14
2.3 Group and representation theory	17
2.3.1 Symmetries and groups	18
2.3.2 Presentations	19
2.3.3 Representations	20
2.3.4 Irreducible representations	22
2.3.5 Equivariant functions	25
2.3.6 Tensor products	26
2.3.7 Learnable tensor products	27
2.4 Machine Learning Force Fields	28
2.4.1 General approach	29
2.4.2 Gaussian Approximation Potentials (GAP)	32
2.4.3 Atomic Cluster Expansion (ACE)	35
2.4.4 Tensor sketching and TRACE	37
2.4.5 Equivariant message passing neural networks	38
2.4.6 Higher order message passing (MACE)	39

3	Active learning for accurate reaction rates in catalysis	43
3.1	Background and motivation	44
3.1.1	Opportunities for catalysis	44
3.1.2	Challenges in computational catalysis	45
3.1.3	The role of machine learning force fields	50
3.1.4	Uncertainty estimation	51
3.2	Active learning for catalysis	53
3.3	Automated training protocol	55
3.4	Validation on known pathway	59
3.4.1	Carbon-dioxide hydrogenation over indium oxide	59
3.4.2	Computational details	61
3.4.3	Leveraging computation acceleration	63
3.4.4	Computational cost comparison	66
3.5	Transferability to unexplored surfaces	68
3.6	Discussion	70
4	Foundation models and fine-tuning	73
4.1	Background and motivation	74
4.1.1	Early foundation models	74
4.1.2	MACE-MP0	75
4.2	Performance on out domain test cases	76
4.2.1	Transition state search for catalysis	76
4.2.2	Molecular dynamics of large catalytic systems	78
4.3	Fine-tuning	79
4.4	Discussion	81
5	Effects of dynamics on NMR spectra across time-scales	83
5.1	Background and motivation	83
5.2	Nuclear magnetic resonance	84
5.2.1	NMR spectroscopy	86
5.2.2	Anisotropic shielding	86
5.2.3	<i>ab-initio</i> NMR spectra	87
5.2.4	Stochastic Liouville equation	88
5.3	Computing dynamic NMR spectra	88
5.3.1	Exploration of dynamic events	88
5.3.2	Modelling NMR spectra across time-scales	89
5.4	Application to metal-organic frameworks for carbon capture	91
5.4.1	MFU-4l MOF	92

5.4.2	Potassium bicarbonate	94
5.4.3	CD-MOF	96
5.5	Discussion	97
6	Matrix Function Neural Networks	99
6.1	Background and motivation	99
6.2	Related work	100
6.3	The Architecture	101
6.3.1	Local graph layer	102
6.3.2	Matrix construction	102
6.3.3	Matrix normalization	106
6.3.4	Matrix functions	107
6.3.5	Update	108
6.3.6	Readout	110
6.4	Resolvent parameterization of matrix functions	110
6.5	Understanding the non-locality of Matrix Function Networks	111
6.5.1	Relationship to infinite layers of MPNNs	111
6.5.2	Combes-Thomas theorem	112
6.5.3	Visualisation of non-locality	113
6.6	Interpretability of the MFN operators	114
6.7	Applications of matrix functions to cumulenes	115
6.7.1	Trends with chain length and rotation	117
6.7.2	Guaranteed Non-Local dataset	117
6.7.3	Computational Details	118
6.8	Relevance for reactions at interfaces	121
6.9	Discussion	121
7	Conclusions and Outlook	123
7.1	Conclusions	123
7.2	Outlook	124
	References	127
	List of figures	147
	List of tables	149
	Appendix A Computational details	151
A.0.1	Visualisation of paths	156

List of abbreviations

ACE Atomic Cluster Expansion. 35–40, 102

AL Active Learning. 60, 61

CatKit Catalysis Toolkit. 55, 57, 62, 63, 66, 152

CD-MOF Cyclodextrin-based Metal-Organic Framework. 95–97

DFT Density Functional Theory. 3, 8, 14, 16, 30, 34, 43, 51–53, 55, 58–62, 64, 66–70, 74–76, 83, 84, 87, 88, 95, 96, 105, 116, 151–153, 155, 156

E(3) Euclidean Group in three dimensions. 29

FPS Farthest Point Sampling. 57

GAP Gaussian Approximation Potential. 16, 17, 32, 33, 39, 51, 52, 58, 60, 61, 66, 68, 153, 155, 156

In₂O₃ Indium(III) Oxide. 54, 57

LAMMPS Large-scale Atomic/Molecular Massively Parallel Simulator. 62

MACE Message Passing with Atomic Cluster Expansion. 29, 39, 70, 76, 77, 102, 109, 115, 118–120, 154

MACE-MP0 MACE Foundation model trained on Materials Project structures. 74–78, 97

MD Molecular Dynamics. 54, 55, 57, 58, 66, 89, 91, 94–96

ML Machine Learning. 30, 34, 90

- MLFF** Machine Learning Force Field. viii, 4, 5, 30, 43, 53, 58–60, 62, 63, 66–70, 79, 84, 95, 121, 147
- MLIP** Machine Learned Interatomic Potential. 3
- MOF** Metal-Organic Framework. 90, 92, 93, 96, 147
- MPNN** Message Passing Neural Network. 39, 102, 112
- NEB** Nudged Elastic Band. 58–60, 62–66, 76, 77
- NMR** Nuclear Magnetic Resonance. viii, 4, 5, 73, 74, 80, 83–97, 123, 147
- ORCA** Quantum Chemistry Software Package. 16, 118
- PAW** Projector Augmented Wave. 61
- PBE** Perdew–Burke–Ernzerhof functional. 61
- SCF** Self-Consistent Field. 61, 118
- SO(3)** Special Orthogonal Group in three dimensions. 18, 19, 23–26, 31
- SOAP** Smooth Overlap of Atomic Positions. 32, 33, 36, 52, 57, 66
- TS** Transition State. 58

Chapter 1

Introduction

...the most remarkable discovery in all of astronomy is that the stars are made of atoms of the same kind as those on the earth...

Richard Feynman

This chapter provides the broader context for the work presented in this thesis by highlighting the importance of atomic-scale processes. It outlines how the presented research contributes to the field of computational atomic-scale modelling and its relevance to society. The chapter concludes with a brief overview of the thesis structure and a summary of publications that resulted from the presented research.

1.1 Atomic processes

Our day-to-day life is shaped by processes occurring at the atomic scale. Molecular machines replicate DNA with remarkable precision, lithium ions diffuse through batteries releasing energy, and toxic gases are converted to less harmful pollutants in catalytic converters [3, 4]. All of these processes are the result of intricate atomic interactions [5].

Interestingly, the large diversity of objects around us emerges from a small number of atomic building blocks and their complex interactions. Atomic elements combine to form vastly different materials and molecules. Iron, for example, is a highly ductile metal, but when combined with a small percentage of carbon, it becomes harder and more brittle. The complex microscopic interactions at the atomic scale give rise to the macroscopic properties of materials and molecules. Indeed, even materials made of a single element may have vastly different observable properties based only on differences in 3D atomic arrangement. For example, carbon exists as graphite, a soft

black conducting material found in pencils, and diamond, a valuable transparent jewel with exceptional hardness. Given that these contrasting materials differ only by their crystal structure, an accurate atomic scale description is critical for understanding our macroscopic reality [6, 7].

Unsurprisingly, humans have a long history of investigating the physical laws that govern smaller and smaller scales. The proposition that the universe is composed of fundamental, indivisible particles dates back to ancient history. Indeed, the word "atom" has its origins in Greek natural philosophy and means uncuttable or indivisible [8]. Even though we know since J. J. Thomson's discovery of the electron in 1897 that atoms are indeed divisible and made of more fundamental particles [9], the atomic scale remains particularly interesting to investigate theoretically. Indeed, the model of the atom has continued to change throughout history, from the discovery of the nucleus by Ernest Rutherford in 1911 to the modern quantum mechanical description of the atom [10–12].

Our ever-growing understanding of atoms and their interactions has allowed us to design new devices, materials, and medications. For example, the discovery of the electron and its negative charge was central to understanding and designing cathode tubes found in early television devices. Atomic scale processes continue to be at the heart of pivotal technologies from semiconductors in processing chips to batteries in electric cars [4, 13].

1.2 Computational modelling and rational design

Our accurate understanding of the physical laws that govern the atomic scale has led to a new scientific field: atomistic computational modelling. We can simulate atomic structures to determine properties and dynamic processes without needing to synthesise a specific molecule or material. Quantum mechanics, specifically the Schrödinger equation, has been remarkably effective at describing and predicting the behaviour of atomic configurations [11, 12, 14–17]. A key limitation is that finding solutions for all but the smallest and simplest of systems is demanding, requiring numerical methods [18]. The solution, or approximate solutions, to the Schrödinger equation allow us to predict properties of molecules and materials, *ab-initio*, that is purely based on physical laws without the need for experimentation [12]. The accuracy of *ab-initio* methods has led to the proliferation of computational quantum chemistry. While observing atomistic structures experimentally is complicated, physical simulations can act as a magnifying glass with atomistic resolution [19].

The utility of computational modelling can mainly be categorised into three prominent use cases. Firstly, computational approaches can be used to understand experimental observations and provide **mechanistic insight**. The advantage of computational tools in this context is that the location and behaviour of atoms can be trivially observed. Experimentation at the atomic scale is often complicated because instruments may modify the experimental outcome, and observation under operational conditions is challenging. Secondly, computational methods can explore simplified, idealised systems that could not be directly observed experimentally. This simplification can help **build fundamental understanding** and develop trends in behaviour that generalise to many systems [20]. Lastly, computational methods can be used for **property prediction** for downstream tasks such as in silico screening. Given an accurate model of the underlying atomistic interactions, computational methods can be used as predictive tools to determine properties without the need for costly and time-intensive synthesis. By leveraging knowledge obtained from computational modelling, materials and molecules can now be *designed* rather than discovered through trial and error [21, 22].

In this thesis, we use computational methods to describe the 3D arrangement of materials and molecules and predict their dynamics. To determine observable macroscopic quantities, such as reaction rates, one can compute the forces on atoms and then evolve them in time. At each simulation step, the forces on the atoms are computed. Assuming the nuclei to behave classically, the time evolution of the atomic positions can be described with Newton's equations of motion. Through computational simulations, we can hence directly observe atomic processes. Furthermore, it is possible to describe rare events, beyond the reach of direct molecular dynamics simulations, by relying on approximations, such as transition state theory or by using enhanced sampling techniques.

One of the key limitations of *ab-initio* methods is their large computational cost. In this thesis, we train machine learning models to predict the forces on atoms without needing to solve the electronic structure problem [2, 23–27]. The total energy of a configuration and the forces on the atoms are predicted directly based on the positions of the atoms and their chemical elements. The model is trained on a representative set of quantum-mechanical calculations. The resulting machine learning force fields (MLFFs) typically scale linearly with system size and are 4-5 orders of magnitude faster than even approximate quantum mechanical approaches like DFT [2, 28–30]. In the context of material modelling, MLFFs are also referred to as machine-learned interatomic potentials (MLIPs). The sizeable computational

speedup allows for much longer simulation times of larger systems, significantly extending the computational capabilities, with minimal trade-off in accuracy [31].

We focus on both applying MLFFs and improving their architecture. We explore modelling rare events in the context of catalysis and porous carbon capture materials. Specifically, we compute energy barriers at catalytic interfaces, which directly relate to reaction rates. Additionally, we model the dynamics in metal-organic frameworks, which are known for their CO₂ adsorption abilities. Here, we compare our computational results with the experimental data and show that dynamics have a significant effect on the experimental spectra. After exploring rare events at complex interfaces with existing architectures, we examine the current limitations of local machine learning force fields. We introduce a new non-local architecture and show that it can describe systems where local models fail qualitatively.

1.3 Outline of the thesis

In this section, we provide an outline of the thesis. The first part of the thesis focuses on how machine learning force fields (MLFFs) can be used to model catalytic reactions and describe the effects of dynamics on NMR parameters. Additionally, we address one of the key limitations of existing MLFF architectures: their locality.

The work has resulted in a number of publications, which are listed below. Here, we provide a brief summary of each chapter, along with the publications on which it is based.

Chapter 2 provides a brief background of relevant previous research and the necessary context for the work of this thesis. We start by giving a brief overview of quantum mechanics and empirical force fields. After a brief section on group and representation theory, we introduce the machine learning architectures which are used in **Chapter 3-5**. We address a key shortcoming of these methods in **Chapter 6** where we investigate non-local effects.

Chapter 3 is the first results section, focused on applying MLFF to model catalytic reactions. This chapter outlines an automated training protocol to systematically improve machine learning force fields for catalytic systems with active learning. We show how the reduced computational cost can be leveraged to find lower energy barriers for extensively explored reactions and used to compute free energy barriers. The work of this chapter is based on Publication [1].

Chapter 4 focuses on the performance of atomistic foundation models. Recent advances in MLFF architectures have enabled single models to be trained on

large and chemically diverse datasets. These so-called foundation models are able to describe most of material chemistry with reasonable accuracy. In this section, we test a model trained only on crystal structures and its accuracy on catalytic tasks. Furthermore, we investigate how the foundation models can be fine-tuned to produce quantitatively accurate results for carbon capture in porous systems. This chapter is based on Publications [2–4].

Chapter 5 addresses how to use MLFF to predict the effects of molecular dynamics on NMR spectra. This section is the result of close collaboration with NMR experimentalists and is based on Publications [3, 4].

Chapter 6 introduces matrix function neural networks, a new machine learning architecture that captures non-local interactions. The results presented in this chapter are the result of collaborative work published in Publication [5].

Chapter 7 concludes the thesis with a summary of the findings and an outlook to future trends in computational catalysis.

Publications

This thesis is based on published work. A list of papers, including the collaborating authors, is provided below.

1. Schaaf, L. L., Fako, E., De, S., Schäfer, A. & Csányi, G. Accurate Energy Barriers for Catalytic Reaction Pathways: An Automatic Training Protocol for Machine Learning Force Fields. *npj Computational Materials* **9**, 180 (2023).
2. Batatia, I. *et al.* A Foundation Model for Atomistic Materials Chemistry arXiv: 2401.00096. (2024). Pre-published.
3. Schaaf, L. L., Rhodes, B. J., Zick, M. E., Pugh, S. M., Hilliard, J. S., Sharma, S., Wade, C. R., Milner, P. J., Csányi, G. & Forse, A. C. ML Force Fields for Computational NMR Spectra of Dynamic Materials across Time-Scales. *Neurips 2024 Workshop on AI for Accelerated Materials Design* (2024).
4. Rhodes, B. J., Schaaf, L. L., Zick, M. E., Pugh, S. M., Hilliard, J. S., Sharma, S., Wade, C. R., Milner, P. J., Csányi, G. & Forse, A. C. 17O NMR Spectroscopy Reveals CO₂ Speciation and Dynamics in Hydroxide-Based Carbon Capture Materials. *ChemPhysChem*, e202400941. ISSN: 1439-7641 (2024).
5. Batatia, I., Schaaf, L. L., Chen, H., Csányi, G., Ortner, C. & Faber, F. A. Equivariant Matrix Function Neural Networks. *International Conference on Learning Representations 2024*. arXiv: 2310.10434 (Oct. 16, 2023).

6. Schaaf, L. L., Batatia, I., Brunken, C., Barrett, T. D. & Tilly, J. *BoostMD: Accelerating Molecular Sampling by Leveraging ML Force Field Features from Previous Time-Steps* arXiv: 2412.18633. (2025). Pre-published.
7. Norwood, S. W., Schaaf, L. L., Batatia, I., Csányi, G. & Bhowmik, A. Enhancing the Local Expressivity of Geometric Graph Neural Networks. *NeurIPS 2023 Workshop: Machine Learning and the Physical Sciences* (2023).

Chapter 2

Background

This chapter provides the context needed to understand the contributions outlined in the subsequent chapters of this thesis. The aim is to introduce past developments in the area of machine-learning accelerated atom scale modelling. Machine learning force fields are trained on energies and forces from quantum mechanical calculations. We therefore start with a brief overview of quantum chemistry methods. These can be used to run atomic simulations, such as molecular dynamics, to predict material and molecular properties *ab-initio*, that is without the need for experimental input.

We then discuss an alternative approach to atom scale simulations, namely classical force fields, which are parametrised to reproduce experimental observables. While classical force fields are orders of magnitude faster compared to quantum mechanical methods, they do not share the same *ab-initio* predictive accuracy.

We finish the chapter by introducing machine learning force field architectures, which obtain near quantum mechanical accuracy at the cost of a classical force field. Additionally, we use this chapter to introduce some background on group and representation theory, which we will refer to in [Chapter 6](#) when introducing matrix function networks.

2.1 Quantum chemistry

There are many textbooks and lecture notes on introductory quantum mechanics. This section serves as a reminder of the core concepts relevant to this thesis. First, we provide an overview of general quantum chemistry methods and introduce the Schrödinger equation. We then introduce the concept of the potential energy surface (PES) and the Born-Oppenheimer approximation. Next, we introduce density functional theory (DFT), an approximation which has significantly lowered the com-

putational cost of *ab-initio* methods. In this thesis, all machine learning force fields are trained on data generated with DFT.

2.1.1 Schrödinger equation

A core postulate of quantum mechanics is that the state of a system is fully described by its wave function $\psi(\mathbf{r}_1, \mathbf{r}_2, \dots, t)$, where \mathbf{r}_i are the spatial coordinates of the system's particles. The wave function evolves in time according to the famous time-dependent Schrödinger equation,

$$i\hbar \frac{\partial \Psi(\mathbf{r}, t)}{\partial t} = H(\mathbf{r}, t) \Psi(\mathbf{r}, t), \quad (2.1)$$

where H is the Hamiltonian operator for the system of interest, corresponding to the total energy, the sum of the kinetic and potential energy. When the Hamiltonian doesn't depend explicitly on time, the partial differential equation of Equation 2.1 has energy eigenstates given by,

$$\Psi_n(\mathbf{r}, t) = \psi_n(\mathbf{r}) \theta_n(t), \quad (2.2)$$

resulting in the separation of a spatial $\psi_n(\mathbf{r})$ and temporal part $\theta_n(t)$. Under this Ansatz, the time-dependent Schrödinger equation can be separated into two ordinary differential equations

$$-\frac{\hbar^2}{2m} \nabla^2 \psi_n(\mathbf{r}) + V(\mathbf{r}) \psi_n(\mathbf{r}) = E_n \psi_n(\mathbf{r}), \quad i\hbar \frac{\partial}{\partial t} \theta_n(t) = E_n \theta_n(t), \quad (2.3)$$

where the $\psi_n(\mathbf{r})$ satisfies the time-independent Schrödinger equation (TISE). The TISE is an eigenvalue equation with eigenstates ψ_n and eigenvalues E_n . In general the solution to the Schrödinger equation with a time independent Hamiltonian is given by a superposition of these eigenstates

$$\Psi(\mathbf{r}, t) = \sum_n c_n \psi_n(\mathbf{r}) e^{-iE_n t/\hbar}, \quad (2.4)$$

where $H(\mathbf{r})$ has discrete eigenvalues E_n . As the probability density given by $|\Psi(\mathbf{r})|^2$ does not depend on time when Ψ is an energy eigenstate, these are often referred to as the stationary states.

In the context, of atomistic modelling, the particles of interest are the electrons and nuclei of a material or molecule. The Hamiltonian for an atomistic system is

hence given as a sum of nuclear (nn), electron (ee) and interacting (ne) terms,

$$H = T_{ee} + T_n + V_{ee} + V_{nn} + V_{ne}, \quad (2.5)$$

where T and V correspond to the kinetic and potential energy operators, respectively. Explicitly writing the individual terms, we obtain the many-body Schrödinger equation,

$$\left[-\sum_i \frac{\hbar^2}{2m_e} \nabla_i^2 - \sum_I \frac{\hbar^2}{2M_I} \nabla_I^2 + \sum_{i<j} \frac{e^2}{4\pi\epsilon_0 |\mathbf{r}_i - \mathbf{r}_j|} + \sum_{I<J} \frac{e^2}{4\pi\epsilon_0 |\mathbf{R}_I - \mathbf{R}_J|} - \sum_{i,I} \frac{e^2}{4\pi\epsilon_0 |\mathbf{r}_i - \mathbf{R}_I|} \right] \psi_k = E_k \psi_k, \quad (2.6)$$

where the wave-function ψ is a function of both the electron \mathbf{r}_i and nuclear \mathbf{R}_I positions, m_e is the electron mass, M_I is the mass of the nucleus I , Z_I and Z_J are the atomic numbers of nuclei I and J , e is the elementary charge, ϵ_0 is the vacuum permittivity, ∇_i and ∇_I are the gradient operators with respect to the coordinates of electrons i and nuclei I , and E is the total energy of the system.

2.1.2 Born-Oppenheimer approximation

The complete many-body Schrödinger equation Equation (2.6) is a function of both the electronic and nuclear coordinates. Due to the stark differences in masses, the electronic and nuclear motions can be approximately decoupled. The so-called Born–Oppenheimer approximation assumes that the electrons respond instantaneously to the relatively slow changes in the nuclear positions.

Within the Born–Oppenheimer picture, one first treats the nuclei as fixed parameters and solves the electronic problem at a given nuclear configuration. We therefore decompose the total molecular Hamiltonian of Equation 2.6 into electronic and nuclear contributions,

$$H(\mathbf{r}, \mathbf{R}) = H_{\text{el}}(\mathbf{r}; \mathbf{R}) + H_{\text{nuc}}(\mathbf{R}), \quad (2.7)$$

where the electronic and nuclear part is given as

$$H_{\text{el}}(\mathbf{r}; \mathbf{R}) = T_e(\mathbf{r}) + V_{ee}(\mathbf{r}) + V_{ne}(\mathbf{r}; \mathbf{R}), \quad (2.8)$$

$$H_{\text{nuc}}(\mathbf{R}) = T_n(\mathbf{R}) + V_{nn}(\mathbf{R}) \quad (2.9)$$

$$\left[-\sum_i \frac{\hbar^2}{2m_e} \nabla_i^2 + \sum_{i<j} \frac{e^2}{4\pi\epsilon_0} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - \sum_{i,I} \frac{e^2}{4\pi\epsilon_0} \frac{Z_I}{|\mathbf{R}_I - \mathbf{r}_i|} \right] \psi_k(\mathbf{r}; \mathbf{R}) \quad (2.10)$$

$$= \epsilon_k(\mathbf{R}) \psi_k(\mathbf{r}; \mathbf{R}),$$

where $\epsilon_k(\mathbf{R})$ are the electronic eigenvalues which parametrically depend on the nuclear positions.

We assume that the eigenstates of the entire molecular wave-function can be written in as a sum of product states,

$$\Psi(\mathbf{r}, \mathbf{R}) = \sum_k \chi_k(\mathbf{R}) \psi_k(\mathbf{r}; \mathbf{R}). \quad (2.11)$$

where ψ_k are the electronic wavefunctions, which depend parametrically on the nuclear positions \mathbf{R} , and χ_k are the nuclear wavefunctions. Substituting the electronic wave-function into the complete Schrödinger equation, introduces terms which couple electronic energy states as the nuclear kinetic energy acts on both the nuclear wavefunctions $\chi_k(\mathbf{R})$ and the parametric dependence of the electronic wavefunctions $\psi_k(\mathbf{r}; \mathbf{R})$. The Born-Oppenheimer approximation neglects any coupling between electronic energy states, assuming that the electrons adiabatically follow a single electronic eigenstate as \mathbf{R} varies. Consequently, the nuclear wave-function satisfies

$$\left[-\sum_I \frac{\hbar^2}{2M_I} \nabla_I^2 + \sum_{I<J} \frac{e^2}{4\pi\epsilon_0} \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|} + \epsilon_k(\mathbf{R}) \right] \chi_k(\mathbf{R}) = E_k \chi_k(\mathbf{R}), \quad (2.12)$$

where E_k is the total energy eigenvalue. Here we assume that changes in nuclear position do not result in transitions between electronic states and remain on one potential energy surface, ie that of the ground-state. This approximation breaks down near avoided crossings, where the electronic energy gaps are small. The nuclei experience an effective potential corresponding to the sum of the Coulombic nuclear–nuclear interaction and the electronic energy,

$$V_{\text{eff}}(\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_N) = \sum_{I<J} \frac{e^2}{4\pi\epsilon_0} \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|} + \epsilon_k(\mathbf{R}). \quad (2.13)$$

This effective potential for the nuclei is commonly referred to as the Born-Oppenheimer potential energy surface (PES). In principle, we can now determine the dynamics of our system by alternately solving the electronic Schrödinger equation and the nuclear Schrödinger equation.

We can make a further approximation that the nuclear dynamics can be described classically. The forces on nucleus I can be computed as

$$\mathbf{F}_I = -\frac{\partial}{\partial \mathbf{R}_I} V_{\text{eff}}(\mathbf{R}_1, \dots, \mathbf{R}_I, \dots, \mathbf{R}_N). \quad (2.14)$$

We can hence describe the dynamics of an atomic system by solving the electronic Schrödinger equation, computing the classical forces on the nuclei, evolving the nuclear positions using Newton's equations of motion, and then solving the electronic Schrödinger equation for the new nuclear positions. This procedure is commonly referred to as *ab-initio* molecular dynamics [32].

2.1.3 Quantum many-body problem

Even after separating the solution to the Schrödinger equation into electronic and nuclear parts, finding the exact solution is unattainable for all but the smallest of systems. The complexity arises because the state of a quantum system is defined by its wave function, which starkly contrasts with classical mechanics. Let a state be all possible information that fully describes a system, such that we can use physical laws to predict how the system changes with time. In classical mechanics, the state of a particle is entirely defined by its position and momentum. Newton's equations of motion can be used to predict all future positions and momenta, ie all future states. In quantum mechanics, the state is defined by a complex-valued function. Imagine that we have N particles where each is restricted to be on one of two sites with positions $(0,0)$ and $(0,1)$. To define the state of the system, we need to specify a complex-valued number for all possible combinations of the N particles between these two sites. For this system, we hence require a total of 2^N complex numbers to define its state. Given a Hamiltonian, we can then predict the time evolution of the 2^N complex numbers using the Schrödinger equation. Given that 2^{300} is more than the number of protons in the known universe, the exponential scaling with system size is a fundamental challenge in solving quantum mechanical problems exactly [33].

2.1.4 Density functional theory

A key observation that has significantly reduced the computational complexity of the many-body electronic Schrödinger equation is that the ground-state energy can be obtained from the electron density only [34, 35]. Let $n(\mathbf{r})$ denote the electron density of the many-body ground-state $\Psi(\mathbf{r}_1, \dots, \mathbf{r}_N)$. For a given ground-state density $n(\mathbf{r})$, Hohenberg-Kohn theorems show that there is a unique external potential $v_{\text{ext}}(\mathbf{r})$ up

to an additive constant [34]. This uniquely defines the electronic Hamiltonian H , with eigenstates Ψ_i and so the many-particle ground-state is a unique functional of $n(\mathbf{r})$.

Hence, if E is the ground-state energy of a system with external potential v_{ext} , it can be expressed as a functional of the electron density n ,

$$E[n(\mathbf{r})] = \int v_{ext}(\mathbf{r})n(\mathbf{r})d\mathbf{r} + F[n(\mathbf{r})], \quad (2.15)$$

where $F[n]$ is a system independent (universal) functional. Furthermore, Hohenberg and Kohn showed that for any trial density not corresponding to the ground-state, $E[n]$ is higher. Solving the electronic Schrödinger equation can hence be reformulated as the variational problem of minimising the energy over all physical densities n [34].

While the Hohenberg-Kohn theory is exact, a universal functional $F[n]$ is not known and difficult to determine explicitly. Kohn-Sham DFT makes the HK theories practical by introducing an auxiliary non-interacting electron reference system chosen to have the same ground-state density as the exact interacting Hamiltonian [36]. First, the total functional is separated into contributions,

$$\begin{aligned} F[n] &= T_s[n(\mathbf{r})] + E_H[n(\mathbf{r})] + E_{xc}[n(\mathbf{r})] \\ &= T_e[n(\mathbf{r})] + \frac{e^2}{8\pi\epsilon_0} \int \int \frac{n(\mathbf{r})n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}d\mathbf{r}' + E_{xc}[n], \end{aligned} \quad (2.16)$$

where the three terms correspond to the kinetic energy T_e of the independent electron system, E_H is the Hartree energy and the exchange-correlation energy [36]. The Hartree potential corresponds to an approximation of the electron-electron interaction, assuming that the electrons interact with a mean field created by themselves. This neglects any electron-electron correlation and also any effects due to the Pauli-exclusion principle and antisymmetry of the fermionic wave-function. The exchange-correlation term E_{xc} accounts for the difference between the E and everything captured by T_s and E_H . In principle, the exchange-correlation functional exists, and if it were known, one could calculate the ground-state energy from the electron density exactly. This rearrangement is useful as the exchange-correlation functional is more tractable and local than the entire universal functional $F(n)$.

To obtain the ground-state density, in practice, Kohn and Sham use the fact that the kinetic energy T_s of an independent electron model is also a functional of an electron density, based on the Hohenberg-Kohn theorems. Hence, to obtain $n(r)$, we

can solve the Schrödinger equation of independent electrons,

$$\left[-\frac{\hbar^2}{2m_e} \nabla^2 + v_s(\mathbf{r}) \right] \phi_i(\mathbf{r}) = \varepsilon_i \phi_i(\mathbf{r}), \quad (2.17)$$

where ϕ_i are single-particle orbitals of the auxiliary system in an effective potential,

$$v_s(\mathbf{r}) = v_{ext}(\mathbf{r}) + v_H(\mathbf{r}) + v_{xc}(\mathbf{r}), \quad (2.18)$$

where the potentials are for a given electron density $n(r)$. As the effective potential depends on $n(\mathbf{r}) = \sum_i |\phi_i(\mathbf{r})|^2$, the KS equations are solved self consistently. Starting from an initial guess of $n(\mathbf{r})$ the starting potentials are computed. One solves for ϕ_i , updates the density and corresponding potentials until self-consistent convergence. The total energy for the interacting KS energy is obtained using the converged densities and single particle orbitals and substituting into Equation 2.16, such that

$$\begin{aligned} E[n(\mathbf{r})] = & \int v_{ext}(\mathbf{r})n(\mathbf{r})d\mathbf{r} - \frac{\hbar^2}{4m_e} \sum_i \int \phi_i^*(\mathbf{r})\nabla^2\phi_i(\mathbf{r})d\mathbf{r} \\ & + \frac{e^2}{8\pi\epsilon_0} \iint \frac{n(\mathbf{r})n(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}d\mathbf{r}d\mathbf{r}' + E_{xc}[n(\mathbf{r})]. \end{aligned} \quad (2.19)$$

The approximation originates from the choice of the exchange correlation functional $E_{xc}[n]$.

Performance In practice various functional forms and parametrisations are used to approximate the exchange-correlation contributions, which result in predictive error. Qualitative failures may arise due a systematic tendency for local functionals to spread electron density [37, 38]. This affects charge transfer situations, charge distributions during reactions and systems containing localized d/f orbitals, such as transition metal oxides [38–40]. Nonetheless, comparisons to more accurate *ab-initio methods* such as CCSD(T) reveal, that modern DFT functional approximations often give quantitatively useful relative energies for many weak and moderately correlated systems [41]. Hence, the optimal choice of DFT functional, is both system and simulation dependent.

The hierarchy of increasingly complex and costly exchange-correlation functional is often displayed as a Jacob’s ladder. At the bottom, local density approximations (LDA) and generalised gradient approximation (GGA) methods depend purely on the local electron density $n(r)$ and its gradient $\nabla n(r)$ respectively. More complex hybrid functionals furthermore incorporate exact Hartree-Fock exchange. Often functionals

are compared to CCSD(T), which is often taken as the gold standard in quantum chemistry. Depending on the system, a different functional obtain better agreement with CCSD(T) and experiment. Additionally, there exist dispersion corrections, which correct for the locality of the previously mentioned exchange-correlation functional [42]. To obtain the ground-state density $n(r)$ it is sufficient to minimise the energy $E[n(r)]$. Making use of the so-called variational principle, the task of finding the ground-state energy is an optimisation task, which can be reformulated as solving an eigenvalue problem.

The computational cost of most DFT codes scales cubically with the number of neighbours. Depending on the choice of functional and other parameters, the overall computational cost can vary significantly. The reduced computational cost compared to solving for the wave function directly allows for routine calculations of up to hundreds or thousands of atoms.

2.2 Empirical force fields

Not all atom-scale simulations start from a quantum-mechanical description. Indeed, predicting the forces and energies directly from 3D nuclear positions using simple analytic functional forms is a large research area by itself [43, 44]. The simplicity of directly predicting forces without explicitly considering electrons means these approaches are orders of magnitude faster. Consequently, it is possible to directly simulate large systems with hundreds of thousands of atoms for long time scales. In the context of small molecules and computational biology, these are typically referred to as empirical force fields, whereas the materials simulation community refers to them as interatomic potentials.

As outlined in the previous section, atomic nuclei can be simulated as classical objects in an effective potential that originates from Coulomb nuclear-nuclear interactions and the interaction of the nuclei with an electron density. Rather than describing the electrons explicitly, the interactions can be described by a simple functional form. A common approach is to incorporate locally bonded terms and long-range interactions separately. A typical functional form includes,

$$E = E_{\text{bonds}} + E_{\text{angles}} + E_{\text{dihedrals}} + E_{\text{Coulomb}} + E_{\text{vdW}}, \quad (2.20)$$

where we explicitly define which atoms are bonded a priori. This is the base form for many popular force fields such as Amber and OPLS. Commonly, the bonded terms are simple quadratic terms, which depend on the distance between two bonded atoms

and their atom type,

$$E_{\text{bonds}}(\mathbf{r}_i, \mathbf{r}_j) = -\frac{1}{2}\kappa_{ij} \left(|\mathbf{r}_j - \mathbf{r}_i| - r_{ij}^{\text{equi}} \right)^2, \quad (2.21)$$

where the spring constant κ and the equilibrium distance r^{equi} depend on the atom types of the two bonded atoms i and j . Note that there typically exist multiple atom types for each chemical element, determined by their bonding to neighbouring atoms. Like the bond connectivity, the atom types are assigned a priori and are typically fixed throughout a simulation. Similar energy terms can be constructed for the angles between atoms and dihedral angles involving 3 and 4 bonded atoms, respectively. Parameters such as κ are then typically adjusted to reproduce experimental or quantum mechanical observables such as the density or two-body dimer curves. This adjustment to reproduce observables gives them the name *empirical* force fields.

The non-bonded interactions commonly contain an electrostatic term

$$E_{\text{coulombic}} = \sum_{i < j} \frac{q_i q_j}{4\pi\epsilon |\mathbf{R}_j - \mathbf{R}_i|}, \quad (2.22)$$

describing the Coulombic repulsion between nuclei, based on their charge. Furthermore, there are terms to capture intermolecular interactions, including dispersion forces. The Lennard-Jones potential is one of the most studied two-body intermolecular potentials,

$$E_{\text{inter}} = \sum_{i < j} \frac{A_{ij}}{|\mathbf{R}_j - \mathbf{R}_i|^{12}} - \frac{B_{ij}}{|\mathbf{R}_j - \mathbf{R}_i|^6}, \quad (2.23)$$

where A and B are atom-type dependent parameters. The $1/R^6$ attractive term captures the effects of instantaneously induced dipoles, which result in the so-called London-dispersion force. More complex dispersion corrections exist, which take into account additional dispersion corrections, such as Debye forces, which arise from the interaction of permanent dipoles with induced dipoles.

The calculation of interatomic distances, angles and dihedrals has been optimised such that empirical force fields can be evaluated more than 10^8 times per day. This is orders of magnitude faster than quantum mechanical calculations, which can take hours to days for single-force evaluation. The significant reduction of the computational cost of force fields comes at a price, however. Often, the parameters are adjusted to reproduce specific experimental observables. These parameters are, however, frequently not compatible, requiring different parameterisations for different observables, such as the density or diffusion coefficients. The observable properties are, hence, not emergent phenomena from accurate atom scale interactions.

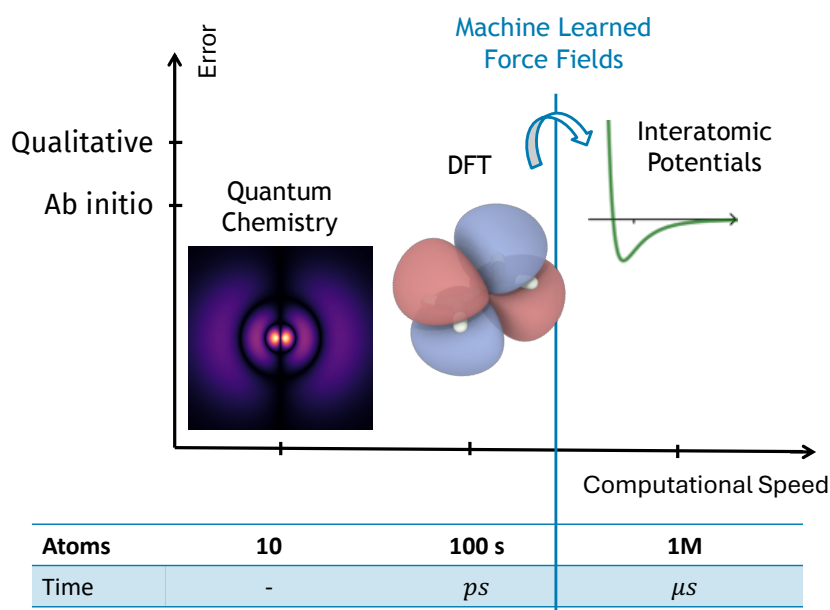


Figure 2.1 Atomistic simulations and the trade-off between accuracy and computational cost. Showing a scale of approaches from wave-function based quantum chemistry methods, through to approximate methods such as density functional theory (DFT) to interatomic potentials also referred to as empirical force fields. The x axis shows the rough length and time scale of routine computational simulations for each approach. Inset illustrations show, from left to right, a solution to the electronic Schrödinger equation for the hydrogen atom, molecular orbitals obtained using DFT (computed using the ORCA package [46]) and the functional form of the Lennard-Jones potential[47] for interatomic interactions. The role of machine learning force fields is to bridge the GAP between *ab-initio* methods and empirical force fields.

This raises the question of how informative dynamics with empirical force fields truly are, given that arbitrary results can be obtained through simple reparametrisation.

Another disadvantage is that empirical force fields typically require defining bond connectivity a priori and keeping it constant throughout a simulation. This makes them unsuited for describing reactions or changes in bonding. There are some exceptions, where bonded terms are dynamically weighted during molecular dynamics, such as ReaxFF[45]. However, their accuracy is not sufficient to match *ab-initio* methods. Furthermore, like other empirical potentials, these are fit to directly reproduce desired observables and hence have limited predictive ability.

Machine learning force fields aim to predict energies and forces with near quantum mechanical accuracy at the cost of empirical force fields. As such, it breaks the traditional trade-off between accuracy and computational cost of atomistic simulations, as visualised in Figure 2.1. For machine learning approaches to outperform their empirical counterparts, we need to generate more expressive functional forms

than simple bond and angle terms. Compared to their empirical counterpart, these can be trained on a large dataset of *ab-initio* energy and forces. The most successful and transferable machine learning architectures are intrinsically reactive as they do not require any bond connectivity as input.

Note that all terms in the empirical potential are defined such that the total energy does not change under certain transformations, including translations and rotations. The key is to use internal coordinates, such as bond distances and angles, which remain constant under these transformations. When trying to fit more complicated functional forms compared to those above using machine learning approaches, a key challenge is to maintain this so-called equivariance. Before introducing machine learning architectures, we therefore provide a brief outline of how to handle such symmetries in the next section. The role of machine learning force fields is to bridge the GAP between *ab-initio* methods and empirical force fields.

2.3 Group and representation theory

Many physical quantities remain unchanged under certain transformations. The total energy of an atomic system, for example, does not change, or is invariant, with respect to rotations, translations, and reflections. These symmetries have fundamental consequences; for example, translation invariance leads to the conservation of momentum.

In many areas of machine learning, restricting the architecture to obey the underlying symmetry of the data is seen as an inductive bias. A restriction is imposed on the architecture to improve performance, for example, for out-of-domain tasks or in low-data regimes. In the context of atomistic modelling, for every continuous symmetry, there exists a corresponding conservation law. Violations of such conservation laws can lead to unphysical simulations and reduce the trustworthiness of calculated observables. Ensuring that learned functions are invariant to certain transformations takes on much greater importance compared to other machine learning domains. A common approach in vision applications is to augment the training data by randomly rotating the images and thereby learning functions that are close to being invariant. This is very effective and partially outperforms architectures that are restricted to be equivariant. However, if the confidence of an image classifier is marginally different for a horse translated to the left, this is of little importance to the application. If, on the other hand, a molecule has a lower energy due to global rotation or translation, it will exhibit some non-physical preferential alignment.

In this section, we look at the key ingredients needed to learn equivariant functions of 3D point clouds. We provide a brief background on the 3D rotation group, the irreducible representations of the $SO(3)$ group and their basis functions. Furthermore, we examine tensor products and how the results can be decomposed as a sum of spherical basis vectors.

2.3.1 Symmetries and groups

We briefly introduced the importance of symmetries when modelling atomistic systems in the previous section. Group theory is the study of such symmetry operations. We will give some brief background to motivate the historical choices made in designing machine learning architectures. This section provides a brief overview to introduce the language needed to understand equivariant machine learning architectures. For an in-depth background and overview, please see the extensive literature and lecture notes on the subject [48, 49]

Symmetry operations, such as 3D translations, are represented by group elements. A group G consists of a set of group elements g_λ and a group operation \cdot , which composes two group elements together. The composition is often referred to as the product or multiplication of two group elements. The group elements and composition satisfy the following conditions:

Closure The product of two group elements is also a group element.

$$\forall g_1, g_2 \in G, \quad g_1 \cdot g_2 \in G \quad (2.24)$$

Associativity The way in which elements are grouped during the product operation does not matter.

$$\forall g_1, g_2, g_3 \in G, \quad (g_1 \cdot g_2) \cdot g_3 = g_1 \cdot (g_2 \cdot g_3) = g_1 \cdot g_2 \cdot g_3 \quad (2.25)$$

Identity There exists an identity element e in G , for which any element in the group is unchanged under multiplication. The element is the same for left and right multiplication.

$$\exists e \in G, \quad \text{such that } e \cdot g_\lambda = g_\lambda = g_\lambda \cdot e \quad (2.26)$$

Inverse Each group element has an inverse which is also a group element of G and this element is both the left and right inverse.

$$\forall g_\lambda \in G, \exists g_\lambda^{-1} \in G, \text{ such that } g_\lambda \cdot g_\lambda^{-1} = e = g_\lambda^{-1} \cdot g_\lambda \quad (2.27)$$

Note that the label λ , which distinguishes group elements, can be continuous (e.g. 3D rotations) or discrete (e.g. permutations). In addition, the composition does not necessarily need to commute, which means that $g_1 \cdot g_2$ is not necessarily the same as $g_2 \cdot g_1$.

Examples of Groups

We now give some examples of groups relevant in the context of machine learning force fields. The set of all possible rotations in three-dimensional Euclidean space, for example, forms a group, the special orthogonal group $SO(3)$. If this thesis is rotated around the x-axis by 90 degrees and then around the y-axis by 90 degrees, the result will be another rotation, which is also part of the special orthogonal group, illustrating closure. A rotation can be trivially reversed, and so there exists an inverse. Note that if we had first rotated the thesis around the y-axis and then the x-axis, the thesis would be in a different orientation. In other words, the group elements of the $SO(3)$ group do not commute, and so it is a nonabelian group.

The set of all possible group elements of $SO(3)$ combined with all possible translations and inversions forms the Euclidean group $E(3)$. In the remainder of the section, we will focus on the $SO(3)$ group, while we will focus on how to treat the other symmetries in [Section 2.4.1](#) in the context of introducing machine learning force field architectures.

2.3.2 Presentations

A specific group is defined by the properties of its group elements and how they relate to each other through composition. For finite groups, this can be presented as a multiplication table. More generally, they can be defined by their presentations, which consist of a set of essential group elements (generators) and the relations these elements satisfy. All other group elements can be obtained by the composition of the essential elements.

To illustrate a specific presentation, we examine the permutation group of three elements S_3 . To represent group elements of the permutations on a set 1, 2, 3 we use cyclic notation. Take the group element, which maps $3 \mapsto 2$, $2 \mapsto 1$ and $1 \mapsto 3$, this

Table 2.1 Multiplication table of the rank 3 permutation group, S_3 . Showing how different group elements of the S_3 group compose with each other, where e is the identity and other elements are presented in cycle notation, such that the group element, which maps $3 \mapsto 2$, $2 \mapsto 1$ and $1 \mapsto 3$, is given as (321) .

	e	(12)	(13)	(23)	(123)	(132)
e	e	(12)	(13)	(23)	(123)	(132)
(12)	(12)	e	(132)	(123)	(23)	(13)
(13)	(13)	(123)	e	(132)	(12)	(23)
(23)	(23)	(132)	(123)	e	(13)	(12)
(123)	(123)	(13)	(23)	(12)	(132)	e
(132)	(132)	(23)	(12)	(13)	e	(123)

can be expressed as (321) . Group elements, or specific permutations, compose with each other to form a different set of permutations, for example $(12) \circ (23) = (123)$. The group has $3 \cdot 2 \cdot 1 = 6$ elements and the complete set of all possible compositions is shown in the multiplication [Table 2.1](#).

The presentation, showing only the essential elements and their relations, can be written as,

$$\langle g_1, g_2 \mid g_1^2 = g_2^2 = (g_1 g_2)^3 = e \rangle, \quad (2.28)$$

from which all other elements can be obtained. Taking g_1 to be (12) and g_2 to be (13) , we quickly see that $g_1 \circ g_2 = (132)$ and $(132)^3 = e$ and so the relation in [Equation \(2.28\)](#) holds. Furthermore, we can use the multiplication table to see that all other elements can be obtained by compositions containing only g_1 and g_2 .

2.3.3 Representations

As seen in the previous section, specific groups are defined by the relations between group elements. It is possible to represent the abstract group elements as specific invertible linear transformations (matrices) acting on a vector space. The group elements, represented as invertible matrices, can then be composed by matrix multiplication.

Definition 2.3.1. *Formally, a representation of a group G on a vector space V over a field \mathbb{F} is a group homomorphism,*

$$\rho : G \rightarrow GL(V), \quad (2.29)$$

where $GL(V)$ is the general linear group of V , consisting of all invertible linear transformations of V to itself.

More simply, each group element $g \in G$ can be represented by $\rho(g)$ as an invertible linear transformation on V , where the map ρ preserves the group structure,

$$\rho(g_1 \circ g_2) = \rho(g_1) \cdot \rho(g_2). \quad (2.30)$$

As the map ρ preserves the multiplicative structure of G , it is a homomorphism. Note, that when V has finite dimensions n with a basis B , the linear transformations on V can be represented by a set of $n \times n$ invertible-matrices. In Equation (2.30), we have explicitly differentiated between, \circ , the composition of two group elements and \cdot , the multiplication of two matrices. The core idea of representation theory is, therefore, to associate each group element g with a corresponding matrix representing a linear transformation on a vector space.

Example representations

As previously for presentations, we illustrate the concept of a representation on the permutation group of three elements, S_3 . Let us represent S_3 on a vector space \mathbb{R}^3 , spanned by the basis vectors

$$\mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad \mathbf{e}_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}. \quad (2.31)$$

Each basis vector can be thought of as corresponding to the first, second and third objects, respectively. The group element (12) can now be represented as a matrix,

$$\rho(12) = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (2.32)$$

If we apply $\rho(12)$ to e_1 we obtain e_2 . We can also compose group elements using simple matrix multiplication, for example,

$$\rho(12) \cdot \rho(23) = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} = \rho(123) = \rho((12) \circ (23)) \quad (2.33)$$

2.3.4 Irreducible representations

The same group can have many representations. In this section, we define and introduce equivalent, reducible and irreducible representations.

Equivalent representations

Taking the example from the previous section, assume that the representation $\rho(g)$ is a 3×3 matrix, then

$$\rho'(g) = S^{-1}\rho(g)S,$$

is also a representation where S is a 3×3 invertible matrix. The fact that ρ' is also a representation is easily seen,

$$\rho'(g_1)\rho'(g_2) = S^{-1}\rho(g_1)S S^{-1}\rho(g_2)S = \rho'(g_1g_2). \quad (2.34)$$

Here, the change in representation corresponds to a simple change of basis. The two representations are considered equivalent.

Reducible representations

Furthermore, it is possible for a representation to be reducible, meaning it can be expressed as a direct sum of other representations.

Definition 2.3.2. *Assume, as previously, that ρ is a representation of a group G on a vector space V . The representation is said to be reducible if there exists a non-trivial subspace $W \subset V$ such that W is invariant under the action of G . A linear subspace $W \subset V$ is G -invariant if $\rho(g)w \in W$ for all $g \in G$ and all $w \in W$.*

In this context, the trivial subspaces of a vector space V are the zero vector $W = \{0\}$ and the entire vector space $W = V$. If a representation cannot be reduced, it is called an irreducible representation.

Example of a reducible representation

We illustrate how a reducible representation can be decomposed into irreducible subrepresentations with an example of the permutation group S_3 . The representation given in the previous section of S_3 on the vector space $V = \mathbb{R}^3$ with basis vectors e_1, e_2, e_3 defined in Equation 2.31 is reducible.

To find a relevant subspace, first note that a vector $v \in V$, with $\mathbf{v} = a\mathbf{e}_1 + b\mathbf{e}_2 + c\mathbf{e}_3$ is invariant under any permutation if and only if $a = b = c$. As a consequence, the subspace $W^{(1)}$ defined as,

$$W^{(1)} = \{(a, b, c) \in \mathbb{R}^3 : a = b = c\}, \quad (2.35)$$

is S_3 invariant as for all $w^{(1)} \in W^{(1)}$ any group action $g \in S_3$, $\rho^{(1)}(g)w^{(1)} = w^{(1)} \in W^{(1)}$. Taking the basis as $\mathbf{f}^{(1)} = \{[1, 1, 1]^T\}$, the representation of S_3 on $W^{(1)}$ is the trivial identity.

Note also that for any permutation, $\rho(g)$ acting on a vector $\mathbf{v} = a\mathbf{e}_1 + b\mathbf{e}_2 + c\mathbf{e}_3$, the resulting vector $\mathbf{v}' = \rho(g)\mathbf{v} = a'\mathbf{e}_1 + b'\mathbf{e}_2 + c'\mathbf{e}_3$, the sums $a + b + c = a' + b' + c'$. As a consequence, the subspace $W^{(2)}$,

$$W^{(2)} = \{(a, b, c) \in \mathbb{R}^3 : a + b + c = 0\}, \quad (2.36)$$

is also S_3 invariant. We can specify a new basis,

$$\mathbf{u} = \mathbf{e}_1 + \mathbf{e}_2 + \mathbf{e}_3, \quad \mathbf{f}_1 = \mathbf{e}_1 - \mathbf{e}_2, \quad \mathbf{f}_2 = \mathbf{e}_2 - \mathbf{e}_3, \quad (2.37)$$

which spans \mathbb{R}^3 . The irreducible representation on the vector space $W^{(2)}$ consists of set of 2×2 matrices. The set of generators,

$$\rho^{(2)}(12) = \begin{bmatrix} -1 & 1 \\ 0 & 1 \end{bmatrix}, \quad \rho^{(2)}(23) = \begin{bmatrix} 1 & 0 \\ 1 & -1 \end{bmatrix}, \quad \rho^{(2)}(123) = \begin{bmatrix} 0 & -1 \\ 1 & -1 \end{bmatrix}, \quad (2.38)$$

can be used to construct all group elements of the irreducible representations of S_3 on \mathbb{R}^2 . It is, hence, possible to reduce the representation ρ of S_3 on vector space \mathbb{R}^3 into a sum of irreducible representations

$$P^{-1}\rho(g)P = \rho^{(1)}(g) \oplus \rho^{(2)}(g), \quad V = W^{(1)} \oplus W^{(2)}, \quad (2.39)$$

where P is a change-of-basis matrix corresponding to expressions 2.37.

Irreducible representations of the SO(3) group

We now shift our focus to the main group of interest, the group of three-dimensional rotations, $SO(3)$. The composition of two rotations around the origin is another rotation and for every rotation there exists a unique inverse rotation. Since rotations

are associative and satisfy all previously mentioned group properties, 3D rotations around the origin form a group.

To introduce the irreducible representations of $SO(3)$, the Wigner-D matrices, we start by defining the relevant vector space and its orthogonal basis. Let S^2 denote a two-dimensional unit sphere in \mathbb{R}^3 . Then consider the set of square-integrable functions on S^2 ,

$$L^2(S^2) = \left\{ f : S^2 \rightarrow \mathbb{C}, \int_{S^2} |f(\theta, \phi)|^2 \sin(\theta) d\theta d\phi < \infty \right\}, \quad (2.40)$$

where θ and ϕ are the polar and azimuthal angles, respectively. The addition and scalar multiplication in $L^2(S^2)$ are defined pointwise and satisfy all vector space axioms. The actions of group elements of $SO(3)$ correspond to norm-preserving rotations of these functions. The $SO(3)$ group can, hence, be represented on the set of square-integrable functions on the unit sphere.

The spherical harmonics form an orthonormal basis for $L^2(S^2)$.

Definition 2.3.3. *The spherical harmonic $Y_{lm} : S^2 \rightarrow \mathbb{C}$ for every non-negative integer l and integer m satisfying $-l \leq m \leq l$ is defined as,*

$$Y_{lm}(\theta, \phi) = (-1)^m \sqrt{\frac{(2l+1)(l-m)!}{4\pi(l+m)!}} P_l^m(\cos \theta) e^{im\phi}, \quad (2.41)$$

where P_l^m are the associated Legendre polynomials, where the index l is referred to as the degree and index m as the order of the spherical harmonic.

The set of spherical harmonics forms a complete orthonormal basis for $L^2(S^2)$, such that any square-integrable function can be expressed as a linear combination of the spherical harmonics,

$$f(\theta, \varphi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l f_{lm} Y_{lm}(\theta, \varphi), \quad (2.42)$$

where f_{lm} corresponds to the expansion coefficients. In the context of machine learning, \mathbf{f} is often referred to as spherical tensor with elements f_m^l .

We now define the irreducible representations of $SO(3)$ in the spherical harmonic basis. Let $R \in SO(3)$ be a rotation about the origin that rotates a function f in $L^2(S^2)$ to a function f' . Both functions can be expanded in the spherical harmonic basis with coefficients f_{lm} and f'_{lm} respectively. Under the group action, a spherical harmonic of degree l transforms to a linear combination of spherical harmonics of

the same degree, such that

$$f'_{lm}(\hat{\mathbf{r}}) = f_{lm}(\mathbf{R}\hat{\mathbf{r}}) = \sum_{m'=-l}^l f_{lm} D_{mm'}^l(\mathbf{R}) Y_{lm'}(\hat{\mathbf{r}}), \quad (2.43)$$

where \mathbf{r} is a unit vector on the sphere and $D_{mm'}^{(l)}(\mathbf{R})$ is the Wigner-D matrix for degree l associated with the rotation \mathbf{R} around the origin. These correspond to $(2l+1) \times (2l+1)$ dimensional matrices. Hence, under rotation, spherical tensors transform as,

$$f'_{lm} = \sum_{m'=-l}^l D_{m'm}^l(R) f_{lm'}. \quad (2.44)$$

Note that under rotations, the expansion coefficients only mix within the same degree l . For a fixed l we can hence define a $(2l+1)$ dimensional subspace of $L^2(S^2)$,

$$\mathcal{V}_l = \text{span}\{Y_{lm}\}_{-l \leq m \leq l} \subset L^2(S^2) \quad (2.45)$$

where \mathcal{V}_l is spanned by the spherical harmonics with fixed l . The representations of $\text{SO}(3)$ on this subspace are the irreducible Wigner-D matrices.

To summarise, the group elements of $\text{SO}(3)$ can be represented on a vector space consisting of the set of all square-integrable functions on the unit sphere $L^2(S^2)$. In the spherical harmonic basis, the group elements are represented by the Wigner-D matrices, which describe how the basis functions transform under rotations depending on the order l .

2.3.5 Equivariant functions

When the positional inputs to ML force fields are rotated, intermediate feature vectors inside the network should transform in a predictable way, such that we can ensure the invariance of the final energy to rotations. This limits the learnable operations we are able to apply to the network inputs. For example, Cartesian coordinates can't be passed into generic transformer blocks without breaking equivariant. Here we define equivariant to motivate the introduction of learnable tensor products in the following sections.

Definition 2.3.4. A function $f : X \rightarrow Y$ is said to be equivariant with respect to the group G , if it commutes with the action of G , such that for all $g \in G$ and all

$x \in X$,

$$\rho_Y(g) \circ f(x) = f(\rho_X(g) \circ x), \quad \forall g \in G, \forall x \in X \quad (2.46)$$

where ρ_Y and ρ_X are the representations of a group element g on the vector space X and Y respectively.

$$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ \rho_X(g) \downarrow & & \downarrow \rho_Y(g) \\ X & \xrightarrow{f} & Y \end{array}$$

The learnable tensor product is the key ingredient in constructing multivariate functions that transform equivariantly with rotations.

2.3.6 Tensor products

We now define the tensor product of two spherical tensors, investigate a sensible change in basis and relate it to the irreducible representations of $SO(3)$.

Definition 2.3.5. *The tensor product of two vector spaces V and W is a vector space $V \otimes W$ which has a bilinear map that maps a pair of (v, w) , $v \in V, w \in W$ to an element of $V \otimes W$.*

Let $v_{l_1 m_1}$ and $w_{l_2 m_2}$ be two spherical tensors. Their raw outer product is given by,

$$k_{l_1 m_1 l_2 m_2} = v_{l_1 m_1} w_{l_2 m_2}, \quad (2.47)$$

where $k \in V \otimes W$ contains all possible pairwise combinations of the components of \mathbf{v} and \mathbf{w} . Under a rotation R , the resulting vector transforms as the product representation

$$\rho(R)v \otimes \rho(R)w = (\rho(R) \otimes \rho(R))(v \otimes w), \quad (2.48)$$

Hence, the rotation of $k_{l_1 m_1 l_2 m_2}$ results in,

$$k'_{l_1 m_1 l_2 m_2} = \sum_{m'_1 m'_2} D_{m'_1 m_1}^{l_1}(R) D_{m'_2 m_2}^{l_2}(R) k_{l_1 m'_1 l_2 m'_2}, \quad (2.49)$$

where the Kronecker product of the Wigner-D matrices $D^{(l)} \otimes D^{(l')}$, is a $(2l_1 + 1)(2l_2 + 1)$ dimensional matrix representation of $SO(3)$ on the $V \otimes W$ vector space. These can

be reduced through a change of basis to the irreducible representations with degree and order l and m , such that,

$$A_{lm,l_1l_2} = \sum_{m_1m_2} C_{l_1m_1l_2m_2}^{lm} v_{l_1m_1} w_{l_2m_2}, \quad (2.50)$$

where $C_{l_1m_1l_2m_2}^{lm}$ are the Clebsch-Gordan coefficients, which are only non-zero for $|l_1 - l_2| \leq l \leq l_1 + l_2$ and $m_1 + m_2 = m$. The resulting tensor A_{lm,l_1l_2} transforms under rotations as $D_{m'm}^l(R)$. The change of basis block diagonalises $D^{(l)} \otimes D^{(l')}$ into irreducible representations. Both \mathbf{A} and \mathbf{k} contain the same information just expressed in a different basis. The tensor product of two irreducible representations can hence be decomposed into a direct sum of irreducible representations through a change of basis. The indices l_1, l_2 indicate that there are multiple ways to obtain tensors in the subspace V_l . For example $l = 1$ tensors may originate from both $l_1 = 1, l_2 = 0$ and $l_1 = 1, l_2 = 1$. In the machine learning literature, these are referred to as different paths. To distinguish different paths η we introduce the notation

$$A_{lm\eta} := A_{lm,l_1l_2}, \quad (2.51)$$

where η enumerates all possible paths that lead to a tensor of degree l .

2.3.7 Learnable tensor products

The resulting tensor $A_{lm\eta}$ can be scaled independently for every l and η without changing how the tensors transform under rotations. We hence define a learnable tensor product as,

$$A_{lm} = \sum_{l_1l_2} W_{l_1l_2}^l \sum_{m_1m_2} C_{l_1m_1l_2m_2}^{lm} v_{l_1m_1} w_{l_2m_2}, \quad (2.52)$$

where $W_{l_1l_2}^l$ are independent weights that can be optimised. Note that the weights \mathbf{W} can be non-linear functions of arbitrary scalar values, $\boldsymbol{\alpha}$, without compromising equivariance,

$$A_{lm} = (\mathbf{v} \overset{\boldsymbol{\alpha}}{\otimes} \mathbf{w})_{lm} := \sum_{l_1m_1,l_2m_2} C_{l_1m_1l_2m_2}^{lm} F_{l_1l_2}^l(\boldsymbol{\alpha}) v_{l_1m_1}^{m_1} w_{l_2m_2}^{m_2}, \quad (2.53)$$

where $F_{l_1l_2}^l$ is a learnable non-linear function, $\boldsymbol{\alpha}$ are a list of scalar features, and we have defined the learnable tensor products as $\overset{\boldsymbol{\alpha}}{\otimes}$. Since the scalar features are invariant under group actions, any function $F(\boldsymbol{\alpha})$ remains invariant. The function could, for example, be a multi-layer perceptron or transformer. The approach of

mapping the tensor product of two spherical tensors back to a spherical basis hence facilitates the construction of learnable multi-variate functions.

Channels and tensor products

In a deep learning context, we often deal with a set of spherical tensors. Consequently, in addition to the spherical l, m -indices, tensors may have a so called channel dimension k . In the atomistic context, these may encode different chemical elements, charge or other properties. Given two vectors $v_{k_1, l_1 m_1}$ and $w_{k_2, l_2 m_2}$ with channel dimensions k_1 and k_2 , a fully connected learnable tensor product is defined as,

$$A_{k, lm}^{\text{FC}} = \sum_{l_1 l_2} W_{k_1 k_2, l_1 l_2}^{k, l} \sum_{m_1 m_2} C_{l_1 m_1 l_2 m_2}^{lm} v_{k_1, l_1 m_1} w_{k_2, l_2 m_2}, \quad (2.54)$$

where $W_{k_1 k_2, l_1 l_2}^{k, l}$ can be a non-linear function of any scalar value as previously in Equation (2.53).

To avoid the size of the weights scaling as $k_1 \times k_2 \times k$, the operations are often separated by channels. Assuming the same channel dimensions for both vectors, this corresponds to,

$$A'_{k, lm} = \sum_{l_1 l_2} W_{k, l_1 l_2}^l \sum_{m_1 m_2} C_{l_1 m_1 l_2 m_2}^{lm} v_{k, l_1 m_1} w_{k, l_2 m_2}, \quad (2.55)$$

where the tensor products are completely independent across channels. To reintroduce some mixing, one can apply a linear transformation to obtain,

$$A_{k, lm}^{\text{CP}} = \sum_k W_k^k A'_{k', lm}, \quad (2.56)$$

which provides an interesting trade-off between expressivity and computational cost. This is a form of tensor decompositions as outlined in Section 2.4.4 on the TRACE architecture[50]. Note that operations across channels can be efficiently parallelised and batched on GPUs.

2.4 Machine Learning Force Fields

We now shift focus from learning general equivariant functions to the concrete example of learning properties of atomistic point clouds. There has been significant development in learning quantum mechanical potential energy surfaces using machine learning. Frequently, inductive biases, such as the smoothness and equivariance

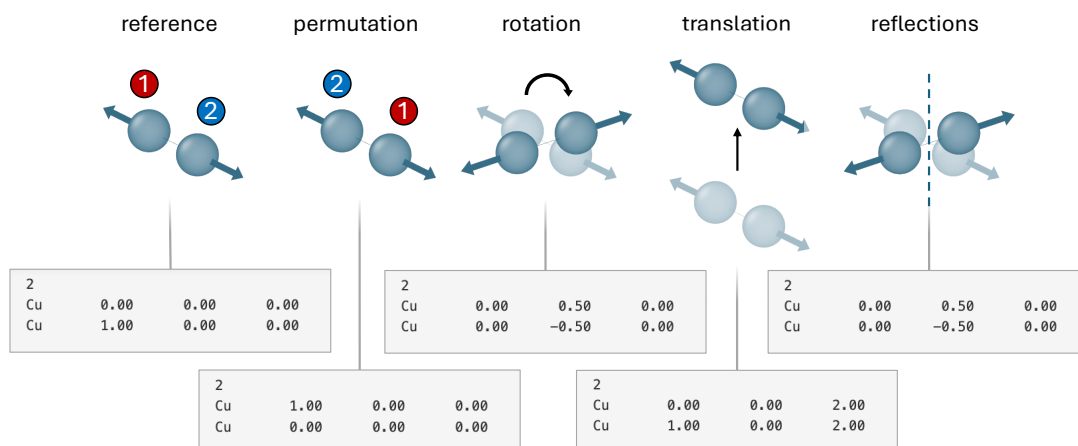


Figure 2.2 Machine learning force fields and relevant symmetries. The aim of machine learning force fields is to learn functions from 3D atomic point clouds to certain properties, in such a way that the outputs transform equivariant with respect to certain transformations. The figure shows the action of various transformations pictorially on a two atom dimer (top) and the effect on a list of 3D positions. The transformation include permutations of like atoms, 3D rotations, translations and inversions.

of the potential energy surface (PES), are directly incorporated into the architecture. A review of all relevant machine learning force field (MLFF) architectures is beyond the scope of this thesis. In this chapter, we focus on the developments which led to the MACE framework. At the time of writing this thesis, it is the state-of-the-art architecture for learning local $E(3)$ equivariant functions with energy-conserving forces.

2.4.1 General approach

There are various approaches to regressing properties of three-dimensional point clouds. Providing a comprehensive review of all relevant architectures is beyond the scope of the thesis. As such, we restrict ourselves to local machine learning approaches, which are extensive and, as such, can be trained on small configurations and predict features for systems of arbitrary size. We first quickly outline the different approaches to learning equivariant functions and then focus on the specific architectures which are relevant for the thesis in the context of catalysis (Chapter 3) and CO_2 capture (Chapter 5) as well as those on which the proposed long-range architecture (Chapter 6) are built.

The core aim of machine learning force fields is to learn functions on 3D point clouds that are equivariant with respect to permutations of like atoms, rotations,

translations and reflections. Figure 2.2 illustrates how these various transformations affect the list of atom positions. While the energy is constant under all transformations, they lead to a change in the atom positions. This illustrates, why it is not appropriate to feed lists of 3D positions directly into universal approximators, such as a multi-layer perceptron. The success of MLFF has led to the proliferation of various architectures. These different ML architectures can be categorised by the order in which they ensure equivariance to the relevant symmetries [51]. In this section, we show the evolution of architectures that ensure permutational, then translation, and lastly, rotational equivariance.

Permutational equivariance

To incorporate permutational equivariance, we move away from representing atomic environments as ordered lists to representing atom configurations as atom density fields. More specifically, we define a function $\rho : \mathbb{R}^3 \rightarrow R$ of N atoms as,

$$\rho(\mathbf{r}) = \sum_i^N g(\mathbf{r} - \mathbf{r}_i), \quad (2.57)$$

where g is some local function like a Gaussian or Dirac-delta function and r_i is the nuclear position of the i th atom. As the summation operation is invariant with respect to permutation, the atom density is also permutationally invariant.

Locality approximation

One key approximation made by many models is that the total energy of an atomic system can be expressed as the sum of local energies. Assuming a system of N atoms, the total energy E is a function of all positions r_k and elements z_k in the structure. Under the locality approximation, this total energy is expressed as a sum of atom energies ϵ_i ,

$$E(\{\mathbf{r}_k, z_k\}_{k \leq N}) = \sum_i^N \epsilon_i(\{\mathbf{r}_j, z_j\}_{j \in \mathcal{N}(i)}), \quad (2.58)$$

where ϵ_i is only a function of the positions and element within a neighbourhood $\mathcal{N}(i)$. The atom-centered energy can then be restricted to depend only on atoms within a certain cut-off distance. This is an approach that ML force fields inherit from many empirical force field approaches.

Note that this approximation has many consequences. Firstly, it allows the computational cost to scale linearly with system size, compared with N^3 for DFT methods or N^7 for coupled cluster approaches. Furthermore, it allows the application

of a model trained on N atoms to be applicable to a system of M atoms where $M \neq N$. However, there is no guarantee that the underlying ground truth quantum mechanical potential energy surface follows this locality approximation. Various papers investigate the locality of various reference methods[52]. For some simplified semi-empirical methods, such as tight-binding models, it can be shown that non-local interactions decay exponentially[52]. Empirical evidence suggests that this approximation is valid for many systems[31, 53, 54]. Addressing the limitations of local models is the main focus of [Chapter 6](#), which introduces an approach to capture non-local interactions in a cost-effective manner.

Translational equivariance

It is possible to shift the reference frame of our coordinate system to ensure translation equivariance. For global translations, one could, for example, shift the origin to the centre of mass of the system. In the context of regressing functions of local environments, we can shift the reference frame to the position of atom i , such that the local energy is a function of local displacements only

$$\epsilon_i(\rho(\mathbf{r} - \mathbf{r}_i)), |\mathbf{r} - \mathbf{r}_i| < R_{\max}, \quad (2.59)$$

where R_{\max} is a chosen cut-off distance [23].

Rotational equivariance

To ensure rotational equivariance, many architectures leverage the basis of irreducible representations of the $SO(3)$ group. As detailed in [Section 2.3](#) these can be used to construct invariant and equivariant functions that transform predictably under rotations. More details will be provided in the descriptions of the detailed architectures.

Recently, approaches have also become popular that rotate the atomic environment into a relevant reference frame [55]. The challenge is to create a reference frame that depends only on the local environment and remains smooth in response to changes in the local environment [56].

Other approaches

One architecture worthy of mentioning that does not first construct permutationally invariant and then rotationally equivariant features are the Behler-Parrinello neural network potentials. These are arguably the first serious attempt at creating surrogate

models for quantum chemistry. The first step is to represent the atomic structures with translationally invariant internal coordinates. The resulting descriptors are averaged over permutations resulting in the so called atom centred symmetry functions [1]. The computational cost of constructing permutationally invariant symmetry functions grows with the body order. To construct functions of angles and distances, these functions need to be summed over all possible atomic pairs and triplets within a local environment.

2.4.2 Gaussian Approximation Potentials (GAP)

A prevalent approach is the GAP framework[2, 57], which is used throughout [Chapter 3](#). First, smooth rotationally invariant descriptors, called smooth overlap of atomic positions (SOAP) are constructed [58]. An arbitrary architecture could then be used to predict invariant features such as the energy, but most commonly, kernel regression is used.

The SOAP descriptor is constructed for every atom i based on its neighbouring atoms j . First, we shift the positions of the atoms to the origin

$$\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i. \quad (2.60)$$

The atom density can then be constructed for every chemical element z .

$$\rho^{i,z}(\mathbf{r}) = \sum_{j \in \mathcal{N}(i)} \delta_{zz_j} \exp\left(-\frac{|\mathbf{r} - \mathbf{r}_{ij}|^2}{2\sigma_z^2}\right) f_{\text{cut}}(r_{ij}), \quad (2.61)$$

such that the sum goes over all atoms j within the neighbourhood \mathcal{N} of i . The cut-off function f_{cut} , smoothly goes to zero as the displacement $|\mathbf{r}_{ij}|$ approaches R_{max} . This ensures that the potential energy is smooth as atoms enter and exit the cut-off distance. The use of Gaussian smoothing with element-dependent variance σ_z is motivated by the observation that the Born-Oppenheimer potential energy surface is smooth.

The permutational invariant atom density ρ is then expanded in a set of radial $R_n(r)$ and angular basis functions Y_{lm} ,

$$\rho^{i,a}(\mathbf{r}) = \sum_{nlm} c_{nlm}^{i,a} R_n(r) Y_l^m(\hat{\mathbf{r}}), \quad (2.62)$$

where the coefficients are obtained from,

$$c_{nlm}^{i,a} = \int d\mathbf{r} R_n(r) Y_l^m(\hat{\mathbf{r}})^* \rho^{i,a}(\mathbf{r}) \quad (2.63)$$

To construct a higher-order descriptor one can take a tensor product between coefficients $c_{nlm}^{i,a}$,

$$P_{nn'l_3m_3}^{i,aa'} = \sum_{l_1m_1, l_2m_2} C_{l_1m_1, l_2m_2}^{l_3m_3} c_{nl_1m_1}^{i,a} c_{n'l_2m_2}^{i,a'} \quad (2.64)$$

which is the tensor product introduced in Equation (2.50) using different variable names to coincide with the original GAP paper [2] and where, as previously, $C_{l_1m_1, l_2m_2}^{l_3m_3}$ are the Clebsch-Gordan coefficients and map the output to the basis for the irreducible representation of SO(3). Note that we also take an outer product across the radial basis n and chemical element a , meaning the tensor p scales quadratically with respect to the number of radial basis functions and the number of elements.

For the application of force fields, the task is to regress energies, and hence we are only interested in the invariant part ($l_3 = 0$). For $l_3 = 0$ all Clebsch-Gordan coefficients are zero except when $l_1 = l_2$. The above equation can hence be expressed in terms of only one l -index. As the Clebsch-Gordan coefficients are trivial for this combination, one can explicitly include them in the equation for the power spectrum, resulting in the expression for the power spectrum commonly found in literature,

$$P_{nn'l}^{i,aa'} = \frac{1}{\sqrt{2l+1}} \sum_m (c_{nlm}^{i,a})^* c_{n'lm}^{i,a'}. \quad (2.65)$$

The resulting invariant tensor was introduced as the SOAP (smooth overlap of atomic positions) descriptor [58]. This description of the atomic environment can now be used as an input to an arbitrary machine learning architecture to learn invariant functions of 3D-point clouds.

Historically, the most commonly used architecture has been kernel regression. In particular, the use of low-order polynomial kernels is popular [54, 59, 60]. Given two atom densities a polynomial kernel of order ζ is,

$$k(\mathbf{p}, \mathbf{p}') = \left(\sum_{\alpha\beta} \sum_{nn'l} P_{nn'l}^{\alpha\beta} P_{nn'l}'^{\alpha\beta} \right)^\zeta \quad (2.66)$$

The resulting kernel can then be used to regress the total energies obtained by DFT using Gaussian process regression. The local energy is hence given by,

$$\epsilon_i(\rho) = \sum_m c_m k(\mathbf{p}, \mathbf{p}_m), \quad (2.67)$$

where c_m are learnable parameters and \mathbf{p}_m is a representative set of possible power spectra. This is the architecture used in [Chapter 3](#) to explore the use of ML force fields for heterogenous catalysis.

Body order and completeness

One measure of how expressive various architectures are is to categorise them by their body order. As seen previously, empirical force fields often contain terms that depend on bond distances (2-body), angles (3-body), and dihedrals (4-body) terms. In general, a multivariate function E has a body-order expansion such that,

$$\begin{aligned} E(\{r_i\}_{i=1}^N) &= e^{(0)} + \sum_{i=1}^N e^{(1)}(r_i) + \sum_{i_1 < i_2}^N e^{(2)}(r_{i_1}, r_{i_2}) + \\ &\dots + \sum_{i_1 < \dots < i_N}^N e^{(N)}(r_{i_1}, r_{i_2}, \dots, r_{i_N}), \end{aligned} \quad (2.68)$$

where the sums ensure permutational invariance and $e^{(n)}$ explicitly defines the n^{th} body order term. It is possible to approximate the multivariate function by truncating the body-order expansion at a maximum N . This approach is used in the ACE framework, which will be introduced shortly. The body order of an approximation \tilde{E} to the entire expansion is defined as the highest derivative order \mathcal{T} such that,

$$\frac{\partial^{\mathcal{T}} \mathbf{m}^{(t)}}{\partial r_{i_1} \dots \partial r_{i_{\mathcal{T}}}} \neq 0, \quad i_k \neq i_l \text{ for all } k \neq l, \quad (2.69)$$

while all higher-order derivatives are zero. Note that for any non-linear function σ applied to a body ordered terms, such as

$$\sigma \left(\sum_{i_1 < \dots < i_N} e^{(N)}(r_{i_1}, r_{i_2}, \dots, r_{i_N}) \right), \quad \text{for } N > 2 \quad (2.70)$$

the body order, is technically infinite as the Taylor expansion of σ is infinite. However, not all coefficients in the body-order expansion of [Equation \(2.68\)](#) independent. Models which explicitly construct the expansion of [Equation \(2.68\)](#) are

referred to as body-ordered and are hence restricted to linear activation functions σ [61].

In the context of GAP, the power spectrum p in Equation (2.65) is body ordered and goes to the maximum body order of 3. When using a polynomial kernel order ζ , the energy function is no longer body-ordered and has terms up to body order 3ζ in its expansion. Changes in dihedral angles, which conserve all distances and angles in a structure, cannot be captured by the power spectrum of the SOAP ($\zeta = 1$) descriptor. As a consequence, different atomic structures can be mapped to the same descriptor. Reference [62] shows that it is trivial to find such example structures.

2.4.3 Atomic Cluster Expansion (ACE)

An approach which remains body-ordered but extends to higher body orders is the ACE framework. Reference [25] introduces the Atomic Cluster Expansion, which can be used to construct polynomial basis functions for functions that are invariant with respect to rotations, reflections and permutations. The ACE framework can be used to construct equivariant functions of arbitrary body order systematically.

The atom-centred energy ϵ_i can be written as a many-body expansion in terms of the relative atomic positions $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i$ and the atomic elements z_i of atoms in the neighbourhood \mathcal{N} ,

$$\begin{aligned} \epsilon_i(\{\mathbf{r}_j, z_j\}_{j \in \mathcal{N}(i)}) &= V_0(z_i) + \sum_{j_1} V_1(\mathbf{r}_{ij_1}, z_{j_1}; z_i) \\ &+ \sum_{j_1 < j_2} V_2(\mathbf{r}_{ij_1}, z_{j_1}, \mathbf{r}_{ij_2}, z_{j_2}; z_i) + \dots \\ &+ \sum_{j_1 < \dots < j_N} V_N(\mathbf{r}_{ij_1}, z_{j_1}, \dots, \mathbf{r}_{ij_N}, z_{j_N}; z_i), \end{aligned} \quad (2.71)$$

where V_n , denotes the $(n+1)$ -body contribution involving atom i and n distinct neighbours and the sums run over all n -tuples of neighbours. The combinatorial scaling of the sums with body order makes this approach computationally expensive. Rather than expanding the atom energy as a many-body expansion ϵ_i in the ACE framework is given by,

$$\begin{aligned} \epsilon_i(\{\mathbf{r}_j, z_j\}_{j \in \mathcal{N}(i)}) &= U_0(z_i) + \sum_{j_1} U_1(\mathbf{r}_{ij_1}, z_{j_1}; z_i) \\ &+ \sum_{j_1, j_2} U_2(\mathbf{r}_{ij_1}, z_{j_1}, \mathbf{r}_{ij_2}, z_{j_2}; z_i) + \dots \\ &+ \sum_{j_1, \dots, j_N} U_N(\mathbf{r}_{ij_1}, z_{j_1}, \dots, \mathbf{r}_{ij_N}, z_{j_N}; z_i), \end{aligned} \quad (2.72)$$

where the summations are not restricted to unique atom pairs $j_N < j_{N+1}$ but go over all possible combinations j_N, j_{N+1} , including unphysical self-interacting terms $j_N = j_{N+1}$. This choice makes it possible to express U_N in a tensor product polynomial basis, where first, a two-body basis is created, and all higher body-order basis functions can be expressed as tensor products of the two-body basis. As a consequence, the cost of constructing the basis scales linearly with the number of neighbours of an atom, rather than combinatorially as the body-order expansion in Equation (2.71).

We now detail the construction of the ACE basis. First, as with the SOAP descriptor, the relative positions r_{ij} are expanded in a radial and angular basis,

$$\phi_{znlm}(\mathbf{r}_{ij}, z_i, z_j) = R_n(r_{ij}, z_i, z_j) Y_l^m(\hat{\mathbf{r}}_{ij}) \delta_{zz_j} \quad (2.73)$$

Note that here we use the same notation as in the previous section when introducing the SOAP vector. One difference here is that there is no Gaussian smoothing. The one-particle basis is then summed over all neighbours of atom i ensuring permutational invariance,

$$A_{i,znlm} = \sum_{j \in \mathcal{N}(i)} \phi_{znlm}(\mathbf{r}_{ij}, z_j) \quad (2.74)$$

The A basis remains a two-body basis. To construct higher-order features, we need to tensor product the A basis with itself,

$$\mathbf{A}_{i,znlm}^\nu = \prod_{t=1}^{\otimes \nu} A_{i,z_t n_t l_t m_t} \quad (2.75)$$

where \otimes indicates a tensor product, resulting in features of body order $\nu + 1$. The bold indices here correspond to tuples of previous indices, for example, $\mathbf{l} = (l_1, \dots, l_\nu)$. The basis A is not a basis of a reducible representation of $\text{SO}(3)$. As detailed in Section 2.3 on group theoretic background we hence perform a change of basis,

$$\mathbf{B}_{i,\mathbf{zn}}^{LM\eta\nu} = \sum_{\mathbf{m}} C_{\mathbf{lm}}^{LM\eta\nu} \mathbf{A}_{i,\mathbf{znlm}}^\nu \quad (2.76)$$

where $C_{\mathbf{lm}}^{LM\eta\nu}$ are the generalised Clebsch-Gordan coefficients and η enumerates the possible paths that lead to degree L features as introduced Section 2.3 on the group theoretic background. In the limit of infinite l, ν and n the ACE framework can be used to construct a complete linear basis. The B -basis can now be used to learn equivariant functions, for example using linear models or Finnis-Sinclair type models.

The atom-centred energies are given as,

$$\epsilon_i = \sum_{\mathbf{zn}\eta\nu} W_{\mathbf{zn}}^{\eta\nu} \mathbf{B}_{i,\mathbf{zn}}^{L=0,\eta\nu}, \quad (2.77)$$

where the weights $W_{\mathbf{zn}}^{\eta\nu}$ can be regressed using Bayesian ridge regression or regularised least squares. Note that the self-interaction terms can be removed from the basis[63].

2.4.4 Tensor sketching and TRACE

One of the key limitations of the ACE basis is its combinatorial scaling with the number of radial basis functions and the number of elements. This makes the approach infeasible for fitting diverse datasets, such as the Materials Project. In the ACE framework, the standard approach is to construct functions for each element separately and then take tensor products between them. One of the key innovations that has led to the fitting of foundation models is that it is possible to sketch the tensor products and thereby avoid the combinatorial scaling [27, 50, 61].

The core idea is to embed the chemical elements and radial basis functions into a fixed-dimensional space. For a full treatment of possible ways to reduce the combinatorial scaling, please see Reference [50]. Let us focus on one way of tensor sketching, such that the architecture scales linearly with the number of radial basis functions n_{\max} and embedding dimension k_{\max} . First, the single-particle basis ϕ from Equation (2.73) is embedded in a channel dimension k , such that

$$\phi_{klm}(\mathbf{r}_{ij}, z_i, z_j) = \sum_{nz} W_{zn}^k \phi_{znlm}(\mathbf{r}_{ij}, z_i, z_j), \quad (2.78)$$

where W_{zn}^k mixes the radial and element dimensions to a single channel dimension. This embedding can be learnt. The remaining part of the architecture remains the same, except that the channel dimensions are not "coupled" to one another. Previously, when the body order of the descriptor was increased by taking tensor products of the vector space spanned by the A -basis, outer products were also taken across z and n . Here, the B -basis is constructed for each k -channel individually,

$$A_{i,klm} = \sum_{j \in \mathcal{N}(i)} \phi_{klm}(\mathbf{r}_{ij}, z_j), \quad \mathbf{A}_{i,klm}^\nu = \prod_{t=1}^{\otimes \nu} A_{i,kl_t m_t}, \quad \mathbf{B}_{i,k}^{L\eta\nu} = \sum_{\mathbf{m}} C_{\mathbf{lm}}^{L\eta} \mathbf{A}_{i,klm}^\nu, \quad (2.79)$$

Only at the final regression step are the channels "mixed" when predicting the final output.

$$\epsilon_i = \sum_{k\eta\nu} W_k^{\eta\nu} \mathbf{B}_{i,k}^{L=0,\eta\nu}, \quad (2.80)$$

Clearly, the weights as well as the B -basis scale linearly with channel dimension. This flavour of the model was shown to retain significant performance, even when the embedding weights were initialised randomly and not learnt [50]. This allows the architecture to learn more transferability between elements and is a key ingredient in equivariant graph neural networks.

2.4.5 Equivariant message passing neural networks

Current state-of-the-art methods are based on many-body equivariant message-passing neural network architectures [26, 27, 64]. In this section, we first introduce the language of message passing and then focus on how atomic point clouds can be represented by a graph embedded in three dimensions. We then focus on developing equivariant graph neural networks. Finally, we provide some background on MACE [27], which can be seen as a fusion of the ACE basis with message passing and is used throughout Chapters [Chapter 4-6](#). In this section, we will use the same variable names and nomenclature as the original Mace paper [27].

Message passing

Many machine learning tasks can be approached by representing data on a graph. Let a graph G be defined as a collection of nodes V and edges E ,

$$G = (V, E), \quad (2.81)$$

where for each node $v_i \in V$ there exists a feature vector $h_i^{(0)}$, and for each edge $(v_i, v_j) \in E$ there is an associated edge feature e_{ij} . One of the most prevalent architectures for learning node-permutation equivariant functions on graph structures are message-passing neural networks[65–67]. At each layer, nodes are able to pass a message to nodes with which they share an edge. MPNNs can be thought of as a generalisation of convolutional neural networks to graph structures, where after t layers, information from node i can propagate to its t nearest neighbours.

In the context of machine learning force fields, an atomic configuration can be represented by 3D graphs, where the nodes represent nuclei and edges connect atoms that lie within a certain cut-off distance from each other. This choice in architecture has been successful due to the empirical observation that energies and forces are

local in many systems. We can think of all previously mentioned methods, including GAP and ACE, as single-layer message-passing architectures.

Using the same notation as MACE [27], we represent the state of the i th node at layer t of the MPNN by a tuple

$$\sigma_i^{(t)} = (\mathbf{r}_i, z_i, \mathbf{h}_i^{(t)}), \quad (2.82)$$

where $\mathbf{r}_i \in \mathbb{R}^3$ is the position of atom i , z_i the chemical element, and $\mathbf{h}_i^{(t)}$ are its learnable features. Throughout the MPNN, the learnable node features are updated while the properties r_i z_i stay constant.

There are parts to the network: the *message passing* phase, where nodes exchange information and their features are updated, and the *readout* phase, where the node features are combined to predict node or graph-level properties.

At each layer, the node feature is updated,

$$\mathbf{h}_i^{(t+1)} = U_t(\sigma_i^{(t)}, \mathbf{m}_i^{(t)}), \quad (2.83)$$

based on a learnable update function U_t and message $\mathbf{m}_i^{(t)}$. The message depends on the state of its neighbours,

$$\mathbf{m}_i^{(t)} = \bigoplus_{j \in \mathcal{N}(i)} M_t(\sigma_i^{(t)}, \sigma_j^{(t)}), \quad (2.84)$$

where M_t is a learnable message function, $\mathcal{N}(i)$ are all nodes which share an edge with node i , and $\bigoplus_{j \in \mathcal{N}(i)}$ is a permutation invariant pooling operation. Commonly, the pooling operation is simply a sum, but in principle, it can also be learnable [68].

The final step is the readout, where, after T node updates, the learnable readout functions \mathcal{R}_t map the node states $\sigma_i^{(t)}$ to the specified target. For MLFFs, this is simply the site energy of atom i ,

$$E_i = \sum_{t=1}^T \mathcal{R}_t(\sigma_i^{(t)}), \quad (2.85)$$

the sum of which gives the total energy. The message-passing structure turns out to help models be transferable.

2.4.6 Higher order message passing (MACE)

In the MACE architectures, messages are constructed using the ACE framework. Not only are the messages equivariant, but they also contain information about many-body interactions.

Message construction

To construct the message, the authors first build the ACE basis,

$$A_{i,kl_1m_1}^{(1)} = \sum_{j \in \mathcal{N}(i)} R_{kl_1}^{(1)}(r_{ji}) Y_{l_1}^{m_1}(\hat{\mathbf{r}}_{ji}) W_{kz_j}^{(1)}, \quad (2.86)$$

where the elements and radials are embedded into k dimensional channels and all other variables are defined as previously. As in the ACE architecture, the A basis can be tensor producted with itself to create the B basis,

$$\mathbf{B}_{i,\eta_\nu kLM}^{(t)} = \sum_{lm} \mathcal{C}_{\eta_\nu, lm}^{LM} \prod_{\xi=1}^{\nu} \sum_{\tilde{k}} w_{k\tilde{k}l\xi}^{(t)} A_{i,\tilde{k}l\xi m_\xi}^{(t)}, \quad \mathbf{lm} = (l_1 m_1, \dots, l_\nu m_\nu) \quad (2.87)$$

where ν is a given correlation order, $\mathcal{C}_{\eta_\nu}^{LM}$ corresponding to the generalised Clebsch-Gordan coefficients ensuring that $\mathbf{B}_{i,\eta_\nu kLM}^{(t)}$ transforms as a spherical tensor with degree L and the weights $w_{k\tilde{k}l\xi}^{(t)}$ mix the channels (k) of $\mathbf{A}_i^{(t)}$. As previously the η_ν enumerated all possible couplings of l_1, \dots, l_ν features that yield the selected spherical tensor degree L . The B basis is now used to create the messages,

$$m_{i,kLM}^{(t)} = \sum_{\nu} \sum_{\eta_\nu} W_{z_i kL, \eta_\nu}^{(t)} \mathbf{B}_{i,\eta_\nu kLM}^{(t)}, \quad (2.88)$$

where $W_{z_i kL, \eta_\nu}^{(t)}$ depends on the chemical element z_i of the receiving atom and message degree L .

Update

The node features are updated as a linear function of the message and the residual connection [69]:

$$h_{i,kLM}^{(t+1)} = U_t^{kL}(\sigma_i^{(t)}, \mathbf{m}_i^{(t)}) = \sum_{\tilde{k}} W_{kL, \tilde{k}}^{(t)} m_{i,\tilde{k}LM}^{(t)} + \sum_{\tilde{k}} W_{z_i kL, \tilde{k}}^{(t)} h_{i,\tilde{k}LM}^{(t)}, \quad (2.89)$$

Subsequent layers

For subsequent layers, the construction of the A-basis depends on the equivariant node features. Specifically, they are constructed using a learnable tensor product as,

$$A_{i,kl_3m_3}^{(t)} = \sum_{l_1 m_1, l_2 m_2} C_{l_1 m_1, l_2 m_2}^{l_3 m_3} \sum_{j \in \mathcal{N}(i)} R_{kl_1 l_2 l_3}^{(t)}(r_{ji}) Y_{l_1}^{m_1}(\hat{\mathbf{r}}_{ji}) \sum_{\tilde{k}} W_{k\tilde{k}l_2}^{(t)} h_{j,\tilde{k}l_2 m_2}^{(t)}, \quad (2.90)$$

where $C_{l_1 m_1, l_2 m_2}^{l_3 m_3}$ are the standard Clebsch-Gordan coefficients ensuring that $A_{i,kl_3 m_3}^{(t)}$ maintain the correct equivariance, r_{ji} is the (scalar) interatomic distance, and $\hat{\mathbf{r}}_{ji}$ is the corresponding unit vector. $R_{kl_1 l_2 l_3}^{(t)}$ is obtained by feeding a set of radial features that embed the radial distance r_{ji} using Bessel functions multiplied by a smooth polynomial cut-off to a multi-layer perceptron (MLP). The remaining architecture is the same as previously. Note that the MLP combines various elements and radials, allowing for significant weight sharing between vastly different atomic environments. In Chapter 3, we will see that this choice in architecture allows for transferable learning to entirely new elements with minimal additional data.

Readout

In the readout phase, the invariant part of the node features is mapped to site energies via readout functions,

$$E_i = E_i^{(0)} + E_i^{(1)} + \dots + E_i^{(T)}, \quad \text{where}$$

$$E_i^{(t)} = \mathcal{R}_t(\mathbf{h}_i^{(t)}) = \begin{cases} \sum_{\tilde{k}} W_{\text{readout},\tilde{k}}^{(t)} h_{i,\tilde{k}00}^{(t)} & \text{if } t < T \\ \text{MLP}_{\text{readout}}^{(t)}(\{h_{i,k00}^{(t)}\}_k) & \text{if } t = T \end{cases} \quad (2.91)$$

The readout only depends on the invariant features $h_{i,k00}^{(t)}$ to ensure that the total energy doesn't change with rotations. Note that the readout function is linear for all but the last layer, where it is a one-layer multi-layer perceptron.

Performance

The mace architecture was first published in 2022 and remains one of the most performant equivariant MLFFs. The innovation of message passing many-body features allows for a much smaller number of message-passing steps compared to architectures which only pass two body features [26, 27, 70]. Indeed, many of the published models only contain two message passing layers, striking a competitive trade-off between computational cost and performance, while Nequip [26] and Equiformer [71] models typically require 6 and 10 layers, respectively. The architecture has subsequently been trained on large datasets, resulting in one of the first foundation models for atomistic material chemistry [72].

Chapter 3

Active learning for accurate reaction rates in catalysis

We now focus on modelling catalysts with machine learning force fields (MLFFs). A major challenge in using MLFFs in catalysis is curating a relevant dataset that ensures reliable simulations. Consequently, researchers have relied on direct *ab-initio* calculations, such as geometry DFT optimisations [71] to collect relevant catalytic configurations. As an alternative, MLFFs can be trained in an iterative manner, where novel configurations are sampled from MLFF simulations and added back into the training set. The relevant geometries can be selected manually or based on explicit criteria, such as changes in density or bonding topology [1, 31, 64, 73–77]. While successful, this approach relies on human expert knowledge, inherently limited to a low-dimensional understanding of force field failures. More recently, active learning has been used to automate this process using a model’s uncertainty metric [78–85]. Note that for pure material or molecular systems, a common approach is to use empirical or semi-empirical methods to firstly sample relevant configurations. The generated samples can then be re-evaluated with highly accurate *ab-initio* methods [86] for training.

In this chapter, we explore how active learning can be used to train accurate MLFFs for modelling catalytic reactions. We introduce an automated training protocol that can be applied to any catalytic system with minimal human supervision. We validate the protocol on, CO₂ hydrogenation on Indium oxide, as a well-studied catalytic reaction to obtain accurate adsorption energies and reaction barriers [87]. The reaction barriers can then be used for downstream tasks to compute reaction rates. With the help of active learning, the final force field obtains energy barriers within 0.05 eV of Density Functional Theory. Thanks to the computational speed-

up, not only do we reduce the cost of routine in-silico catalytic tasks, but we also find an alternative path for the previously established rate-limiting step, with a 40% reduction in activation energy, changing our understanding of a well-explored reaction mechanism [87–99].

Furthermore, we illustrate the importance of finite-temperature effects and compute free energy barriers. After validating the protocol on the extensively explored hydrogenation of carbon dioxide to methanol over indium oxide, we demonstrate its transferability on the experimentally relevant, yet unexplored, top-layer reduced indium oxide surface. This chapter is based on work published in Reference [54]. The author of this thesis was responsible for training all the models presented in this chapter and exploring different training approaches to arrive at the automated protocol presented in this chapter.

3.1 Background and motivation

We begin this section with a brief outline of why catalysis is important to society and why, in the context of the green energy transition, the need for innovations in catalytic discovery is urgent. We also explore the current limitations in catalytic *ab-initio* modelling to motivate the need for faster but equally accurate methods.

3.1.1 Opportunities for catalysis

Heterogeneous catalysts play a central role in the world economy, from the agricultural industry to the production of garments, cosmetics and pharmaceuticals. Indeed, 80% of all manufactured goods rely on processes that are catalysed by heterogeneous catalysis [100]. Catalysts reduce the energy required for specific reactions to occur without being consumed during the reaction and remain unchanged except for slow degradation. A catalyst effect on energy barriers between elementary steps changes the entire reaction network. Consequently, not only the activity, how much of A is made per unit of time, but also the selectivity, how much of substance A is produced compared to substance B, is significantly impacted. A famous example of this is the Lindlar catalyst, which partially hydrogenates organic molecules [101]. Normally, when compounds containing carbon-carbon triple bonds $R - C \equiv C - R$ are hydrogenated under reductive conditions, the carbons gain two hydrogens $R - CH_2 - CH_2 - R$. However, in the presence of a palladium catalyst with deposited calcium carbonate, the molecule can be partially hydrogenated to $R - CH = CH - R$, without breaking the carbon-carbon double bond. Partial

hydrogenation reactions are critical in Vitamin A synthesis, for example. As such, catalysts not only reduce the heat and pressure conditions needed for a desired reaction to occur but also change the possible products entirely.

Many parts of everyday life are affected by the historic discovery of new catalysts. The Haber-Bosch process is responsible for the mass synthesis of ammonia needed for fertilisers [102]. Furthermore, catalytic converters in cars transform toxic exhaust gases into less toxic pollutants. Additionally, catalysts play a crucial role in renewable energy storage, particularly in the form of liquid fuels. Electro-catalysts are required to split water into hydrogen and oxygen for green hydrogen production [103, 104]. The hydrogen can then be used to hydrogenate captured CO₂, to produce methanol.

In comparison to hydrogen, methanol can be transported in the same way as crude oil and has a significantly higher energy density per volume [105]. Making catalysts more efficient is critical to making liquid renewable energy storage economically viable. When investigating automated workflows in the subsequent sections of this chapter, we take the production of methanol on indium oxide as a complex trial system.

One of the key reasons catalysis is currently receiving significant attention is our global transition away from fossil fuels [106]. For the chemical industry, the role of fossil fuels is twofold: firstly, they are used as a source of heat, and secondly, fossil fuels serve as direct chemical feedstocks and the source of hydrocarbons as reactants. Fossil fuels, hence, form the starting material for many fundamental chemicals. Replacing the dependents on non-renewable fuels requires an alternative renewable carbon source, such as biomass. Reactions starting from sugars and aromatics have not been optimised to the same extent as fossil fuels. This sudden change means that we need to discover new catalysts for critical reactions in only a few years, a process which took decades for the use of fossil fuels.

3.1.2 Challenges in computational catalysis

To accelerate discovery, one of the most promising techniques is rational design. Rather than relying on serendipitous discoveries of catalytic materials, we can use computational screening to accelerate the discovery process. Indeed catalysts have been designed based on computational trends [21, 22]. However, significant discrepancies remain between simplified, computationally modelled systems and true catalysts under operando conditions, potentially reducing the validity of computational screening. In this section, we outline some of the current challenges in *in-silico* catalytic design which are visually represented in [Figure 3.1](#).

Complex surface structure The utility of catalysts originates from the fact that the barriers of the elementary steps change significantly depending on the catalytic surface. Not only does the chemical composition matter, but the morphology and 3D arrangement can also have a strong impact. Significant surface changes, such as steps, edges, and vacancies, significantly increase the system size that needs to be simulated. Traditional *ab-initio* methods, however, are restricted to only significantly simplified surfaces.

Indeed, for many systems, it is challenging to determine the atomistic structure of a surface. Without the correct surface, calculations of barriers can be significantly different compared with the experiment. MLFFs can be used to determine surface structure. Timmermann et al. [107], for example, used long molecular dynamics simulations of simulated annealing to investigate the surface of iridium oxide. During simulated annealing, the temperature of the dynamics is set very high and then continuously decreased. Typically, the most energetically favourable surface for rutiles has miller index (110). The authors, however, found computationally that under experimental conditions, the (101) surface is energetically favourable due to a top-layer reconstruction. Independent of these computational investigations, the (101) facet was observed experimentally, validating the formation of complexions on the (101) surface. This illustrates the predictive ability of machine learning atomistic simulations.

Finite temperature effects We previously saw that it is often challenging to find relevant low-energy surfaces. This is significantly complicated when considering that under operando conditions, the surface is not static and dynamically changes. Dattila et al. used *ab-initio* molecular dynamics to investigate common copper surfaces, which constitutes one of the most explored catalytic surfaces both computationally and experimentally [108]. They found that the surface restructures with specific recurring motifs, which impact catalytic performance.

Accurate energy barriers To determine the rate of reaction steps, we can compute the *ab-initio* energy barriers. The rate of a reaction is related to the energy barrier by the Arrhenius equation,

$$k = Ae^{\Delta E_a/k_b T}, \quad (3.1)$$

where T is the temperature, k_b is the Boltzmann constant, ΔE_a is the energy barrier, also referred to as the activation energy and A is the pre-exponential factor [109, 110]. First-order approximations to the pre-exponential can be made using vibrational analysis of the initial and transition state [110]. As the rate depends on the




Figure 3.1 Challenges in computational catalysis. Accurately describing catalytic processes is challenging due to the compounding complexity of diverse surfaces, molecule-molecule interactions and the sheer size of the design space. Showing a simplified catalytic surface in the foreground surrounded by illustrations of the true complexity. The role of MLFFs is not limited to speeding up existing conventional computational methods but also enables a completely different approach to tackling these challenges.

exponential of the energy barrier, small errors in the potential energy surface can lead to significant errors in predicted reaction rates. Computing a single barrier to determine the activation energy E_a requires hundreds to thousands of single-point evaluations. For individual barriers, this is possible at *ab-initio* accuracy. However, especially for more complex surfaces, there are multiple reaction progressions and active sites. An active site is a location on a catalytic surface where the molecule can adsorb and react. This could be on top of a surface atom, in between two (bridge) or three surface atoms. Furthermore, various defects and vacancies can affect the barriers of reaction steps. In principle, we need to investigate all possible barriers at every active site for each reaction step. Due to the large computational cost, even of approximate methods like density functional theory (DFT), often only single or very few active sites are investigated. We will see the limitations of these approximations in [Chapter 3](#).

Accurate reaction rates The previously introduced Arrhenius equation originated from empirical observations. From statistical mechanical arguments, one can derive the Eyring-Polanyi equation, which has the same functional form,

$$k = \frac{\kappa k_b T}{h} e^{-\frac{\Delta G_a}{k_b T}} = \frac{\kappa k_b T}{h} e^{-\frac{\Delta H_a - \Delta S_a T}{k_b T}}, \quad (3.2)$$

where κ is the transmission coefficient, h is the Planck constant, ΔG is the Gibbs energy, ΔH is the enthalpy of activation, ΔS is the entropy of activation and all other variables are defined as previously [[60](#), [110](#)]. The transmission coefficient quantifies the fraction of the reactants which go directly to the product without recrossing the barrier. Under the transition state theory approximation, this is often assumed to be one. To obtain a more accurate rate, one would have to determine it through simulation. One option is starting molecular dynamics simulations at the transition state with different initial conditions and observing whether the system evolves to the reactant or product state.

The entropic contributions can be calculated using various approximations. A widely used assumption is the harmonic approximation, where a simple vibrational analysis can be used to determine the prefactor. Alternatively, we can calculate free energy differences explicitly. Free energy calculations, however, require on the order of 10^6 evaluations even for simple reactions. This is beyond the capabilities of direct *ab-initio* simulation. We show in [Chapter 3](#) that using MLFFs, we can determine free energy barriers using enhanced sampling approaches.

***ab-initio* accuracy** As the reaction rate depends exponentially on the calculated energy barriers, the accuracy of conclusions drawn from the computational modelling strongly relies on the precision of the underlying *ab-initio* method. Indeed, studies indicate that for different types of reactions, different reparametrisations of approximate (density functional theory) methods are optimal[111]. This makes it necessary to validate observations with more accurate and hence costly methods.

Reaction network Until now, we have discussed accurately determining the rate of individual elementary reactions. In reality, reactions occur in multiple steps, forming intermediate compounds before the final product is formed, which then desorbs from the surface. Furthermore, a given set of reactants can be converted to a multitude of products. A desired catalyst not only produces products at a high rate but also selectively produces only the desired compounds.

It is helpful to think of a reaction network as analogous to electric circuits with resistors. Mathematically, there are strong parallels between the current flow in a circuit and the turnover steady-state frequency of chemical conversion. This analogy, known as the energy span model [112], illustrates that reaction dynamics are complex and require a complete treatment of the entire network.

In electric circuits, the overall current flow depends on the entire network of wires and resistors. However, the current can be particularly sensitive to only a few components, which are not necessarily the largest resistors. Similarly, in a catalytic reaction, the overall efficiency of the entire reaction network depends on the rates of each intermediate step. Determining which intermediates and barriers dominate the resulting catalytic turnover frequency is a complex task that depends on the entire reaction network. The conversion frequency-determining reaction step, hence, does not necessarily correspond to the highest barrier. As the number of atoms of the reactants increases, the number of possible intermediate and paths grows exponentially [113]. This places strong constraints on the computational cost of a single-barrier calculation.

Surface kinetics Molecules on top of a surface are not stationary under operational conditions. A catalyst's efficiency is often strongly influenced by the availability of reactants on the surface, which is controlled by adsorption, desorption and surface diffusion rates. The effect of surface kinetics can be illustrated with the example of single-atom alloys, a class of catalysts where the active sites consist of isolated atoms. The ability for intermediates to diffuse on the surface significantly impacts

the feasibility of tandem reactions occurring at different active sites consisting of different isolated dopant atoms[22].

Coverage effects One factor that significantly affects surface kinetics is the interaction between adsorbed molecules. Famously carbon dioxide can occupy active sites and thereby "poison" the catalyst [114, 115]. The interaction between the two reactants at an active site is particularly important, determining if they effectively repel or attract each other [116].

Large design space A further combinatorial complexity is the large number of possible catalysts that could be investigated. Catalytic properties depend on the elemental composition of a surface, its atomic arrangement and the operating conditions. Effective catalysts range from single-element surfaces to high-entropy alloys. Moreover, the 3D structure of catalysts can be designed to improve selectivity for specific compounds. For example, the pore size of porous structures like zeolites and metal-organic frameworks can be tuned to increase product selectivity[117, 118]. When optimising catalytic performance, the search space scales combinatorially.

Summary

After this section, the reader may be overwhelmed by the seemingly impossible task of designing and improving catalysts. Indeed, the large design space is one of the main reasons computational property screening has received so much attention. This section illustrates that there is a strong need for accurate and computationally efficient atom scale simulation. Many of the above-mentioned challenges can be addressed with MLFFs, which provide quantum mechanical accuracy at orders of magnitude reduced cost. In this thesis, we show that the role of MLFFs is not limited to speeding up existing computational workflows but enables a completely different approach to investigating catalysts computationally.

3.1.3 The role of machine learning force fields

Machine learning accelerated *ab-initio* simulations are set to have a significant impact on computational catalysis. The reduced computational cost will allow modelling of more realistic systems and more systematic exploration reaction networks. At the start of this PhD there was minimal work on using MLFFs for catalysis. The investigations were restricted to simple reactions [119] or purely modelling the catalytic surface[107]. The main reason for this complication was the difficulty of

curating relevant training datasets. In this chapter, we tackle this problem by showing that active learning can be used to systematically curate training data, which results in stable and accurate MLFFs. Furthermore, we show that the models can be used to compute accurate energy barriers. We start by outlining how energy uncertainties can be computed from Gaussian process regression.

3.1.4 Uncertainty estimation

In this chapter, we use the Gaussian Approximation Potential (GAP) framework [2] to fit DFT energies and forces obtained using Quantum Espresso [120]. GAP is chosen due to its maturity, past success in describing a wide variety of systems, and because it provides an analytical uncertainty estimate [59, 73, 121]. As a many-body descriptor, we use the smooth overlap of atomic positions (SOAP). To assess the transferability of the training set, we employ MACE [27], a higher-order equivariant message-passing neural network that leverages recent advancements in graph neural networks for atomic simulations [26, 27, 122].

Various approaches exist to predict a model’s error for active-learning [84, 85, 123, 124]. The Gaussian Approximation Potential framework allows for a rigorous estimate of the local energy uncertainty from the underlying Gaussian Process Regression (GPR). In the GAP framework, the total energy of a configuration with N atoms is given by the sum of atomic energies,

$$E = \sum_i^N \epsilon(\boldsymbol{\rho}_i), \quad (3.3)$$

where $\boldsymbol{\rho}_i$ is a descriptor of an atom’s local atomic environment that only depends on the positions and elements of its neighbours within a specified distance cut-off r_{cut} . Thanks to this locality approximation, the computational cost of the force field scales linearly with the number of atoms.

To describe the many-body atomic environment $\boldsymbol{\rho}$ we use the smooth overlap of atomic positions (SOAP) descriptor, which is invariant with respect to rotations, permutations, translations, and reflections [58].

The local energy function ϵ is modelled using sparse GPR. To quantify the similarity between the two environments, we use a polynomial kernel,

$$k(\boldsymbol{\rho}_n, \boldsymbol{\rho}_m) = (\boldsymbol{\rho}_n \cdot \boldsymbol{\rho}_m)^\zeta, \quad (3.4)$$

where ζ is the polynomial degree and $\boldsymbol{\rho}$ are SOAP vectors. For a set of N and M local environments, we can define the kernel matrix $[\mathbf{K}_{\text{NM}}]_{nm} \equiv k(\boldsymbol{\rho}_n, \boldsymbol{\rho}_m)$ [53]. If one were to perform GPR conditioned on a training set of N local environments $\boldsymbol{\rho}_n$ with ground truth energies \mathbf{y} , the predicted energy distribution for unseen atomic environments $\boldsymbol{\rho}$ would be Gaussian

$$\epsilon(\boldsymbol{\rho}) \sim \mathcal{N}(\mu(\boldsymbol{\rho}), \Sigma^2(\boldsymbol{\rho})), \quad (3.5)$$

with mean and covariance,

$$\mu(\boldsymbol{\rho}) = \mathbf{k}_N^T [\mathbf{K}_{NN} + \sigma_n^{-2} \mathbb{1}]^{-1} \mathbf{y} \quad (3.6)$$

$$\Sigma^2(\boldsymbol{\rho}) = k - \mathbf{k}_N^T [\mathbf{K}_{NN} + \sigma_n^{-2} \mathbb{1}]^{-1} \mathbf{k}_N, \quad (3.7)$$

where k and \mathbf{k}_N are the kernels between the environment $\boldsymbol{\rho}$, itself and the training set $\boldsymbol{\rho}_n$ respectively. While the time complexity of obtaining the true mean μ and covariance Σ functions would scale cubically with training set size, they can be approximated with a sparse method. Rather than quantifying the similarity of novel environments with all N training data points, we use a sparse set of M environments, where $M \ll N$. Modifications to Equation (3.6) can be found by minimizing the Kullback-Leibler divergence between the sparse method and the full posterior in Equation (3.5) [125]. The approximate mean and covariance functions are,

$$\mu(\boldsymbol{\rho}) = \mathbf{k}_M^T \mathbf{C}_{MM} \mathbf{K}_{MN} \mathbf{y} \quad (3.8)$$

$$\Sigma^2(\boldsymbol{\rho}) = k - \mathbf{k}_M^T \mathbf{K}_{MM}^{-1} \mathbf{k}_M + \mathbf{k}_M^T \mathbf{C}_{MM} \mathbf{k}_M, \quad (3.9)$$

where

$$\mathbf{C}_{MM} = [\mathbf{K}_{MM} + \sigma_m^{-2} \mathbf{K}_{MN} \mathbf{K}_{MN}^T]^{-1}, \quad (3.10)$$

as found using variational inference [125]. The training cost of the resulting GAP model scales as $\mathcal{O}(NM^2)$. During simulation, we use the mean of the sparse posterior distribution, while the variance Σ^2 provides the local energy error estimate for active learning. As is commonly done within the GAP framework, we select a subset of M representative environments from the training set using CUR matrix decomposition [53].

Note that in addition to using the uncertainty to sample new configurations, it can be used to monitor the accuracy of a simulation without having to run DFT

calculations. If the predicted uncertainties are high, the simulation has reached a part of configuration space which is not represented in the training set. This consequently means that the resulting observables are less trustworthy. Indeed, uncertainty prediction can be used to propagate uncertainties to the observables. This is commonly done by producing a committee of models and running multiple simulations [126, 127]. In the case of a Gaussian process, regression committees can be sampled from the posterior distribution of the weights. Alternatively, when computing thermodynamic averages, individual samples can be reweighed based on their uncertainty [127]. However, in this chapter, we focus on the use of uncertainty for active learning.

3.2 Active learning for catalysis

When interrupting simulations for sampling it is useful to look at per-atom uncertainties. In a large simulation, high uncertainty configurations can occur in only a small fraction of the configuration, such as bonds breaking or atoms colliding. If we were to take the mean uncertainty over the entire configuration, such changes would be unobservable. Hence we employ a sampling strategy based on the local energy uncertainty of individual atoms [84, 128, 129]. When the uncertainty of any atom exceeds a predefined threshold, new configurations are sampled from the ongoing simulation. These new geometries are subsequently evaluated with Density Functional Theory (DFT) and incorporated into the training set. In cases where the threshold is not reached during the simulation, we sample either the configuration with the highest uncertainty (for molecular dynamics simulations) or the final configuration (for optimization tasks). The MLFF is then re-trained using the expanded dataset, and the aforementioned steps are repeated. This active learning loop continues until a desired level of accuracy is achieved.

Our specific active learning approach, as depicted in Figure 3.3b, incorporates two distinct stopping criteria to guide the automatic iterative training of the MLFF. The first criterion, referred to as the uncertainty threshold (σ_{thr}), determines when to interrupt a simulation and sample a configuration. Throughout the training protocol, the uncertainty threshold is set to 50 meV, a choice that is justified in Figure 3.2. The second criterion, referred to as the termination criterion, determines when to conclude the active learning loop and considers the MLFF to be sufficiently accurate for the intended task. The termination criteria are tailored to the specific simulation type (see Section 3.3) and can be determined based on the observed DFT-MLFF error in the sampled configurations. The exact termination criteria depend on the user's

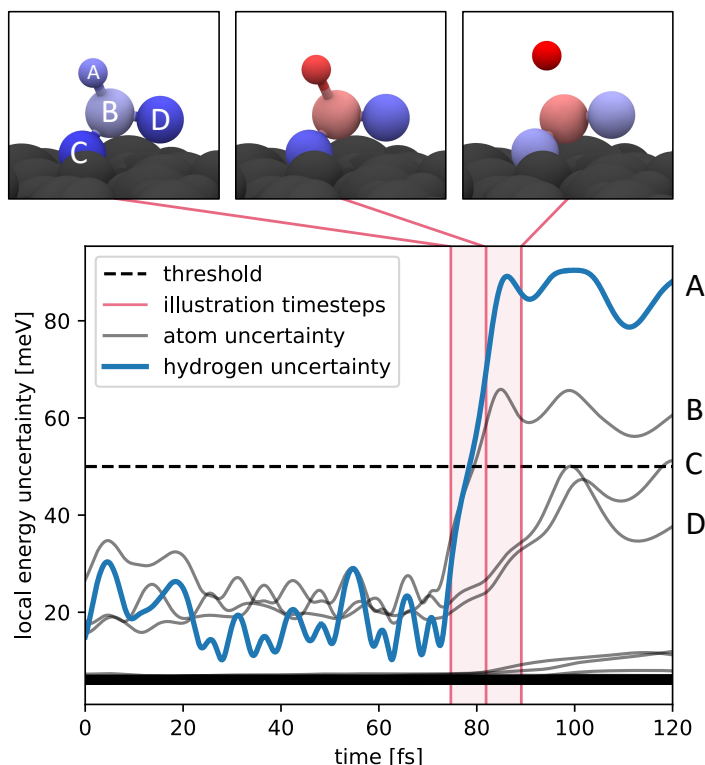


Figure 3.2 Predicted energy uncertainty during molecular dynamics (MD). This figure demonstrates the initial iteration of active learning with molecules, highlighting the occurrence of a non-physical bond break when the energy uncertainty surpasses a threshold of 50 meV. The top panel displays MD snapshots of formate (HCOO) on In₂O₃, where molecular atoms are colour-coded by the local energy uncertainty, while the surface atoms are depicted in grey.

desired level of accuracy. This criteria automates the entire workflow and allows the model to be trained without significant supervision.

In this manner, the model's accuracy at the desired task converges to the specified accuracy without prior knowledge of the reaction path energetics. Active learning reduces the accumulation of uninformative training data while systematically improving the model's accuracy in relevant parts of the potential energy surface. The method section contains more detail on how the underlying sparse Gaussian process provides a cheap metric for the model's predicted error.

We demonstrate the effectiveness of using local energy uncertainty to determine when to interrupt simulations by running molecular dynamics (MD) with an adsorbate on indium oxide. Figure 3.2 depicts the evolution of energy uncertainty for each atom during the MD. Notably, the four molecular atoms exhibit significantly higher uncertainty compared to the 79 bulk atoms. Throughout the MD trajectory, the local

uncertainties remain relatively low until the energy uncertainty of the hydrogen atom exceeds the predetermined threshold of 50 meV at 78 fs.

The substantial increase in uncertainty for the hydrogen atom suggests that the simulation has encountered an atomic environment that greatly differs from any seen in the training set. This indicates the need to sample a new configuration for the next iteration of the potential. It is noteworthy that only a few atoms exhibit a pronounced increase in uncertainty, highlighting the importance of monitoring individual atom energy errors rather than total errors. The specific threshold value becomes less critical due to the sharp rise in uncertainty observed. In this particular system, a threshold of 50 meV ensures that the sampled configurations remain chemically relevant and informative for further training iterations.

3.3 Automated training protocol

We now go into detail about how we can combine active learning of different simulation types to produce accurate force fields for computational catalysis. The training protocol consists of six active learning blocks which sample configurations from MD, geometry optimization, or nudge elastic band calculations. The flow chart in [Figure 3.3a](#) shows how these blocks link in series to obtain the final training set. The only inputs to the protocol are the bulk configuration of indium oxide and the set of approximate geometries corresponding to the intermediates along the reaction path. These intermediate geometries can come from force field minimization, manual placement, or adsorption prediction tools like CatKit.

Throughout the protocol, we only ever perform single-point calculations with DFT, while all dynamical simulations are performed with the MLFF. Below we outline the protocol's six active learning blocks and their termination criterion. Once a sampled configuration satisfies the termination criterion, the training protocol moves on to the next block. While we use very strict termination criteria, these will depend on the user's desired accuracy.

The first two blocks model the surface itself, while the remaining four are used to capture inter-molecular and molecule-surface interactions. The first four blocks consist of running active learning with MD to gain a stable potential, and the last two correspond to geometry optimization and nudge elastic band calculations to obtain accurate adsorption energies and barriers. We find that running active learning with MD before geometry optimization helps sample repulsive behaviour effectively. The number of configurations in the training set at each stage of the protocol is tabulated in [Appendix A Supplementary Note 1](#).

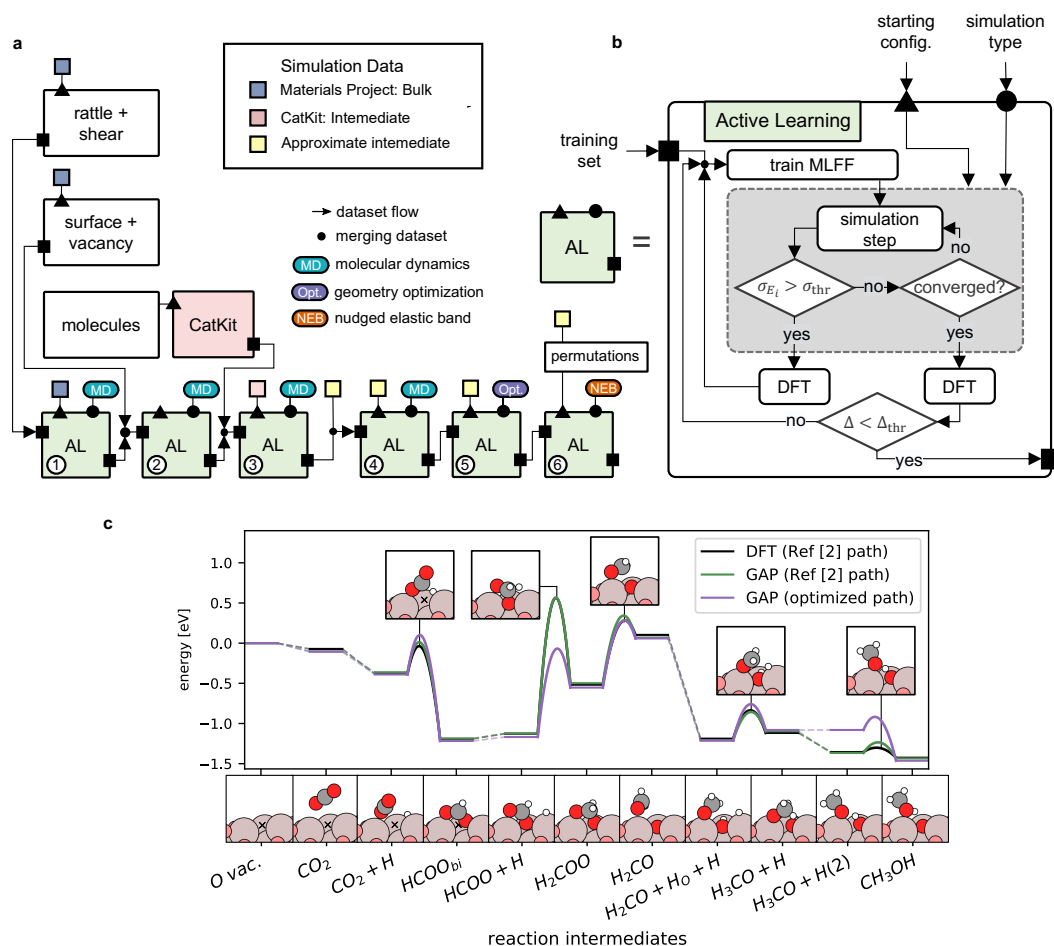


Figure 3.3 Automated training protocol. A flow chart (a) of the six active learning (AL) blocks that make up the protocol. Each active learning block (b) iteratively samples new training configurations when the model uncertainty surpasses a threshold (σ_{thr}) until a sampled configuration satisfies the termination criterion (Δ_{thr}). Simulations include running MD, geometry optimization, and nudged elastic bands. The protocol is validated by replicating the Dang et al. [87] path (green) and used to optimize the path (purple), finding a 40% reduction in the rate-limiting energy barrier (c). The location of the oxygen vacancy is marked with a cross, and the opacity of the surface atoms is reduced.

0. Initial training set The initial training set consists of the bulk, conventional unit cell of cubic In_2O_3 , taken from the Materials Project [130]. Additionally, we included ten configurations obtained by rattling all atom positions with random displacements drawn from a Gaussian with standard deviation 0.02 \AA and applying random deformation to the unit cell, with a standard deviation of 0.01. This ensures that we sample repulsive behaviour immediately. Additionally, we obtain a set of initial geometries for the intermediates, referred to as approximate intermediates, which enter the protocol in blocks 3-5.

1-2. Bulk and surface dynamics In the first two active learning blocks, we iteratively increase our training set until the potential can run stable dynamics. The potential is deemed stable if, during a 10 ps molecular dynamics (MD) run at 300K, the uncertainty threshold is never breached for all intermediates. We repeat the same procedure for the slab and oxygen vacancy of interest, requiring one additional iteration, using the same termination criterion.

3-4. Adsorbate dynamics We now extend our model to capture molecule-surface interaction, creating a four-element force field (H, C, O, In). To reduce the number of iterations needed to obtain stable MD we enhance the training set before active learning with:

- **Dimers:** configurations consisting of two atoms spaced such that their repulsive forces never exceed 20 eV\AA^{-1} . We include four such dimer configurations for each possible combination between carbon, oxygen, and hydrogen, totalling 24 configurations. We set the energy and force weights of all dimer configurations at a factor of five lower than the rest of the training set.
- **CatKit configurations:** obtained by placing the molecules of the intermediates around the active site using CatKit's adsorbate placement tool [131]. Rather than placing molecules directly on the relaxed surface, we use the potential from the second block to run MD on the active site first. This avoids sampling unnecessarily similar atomic environments from the surface. After placing the intermediates within 3 \AA of the oxygen vacancy with CatKit, we select three adsorbate structures per intermediate using farthest point sampling (FPS) of all atomic environments. Here we use the distances between SOAP vectors of the local environments to measure dissimilarity, following the same procedure as De et al. [132].

We now run MD with the MLFF on all approximate intermediates for 1 ps. As previously we terminate active learning if, during MD, the predicted uncertainty stays below the threshold for all intermediates. As mentioned in [Section 3.2](#), using the uncertainty estimate is crucial for sampling new configurations in this block.

5. Adsorbate Geometry optimization Using geometry optimization, we now try to find the DFT adsorbate structure starting from the approximate intermediates. During active learning, we monitor the performance of the potential by evaluating the GAP minima configurations with single-point DFT. If the DFT forces are sufficiently small for all intermediates, the MLFF's minima correspond to true minima in the DFT potential energy surface, and we terminate the active learning loop. However, if the DFT forces are above 0.2 eV\AA^{-1} for any intermediate, we add the configurations to the training set and continue the active learning loop. Note that the adsorbate structures further improve during the last active learning block.

6. Minimum Energy Path Search Ultimately we want to obtain accurate energy barriers with our purpose trained MLFF. To find transition states (TS) we use the nudged elastic band (NEB) method [[133](#), [134](#)]. This two-sided search connects the initial and final states with a discrete chain of configurations called images. The atoms of the images are connected by artificial springs that spread the images along the reaction path. The initial path is often generated by interpolating the atomic positions of the images between the initial and final state. During optimisation the images experience the forces from the unbiased potential energy perpendicular to the path and the artificial spring forces along the path. The forward energy barrier is given by the energy difference between the highest and initial image. After initial convergence we use the climbing image NEB to refine the barrier further. Here image with the highest energy is treated differently and its energy is optimised along transition path direction and minimised perpendicular to the reaction path. Alternatively one could also initialise a single sided transition state search, such as the dimer method [[135](#)] following, from the initial guess structure.

Before starting active learning, we add a set of linearly interpolated images between the reaction intermediates to the training set. Running active learning, we then perform NEB calculations, sample the resulting minimum energy paths, and add them to the training set. The active learning loop is terminated when the total energy error (DFT vs. MLFF) on all sampled configurations is less than 50 meV. Note that for each reaction, we run NEBs for all possible permutations of like atoms of the

molecule (see Section 3.4.3 for more detail). During sampling, we only DFT evaluate the NEB path with the lowest energy.

3.4 Validation on known pathway

We validate our training protocol using a well-studied reaction mechanism, namely the carbon dioxide hydrogenation to methanol on an indium oxide surface [87–98, 136–140]. In the context of reducing carbon dioxide emissions and carbon recycling, methanol synthesis from waste carbon oxide (CO and CO₂) streams is of significant industrial importance. Indium oxide has emerged as a promising catalyst, particularly for operation under CO₂-rich conditions, offering improved selectivity compared to copper-based catalysts that currently dominate industrial practice [141, 142]. This reaction pathway, comprising ten intermediates and five distinct reactions, presents significant challenges for training a single MLFF due to the diversity of the adsorbate species and the complexity of the surface structure. As a benchmark, we investigate the oxygen vacancy on cubic-In₂O₃(110) with the lowest formation energy, as reported by Dang et al. [87]. We use the same oxygen vacancy as in Reference [87]. Each configuration for a different molecule has a different number of atoms, based on the adsorbed molecule. To compare energies, we assume a chemical potential that corresponds to H₂, CO₂ and O₂ adsorption energies.

Additionally, we showcase the transferability of the protocol on a previously unexplored surface. Recent experimental and computational investigations have revealed that, under experimental conditions, the top surface of indium oxide undergoes complete reduction [143, 144]. By applying our protocol to this reduced surface, we demonstrate its capability to describe a more realistic and experimentally relevant surface.

We illustrate the utility of a fast yet accurate MLFF by exploring multiple minimum energy paths for each reaction and calculating free energy barriers. This systematic investigation uncovers notable alterations in the reaction pathway. Specifically, we identify a preferred minimum energy path for the rate-limiting step, characterized by a substantially lower barrier. Additionally, investigations at finite temperatures reveal that the true rate-limiting step corresponds to the production of formaldehyde.

3.4.1 Carbon-dioxide hydrogenation over indium oxide

After applying the outlined protocol, we find close alignment between DFT (black) and the final MLFF (green) throughout the entire reaction path as visible in Figure 3.3.

The final training set consists of 622 configurations. Appendix A Supplementary Note 1 outlines the number of training configurations needed at each stage of the protocol.

Adsorption Energies

The adsorption energies converge throughout the training protocol. Adsorbate-surface interactions are first sampled in active learning (AL) blocks three and four, after which the mean adsorption energy error is 82 meV. We improve upon this with AL and geometry optimization in block five. To allow a direct comparison, relaxations are initialized from configurations similar to Reference [87] (see Appendix A Supplementary Note 2 for further details on starting configurations). After eleven iterations, we reach the final termination criterion outlined in Section 3.3, at which point the mean adsorption energy error is 23 meV across intermediates. These errors further reduce to 12 meV throughout NEB active learning. See Appendix A Supplementary Figures 1 and 2 for more details on the adsorption energy convergence across iterations.

Minimum Energy Paths

Throughout the protocol, the energy barriers converge to those obtained with DFT. To validate the final force field, we initialize DFT-NEB calculations with the MLFF's minimum energy paths (MEPs). We find that the DFT-NEBs are converged without a single optimization step and have a projected force below 0.05 eV\AA^{-1} for all five reactions. The MLFF has hence found true MEPs on the DFT potential energy surface. Note that as we aim to validate the protocol against Reference [87] in this section, we don't permute all atoms with identical elements, as detailed in the protocol, but choose a permutation that is similar to the previously published literature.

Figure 3.4, shows the convergence of the MEP throughout active learning iterations. We find that without any NEB active learning, it is possible to draw qualitative conclusions with a mean barrier error of 50 meV. At this stage of the protocol, the training set only contains 325 configurations. After six iterations of NEB active learning, the GAP MEPs converge to those of DFT, providing an energy barrier estimate within 45 meV for all five reactions. As highlighted in Section 3.4.4, training an MLFF following this automated protocol requires approximately eight times fewer DFT single-point calculations than running DFT-NEBs directly.

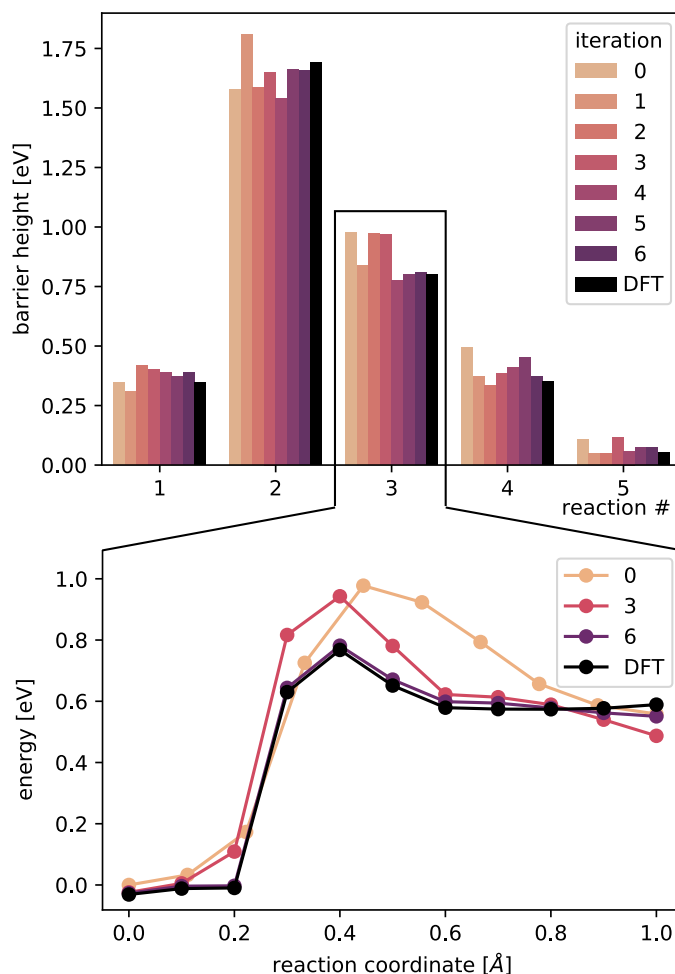


Figure 3.4 Energy barrier accuracy throughout active learning. Barrier height comparison between DFT and GAP models at different AL iterations for all energy barriers (top panel). The bottom panel gives a detailed view of the different minimum energy paths throughout active learning for the third reaction ($\text{H}_2\text{COO} \rightarrow \text{H}_2\text{CO} + \text{O}$).

3.4.2 Computational details

DFT All density functional theory calculations, including single point evaluations, geometry optimizations as well as NEB [145] calculations were done with QuantumEspresso [120]. We use the Perdew-Becke-Ernzerhof (PBE) functional [146], core electrons are described by projector augmented-waves (PAW) [147, 148], the valence mono-electronic states are expanded as plain waves with a maximum kinetic energy of 884 eV (65 Ry). Core electrons are represented by pseudo-potentials [149] while the valence shell is represented with 4, 6, 13, and 16 electrons for C, O, In, and Pt. We use a dense reciprocal-space mesh, with a maximum spacing of 0.25 \AA^{-1} and Gaussian smearing of 0.1 eV with an SCF convergence criterion of 10^{-7} eV, to ensure sufficient

convergence of gradients, which is essential for training MLFFs. Reference [87] used the Vienna Ab initio Simulation Package (VASP). A calculation on the reduced energy barrier for the rate-limiting step shows that QuantumEspresso calculations correlate well with those obtained from VASP (see Appendix A Supplementary Method 1).

GAP We used the Gaussian Approximation Potential (GAP) [2, 53] in combination with a double turbo SOAPs [150] as our many body descriptor, with 4 Å and 6 Å cut-offs. To evaluate the potential, we use the quippy python interface [151, 152]. A complete list of hyper-parameters is included in Appendix A Supplementary Method 2.

MACE Mace [27] is a higher-order equivariant message-passing neural network. We use default settings, with two message passing layers, 256 channels, and equivariant messages of order $L=1$. All further settings are tabulated in Appendix A Supplementary Method 3.

Geometry optimization and NEB calculations All slabs are centred between 8Å of vacuum with the lowest two layers (40 atoms) fixed. Relaxations were converged to 0.001 eVÅ^{-1} and 0.01 eVÅ^{-1} when using the MLFF and DFT, respectively, unless explicitly stated otherwise. For NEBs calculations, the convergence criteria were set to 0.01 eVÅ^{-1} and 0.05 eVÅ^{-1} respectively. The criteria were set lower for the MLFF due to its smooth potential energy surface and minimal computational cost. All NEBs were calculated using the atomic simulation environment [153] using splines and exponential preconditioning [154]. The initial image guesses were made using the IDPP method [155]. We use ten and twenty images per NEBs for the oxygen vacancy and reduced surface, respectively.

Umbrella sampling The free energy barriers are obtained using the Colvar [156] package in combination with LAMMPS [157]. We run 110 ps of molecular dynamics for 32 bins along the collective variable, with a 1 fs time step and a spring constant of 75 eVÅ^{-2} for all temperatures. Starting configurations are taken from the relevant NEB path and allowed to equilibrate for 10 ps. Using umbrella integration [158], the samples are mapped to the free energy barrier.

CatKit CatKit [131] was used to create initial adsorbate structures to start active learning and to initialize the minima-hopping configurations. Here all surface oxygen atoms were explicitly defined.

Active site We investigate the same active site as Reference [87]. Specifically the oxygen vacancy on the 110 surfaces of cubic indium oxide, with the lowest formation energy. The vacancy configuration itself has 79 atoms per unit cell.

Data availability The datasets generated throughout the active learning protocol and used to train the MLFFs are available in the Zenodo repository, <https://doi.org/10.5281/zenodo.8268725>.

3.4.3 Leveraging computation acceleration

Lower rate-limiting step

We re-initiate the training protocol, assuming the adsorbate configurations are unknown for this reaction. To generate initial guesses for the reaction intermediates, we employ CatKit and select the adsorbate configuration closest to the position of the oxygen vacancy. After successful termination of the protocol, the acquired training set contains 696 configurations. **Figure 3.3c** compares the reaction pathway of the resulting MLFF (purple) with the pathway explored by Dang et al. [87] (green). The paths align closely, except for a significant reduction in the rate-limiting barrier and the last reaction path. The difference in the last reaction step originates from a subtle rearrangement of the $\text{H}_3\text{CO} + \text{H}$ intermediate for the path found in the literature. As we start from CatKit configurations, in the training protocol these two intermediates are identical. In the following, we discuss in greater detail how the automatic protocol finds an optimized reaction path for the rate-limiting step.

When performing NEB calculations in an automatic fashion, a crucial consideration arises regarding the ordering of like atoms. Since the atoms in the initial image are connected to atoms in the final image along a smooth path, any alteration in the index of atoms in the initial image leads to a distinct reaction path. Traditionally, the selection of the most plausible arrangement has relied on chemical intuition. However, given the computational efficiency of NEB calculations with the MLFF, we adopt a comprehensive approach by performing all possible permutations of atoms with the same element of the molecule. Subsequently, we identify the path that exhibits the lowest energy as the preferred path.

This comprehensive approach unveils a preferential reaction path for the rate-limiting step ($\text{HCOO} + \text{H} \rightarrow \text{H}_2\text{COO}$). This optimized path, depicted in purple in **Figure 3.5**, has a significantly reduced barrier of 1.0 eV. The insets reveal a key difference between the two paths: throughout the optimized path, one of the intermediate's oxygen atoms occupies the oxygen vacancy. By screening permutations, we achieve a

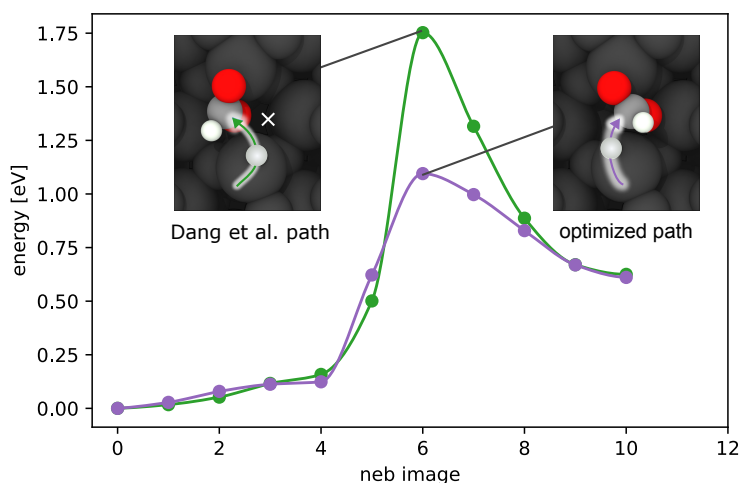


Figure 3.5 Lower rate-limiting step. Comparison of two NEBs for the rate-limiting step, with different permutations of like atoms. The protocol investigates all possible permutations, finding a significantly lower barrier (green) than Reference [87] (blue). Visualizations of the transition states (insets) show the the hydrogen atom's trajectory and the location of the oxygen vacancy (white cross).

systematic approach to investigate barriers, avoiding the potential pitfalls of possibly misleading chemical intuition. We note that for larger intermediates, screening all possible permutations may become unfeasible. Here we suggest pre-selecting the most relevant permutations by comparing the geometric distance between the initial and final NEB configurations.

Finite temperature effects

Until now, we have assumed that the zero kelvin potential energy surface can be used to infer properties of the catalyst. This approximation is often made due to the computational cost of DFT. We now illustrate the importance of incorporating finite temperature effects by investigating the free energy barrier of dioxymethylene to formaldehyde conversion ($\text{H}_2\text{COO} \rightarrow \text{H}_2\text{CO} + \text{O}$). In doing so, we find that the energy barrier is larger than previously assumed and that formaldehyde production is the rate-limiting step.

We perform umbrella sampling at ambient and operating (500K) temperatures, using the distance between the oxygen atom and the carbon atom as a collective variable. Within the first picosecond of molecular dynamics, the starting configuration, taken as the local minima found in literature, moves to a lower energy state with different geometry. The insets in Figure 3.6 illustrate that the H_2COO intermediate preferentially binds to the indium atom with lower oxygen coordination.

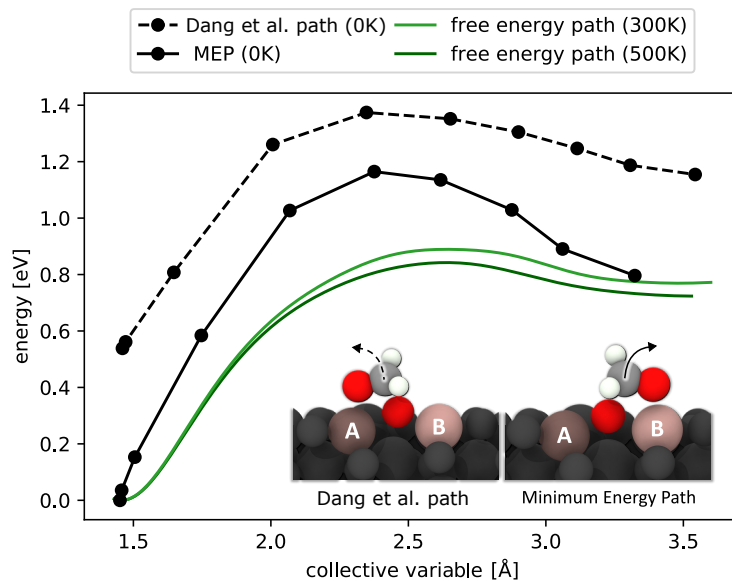


Figure 3.6 Free energy barriers for formaldehyde production. Comparison of the Reference [87] path (dotted) and the minimum energy path (MEP, bold). The black lines represent NEB energies, while the green lines depict the free energy barriers along the MEP, calculated by umbrella integration. The insets display the transition states of Ref. [87] and the MEP, highlighting the distinct oxygen-coordination of the two indium atoms (atom A is four-fold coordinated, atom B is three-fold coordinated).

Consequently, the minimum energy path (depicted by the bold line) differs from the path investigated in Reference [87] (represented by the dotted path). With a barrier height of 1.1 eV, the production of formaldehyde (H_2CO) rather than H_2COO is rate-limiting. For a more detailed visual comparison of the paths, please refer to Appendix A Supplementary Figures 3-7. This exploration of the potential energy surface, facilitated by the computational speed of MLFF, reveals that although there may exist some high-energy local minima, their relevance to the overall reaction progression is limited. Finite temperature simulations promptly reveal the thermodynamically relevant intermediate geometries.

By running umbrella sampling, we circumvent the approximation that reaction rates are solely determined by a single transition state and local energy minima on the potential energy surface. The temperature dependence of the free energy barrier can have a substantial impact on the predicted selectivity and activity of catalysts, as well as the behaviour of micro-kinetic models employed to analyse the progression of reactions [159–162].

Exploration of adsorbate structures

We showed in Section 3.4.3 that thermodynamically relevant adsorption configurations can elude DFT investigations, as adsorbates may get stuck in local energy minima. To investigate whether this is a more general property of these complex surfaces, we use the MLFF to sample adsorption configurations of the very first intermediate: CO₂. A thorough sampling of the surface reveals a plethora of local minima as illustrated in Figure 3.7a. Each such minimum is a potential starting structure for studying the subsequent reaction step. This suggests that running a small number of geometry optimization with expensive *ab-initio* methods is unsuited for conclusively finding adsorbate configurations.

We systematically screen local minima using minima hopping, which consists of alternating high-temperature MD with geometry optimization [163, 164]. We use CatKit to sample ten starting configurations for each intermediate and then run ten iterations of active learning with minima hopping. The final potential finds 236 unique minima. All non-unique minima are removed by comparing the local environment of the carbon atoms using SOAP vectors (see Appendix A Supplementary Method 2 for details). To validate the new adsorbate geometries, we relax 25 MLFF-found minima with DFT to a maximum force of 0.01 eVÅ⁻¹. We find that the MLFF (Figure 3.7c) and DFT minima (Figure 3.7d) are extremely similar. Note that the top two layers of the surface are not fixed and, hence, differ for all adsorbate configurations.

Figure 3.7 shows all 236 unique adsorption geometries placed on top of each other. For visualization purposes, only the CO₂ molecules are included and color-coded by adsorption energy, while the geometry-optimized oxygen vacancy in dark grey acts as a backdrop. The lowest energy configurations are found when CO₂ reacts with a surface oxygen atom to form carbonate. Previous studies that also identified the thermodynamic stability of carbonate found that it doesn't partake in the conversion to methanol [87]. Carbonate configurations are hence excluded from the figure. As visible from the histogram, a significant number of local minima have negative adsorption energies.

3.4.4 Computational cost comparison

The findings summarised above rely heavily on the computational speed-up gained by the MLFF. To compare the computational cost, we calculate the first energy barrier, starting from a simple interpolation of images with DFT. After 38 steps with BFGS, requiring 1092 single-point DFT evaluations, the NEB converges to 0.08 eVÅ⁻¹. As a comparison, the final GAP model requires only 622 DFT training




Figure 3.7 Adsorbate configuration exploration. Displaying 236 unique MLFF adsorption geometries (b) of carbon dioxide found with minima hopping and their adsorption energy distribution (a). Figures (b-d) contain only the molecular atoms of the adsorbate states, colour-coded by adsorption energies, and the relaxed oxide surface in dark grey as a backdrop. All carbonate configurations are excluded, and their energies are shaded in grey in the histogram (a). A subset of MLFF minima (c) is relaxed with DFT (d) for validation.

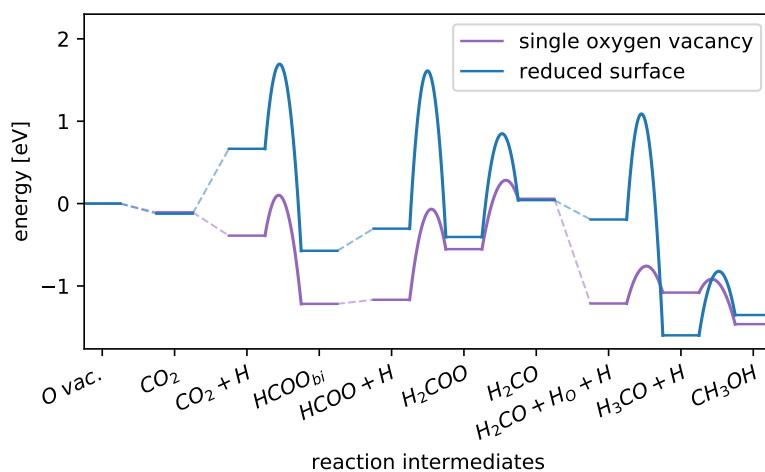


Figure 3.8 Reduced surface reaction pathway. Comparing the reaction pathway for the single oxygen vacancy and the top-layer reduced surface.

configurations and finds NEBs for all five reaction steps that converged to below $0.05 \text{ eV}\text{\AA}^{-1}$ when evaluated with single point DFT. Running all DFT-NEBs would hence require approximately five thousand single-point evaluations, eight times more than the size of the training set.

Beyond its computational efficiency for routine catalytic tasks, the MLFF offers extensive exploration capabilities of the potential energy surface. Considering all feasible atom permutations within the same chemical element helps us obtain a 40% reduction in the activation energy of the rate-limiting step. Furthermore, determining the free energy barriers required 6.4 million samples of the potential energy surface. For an 84-atom configuration, the GAP model runs 9.2 steps per second on the dual AMD EPYC™ 7742 64-core processors at 2.25GHz. In contrast, a single DFT calculation requires an average of 33 minutes on the same hardware, resulting in a speed-up of more than 18 thousand times.

3.5 Transferability to unexplored surfaces

To demonstrate the transferability of the automatic training protocol, we apply it to a previously unexplored surface. Additionally, we illustrate that the obtained training set can be effectively utilized by different machine learning methods, indicating its versatility beyond the specific GAP framework.

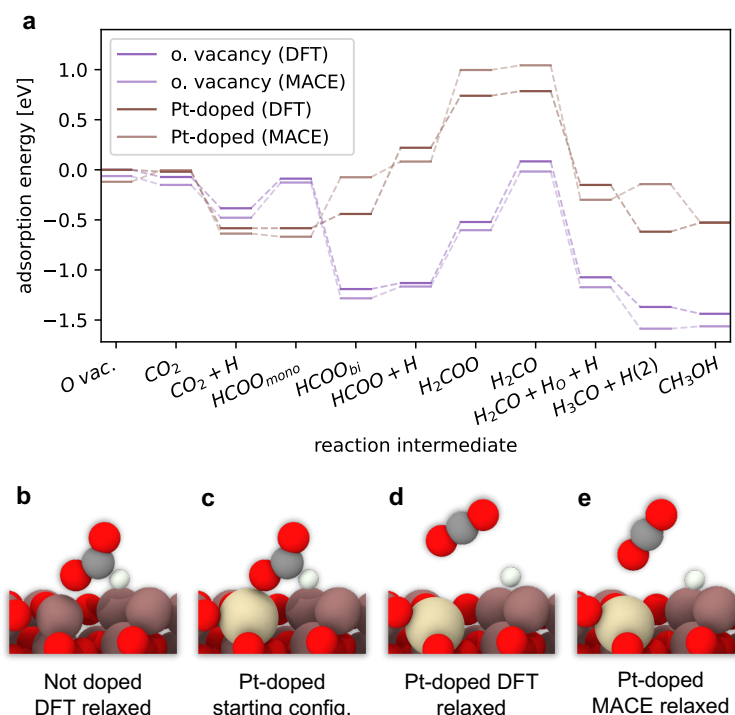


Figure 3.9 Training set transferability to other MLFF frameworks. A comparison between the MACE [27] model and DFT for both undoped and doped surfaces (a). The intermediates on the Pt-doped surface differ significantly, such as formate splitting into CO₂ + H (b-e).

Reduced surface

Recent computational investigations have revealed the thermodynamic preference for the reduced surface under experimental conditions [90]. To evaluate the transferability of our protocol, we apply it to the top-layer reduced surface, yielding the reaction pathway depicted in Figure 3.8. The construction of the training set for this force field required less than one thousand single-point DFT calculations, with 70% of the calculations dedicated to the final active learning phase to accurately determine energy barriers. Notably, the pathway reveals a large barrier for the hydrogenation of carbon dioxide.

Doped surfaces

We examine the transferability of our training dataset to other MLFF frameworks. Specifically, we use MACE [27], a higher-order equivariant message-passing neural network, to train an MLFF on the dataset acquired in Section 3.4.1 for the single

oxygen vacancy surface. Without any additional training data, it achieves a mean adsorption energy error of 38 meV.

Furthermore, we investigate the behaviour of doped surfaces. Dopants play a crucial role in modifying the adsorbate structure and activation of intermediates, leading to significant changes in catalytic activity and selectivity [22]. In the case of indium oxide, platinum, and palladium dopants have been demonstrated to enhance catalytic performance by promoting H₂ splitting [165, 166]. These dopants not only increase the abundance of adsorbed hydrogen, needed for CO₂ hydrogenation and vacancy formation but also influence the entire reaction pathway.

To avoid repeating the entire training protocol, which would be inefficient, we explore the possibility of enhancing the training set without additional active learning steps. We find that with sixty additional configurations, the MACE framework captures general trends in the adsorption energies, as visible in Figure 3.9. The additional training configurations originate from the oxygen vacancy training set, where the Indium atom closest to the active site is replaced by a platinum atom and randomly rattled with a standard deviation of 1Å. It is worth noting that improving the adsorption energies through active learning with the MACE model could be pursued, but is beyond the scope of this thesis. Application to a different MLFF framework demonstrates the transferability of the acquired training data.

3.6 Discussion

We present an automatic training protocol for MLFF capable of accurately capturing the energetics of a given heterogeneous reaction path. We validate our approach on the extensively explored conversion of CO₂ to methanol on indium oxide with a single oxygen vacancy. The minimum energy paths (MEPs) for all five reactions found using the final MLFF correspond to true MEPs on the DFT potential energy surface. A single model is hence able to accurately describe a reaction with ten intermediates on a complex oxide surface. Additionally, we showcase the transferability of the protocol by investigating the same mechanism on a new and previously unexplored surface: top-layer reduced indium oxide.

We show that MLFFs offer more than just computational cost reduction for routine computational catalysis tasks; they emerge as an essential tool for in-depth mechanistic catalytic investigations. By running multiple nudged elastic band calculations for each reaction, we find a preferred MEP for the rate-limiting step with a 40% lower energy barrier. Moreover, through finite temperature sampling, we unveil lower minima for the intermediates of the subsequent largest barrier, identifying it as the

true rate-limiting step. Finally, we illustrate the power of MLFFs by computing free energy barriers at ambient and operating temperatures, requiring over six million samples of the potential energy surface.

Accurately describing the true minimum energy path of individual reactions significantly influences the relevance of computational studies to experiment. Predictions from micro-kinetic or energetic span models can depend strongly on specific barriers, resulting in notable differences in observables, such as the turnover frequency or selectivity [162, 167].

The work presented is a significant step in the direction of leveraging MLFFs in automated computational catalysis exploration. Currently the selection of the reactions and surface is manual and requires some supervision. We are currently working on interfacing the automated protocol of training MLFFs with reaction exploration packages, specifically Pynta [168]. This will help automatically build the relevant reaction network, initial states and reaction paths needed as an input. Currently Pynta requires exascale super computers and uses a pure DFT approach. By combining this work with other high throughput approaches we hope to tackle a significant part of the current challenges in computational catalysis outlined in [Figure 3.1](#). We anticipate that active learning protocols will facilitate the adoption of MLFFs in the wider catalysis community, enabling more comprehensive mechanistic explorations of catalytic cycles.

Chapter 4

Foundation models and fine-tuning

The approach to training machine learning models has undergone a significant transformation in recent years. Traditionally, models were trained from scratch for a specific task, requiring significant amounts of data and computational cost for each application. However, many machine learning domains, from natural language processing to vision, now benefit from training large-scale models on large and diverse datasets. These foundation models can then be fine-tuned for a desired application with minimal data.

In atomistic modelling, the recent emergence of foundation models is set to significantly change the accessibility and utility of machine learning force fields [72, 169]. In this chapter, we give an overview of recent developments and how they change the way we will train machine learning models in the future. We will mainly focus on the MACE-MP0 [72] model, which was trained on the materials project [130], a large dataset of inorganic and organic crystal structures.

In this brief chapter, we focus on the model's performance on molecule-surface interactions. We show stable molecular dynamics for a 7 atomic element surface-molecule interface. Furthermore, we test the model's performance at chemisorbing carbon dioxide within a metal organic frameworks (MOFs). We find that the model performs surprisingly well, given that the training dataset does not include surfaces or molecules. We conclude the chapter with an example of fine-tuning a foundation model, starting from the MACE-MP0 [72] foundation model for CO₂ adsorption in a metal-organic framework. The performance of this fine-tuned model will be further investigated in Chapter 5, where it is used to incorporate the effects of dynamics in computational NMR spectra.

This chapter is based on work published in multiple papers [72, 170, 171]. The performance analysis of the MACE-MP0 model for catalysis was published as part of

Reference [72], which introduces MACE-MP0. The paper contains multiple examples analysing the foundation model’s performance in the greater context of molecule-material interactions. This chapter focuses on the contributions by the author of this thesis in the context of catalysis. The fine-tuning work was performed to help elucidate an experimental NMR spectrum. The models are published as part of References [170, 171]. The comparison to wetlab experiment is the focus of Chapter 5, where all experiments were performed by Benjamin Rhodes.

4.1 Background and motivation

Machine learning force fields are now more accessible than ever, primarily because of large general-purpose models. This section provides an overview of the historical developments and introduces the MACE-MP0 model.

4.1.1 Early foundation models

The first general-purpose neural network potentials focused on training force fields for organic chemistry. The most widely used and notable models include ANI [129, 172, 173] and AIMNet [174–176]. Trained on small-molecule data at the DFT level of theory, they have been shown to be transferable to larger unseen molecules. Furthermore, after pre-training on a large dataset of DFT data, they have been fine-tuned to a subset of coupled cluster data [177].

Another domain that was early to fit large, diverse datasets is catalysis. In particular, the Open Catalysis initiative from Meta [71]. The initial data set contains more than 1.2 million DFT geometry optimisations with more than 200 million single-point DFT evaluations. The configurations are all geometry optimisations of molecules adsorbed to surfaces of up to three different elements. The dataset has been used to benchmark various architectures. The initial DimeNet++ [178] obtaining more than 0.5 eV mean-absolute-error, while the current state of the art equiformerV2 [179] is more than double as accurate [180]. The open catalysis datasets have subsequently been extended to include more than 900 transition state searches [181]. Models previously trained on the non-reactive OC20 dataset were able to obtain energy barriers within 0.1 eV for 70% of the tested transition state searches, showing their transferability.

In the context of material chemistry, MEGNet [182] was a first notable attempt at training on the Materials Project [130] to predict properties such as band gaps and elasticity directly. The dataset aims to provide computational *ab-initio* analysis of all

known inorganic materials. With the development of machine learning force field architectures, which don't suffer from combinatorial complexity with a number of elements, models were trained on the DFT relaxation trajectories. The initial models include M3GNet [182] and ChGNet [183], the second of which introduced a training dataset MPtrj which has subsequently been used to train numerous foundation models. Additionally, the Nequip architecture has been trained on a large diverse dataset, resulting in the GNOME [184] model, which was used to search for new crystal structures.

More recently, models have been trained on various datasets simultaneously. Previously, large-foundation models were focused on either materials, catalysis, or molecular applications. The choice of DFT functional and basis set depends strongly on the application. While periodic calculations are frequently performed on a plane wave basis with pseudo-potentials for the core electrons, molecular calculations profit from an atom-centred basis. The datasets are incompatible, making it impossible to directly train on publicly available DFT calculations. Reference [185] introduces an approach to simultaneously training on heterogeneous atomistic data. They used multi-head training, an approach where the majority of the weights are shared between the different datasets, except a small subsection, typically the readout. The authors used this architecture to train on molecular (SPICE, QM9, MD17, MD22, ANI1x), reactive (Transition1x), material (QMOF, Matbench), and catalysis (OC20, OC22) datasets. The idea is that the shared part of the model learns complex representations of the 3D environment from which all applications profit.

4.1.2 MACE-MP0

A significant recent development is the MACE-MP0 model. Trained on a public database of 150 thousand crystal structures, the model can run sensible simulations of completely out-of-domain systems ranging from molecular liquids to proteins and catalysts. Reference [72] tests the model by running hot molecular dynamics, geometry optimisations, free energy calculations and transition state searches on over sixty test systems. Overall, the authors found that the model was able to reproduce recent computational discoveries that previously required direct *ab-initio* simulation or the training of a purpose-made force field.

4.2 Performance on out domain test cases

A large group of researchers tested the MACE-MP0 model on various material and molecular tasks. For details on the 60+ test cases, from nano-confined water to simulating entire battery systems, please see Reference [72]. This section will focus on the work done by the author of this thesis on testing and fine-tuning MACE-MP0, which is published as part of references [72, 170].

4.2.1 Transition state search for catalysis

We now test the MACE-MP0 foundation model on computing reaction barriers for catalysis. Note that the training set does not include surfaces or molecules. As a test system, we investigate a key step ($\text{CH}_2\text{O}_2 \rightarrow \text{CH}_2 + \text{O}$) in carbon dioxide hydrogenation to methanol over indium oxide. This reaction has been extensively explored in Chapter 3 and directly with DFT [87]. Firstly, we optimise the reactant’s global geometry near an oxygen vacancy. As shown in Figure 3.6 found using a purpose-trained MLFF, the Indium coordinated with three oxygens is the active site. Previous DFT investigations had focused on a neighbouring four oxygen-coordinated indium significantly deviating from the relevant energy profile. MACE-MP0 correctly identifies the Indium coordinated with three oxygens as the active site [54], which previous DFT investigations had missed. This indicates the utility of the pre-trained model even before fine-tuning. In the future, *ab-initio* investigations can start with a thorough exploration of the possible adsorption sites before detailed investigations of specific sites. Finally, we perform a NEB calculation and find a reaction barrier within 10% of that investigated with DFT, as visible in Figure 4.1.

Similarity Investigation

To understand how MACE-MP0 can perform with reasonable accuracy for out-of-domain tasks, we look to investigate structures in the training set, which are mapped to similar MACE features to the catalytic ones. As detailed in Section 2.4.6, the MACE model assigns features to each atom based on its element and local environment. By comparing these high-dimensional features, it is possible to gain some qualitative insight into which training configurations influence the predictions. We see in Figure 4.1b, a projection of the features to two dimensions using a UMAP[186].

This illustrates the underlying reason why MACE-MP0 can extrapolate. While the configuration may be different globally, locally, the environments may be similar,

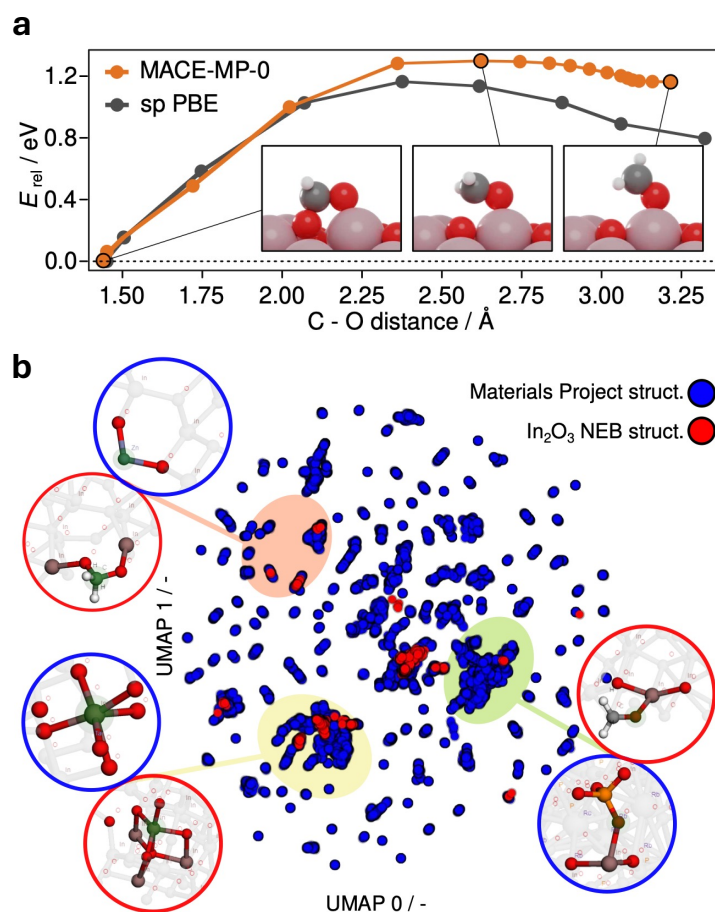


Figure 4.1 Transition state searching using foundation model. (a) Showing the reaction barrier of $\text{CH}_2\text{O}_2 \rightarrow \text{CH}_2 + \text{O}$ on indium oxide from Chapter 3. (B) A Comparison of the atomic environments in the training data (blue) and in the In_2O_3 NEB images (red) in the form of a UMAP plot[186]. Insets show local environments with similar MACE features (inset frames in blue for training data and in red for NEB configurations), exemplifying which bulk training environments influence predictions for the out-of-domain catalytic test case. Figure (a) was reformatted by Edvin Fako to match the colour scheme of the MACE-MP0 paper [28], Figure (b) was created in collaboration with Rokas Elijošius.




Figure 4.2 Molecular dynamics snapshot of an indium oxide slab covered with diverse molecules. An illustration of the simplicity and accessibility of atomistic simulations to the wider computational science community. A foundation model (MACE-MP0) trained only on crystal structures is able to perform dynamic simulations out of the box without the need for *ab-initio* calculations or retraining.

allowing for extrapolative emergent behaviour. One key component is the embedding of elements [27, 50], allowing for transferable learning between chemical elements.

The application of catalysis is a particularly stark illustration of this, as the data set contains no slab surfaces or molecules on top of the surfaces. Figure Figure 4.1b shows the most relevant bulk structures. Notably, the model detects similarities between a carbon of the formaldehyde test structure and a carbon bonded to indium and a phosphorus group in the material project (green). Indeed, the mace model detects similarity between different elements, such as carbon and Zinc (pink), because of the element embedding. As expected, bulk indium oxide mace features are closely mapped to some of the surface atoms in the test structures (yellow).

4.2.2 Molecular dynamics of large catalytic systems

To showcase the versatility of the MACE-MP0 model, we place random molecules on top of a large indium oxide slab as visible in Figure 4.2. The molecules were randomly chosen from PubChem database[187] and contain seven different chemical elements (C, H, Cl, F, N, O, S, Si). The molecular dynamics runs stable for the entire test duration of five picoseconds. Previously, obtaining a machine learned interatomic potential that can run reactive molecular dynamics on even small surfaces with a single CO₂ molecule required multiple iterations of costly active learning.

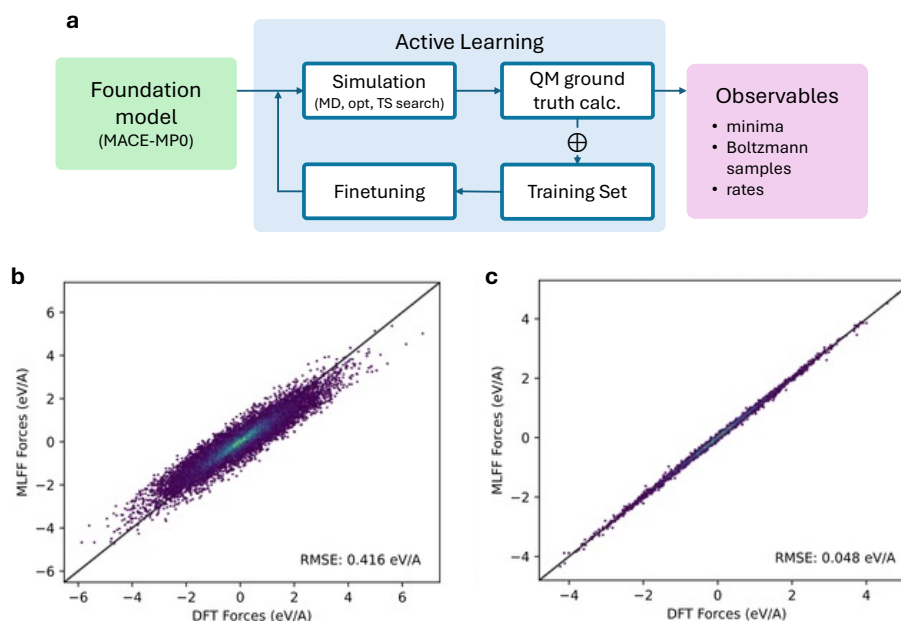


Figure 4.3 Fine-tuning foundation models. (a) Showing a workflow of iterative training outlined consisting of using the foundation model to generate structures, subsequent evaluation with a quantum mechanical ground truth and fine-tuning of the foundation model. (b) Show the force accuracy of the foundation model on a dataset of sampled molecular dynamics of the MFU-4l metal-organic framework. (c) Shows the performance of a fine-tuned model on a test set of molecular dynamic runs. Rather than showing a scatter plot that would hide overlapping observations, these plots show a 2D heat map.

While the dynamics and rates will not have quantitative accuracy, there is significant value in a versatile tool that can sample sensible configurations across a wide variety of atomistic systems. The MLFF can be used to sample relevant configurations optimised with a ground truth *ab-initio* method. If high accuracy is required, the model can be fine-tuned.

4.3 Fine-tuning

Fine-tuning is a type of transfer learning where weights are taken from a large pre-trained model and further adjusted by training on a new small dataset. There is extensive literature in natural language processing (NLP), where different fine-tuning approaches are investigated to reduce the computational cost or allow for training private information. These involve a method such as training only a subset of weights.

In the context of fine-tuning MACE-MP0, the requirements are less stringent, as the model is relatively small compared to NLP models, and there are no privacy

concerns. The main requirement is that the model improves on the dedicated task without losing its stability or general-purpose performance.

We explore the simplest form of fine-tuning on a metal-organic framework. We allow all weights in the model to change and train on less than 60 configurations sampled from a molecular dynamics simulation of the metal-organic framework. [Figure 4.1](#) compared the accuracy of the pre-trained foundation model (b) with that of a fine-tuned model after a single iteration. We see that after one iteration, the force RMSE on unseen configurations drops by more than a factor of eight.

Looking at the force correlation plot of the pre-trained model, we see that forces are systematically under-predicted. Others have also observed this in the context of predicting vibrational properties of materials and described as a general "softening" of the potential energy landscape [188]. The origin of this is hypothesised to be the abundance of low-force configurations in the training set. In practice, this leads to a reduction in the barrier height of dynamic events. In the context of sampling configurations for fine-tuning, this is a valuable property, as it allows sampling relevant configurations with shorter simulations. In the case of the metal-organic framework, this leads to a faster rotational motion of adsorbed bicarbonates in the metal-organic framework. The precise system and its dynamics are discussed in greater detail in [Chapter 5](#) in the context of computationally predicting NMR spectra when dynamic processes significantly affect the experimentally observed spectra.

The ultimate consequence of the efficacy of fine-tuning is a significant speed-up in active learning. As described in [Chapter 3](#), it takes several iterations of active learning for models trained from scratch to allow stable molecular dynamics of complex systems. With the help of foundation models, the first iteration can sample uncorrelated samples of relevant configurations from long-time-scale molecular dynamics. This significantly accelerates the active learning procedure. Indeed, having a threshold for interrupting molecular dynamics is almost superfluous. Instead, the most relevant configurations can be selected as a post-processing step, selecting the configurations with the highest uncertainty. This significantly changes the requirements for uncertainty predictions. Previously, a cost-effective metric with high accuracy was needed to be run in series with the simulation. The uncertainty on potential samples can now be efficiently evaluated in parallel, placing less stringent requirements on the uncertainty metric. The proposed iterative learning approach is graphically summarised in [Figure 4.3](#).

4.4 Discussion

The ability of foundation models to capture the interactions between atoms in a transferable way surprised many in the community. The existence of models trained on small molecular and inorganic crystals that can predict the behaviour of proteins with thousands of atoms or the reaction of molecules on surfaces is remarkable. These developments are a testament to the potential of machine learning force field architectures and reveal something about nature itself, that purely local interactions can lead to the emergence of complex macroscopic phenomena.

The development of accurate and transferable foundation models fundamentally transformed the accessibility of accurate atomistic modelling. Computational chemists no longer need to spend time running direct *ab-initio* dynamics or curating large datasets. Instead, a reliable and stable force field can be used to sample relevant equilibrium configurations from molecules to materials. The foundation model can be fine-tuned for applications requiring high accuracy on a small number of targeted configurations, as demonstrated in this chapter.

Before the advent of foundation models, active learning was crucial for dataset curation, especially for applications such as catalysis, where reliable empirical force fields were lacking. However, such challenges have been mitigated with the proliferation of universal machine learning force fields. MLFFs can now be used directly to generate relevant datasets without avoiding unphysical bond dissociations. Previously, as presented in [Chapter 3](#), choosing a low uncertainty threshold during active learning was critical to avoid catastrophic failure during initial simulations. As such, the role of uncertainty quantification will presumably shift from avoiding simulation failures towards selecting the most relevant configurations for fine-tuning as a post-processing step. This marks a significant advancement in machine learning for atom scale modelling.

Chapter 5

Effects of dynamics on NMR spectra across time-scales

The ultimate test of whether computational modelling accurately represents true dynamic motion is the comparison to experimental observables. An experimental technique with atomic-scale resolution is nuclear magnetic resonance (NMR). In this chapter, we show the utility of using machine learning force fields (MLFFs) to help interpret NMR spectra. This chapter is based on the work done in collaboration with an experimental group published in Reference [170, 171]. While this chapter focuses on the computational work done by the author of this thesis any experimental work included for context was conducted by Benjamin Rhodes.

We start by presenting a workflow for capturing the effects of molecular motion on NMR spectra. We then focus on applying the workflow to carbon dioxide capture in a class of porous materials, namely metal-organic frameworks (MOFs). In the previous chapter, we introduce a workflow for fine-tuning atomistic foundation models. In this chapter, we show how with minimal additional DFT calculations, we can accurately capture the dynamics of metal-organic frameworks not present in the foundation models dataset.

5.1 Background and motivation

Accurately differentiating materials and determining their atomic structures is crucial for mechanistic insight into materials and correct material characterization. Nuclear magnetic resonance (NMR) spectroscopy provides high-resolution insights into a material's atomic structure, making it essential for understanding processes from carbon capture to catalysis. However, solid-state NMR presents unique challenges

due to its complexity in interpretation. Expensive quantum mechanical methods, such as density functional theory (DFT) calculations, are often required to match the experimental result with a material's atomic arrangement. A single DFT NMR prediction represents a spectrum of a frozen configuration, neglecting all dynamical effects.

Machine learning force fields (MLFFs) have been applied to simulate dynamic effects in NMR spectra [189–191]. In these approaches, configurations from molecular dynamics simulations with a MLFF are evaluated with DFT NMR calculations and averaged to obtain effective spectra. Although MLFFs provide computational speed-up, they typically capture only fast dynamics (nanosecond timescales). However, longer time-scale dynamics (e.g. microseconds) can also influence NMR shifts, limiting the effectiveness of traditional approaches. Alternatively, packages allow modelling of stochastic hops between a small number of configurations and associated chemical shifts at arbitrary rates [192]. As the number of relevant minima increases, the number of required barrier calculations grows combinatorially. Furthermore, the procedure doesn't correspond to a true Boltzmann sampling of relevant equilibrium configuration and, hence, often fails to capture complex short-term processes accurately. Furthermore, obtaining rates for these slow processes requires expensive *ab-initio* transition state searching and vibrational analysis.

In this chapter, we present a combined approach that integrates fast- and long-timescale dynamics, enabling agreement between computational and experimental spectra. Furthermore, we present a comprehensive workflow on how to fine-tune atomistic foundation models and use them to both sample equilibrium configurations at short time scales (ns) as well as calculate transition rates for more rare events. We show that it is possible to obtain accurate computational NMR spectra even for materials with complex dynamics.

5.2 Nuclear magnetic resonance

The atomic structure of a material can be probed by analysing how atomic nuclei interact with an external magnetic field. Atomic nuclei have a quantized intrinsic angular momentum, referred to as nuclear spin. The spatial and spin properties of a nucleus are predominately uncoupled, and so the total nuclear wave-function, therefore can be separated into the spatial and spin wave-function,

$$\Psi = \psi_{\text{spatial}}\psi_{\text{spin}}. \quad (5.1)$$

In the presence of an external static magnetic field, \mathbf{B}_0 , the system's energy levels become spin dependent through the Zeeman interaction,

$$\hat{H}_Z = -\hat{\boldsymbol{\mu}} \cdot \mathbf{B}_0 = -\gamma \hbar \hat{I}_z B_0, \quad (5.2)$$

where $\boldsymbol{\mu}$ is the nuclear magnetic moment operator, γ is the gyromagnetic ratio of the nucleus, and $\hbar = h/2\pi$, is the Planck constant. Here we have arbitrarily chosen the time-independent field \mathbf{B}_0 to align with the z -axis. The Zeeman Hamiltonian has the same eigenstates as the operator for the z component of the nuclear spin. The eigenstates, like other angular momentum eigenstates are indexed by I and m , where I depends on the number of neutrons and protons in the nucleus. For a given I , there exist $2I + 1$ spin states indexed by $-I < m < I$. The energies of the eigenstates $|I, m\rangle$ are given by,

$$\hat{H}_Z |I, m\rangle = E_{lm} |I, m\rangle = -\gamma \hbar B_0 m |I, m\rangle, \quad (5.3)$$

where m is the magnetic quantum number corresponding to the projection of \mathbf{I} along the z -axis. The transition energy between neighbouring spin states for a given I is equal to $\gamma \hbar B_0$. The nuclei can transition between these spin states by adsorbing or emitting photons with the corresponding energy. The energy difference corresponds to a frequency of

$$\omega_0 = \gamma B_0, \quad (5.4)$$

which is commonly referred to as the Larmor frequency. The magnetic field originating from the resulting electromagnetic radiation can be experimentally measured.

The Zeeman interaction in Equation (5.2), appears to be equal for all nuclei with the same number of neutrons and protons. However, we have for now neglected the effect of surrounding electrons, electric field gradients and neighbouring nuclei. The most substantial perturbation to the observed NMR frequency originates from the shielding of the external magnetic field, \mathbf{B}_0 by surrounding electrons. Electrons also interact with the external field and effectively reduce the magnetic field experienced by the nuclei. This reduces the energy splitting and the experimentally observed frequency. These secondary effects strongly depend on the nuclear arrangement and, therefore, provide insight into the 3D environment of materials and molecules. Details about all of the relevant interactions are beyond the scope of this thesis and can be found in relevant textbooks [193].

5.2.1 NMR spectroscopy

In an NMR experiment, the sample is first placed in an external magnetic field, and nuclear spins partially align with the external magnetic field. Note that, due to thermal fluctuations and the weakness of the Zeeman interaction, approximately only 1 in 10^6 nuclei are aligned with the magnetic field. The nuclear spin states are then excited by a radio frequency pulse. During the electromagnetic pulse, the nuclear spins interact with the time-dependent magnetic field, leading to a mixing of the spin states. As the spin states transition to lower energy states, the individual nuclear transitions result in a measurable electromagnetic signal. The macroscopic oscillating magnetic field is measured using a receiver coil. A Fourier transform of the observed signal allows one to distinguish radiation originating from different nuclei with different Larmor frequencies. The spectrum for a given material or molecule is highly dependent on the local 3D environments present in the structure.

5.2.2 Anisotropic shielding

As mentioned previously, the magnetic field experienced by the nuclei depends on the surrounding electrons. As the electron density around the nuclei is not necessarily spherically symmetric, these effects depend on the orientation of the atomic structure with respect to the external magnetic field.

The predominant source of magnetic shielding is due to the charge flux caused by the electrons interacting with the external field. The electron motion causes a local field, which opposes the applied field. To first-order approximation, this secondary field is proportional to the external magnetic field, such that,

$$\mathbf{B}_s = \sigma \cdot \mathbf{B}_0, \quad (5.5)$$

where σ is the chemical shielding tensor. As the local B_s field depends on the orientation of the molecule or material, the shielding tensor is a rank-two tensor and captures how the shielding varies with orientation. The shielding Hamiltonian is therefore given by

$$H_s = -\gamma\hbar\hat{\mathbf{I}} \cdot \mathbf{B}_s = -\gamma\hbar\hat{\mathbf{I}} \cdot \sigma \cdot \mathbf{B}_0, \quad (5.6)$$

where B_s is not necessarily oriented in the z-direction, due to the possibly anisotropic electron density. Consequentially, the eigenstates aren't purely eigenstates of I_z . The total Hamiltonian of the system is given by $H_Z - H_s$, where the Zeeman term dominates. Due to the different scales of the interactions, we can solve the equation using perturbation theory. For the first-order energy contributions, the eigenstates

continue to be $|I, m\rangle$, the eigenstates of the Zeeman Hamiltonian. The only contributions to the first-order energy are those that commute with H_Z and therefore I_z . The largest contribution, therefore, originates from z -component of B_s . The total Hamiltonian is hence given by,

$$H_Z - H_s = -\gamma\hbar\hat{\mathbf{I}} \cdot (\mathbb{I} - \sigma) \cdot \mathbf{B}_0, \quad (5.7)$$

where all variables are defined as previously. Considering only the first-order energy contributions, the Larmor frequency is offset,

$$\omega = \gamma(1 - \sigma_{zz})B_0, \quad (5.8)$$

where only the σ_{zz} part of the tensor affects the offset as described earlier.

Note that the shielding tensor σ rotates equivariantly with the underlying nuclear structure. The observed energy levels and observed NMR spectrum, therefore, depend on the molecule's or material's orientation with respect to the applied magnetic field. In solid-state NMR, the material is often in powder form, where the crystals are orientated randomly, and so the observed spectrum is an average over 3D rotations. This leads to broad line shapes, which depend on the entire shielding tensor σ , including its anisotropy.

NMR spectra are normally not displayed directly as a function of ω . Instead the chemical shift, δ is computed with respect to a reference substance defined as,

$$\delta = \frac{\omega - \omega^{\text{ref}}}{\omega^{\text{ref}}} = \frac{\sigma_{zz}^{\text{ref}} - \sigma_{zz}}{1 - \sigma_{zz}^{\text{ref}}} \approx \sigma_{zz}^{\text{ref}} - \sigma_{zz}, \quad (5.9)$$

where ω^{ref} and σ_{zz}^{ref} are the reference Larmor frequency and shielding tensor respectively.

5.2.3 *ab-initio* NMR spectra

NMR spectra are often described as a type of fingerprint for a given material or molecule. For materials, interpreting this fingerprint is often challenging and may require *ab-initio* calculations to elucidate the corresponding 3D structure. *ab-initio* methods can be used to calculate the electron density under some approximations such as density functional theory (DFT). The shielding tensor can then be determined by computing the electronic current induced by the applied field [194–198]. Further details about *ab-initio* magnetic shielding tensors are easily accessible from review papers [199, 200]. Particularly, integrating NMR calculation into plain wave DFT

codes with pseudo-potentials has allowed for a cost-effective determination of *ab-initio* shielding tensors [195]. In addition to the magnetic shielding by the electrons, NMR shifts also depend on electric field gradients, as well as coupling between nuclear spins. These effects can also be computed with DFT codes like CASTEP [197].

5.2.4 Stochastic Liouville equation

As NMR spectra depend on the local electronic environment, molecular motion can affect the spin Hamiltonian and thereby alter observed spectra. Dynamics at the atomic scale can lead to the merging of distinct peaks or the narrowing of individual peaks. The effects of dynamics are often referred to as chemical exchange. We model these dynamics using a stochastic jump model, where the system occupies one of N discrete environments. The Hamiltonian is piecewise constant and switches between environments via Markovian, memoryless jumps. Each environment has an associated chemical shielding tensor and therefore Hamiltonian. The jumps and their effects on ensemble-averaged dynamics are described by the stochastic Liouville equation. At thermal equilibrium, the rates also determine the relative occupation of various environments due to detailed balance. The package EXPRESS [192], explicitly models the density matrix and its time evolution assuming a stochastic change Hamiltonian. The time evolution of the density matrix is described by the Liouville–von Neumann equation and can be derived directly from the Schrödinger equation. We use the stochastic hop approximation outlined above when capturing dynamics with high barriers, which occurs rarely.

5.3 Computing dynamic NMR spectra

We now outline a workflow for using machine learning force fields to incorporate dynamics into computational NMR spectra.

5.3.1 Exploration of dynamic events

Firstly, we explore the potential energy landscape of our material using a mixture of molecular dynamics and global optimization methods. Specifically, we perform minima hopping [201], which consists of alternating NVE molecular dynamics with geometry optimization to find relevant minima. During the NVE molecular dynamics, the number of atoms, volume, and total energy are constant. The same method was used in Section 3.4.3 to explore carbon dioxide adsorption configurations. Individual NVT molecular dynamic runs are performed at each minimum to capture complex

short-term motion. The extensive minima exploration allows a thorough sampling of relevant equilibrium configurations beyond the time-scale limits of direct molecular dynamics (MD). Transition-state searches are used to compute rates between minima, enabling us to model rare dynamic events. This extends the applicability to dynamic events beyond nanoseconds, which still affect NMR spectra. The computational cost of capturing the relevant dynamics is significantly reduced by leveraging a fine-tuned machine learning forcefield. This can be trained iteratively, as outlined in [Chapter 4](#), where samples are taken from the desired simulations.

5.3.2 Modelling NMR spectra across time-scales

The effect of dynamic processes on the observed NMR spectra significantly depends on the relationship of the process' exchange rate, k , and the absolute resonant frequency difference, $\Delta\nu$, between the initial and final states. This separates the effect of chemical exchange into three regimes: slow exchange ($k \ll \Delta\nu$), where each environment is well resolved in the NMR spectrum; fast exchange ($k \gg \Delta\nu$), where rapid dynamics lead to a complete averaging of individual NMR environments; or intermediate exchange ($k \approx \Delta\nu$) when NMR environment coalescence is observed but averaging is only partial. See [Figure 5.1](#) for examples of coalescence (right).

Using the molecular dynamics samples and transition states obtained as described in the previous section, we can compute NMR spectra across chemical exchange regimes. [Figure 5.1](#) visually outlines the steps on the left while showing the effects on the computational spectra compared to the experimental ground truth for each step. The example system is a metal-organic framework dosed with CO_2 . [Section 5.4.1](#) goes into further detail about this specific system and the observed dynamics.

Assume that we have a set of minima $\{\mathbf{R}_i\}$, where \mathbf{R}_i corresponds to a distinct minimum configuration. For each minimum, we have a set of molecular dynamics samples. For every MD sample we calculate the corresponding *ab-initio* chemical shift tensors $\sigma_j^{(i)}$ and electronic field gradients $\mathbf{V}_j^{(i)}$, corresponding to the j th MD sample of minimum i .

Fast regime Fast dynamic motion results in a direct averaging of the shielding tensor. The effective shielding tensor $\sigma_{\text{eff}}^{(i)}$ is hence given by,

$$\sigma_{\text{eff}}^{(i)} = \sum_j \sigma_j^{(i)}, \quad (5.10)$$

where we average the tensor pointwise over all MD samples for a specific minimum. The effect on the spectrum is visualised in [Figure 5.1C](#),

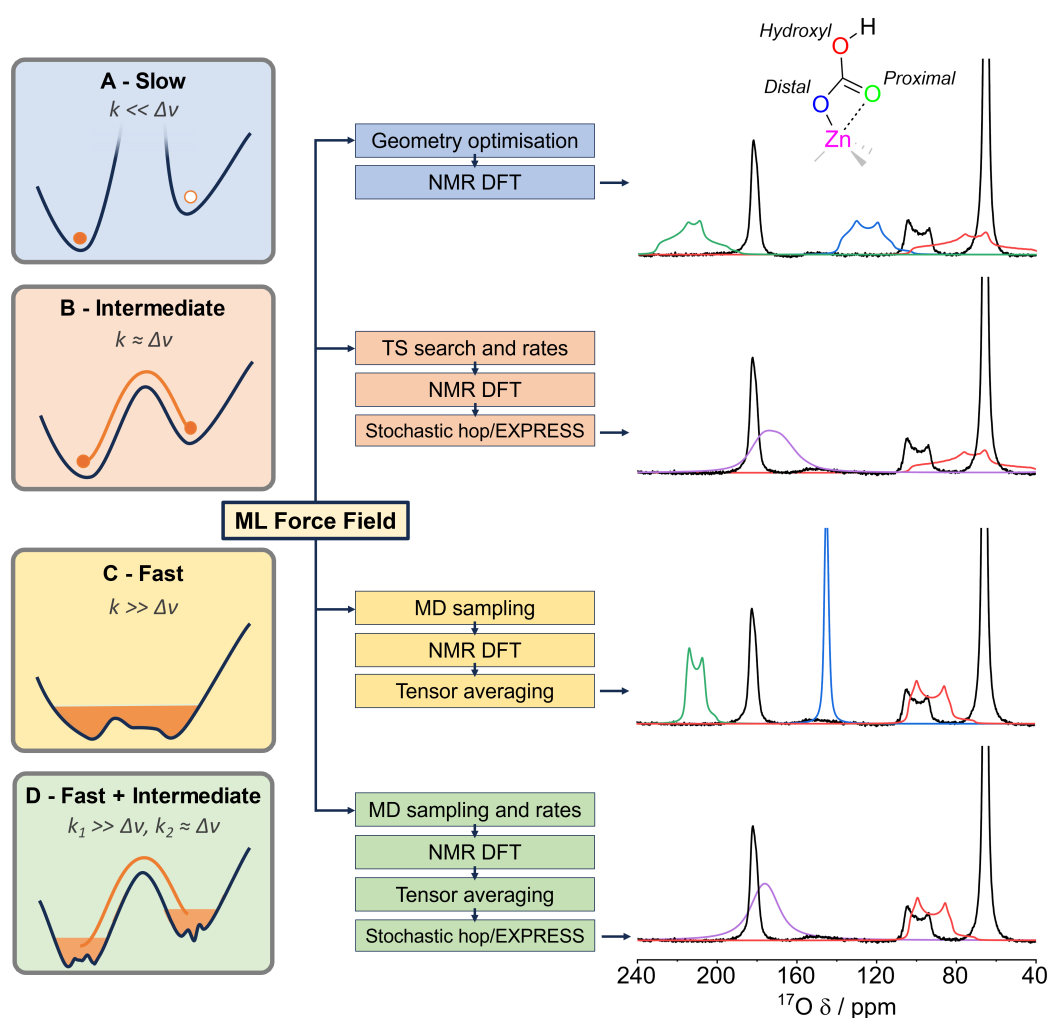


Figure 5.1 Capturing dynamic effects of NMR spectra across timescales. The ground truth experimental ^{17}O NMR spectrum (right, black) of MFU-4l MOF dosed with carbon dioxide. Computational NMR spectra associated with the individual oxygen atoms are shown in colour. Schematics of the exchange regimes and workflows utilising ML force fields are displayed (left). Panel D demonstrates the most accurate model combining fast and intermediate exchange modelling. Note that the peak on the right at <80 ppm is unadsorbed carbon dioxide and is hence not relevant for elucidating the adsorption mechanism.

resulting in a narrowing of the line shape, as the anisotropic parts of the tensor decay.

Intermediate regime To capture longer timescale events we combine the shielding tensors of distinct local minima $\sigma^{(i)}$ and $\sigma^{(i')}$. The minima configurations are found during minima hopping [201]. Using the exchange rates between the minima calculated using a transition-state search, we use the stochastic hopping package EXPRESS, [192] to compute the effective NMR shielding tensors.

Fast + Intermediate regime To combine the effects of fast and intermediate exchange processes, we compute the two MD-averaged tensors as in Figure 5.1C. We then use EXPRESS to model a stochastic hopping between MD averaged tensors $\sigma_{\text{eff}}^{(i)}$ and $\sigma_{\text{eff}}^{(i')}$ as illustrated in Figure 5.1D.


5.4 Application to metal-organic frameworks for carbon capture

We now apply this approach to understand CO₂ adsorption in promising metal-organic frameworks (MOFs). Metal-organic frameworks are highly tunable porous materials that consist of metal clusters connected by organic ligands, forming a 3D lattice structure. They are characterised by large surface areas and large pore volumes, which makes them interesting for gas storage, separation, and catalysis. Furthermore, MOFs can be modified post-synthesis to improve selective interaction with specific molecules.


In this chapter, we focus on the carbon-dioxide adsorption properties of selected MOFs, namely MFU-4l and CD-MOF. NMR spectroscopy is an established tool for exploring the mechanisms of CO₂ capture in a large variety of materials [202–204]. Most NMR studies focus on ¹³C NMR measurements, with ¹⁵N NMR also utilised in amine-based systems [205, 206]. In hydroxide systems, however, such ¹³C NMR data can be ambiguous due to the similar chemical shifts of certain products (e.g. bicarbonate and carbonate) or require further NMR measurements which can introduce their own ambiguity [207, 208]. Recently, *Berge and Pugh et al.* demonstrated ¹⁷O NMR spectroscopy as a novel and effective tool for deconvoluting CO₂ capture mechanisms in amine-functionalised metal-organic frameworks [209]. Here, the material is dosed with special CO₂, where the oxygen is enriched ¹⁷O to be detectable by NMR. As the

natural abundance of the ^{17}O isotope is low (less than 0.04%), the peaks correspond almost exclusively to oxygen atoms originating from enriched CO_2 [209]. This gives us a unique way to track what happens to CO_2 as the MOF adsorbs it.

Both MOFs under investigation are hydroxide-based, where the active site is a metal-bound OH^- group. This class of MOF has been demonstrated to achieve high capacities [210, 211] and high stabilities [212–215]. Carbon dioxide can react with the metal-bound hydroxide to form a metal-bicarbonate species.



However, in systems with unreacted adjacent and/or labile hydroxide anions, successive capture (Reaction 5.12) and deprotonation (Reaction 5.13) steps may occur to form free bicarbonate/carbonate chemisorbed species.



Note that the bicarbonate product mechanism captures twice the number of CO_2 molecules per hydroxide molecule, compared to a carbonate product process.

Experimentally, distinguishing between these mechanisms is challenging. We show that ^{17}O is an effective tool for distinguishing between the bicarbonate and carbonate mechanisms. However, for the interpretation of the experimental spectra, a computational treatment of the dynamics is required. Firstly, we apply the protocol introduced in this chapter to the well-studied MFU-4l MOF [213, 214]. This MOF acts as a validation of the method, as the bicarbonate adsorption mechanism of the MFU-4l MOF is well established. Here the combined NMR-MLFF approach reveals details about the dynamic nature of the bound bicarbonate species. Secondly, we investigate the ^{17}O NMR spectrum of KHCO_3 and K_2CO_3 crystals. This helps us draw conclusions about the mixed bicarbonate-carbonate capture mechanism for a more complex KHCO_3 -CD-MOF.

5.4.1 MFU-4l MOF

We now apply this approach to understand CO_2 adsorption in a promising metal-organic framework (MOF) as illustrated in Figure 5.1. Here, the material is dosed with special CO_2 , where the oxygen is enriched ^{17}O to be detectable by NMR. As the natural abundance of the ^{17}O isotope is low (less than 0.04%), the peaks correspond almost exclusively to oxygen atoms originating from enriched CO_2 [209]. This gives




Figure 5.2 Dynamics of adsorbed CO₂ in MFU-4l MOF across chemical exchange regimes. (A) A cartoon representation of the adsorption process from carbon-dioxide to bicarbonate, rotational tumbling in the fast regime and the rotation of a hydroxyl group in the intermediate regime. (B) A representation of the crystal structure of MFU-4l with the Zn-OH moieties pointing into the pore cavity (Zn - light blue, O - red, C - grey, N - dark blue). (C) The experimental ¹⁷O NMR spectrum (23.5 T, 20 kHz MAS, 10.8 h, hahnecho) of MFU-4l dosed with 0.8 bar C¹⁷O₂. * Denotes spinning side bands. The effect of molecular dynamics on the observed spectrum is the focus of Figure 5.1. The experimental work and figure were created by Benjamin Rhodes.

us a unique way to track what happens to CO₂ as the MOF adsorbs it. When the MFU-4l MOF is dosed with carbon dioxide, it reacts with the MOF forming bicarbonate. This is immediately visible in the experimental spectrum, to the right of Figure 5.1 in black.

The experimental spectrum has three peaks; the far right peak corresponds to carbon dioxide, which has not reacted with the MOF and is physisorbed. We focus only on modelling the two interesting peaks, which correspond to CO₂ strongly chemisorbed by the MOF. As visible from the **static picture** (Figure 5.1A), there are 3 distinct chemisorbed peaks in the computed NMR spectrum (in colour), corresponding to the three oxygen atoms of the bicarbonate. The **intermediate regime** (Figure 5.1B) reveals why the experimental spectrum only has two chemisorbed peaks. We find that the bicarbonate proton can rotate around its bonded oxygen atom, effectively causing the two oxygen environments to be identical, as illustrated in Figure 5.2. A transition-state search finds the barrier to be 0.49 eV. Combined

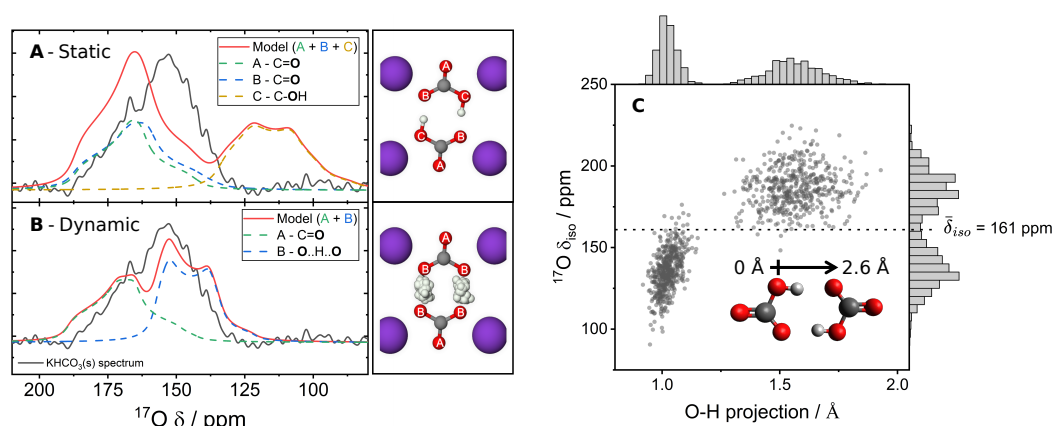


Figure 5.3 Fast regime dynamic modelling of $\text{KHCO}_3(\text{s})$. Comparison of computational static (**A**) and dynamic (**B**) ^{17}O NMR models compared to experimental spectra of $\text{KHCO}_3(\text{s})$ (23.5 T, 20 kHz MAS, 33.0 h, hahnecho). To the side, illustrations of the crystal structure of $\text{KHCO}_3(\text{s})$ show the unique oxygen (red) environments, as well as samples of the hydrogen (white) positions in the dynamic case (grey = carbon, purple = potassium). (**C**) A plot of ^{17}O isotropic shift of all the now equivalent (under exchange) dynamic oxygen environments against the projection (onto the oxygen-oxygen axis) of oxygen-proton distance of all the MD samples, the mean δ_{iso} is marked.

with a vibrational analysis of the start and transition state configuration, the predicted reaction rate corresponds to 2.4×10^5 Hz under the harmonic approximation at room temperature. The EXPRESS [192] calculation shows that this rate causes a coalescence of the two peaks. However, the computational peaks are too wide, especially the peak at 90 ppm. Secondly, a nanosecond-long MD simulation reveals a rotational tumbling of the OH-group of the bicarbonate. Including the effects of these **fast regime** (Figure 5.1C) dynamics leads to a reduction in the average anisotropic chemical shielding resulting in narrower line widths, particularly visible at 90 ppm. Only by combining the **fast and intermediate approach** (Figure 5.1D), we obtain strong agreement with the two chemisorbed experimental peaks.

5.4.2 Potassium bicarbonate

In preparation for understanding CO_2 capture of the more complex CD-MOF, we investigate potassium bicarbonate (KHCO_3) crystals. Figure 5.3, shows its crystal structure and the observed NMR spectrum. Note that in this case, the Oxygen atoms of the structure do not originate from adsorbed CO_2 . The percentage of NMR-active nuclei is just that of the natural abundance of ^{17}O of approximately 0.037%. Consequently, the NMR measurements had to be performed over a longer duration, and the spectrum is noisier.

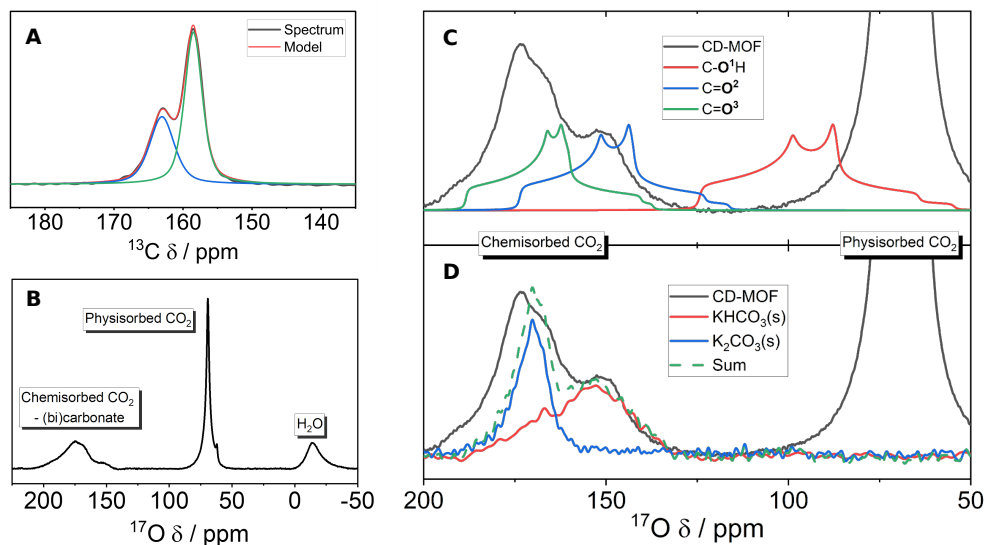


Figure 5.4 The NMR spectra and assignments of CO_2 dosed KHCO_3 CD-MOF. Showing both ^{13}C (A) and ^{17}O (B) NMR spectra of the KHCO_3 CD-MOF dosed with 1 bar enriched CO_2 . We compare the chemisorbed part of the spectrum with single-point DFT predictions (C). Furthermore, we show the ^{17}O spectrum overlaid with the NMR experiments of $\text{KHCO}_3/\text{K}_2\text{CO}_3$ solid powders. Y-axis scaling in arbitrary units.

As visible in panel A, there are three unique oxygen atoms in the crystal structure. From purely static DFT calculations, we find two of these to be around 160 ppm and one at 110 ppm, leading to a bimodal NMR spectrum. The bicarbonate dimer motif within the crystal structure is known to facilitate a rapid proton hop between the hydroxyl-carbonyl pairs [216–220]. Using a purpose-trained MLFF to perform a transition state search, we determine a proton hop frequency of 140 GHz. This corresponds to the fast chemical exchange regime.

We investigate in further detail how the proton hop dynamics affect the observed spectrum. We sample configurations from nanosecond MD simulations and perform single-point DFT-NMR calculations on the sampled structures. Figure 5.3C shows the isotropic chemical shift as a function of O-H distance of the NMR active oxygen. The symmetric proton hopping (Figure 5.3C) produces a fast exchange average value of $\delta_{iso} = 161$ ppm, and, through electric field gradient (EFG) tensor averaging, a $C_Q = 7.0$ MHz and $\eta_Q = 0.2$ for this new ‘dynamic hydroxyl’ peak. The resulting dynamic peak displayed in red in Figure 5.3B shows strong agreement with the experiment in black.




Figure 5.5 Proposed CO₂ binding mechanism for KHCO₃-CD-MOF. (A) An assigned $\{^1\text{H}\}^{17}\text{O}$ NMR spectrum of C¹⁷O₂ dosed (0.8 bar) CD-MOF with a three environment fit (blue, purple, green) derived from MQMAS and multi-field fitting data. Comparing experiment (black) and summed fitted spectrum (red). (B) A snapshot from an MLFF-MD simulation of the CD-MOF demonstrating an example proton-hopping pathway for the ‘dynamic hydroxyl’ environment.

5.4.3 CD-MOF

One of the key promises of ¹⁷O NMR is the ability to differentiate between the bicarbonate and the carbonate adsorption mechanism outlined in reactions 5.11 - 5.13. Therefore, we now look at γ -cyclodextrin metal-organic frameworks (CD-MOFs), where previous studies have not identified the mechanism [215, 221, 222].

CD-MOFs have recently been demonstrated to leverage labile hydroxide-based chemistry to achieve promising CO₂ capture performance under flue gas conditions [215]. The best performing MOF in this family, KHCO₃-CD, demonstrated thermal, oxidative and cycling stabilities along with reasonable capacities for post-combustion CO₂ capture (1.43 mmol g⁻¹ at 15% CO₂/85% N₂ atmosphere at 30°C and 0.06 mmol g⁻¹ at 0.4 mbar CO₂, 25°C). A bicarbonate capture mechanism involving non-metal bound hydroxide counterions within the framework pores (Reaction 5.12) was proposed [215]. The free nature of the hydroxide results from the dissociation of the K-OH bond due to the significant hydrogen bonding from the cyclodextrin sugar framework [222]. However, ambiguity remains in the adsorption mechanism with IR, ¹³C NMR and heats of adsorption data that are unable to fully eliminate the possibility of carbonate formation (Reaction 5.13).

Due to the size of the CD-MOF structure, applying the protocol introduced earlier is computationally too costly. With more than a thousand atoms, creating a sufficient large dataset and calculating DFT-NMR on multiple MD samples is unfeasible. A single-point DFT-NMR calculation on the entire MOF structure reveals that the static

computational prediction is in strong disagreement with the experiment. Similarly to the case of KHCO_3 -crystals of the previous section, the computational spectrum contains three distinct peaks as visible in [Figure 5.4](#).

Therefore, we leverage our understanding from reference systems, namely KHCO_3 and K_2CO_3 , which contain typical bicarbonate and carbonate crystal atomic environments. We first investigated the NMR spectrum experimentally by comparing the CD-MOF spectrum with that of the two crystals. [Figure 5.4](#), shows that the CD-MOF spectrum is similar to the sum of the two experimental spectra.

Our understanding of proton hopping in the case of KHCO_3 indicates that a similar dynamic process occurs in the CD-MOF, leading to the averaging of the two bicarbonate peaks. To investigate this process, we use the MACE-MP0 foundation model to run molecular dynamics on the CD-MOF crystal structure. We find an extensive hydrogen bond network surrounding the active sites, as visible in [Figure 5.5](#). The carbonyl-hydroxyl proton transfer provides a plausible mechanism for a dynamic hydroxyl group within the bicarbonate.

5.5 Discussion

The correct characterisation of complex materials and mechanistic understanding of reactions within porous materials requires experimental techniques that capture atomic-level detail. In this chapter, we showed that for many systems, the use of nuclear magnetic resonance (NMR) spectroscopy may be challenging because of the difficulty in interpreting experimental spectra, primarily due to the dynamic nature of materials. We presented an approach which automates the incorporation of dynamic effects into computational NMR spectra, utilizing cost-effective fine-tuning of expressive atomistic foundation models to capture dynamics across multiple timescales accurately. By combining tensor averaging with stochastic hopping methods, we achieve strong agreement between computational and experimental results. Although demonstrated on a metal-organic framework for carbon capture, this method offers a generalisable solution for dynamic materials. The workflow opens new opportunities for automated characterization, enabling large-scale material screening and advancing automated material design. To reduce the computational cost of this approach, steps could be taken to learn the magnetic tensorial properties [223], potentially as a future foundation model. This would also allow for a more in-depth treatment of the CD-MOF.

Chapter 6

Matrix Function Neural Networks

So far this thesis has explored the use of local machine learning force fields. However, the approximation that the total energy can be separated into atom site energies that depend only on their local environment doesn't hold for all systems and can lead to qualitatively wrong behaviours. In this chapter, we explore the complexity of modelling inherently non-local quantum mechanical interactions and introduce a new machine learning architecture. Matrix function neural networks leverage the non-locality of the spectrum of a learnable matrix that is parametrised based on the geometric atomic configuration.

We start by motivating the need for non-local models and give a brief overview of existing approaches. We then go into detail of the matrix function neural network (MFN) architecture. We illustrate its competitive performance on cumulenes which exhibit stark non local behaviour. We create a new dataset that has guaranteed non-local interaction for the community to benchmark their approaches. Finally, we give motivation as to why architectures that leverage spectral information may allow for the learning of complex non-local interactions in a generalisable manner.

6.1 Background and motivation

The success of machine learning force fields can be attributed to their order of magnitude reduction in computational cost, high accuracy when trained on minimal data, and their transferability in the large data regime. One key ingredient that aids in the transferable nature and linear scaling with system system size is the locality approximation. As introduced in [Section 2.4.1](#), rather than comparing and regressing the energy of entire systems, the total energy is split into atom site energies that depend only on their immediate neighbourhood within a set cut-off.

There are various system classes where the locality approximation fails [224, 225]. An evident example is electrostatic interactions, which decay as $1/r$. Systems that contain unscreened charged or polar molecular clusters are evidently badly described by local force fields. Various approaches have been explored to include long-range interactions by explicitly representing the charge density [226–232]. Such effects are very relevant for biological applications, where localised charges in proteins interact with dissolved ions and effect the surrounding water structure [233, 234]. Furthermore, there are some applications require the modelling of an external electric fields, such as in batteries, supercapacitors and electro-catalysis. This kind of interaction is often referred to as long-range interaction.

Additionally, there are other types of interactions which are more quantum mechanical in nature, such as the behaviour of conjugated molecules. These kinds of interactions are often qualitatively better described by empirical simple models like the Hückel method. One example that we will use is a chain of carbon atoms terminated by two hydrogen atoms on either side. This molecule’s energy changes significantly as a function of the dihedral angle between the 4 hydrogens (see Figure 6.7). Clearly, the distances between atoms change minimally, and this is not a simple $1/r$ classical Coulomb interaction. Throughout this Chapter, we will be referring to this kind of interaction as non-local, in contrast to long-range interactions, which commonly refer to coulomb-type interactions. In this chapter, we present an architecture that is able to accurately model non-local interactions in a transferable way.

This work is published in Reference [235] and is the result of collaborative work. The author of this thesis contributed to the design of the architecture, coding and validation of the models.

6.2 Related work

Previous work has focused on coarse graining ab-initio Hamiltonians [236–238]. Semi-empirical tight binding approaches, for example are often parametrised to replicate DFT hamiltonians. Tight-binding models give good qualitative agreement. Recently, differentiable approaches have been proposed to machine learn larger sets of parameters to improve accuracies and obtain quantitative agreement [236]. Further approaches have been suggested which learn only part of the Hamiltonians spectrum close to the fermi-level [239]. Additionally, Cignoni et al. have explored this design space by training on the Hamiltonian directly, but also on only part of the spectrum using a smaller basis set compare to the reference method [240, 241].

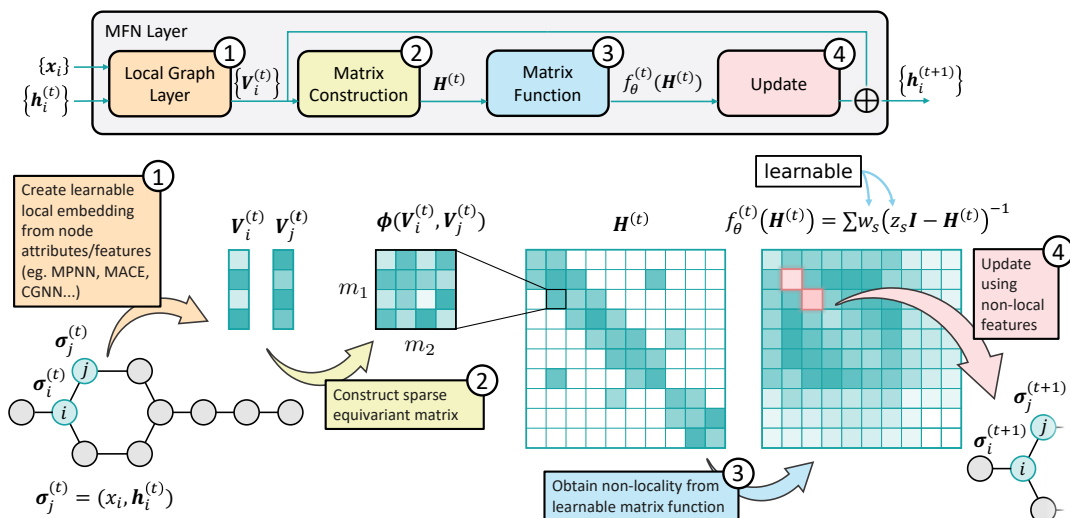


Figure 6.1 Matrix function network architecture. Illustrating matrix construction and non-locality of matrix functions on a molecular graph.

In addition to architectures which parametrise the tight binding Hamiltonian directly, hybrid approaches have been designed. Stöhr et al. have trained repulsive machine learning force fields to correct the force and energy predictions of DFT-TB methods [242].

In comparison to tight-binding or direct Hamiltonian learning, MFN’s are purely trained on energy and forces and not the Hamiltonian directly. Furthermore, while for tight binding approaches the matrix function is commonly the fermi function, in the case of the MFN architecture it is learnable. There is hence no explicit relation between the learned matrix and the hamiltonian of the reference method.

6.3 The Architecture

Matrix function neural networks (MFN) are a framework for learning group equivariant matrix operators of the graph. In principle, the framework applies to any reductive Lie group, such as the Lorentz group, relevant in the context of high-energy physics. We, however, will focus on the use case of atomistic modelling and hence will focus on the rotation and inversion group $O(3)$.

For constructing machine learning forcefields, we use matrix functions in combination with message passing. Assume that, as in Section 2.4.6, we have undirected graph \mathcal{G} embedded in three-dimensional Euclidean space. The architecture operates in four stages at each layer, a **local graph layer** that learns local node features from atomic-neighborhoods, the **matrix construction**, the **matrix function**, and the

update. These are visually outlined in [Figure 6.1](#). The local graph layer is simply a MACE model and is used to construct descriptive many-body node features. The node features are combined to form edge features for all edges in the graph \mathcal{G} . Lastly, the learnable matrix functions introduce non-local interactions. We now look at each of these stages in the t th MFN layer.

6.3.1 Local graph layer

As in the MACE architecture, assume the state of each node i at layer t of the MPNN is given by a tuple,

$$\sigma_i^{(t)} = (\mathbf{r}_i, z_i, \mathbf{h}_i^{(t)}), \quad (6.1)$$

where \mathbf{r}_i is the position in 3D space, z_i are the chemical elements, and $\mathbf{h}_i^{(t)}$ are the learnable features. For each iteration t , the first step in an MFN layer is to form equivariant node features using a local graph layer (abstractly noted \mathcal{L}) as a function of the **local** environment of a node i ,

$$\mathcal{L}^{(t)} : \{\sigma_j^{(t)}\}_{j \in \mathcal{N}(i)} \mapsto V_{i\tilde{c}\tilde{m}}^{(t)}, \quad (6.2)$$

where $V_{i\tilde{c}\tilde{m}}^{(t)}$ are learned **local** equivariant node features and \tilde{m} label the basis vectors of the representation of G acting on \mathbf{V} , such that

$$\mathbf{V}_{\tilde{m}} \circ g = \sum_{\tilde{m}\tilde{m}'} \rho_{\tilde{m}\tilde{m}'}(g) \mathbf{V}_{\tilde{m}'}, \quad (6.3)$$

where ρ is a representation of group element $g \in G$. In the case of $\text{SO}(3)$, \tilde{m} corresponds to the tuple (l, m) of the spherical harmonics and ρ are the Wigner-D matrices $\mathcal{D}_m^{(l)}$.

The specific choice of the graph layer \mathcal{L} depends on the application. A more expressive local layer will result in better coverage of the operator space $\mathcal{H}(\mathcal{G})^G$ and, therefore, better general expressivity. We hence use the MACE architecture in the case of $\text{SO}(3)$, although any such function could also be approximated with arbitrary accuracy using an ACE layer [243]. In the first layer, there are no learnable node features and hence $\sigma_i^{(1)} = (\mathbf{r}_i, z_i)$.

6.3.2 Matrix construction

The second step involves constructing a set of graph matrix operators from the learned node features. The space of graph matrix operators, $\mathcal{H}(\mathcal{G})^G$, corresponds to the space of operators that are (1) **self-adjoint**, (2) **permutation equivariant**, and




Figure 6.2 Edge feature construction. The learnable edge features can be constructed in a variety of ways. Common approaches include leveraging previously computed node features [235] of the edge (a) with a receptive field visualised in (b) or constructing edge features based on the edge position and a spherical cut-off [244].

(3) **G-equivariant.** The off-diagonal entries of the matrix operator correspond to learnable equivariant edge features.

The edge features can be constructed in a variety of ways. An approach used by Reference [244] for learning *ab-initio* Hamiltonians directly is centring a spherical basis on the edge centre. Any atoms within a spherical cut-off are then included in the description of the atomic environment. Alternatively, one can construct learnable equivariant functions of the edge's node features. As this allows the reuse of previously computed features and graph connectivity, we take this approach. The overall receptive field of the two edge feature construction is visually compared in Figure 6.2. Independent of the approach, the edge features are then mapped to a reducible representation to form a square matrix of size $\tilde{M} \times \tilde{M}$ matrix for each channel c and edge, where \tilde{M} is the size of the basis. The off-diagonal square matrices can be combined to form a matrix for the entire configuration of shape $\mathbf{H}_c \in \mathbb{C}^{\tilde{M}n \times \tilde{M}n}$, consistent with $n \times n$ blocks of size $\tilde{M} \times \tilde{M}$, where n are the number of nodes.

$$H_{cij, \tilde{m}_1 \tilde{m}_2}^{(t), \text{off}} = \begin{cases} \phi_{c\tilde{m}_1 \tilde{m}_2}^{(t), \text{off}}(\mathbf{V}_i^{(t)}, \mathbf{V}_j^{(t)}, \sigma_i^{(t)}, \sigma_j^{(t)}) & \text{if } j \in \mathcal{N}(i), i \in \mathcal{N}(j), \text{ and } i \neq j, \\ 0 & \text{else,} \end{cases} \quad (6.4)$$

where i, j denote the indices of the matrix blocks, \tilde{m}_1, \tilde{m}_2 represent the indices within each block, and $\phi_{c\tilde{m}_1 \tilde{m}_2}^{(t)}$ is a learnable equivariant function mapping a tuple of node features to an edge feature. Note that the matrix is sparse and is only non-zero for matrix elements corresponding to node edges.

The on-diagonal elements are simply a learnable function of the node features.

$$H_{cij, \tilde{m}_1 \tilde{m}_2}^{(t), \text{on}} = \begin{cases} \phi_{c\tilde{m}_1 \tilde{m}_2}^{(t), \text{on}}(\mathbf{V}_i^{(t)}) & i = j, \\ 0 & \text{else,} \end{cases} \quad (6.5)$$

$$H_{cij,\tilde{m}_1\tilde{m}_2}^{(t)} = H_{cij,\tilde{m}_1\tilde{m}_2}^{(t), \text{off}} + H_{cij,\tilde{m}_1\tilde{m}_2}^{(t), \text{on}}, \quad (6.6)$$

where $H_{cij,\tilde{m}_1\tilde{m}_2}^{(t)}$ is the final learned matrix for a given configuration. Note that the learned matrix H is entirely local. Any non-locality originates from a learnable matrix function of this matrix.

Matrix construction for SO(3)

Here we detail a concrete example for the matrix construction in the case of $O(3)$ -group. The node features are expanded on the spherical basis, and so the matrix \mathbf{H} is structured as a $N \times N$ block matrix,

$$\mathbf{H} = \begin{pmatrix} \mathbf{H}_{11} & \mathbf{H}_{12} & \cdots & \mathbf{H}_{1N} \\ \mathbf{H}_{21} & \mathbf{H}_{22} & \cdots & \mathbf{H}_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{H}_{N1} & \mathbf{H}_{N2} & \cdots & \mathbf{H}_{NN} \end{pmatrix}, \quad (6.7)$$

where N is the number of atoms (nodes) and each block is a sub-matrix \mathbf{H}_{ij} of shape $N(2l_{\max} + 1) \times N(2l_{\max} + 1)$. We index the sub-matrices using the usual convention of the $\tilde{l}\tilde{m}$ index, such that the tuple (l, m) corresponds to the \tilde{m} index from the previous section. As $-l \leq m \leq l$, the single index translates to

$$\begin{aligned} \tilde{m} = 0 &\rightarrow (l = 0, m = 0) \\ \tilde{m} = 1 &\rightarrow (l = 1, m = -1) \\ \tilde{m} = 2 &\rightarrow (l = 1, m = 0) \\ &\dots \end{aligned} \quad (6.8)$$

Note that each sub-matrix is no longer expressed in a basis for the irreducible representations of SO(3). Instead, the matrix elements transform as

$$\mathbf{D}^{(l)} \otimes \mathbf{D}^{(l')}, \quad (6.9)$$

where \otimes is the Kronecker product and \mathbf{D} are the Wigner D matrices. In practice, we take learnable tensor products of node features as in Equation (2.52). We can then map from the resulting spherical tensors expressed in the basis of irreducible representations to the sub-matrix \mathbf{H}_{ij} . We now look at each step in greater detail in the context of atomistic modeling.




Figure 6.3 Block structure of a Euclidean MFN operator, \mathbf{H} . Each entry in \mathbf{H} corresponds to a different product of representations of the group of rotations (ss, sp, pp, \dots). Example for $L=1$ on a chain-like graph resulting in the block-diagonal matrix \mathbf{H} .

The initial node features are constructed using MACE [27] as the **local graph layer** resulting in the features $V_{i,clm}$ of Equation (6.2). We then construct the edge feature ϵ using a learnable tensor product of node features of atoms i and j ,

$$\epsilon_{cij,\eta l_3 m_3}^{(t)} = R_c(r_{ij})(V_{i,cl_1 m_1} \otimes V_{j,cl_2 m_2}), \quad (6.10)$$

where R corresponds to a function of the distance r_{ij} between i and j where $j \in \mathcal{N}(i)$ and $i \in \mathcal{N}(j)$ and \otimes is the learnable tensor product as defined in Equation (2.53). The index η enumerates multiple features with the same degree l_3 . We provide further details of the number of η required with a concrete example of $l_{\max} = 1$ later in this section.

As the graph is undirected, the matrix is symmetric with respect to i, j . Using the Clebsch-Gordan coefficients, we make a change of basis,

$$H_{cij,l_1 m_1, l_2 m_2}^{(t), \text{off}} = \sum_{\eta l_3 m_3} C_{\eta, l_3 m_3}^{l_1 m_1 l_2 m_2} \epsilon_{cij,\eta l_3 m_3}, \quad (6.11)$$

resulting in c_{\max} matrices with shape $N(2l_{\max} + 1) \times N(2l_{\max} + 1)$. The matrix is sparse, as only the elements corresponding to an edge of the graph are non-zero. We refer to l_{\max} as the maximal value of each spherical index l_1 and l_2 of the matrix ($0 \leq l_{1,2} \leq l_{\max}$), and we refer to this architecture as $\text{MFN}_{l_{\max}}^{(\text{MACE})}$ when comparing performance to local and other non-local models in Section 6.7.

Example matrix for $l_{\max} = 1$

We now explicitly construct the matrix for $l_{\max} = 1$. The $\mathbf{H}_{ij} \in \mathbb{C}^{4 \times 4}$ has parallels to the Hamiltonian in DFT expressed in atom-centred basis. The node features are a concatenation of the scalar (s or $l = 0$) and vectorial features (p or $l = 1$). The direct outer product results in a basis with block structure illustrated in Figure 6.3

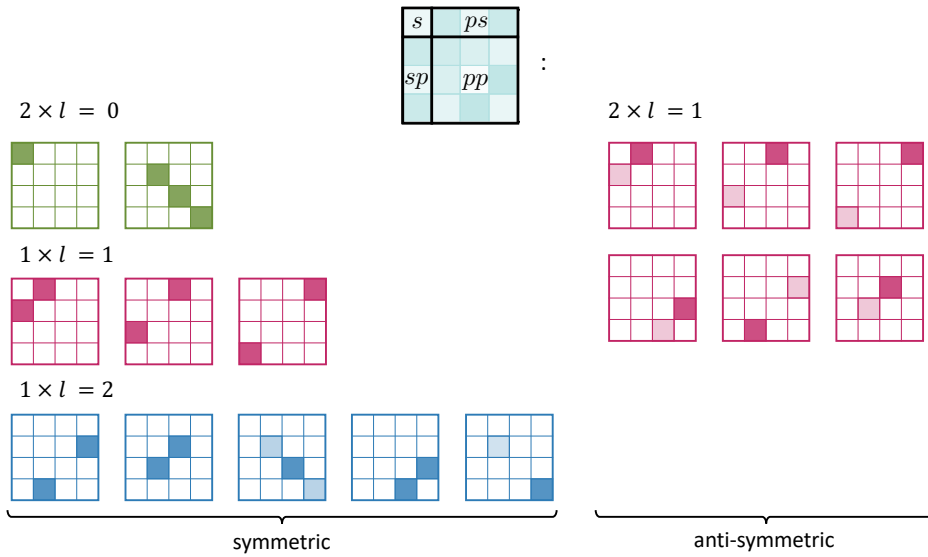


Figure 6.4 Matrix block decomposition for $\text{SO}(3)$. Showing how the tensor product basis of two spherical vectors with maximum degree $l_{\max} = 1$ can be decomposed into a sum of spherical basis vectors.

similar to Hamiltonian matrices in quantum mechanics with ss and sp orbitals. The sub-matrix can be expressed as a sum of matrices that transform like the spherical harmonics as visually illustrated in Figure 6.4. The symmetric part of each sub-matrix can be decomposed into two scalars, one vector and one $l = 2$ tensor. The first scalar corresponds to the ss block in Figure 6.3, and the second can be obtained from the trace of the pp block. The antisymmetric part can be decomposed into: one vector and one pseudo vector. Consequentially when constructing the edge features ϵ of Equation (6.10), η enumerates the required multiplicity for two scalars, two vectors, one pseudo vector and one $l = 2$ spherical tensor. The edge features are then transformed with a change of basis to a square matrix according to Equation (6.11).

6.3.3 Matrix normalization

Batch normalization [245] and layer normalization [246] are two techniques widely used in deep learning to stabilize training, improve convergence, and provide some form of regularization.

In the context of MFNs, a similar normalization strategy is needed to ensure stable training and improved convergence. Instead of normalizing the features directly as conventional normalization layers do, we normalize the eigenvalues of a set of H matrices with batch dimension [batch, channels, n , n], where n is the size of the matrix. The aim is to adjust their mean and variance to 0 and 1, respectively, much

like standardization in traditional batch or layer normalization. For the batch matrix norm, the mean is taken across the batch dimension, and for the layer matrix norm, the normalization is taken across the channel dimension. The mean and the variance of the eigenvalues are computed using the following formulas,

$$\mathbb{E}(\Lambda) = \frac{\text{tr}(\mathbf{H})}{n}, \quad \text{Var}(\Lambda) = \frac{\text{tr}(\mathbf{H}^2)}{n-1} - \frac{\text{tr}(\mathbf{H})^2}{n(n-1)} \quad (6.12)$$

This normalization of eigenvalues ensures that the spectral properties of the graph used for representation in the MFN are effectively standardized, contributing to better training stability and convergence. Note that the batch norm is kept constant during inference and hence does not vary by batch.

6.3.4 Matrix functions

The central operation of the MFN architecture is the matrix function, which introduces long-range many-body effects. Any continuous function $f_\theta : \mathbb{R} \rightarrow \mathbb{R}$, with parameters θ , can be used to define a matrix function that maps a square matrix to another square matrix of equal size. Formally, a matrix function on self-adjoint matrices \mathbf{H} can be defined by its spectral decomposition. The symmetric matrix \mathbf{H} is diagonalizable, if there exists an invertible matrix U and diagonal matrix Λ such that

$$\mathbf{H} = \mathbf{U}\Lambda\mathbf{U}^T, \quad (6.13)$$

where the column vectors of \mathbf{U} are the orthogonal eigenvectors of \mathbf{H} and Λ is a diagonal matrix consisting of the eigenvalues $\{\lambda_i\}$ of H . For any analytic scalar function f defined on the spectrum of H , there exists a corresponding matrix function f defined by,

$$f_\theta(\mathbf{H}) = \mathbf{U}f_\theta(\Lambda)\mathbf{U}^T = \mathbf{U} \begin{bmatrix} f_\theta(\lambda_1) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & f_\theta(\lambda_n) \end{bmatrix} \mathbf{U}^T, \quad (6.14)$$

where the function $f : \mathbb{R} \rightarrow \mathbb{R}$ is applied pointwise to the diagonal elements of Λ . The eigenvalue decomposition is the source of the non-locality. As such, the architecture mimics quantum mechanical calculations. Note that due to numerical stability and computational cost in practice we don't diagonalise the matrix. We use the resolvent expansion to allow for the possibility of linear scaling and allow for training on forces.

Equivariance of matrix function

An essential observation is that any continuous matrix function f_θ preserves equivariance. Let $\rho(g)$ be a matrix representation of group element $g \in SO(3)$ of shape $n \times n$ and $\mathbf{H} \in \mathbb{C}^{n \times n}$ be a self-adjoint matrix. We define matrix \mathbf{H}' as the resulting matrix after the group action g on \mathbf{H} such that,

$$\mathbf{H}' = \mathbf{H} \circ g = \rho(g)\mathbf{H}\rho^{-1}(g) = \rho(g)\mathbf{H}\rho^*(g), \quad (6.15)$$

where $\rho^*(g)$ is the conjugate transpose of g . Given that the symmetric matrix \mathbf{H} is diagonalizable as $\mathbf{H} = \mathbf{U}\Lambda\mathbf{U}^{-1} = \mathbf{U}\Lambda\mathbf{U}^T$, where \mathbf{U} is a unitary matrix, we see that,

$$\mathbf{H}' = \rho(g)\mathbf{H}\rho^{-1}(g) = \rho(g)\left(\mathbf{U}\Lambda\mathbf{U}^{-1}\right)\rho^{-1}(g) = (\rho(g)\mathbf{U})\Lambda(\rho(g)\mathbf{U})^{-1}, \quad (6.16)$$

meaning that \mathbf{H}' has the same eigenvalues Λ as the untransformed matrix \mathbf{H} and has eigenvectors which are transformed as $\rho(g)\mathbf{u}$, where $\mathbf{u} \in \mathbf{U}$ are the original eigenvectors.

A matrix function on the transformed matrix \mathbf{H}' is given by,

$$f(\mathbf{H}') = f\left([\rho(g)\mathbf{U}]\Lambda[\rho(g)\mathbf{U}]^{-1}\right) = [\rho(g)\mathbf{U}]f(\Lambda)[\rho(g)\mathbf{U}]^{-1} \quad (6.17)$$

$$= \rho(g)\mathbf{U}f(\Lambda)\mathbf{U}^{-1}\rho^{-1}(g) = \rho(g)f(\mathbf{H})\rho^{-1}(g), \quad (6.18)$$

and hence analytical matrix functions are $SO(3)$ equivariant such that

$$f(\rho(g)\mathbf{H}\rho^{-1}(g)) = \rho(g)f(\mathbf{H})\rho^{-1}(g). \quad (6.19)$$

As a consequence, after applying the learnable matrix function, we can extract features from the resulting matrix and map these back to irreducible representations. A scalar quantity, for example, in the ss block, will remain invariant with rotations.

6.3.5 Update

There are various ways the resulting matrix $f(H)$ can be used in a message-passing context. We will outline three plausible choices which have consequences on the computational cost and non-locality of the entire architecture: the diagonal, sparse, and dense update. In practice, we restrict ourselves to the diagonal update as it has the most potential for linear scaling implementations.

Diagonal update

The diagonal update updates the state of each node with non-local features extracted from the diagonal blocks of the matrix function. We hence extract the required features

$$h_{icm}^{(t+1)} = V_{icm}^{(t)} + \sum_{\tilde{c}} w_{c\tilde{c}}^{(t)} f_{\theta}^{(t)}(\mathbf{H}_{\tilde{c}}^{(t)})_{ii,m0}, \quad (6.20)$$

where m indexes the relevant basis.

In the case of force fields, we use the MACE architecture as our underlying message-passing framework, where node features are expressed in a spherical basis with a maximum degree l_{\max} . Following the equation the diagonal blocks can be mapped to decomposed into spherical tensors up to the required degree.

This method is the most computationally efficient since selected inversion techniques [247] can be employed to efficiently evaluate the diagonal blocks of a matrix function; cf. [Section 6.4](#).

Dense update

For the dense update, each configuration is assigned a matrix feature. In subsequent message-passing layers, the previous configuration feature is added as a residual in the matrix construction step.

$$H_{cij,l_1m_1,l_2m_2}^{(t), \text{off}} = \sum_{\eta l_3m_3} C_{\eta,l_3m_3}^{l_1m_1,l_2m_2} \epsilon_{cij,\eta l_3m_3}^{(t)} + \sum_{\tilde{c}} w_{c\tilde{c}}^{(t)} f_{\theta}^{(t-1)}(H_{cij,l_1m_1,l_2m_2}^{(t-1), \text{off}}) \quad (6.21)$$

As a consequence, the non-local off-diagonal elements from the previous layer are kept. Note that the $f(\mathbf{H})$, does not necessarily maintain the sparsity of the original matrix H as visually illustrated in [Figure 6.1](#).

Sparse update

Similar to the dense update, each configuration is assigned a matrix feature. However, rather than adding all matrix entries of the previous layer during the matrix construction step, only those that correspond to edges in the graph are added. This insures that the initial matrix H remains sparse, allowing for computational approaches which leverage sparsity for computational speed.

The optimal kind of update is a trade-off between expressivity and computational cost. The node states are then updated using these new non-local node features $\sigma_i^{(t+1)} = (x_i, \mathbf{h}_i^{(t+1)})$ to form the next states. The steps are repeated for T iterations, starting from the local graph layer.

6.3.6 Readout

The readout phase maps the resulting non-local node features to observable properties. Here we simply follow the same approach as MACE [28], as outlined in Section 2.4.6.

6.4 Resolvent parameterization of matrix functions

The evaluation of the matrix function in Equation (6.14) is the practical bottleneck of our method. Directly diagonalising the matrix scales cubically with matrix dimension. Note that the matrix dimension scales quadratically with the maximum degree l_{\max} making this a significant restriction. Furthermore, the second derivatives of the eigenvalue decomposition are numerically unstable in `pytorch`. As we fit on forces, we need to back-propagate through first derivatives, so this approach is not possible.

An alternative to defining a matrix function through its eigenvalue decomposition is through its resolvent expansion. For a general analytic $\tilde{f} : \mathbb{C} \rightarrow \mathbb{C}$, a matrix function can be expressed as

$$\tilde{f}(\mathbf{H}) = \oint_{\mathcal{C}} \tilde{f}(z)(z\mathbf{I} - \mathbf{H})^{-1} \frac{dz}{2\pi i}, \quad (6.22)$$

where \mathcal{C} is a curve encircling the eigenvalues of \mathbf{H} and excluding any poles of \tilde{f} . We can approximate the contour integral with a quadrature rule such that

$$\tilde{f}(\mathbf{H}) \approx \sum_s \tilde{w}_s \tilde{f}(z_s)(z_s\mathbf{I} - \mathbf{H})^{-1}, \quad (6.23)$$

where z_s and \tilde{w}_s are the nodes and corresponding weights respectively. Expressing the coefficients with a single variable $w_s := \tilde{w}_s \tilde{f}(z_s)$ we arrive at the parametrisation,

$$f_{\theta}(\mathbf{H}) = \sum_s w_s (z_s\mathbf{I} - \mathbf{H})^{-1}. \quad (6.24)$$

The inverse is computed by firstly doing a LDL^* factorization, where L is a lower triangular matrix and D is a diagonal matrix and then computing the inverse. This approach is more numerically stable than other approaches such as Gaussian elimination and allows for stable training. Furthermore the LDL^* factorization can benefit from the sparsity of the matrix \mathbf{H} . Depending on the sparsity pattern the number of operations can be reduced from scaling as $O(N^3)$ to $O(n)$ for quasi one-dimensional systems such as nano-tubes. Indeed pole expansion framework is




Figure 6.5 Matrix powers series and non-locality. Showing the matrix function as a power series illustrates its convolutional nature.

the first key ingredient in linear scaling electronic structure methods [248] such as PEXSI [247, 249].

Compared to the standard usage of pole expansions in computing matrix functions [250] both the weights w_s and the poles z_s are learnable parameters. In the limit of infinitely many pole-weight pairs (z_s, w_s) , any analytic matrix function can be expressed in this form. To ensure no divergence for any geometric structure, we can restrict the poles to have non-zero imaginary parts as the matrix \mathbf{H} is self-adjoint, and hence all eigenvalues are real. Additionally, we choose adjoint pole weight pairs (w_s, z_s) and (w_s^*, z_s^*) to ensure that f_θ is real when restricted to real arguments.

6.5 Understanding the non-locality of Matrix Function Networks

We now investigate the non-locality of matrix functions in greater detail. Firstly we show that features obtained from using the MFN architecture would require infinite layers to construct with message passing. This is of interest for modelling causal relationships or non-local interactions by proxy, such as in chemistry or natural language processing. In these cases, the propagation of local effects over long distances results in multiscale effects that are effectively captured by our method. We then focus on the resolvent expansion and investigate the non-locality as a function of pole position. These observations directly impact design choices in the architecture.

6.5.1 Relationship to infinite layers of MPNNs

The expressivity of a one-layer MFN can be related to equivariant message-passing neural networks (MPNN) with linear two-body updates and infinite layers. The

parallels can be seen by expanding the matrix function as an infinite power series. Any analytic matrix function f admits a formal power series expansion, valid in its radius of analyticity,

$$f(H) = \sum_{k=0}^{\infty} c_k H^k \quad (6.25)$$

where c_k are the complex coefficients of the expansion. With each power \mathbf{H} , information propagates to neighbouring nodes, similar to a message passing step. Figure 6.5 shows a block diagonal matrix H and powers of this matrix H . With each power, the matrix becomes more dense, visually showing the propagation of information. The function f_θ can be expanded in this power series.

Let us look at the features that get extracted from $f_\theta(\mathbf{H})$ during the diagonal update. For powers of two and three, we have,

$$(H)_{ii,c00}^2 = \sum_{j \in \mathcal{N}(i)} H_{ij,c0m} H_{ji,cm0}, \quad (6.26)$$

$$(H)_{ii,c00}^3 = \sum_{j \in \mathcal{N}(i)} \sum_{k \in \mathcal{N}(j)} H_{ij,c0m} H_{jk,cm'm'} H_{ji,cm'0}. \quad (6.27)$$

Note that the matrix \mathbf{H}_{ij} is only non-zero if there is an edge between node i and j . It is this sparsity that restricts the sum over indices to the neighbourhood $\mathcal{N}(i)$ of the atoms. This is the same convolutional structure as in message-passing neural networks, where node features are updated based on nearest neighbours. Each coefficient of the matrix function that is **learnable** can act like weights in the linear update of a message passing layer. The power series expansion of our resolvent parametrisation is infinite, and so it would require infinite layers of a two-body MPNN to construct a single-layer MFN feature. This suggests that MFNs are inherently non-local and exhibit a convolution-like structure, similar to message-passing methods with infinite layers and linear updates.

6.5.2 Combes-Thomas theorem

We now investigate the non-locality of the resolvent expansion directly. With the help of the Combes-Thomas theorem [251] we can precisely quantify the limits of the non-locality based on the location of the poles. If \mathbf{H} has a finite interaction range, then the Combes-Thomas theorem [251] implies that

$$|(z\mathbf{I} - \mathbf{H})_{i,j}^{-1}| \leq C e^{-c\gamma z d_{ij}}, \quad (6.28)$$




Figure 6.6 Matrix elements of the resolvent around $z_l = x + iy$. Illustrated on the Laplacian of the path graph with 100 nodes, showing both the diagonal elements (line graph) and all matrix elements (inset).

where $\gamma_z = \text{dist}(z, \sigma(\mathbf{H}))$ quantifies the distance between the spectrum of \mathbf{H} and the pole z , d_{ij} is the length of the shortest path from node i to node j , and C, c are coefficients that may depend on the norm of \mathbf{H} . Note that the norm of the matrix is restricted in our architecture as detailed in Section 6.3.3.

Since we have taken \mathbf{H} to be self-adjoint and it therefore has a real spectrum, an estimate for γ_z is $\text{Imag}(z)$. As a result, if we constrain the poles in the parameterization of f_θ to be at $\text{Imag}(z_s) = \gamma$, then the resulting matrix function will satisfy

$$|[f_\theta(\mathbf{H})]_{i,j}| \leq C e^{-c\gamma d_{ij}}. \quad (6.29)$$

Therefore, the degree of locality can be controlled by constraining $\text{Imag}(z_s)$. The non-locality can also be regularised by adding the magnitude of $\text{Imag}(z_s)$ to the loss.

6.5.3 Visualisation of non-locality

We now present a visualisation of the non-locality for varying pole positions. For simplicity, we take a pure graph not embedded in 3D space. We take a chain graph consisting of nodes in a line connected to its neighbours and examine its Laplacian

matrices. For a weighted graph, the Laplacian matrix is computed as $\mathbf{L} = \mathbf{D} - \mathbf{W}$, where \mathbf{D} is the degree matrix and \mathbf{W} is the weighted adjacency matrix. Laplacian matrices play a crucial role in graph theory, particularly in the analysis of graph structures and properties [252].

As mentioned previously, the non-locality is determined by the distance of the spectrum of \mathbf{H} and the poles. In the case of the linear chain with equal weights, the matrix is symmetric, and hence, the eigenvalues are real and the eigenvectors of the Laplacian matrix exhibit sinusoidal oscillations. The oscillation frequency is proportional to the corresponding eigenvalue. Figure 6.6 shows the diagonal elements of the resolvent term,

$$f_w = (z_w \mathbf{I} - \mathbf{L})^{-1}, \quad (6.30)$$

for various choice of poles $z_w = x + iy$. If we investigate the function f_w by performing a spectral decomposition,

$$f_w(\mathbf{L}) = U \begin{bmatrix} \frac{1}{z_w - \lambda_1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \frac{1}{z_w - \lambda_n} \end{bmatrix} U^T, \quad (6.31)$$

where all variables are defined as previously in Equation (6.31), we see that as z_s approaches λ_i its corresponding eigenvector dominates the resulting matrix $f_w(\mathbf{L})$. In the case where L is the Laplacian of the chain graph, for small y this corresponds to selecting a specific eigenvector. As visible Figure 6.6 bottom, the x position specifies which eigenvector dominates. For poles with non-negligible imaginary parts the f_w is not dominated by a single eigenvector.

6.6 Interpretability of the MFN operators

In the case of the Euclidean group, the matrices learned in MFNs have the same symmetries as Euclidean operators. Euclidean operators play a crucial role in quantum mechanics. They are defined as self-adjoint operators (in the space of square-integrable functions) that are equivariant to the action of the group of rotations, translations, reflections, and permutations. When these operators are expanded on the basis of the irreducible representations of the rotation group, they exhibit a block structure (see Figure 6.4) in which each entry is labelled with representations (ss, ps, pp, dd, \dots).

The most central of these operators is the Hamiltonian operator. The energy of a quantum system is related to the trace of a specific matrix function of the

Hamiltonian,

$$E = \text{Tr} [f(\mathbf{H})\mathbf{H}], \quad (6.32)$$

where f is the Fermi function defined as

$$f(\varepsilon) = \frac{1}{e^{(\varepsilon-\mu)/k_{\text{B}}T} + 1}, \quad (6.33)$$

where μ is the total chemical potential, k_{B} is the Boltzmann constant and T is the absolute temperature. The Hamiltonian is usually computed as a fixed point of a self-consistent loop that minimizes the energy. This loop introduces many-body non-local effects. This motivates us to use many-body functions to parametrise our matrices and to learn the fixed point directly via learning a matrix function of many matrices. This is in opposition to tight-binding methods that usually construct a single low-body Hamiltonian with physics-inspired functional forms but require several self-consistent iterations to converge.

6.7 Applications of matrix functions to cumulenes

We compare the non-locality of MFNs to local MPNNs and global attention MPNNs using linear carbon chains called cumulenes. This example is a notorious challenge for state-of-the-art machine learning force field architectures [253].

Cumulenes are made up of double-bonded carbon atoms terminated with two hydrogen atoms at each end. Cumulenes exhibit pronounced non-local behaviour as a result of strong electron delocalisation, as visualised in Figure 6.7. Small changes in chain length and relative dihedral angle between the terminating hydrogen atoms can result in large changes in the energy of the system, as visually represented in Figure 6.7. These structures are known to illustrate the limited expressivity of local models [253] and are similar to the k-chains introduced by [254] in the context of the geometric WL test. Reference [255] showed that global attention is capable of capturing the angular trends of cumulenes with fixed length. We go beyond and demonstrate that MFNs are capable of accurately extrapolating to longer chains, simultaneously capturing length and angle trends. In contrast, global attention models, such as Spookynet [253], are unable to extrapolate to longer chains, highlighting the benefit of the matrix function formalism. For all MFN models in this section, we use MACE layers [27] to form the matrix at each layer. We refer to the model as MFN(MACE). The reference local MACE model and MFN(MACE) contain two message passing steps each. The receptive field of the local model is displayed in

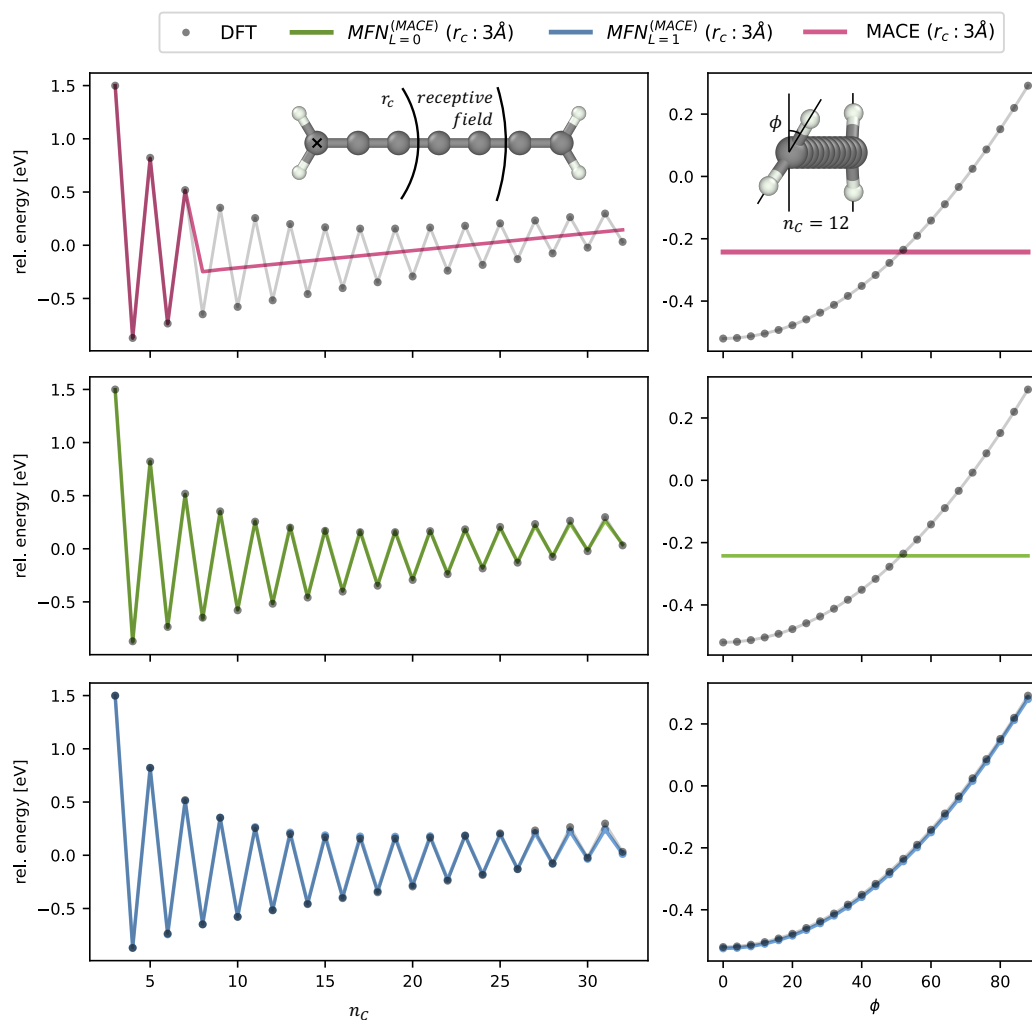


Figure 6.7 Visualizing MFN expressivity on cumulene chains. The left panel depicts energy trends with respect to cumulene chain length at a fixed angle $\phi = 5^\circ$. The right panel shows the DFT (ground truth) and the predicted energy as a function of the dihedral angle ϕ between the hydrogen atoms for a cumulene chain containing 12 carbon atoms. Local many-body equivariant models (MACE) are only able to capture average trends, even though test configurations are included in the training set. Invariant MFNs ($L = 0$) capture only the trends with respect to length, while equivariant MFNs ($L = 1$) capture both non-local effects. All models have a cut-off distance r_c of 3 \AA , corresponding to the nearest neighbours, with two message-passing layers. The cut-off distance, as well as MACE's receptive field for the first carbon atom, is annotated in the left panel.

Figure 6.7. Details of the training set and the specific choice of parametrisation of the matrix entries are included in the [Section 6.7.3](#).

6.7.1 Trends with chain length and rotation

The lowest energy structure of cumulenes alternates between 90- and 0-degree angles for odd and even carbon atom counts, respectively. Consequently, varying the number of carbon atoms at a fixed angle unveils a distinctive zigzag pattern in the ground truth energy ([Figure 6.7](#) left panel). Although the local model, with two message passing layers, is trained on exactly these configurations, this system is beyond its expressivity, as can be seen by the linear trend for $n_c > 7$ in ([Figure 6.7](#) left panel). This is to be expected as the receptive field of the two end atoms doesn't overlap. In contrast, the invariant and equivariant MFN models perfectly reproduce density functional theory (DFT) energies, thanks to their inherent non-locality. To demonstrate the necessity of equivariance, we train models on the energy of a fixed-size cumulene as a function of the dihedral angle between the hydrogen atoms. [Figure 6.7](#) demonstrates that only the equivariant MFN ($L=1$) captures the angular trend. The invariant model on the other hand, $L = 0$, like the mace model, is constant. This illustrates the benefit of having an equivariant basis at the cost of a significantly ($4\times$) larger matrix.

6.7.2 Guaranteed Non-Local dataset

Non-local effects are often difficult to observe from pure RMSE values on the energy and forces. Indeed local models trained on datasets designed to test non-local interactions such as MD22 dataset [256] may yield unexpectedly high accuracy, complicating the assessment of model non-locality [70]. The dataset introduced here is based on cumulenes, whose strong electronic de-localisation results in a directly observable non-locality in the RMSE values. The training set contains geometry-optimized cumulenes with 3-10 and 13, 14 carbon atoms, which are then rattled and rotated at various angles. The test set contains cumulenes created in a similar fashion with the same number of carbons (in-domain) and cumulenes of unseen length, not present in the dataset (out-domain 11,12 and 15,16). [Table 6.1](#) shows that the MFN architecture significantly outperforms both local and attention-based models (Spookynet). In the attention-based model, all nodes can, in principle, communicate with each other. On the training dataset, attention captures some non-locality, resulting in marginally lower errors on the train and in-domain test set compared to the local model. However, the learned non-locality does not generalize to larger molecules.

Dataset	n_C	E (meV/atom)			F (meV/Å)		
		MACE (Local)	SpookyNet (Global attention)	MFN ^(MACE) (ours)	MACE (Local)	SpookyNet (Global attention)	MFN ^(MACE) (ours)
Train	3-10,13,14	41.1	<u>31.4</u>	2.3	179.6	<u>114.1</u>	17.7
Test (In Domain)	3-10,13,14	41.8	<u>30.8</u>	3.1	205.6	<u>162.3</u>	18.4
Test (Out Domain)	11,12	<u>16.5</u>	31.4	0.4	<u>108.5</u>	116.2	8.6
Test (Out Domain)	15,16	<u>12.0</u>	23.4	2.3	<u>77.1</u>	87.6	12.1

Table 6.1 Guaranteed non-local cumulene dataset. Containing rattled cumulene chains, with various chain lengths (n_C) and hydrogen dihedral angles (ϕ). The table compares energy (E) and forces (F) RMSEs between local two-layer MPNNs (MACE), global attention MPNNs (SpookyNet), and equivariant MFNs. Train and in-domain test sets contain cumulenes of lengths 3-10 and 13,14. Two out-domain test sets compare different levels of extrapolation to unseen cumulene lengths, containing cumulenes with 11, 12 and 15, 16 carbon atoms, respectively. Bold is best, and underline second best.

Indeed, when tested on unseen chain lengths, it obtains energy and forces worse than those obtained with a local model. The non-locality is extremely overfit to the training set and doesn’t generalise. Looking at the performance of MFNs, we see orders of magnitude lower errors in the forces and energies on train and test sets. Especially interesting is its competitive performance on unseen chain lengths. The structured non-locality of MFNs enables generalization to larger unseen system sizes. To date, no other alternative architecture has beaten the MFN at this challenging test.

6.7.3 Computational Details

Dataset The cumulene dataset is designed to test the expressivity of graph neural networks and their non-local capabilities. The task is to regress energy and forces from 3D molecular configurations. The ground truth energy is obtained using the ORCA quantum chemistry code at the density functional theory level of accuracy using the *wB97X-D3* functional, the *def2-TZVP* basis set, and very tight SCF convergence criteria.

Length and angle scans The configurations for the length and angle scans consist of chain-like graphs as in [Figure 6.8](#). The length scans range from 3 to 30 atoms fixed at an angle of 5 degrees. Additionally, we scan the angle between the two terminating hydrogen groups for a cumulene with twelve carbon atoms ($n_C = 12$). At increments of 3 degrees, we scan from 0 to 90 degrees. By symmetry, this covers the entire range of unique configurations from 0 to 360 degrees. For the expressivity scan, all internal coordinates (see [Figure 6.8](#)) that uniquely define the cumulene




Figure 6.8 Angles and distances that define a cumulene graph. Used to test expressivity in Figure 6.7. The carbon-hydrogen (r_{CH}), first carbon-carbon ($r_{CC}^{(1)}$), and remaining carbon-carbon distances ($r_{CC}^{(>1)}$) are set to 1.086 Å, 1.315 Å and 1.279 Å respectively. The angle between the hydrogen-carbon-hydrogen (θ_{HCH}) is fixed at 118.71 degrees, and the dihedral angle ϕ depends on the experiment as detailed in the main text.

are kept constant and set to the geometry-optimised configuration of the length 30 cumulenes. The carbon-hydrogen, first carbon-carbon, and remaining carbon-carbon distances are set to 1.086 Å, 1.315 Å and 1.279 Å respectively (see Figure 6.8). For the length 12 cumulene, the distance between the most distant carbons is 14.1Å. Thus, it becomes impossible for a purely local equivariant MPNN, with cutoff 3Å and two message-passing layers, to differentiate between two graphs at different angles.

Note that the relative energies of Figure 6.7 are obtained by subtracting individual atom energies. These are -16.3eV and -1036.1eV for each hydrogen and carbon atom, respectively.

Guaranteed Non-local dataset The GNL dataset tests how well a model can extrapolate to an unseen number of atoms. Furthermore, configurations are rattled to see how well the model can capture both local and non-local effects. Configurations are generated starting from the relaxed geometries of cumulenes with lengths 0-20. Relaxations are carried out with the smaller 6-31G basis set. The cumulenes are subsequently rotated by an equally spaced scan with an increment of 6 degrees, starting from a random angle, and the positions are subsequently randomly perturbed by Gaussian noise with a standard deviation of 0.01 Å. The training and validation set contains cumulenes of length 0-10 and 13-14, with 10 and 3 samples, respectively. The test set contains in-domain configurations with two samples for each of the lengths present in the training set. Furthermore, it contains configurations with 11-12 and 15-16 carbon atoms labelled out-of-domain. In total, the train, validation, and test set contain 200, 50 and 170 configurations.

MACE and MFN Model Both the local MACE and MFN are trained on the same graphs, which means that the cut-off distance is fixed at 3Å, including information from the nearest neighbour. For the MFN model, we use the architecture described

in Section 6.3.2. Both models are trained with two layers, allowing the local model to distinguish changes that are separated by 12\AA . The MFN is trained using 16 matrix channels and 16 poles. All MFN models use the diagonal readout block. For the matrix construction step, we use a MACE [27] layer with a correlation order 3, $l_{\max} = 3$, $L = 1$ and 128 channels to generate node features h_i . We use the diagonal update Equation (6.20) to update the node features of MACE and reiterate. The readout at the first layer is linear, and at the second, it is a one-layer MLP with 16 hidden dimensions.

Spookynet Model The Spookynet architecture was trained on the same dataset as the MFN model. We use a one-layer model with 5.5\AA cut-off to reproduce the receptive field of MACE. We use the global attention and electrostatic interactions. The model has 128 channels.

Training All models underwent initial training on the GNL dataset before transfer learning was applied to the relevant length and angle scan data sets. The saved model corresponds to the epoch that exhibits the minimum loss values. Details on settings such as learning rate and epoch count are disclosed in Table 6.2. The training incorporated both energy and forces, with adjustable weights for each observable. In particular, for length and angle scans, an additional 100 epochs with zero force weight were used after initial training, ensuring the depiction of the minimal possible energy error in Figure 6.7, as the lowest-loss model is saved.

Table 6.2 Model training parameters. For the matrix functions the number of poles (n_p) and matrix channels (c) are indicated.

dataset	model	epochs	lr	E_{weight}	F_{weight}	n_{layers}	r_{max}	other
Length Scan	MFN (L=0)	1240	1e-2	1000	100	2	3	$n_p=16, c=16$
	MFN (L=1)	5628	1e-5	1000	100	2	3	$n_p=16, c=16$
	MACE	5288	0.005	1000	100	2	3	-
Angle Scan	MFN (L=0)	1240	1e-2	1000	100	2	3	$n_p=16, c=16$
	MFN (L=1)	656	1e-5	1000	100	2	3	$n_p=16, c=16$
	MACE	954	0.005	1000	100	2	3	-
GNL	MFN (L=1)	4958	1e-2	100	1	2	3	$n_p=16, c=16$
	MACE	1260	1e-2	100	1	2	3	-
	Spookynet	3500	1e-4	0.10	0.90	1	5.5	attention

6.8 Relevance for reactions at interfaces

MFN's remain a reactive force field architecture. While the matrix is parametrised for the entire structure, off-diagonal elements in the learnable matrix smoothly transition to zero as atoms leave the cut-off distance. The use of a smooth cut-off is similar to any of the MLFF architectures outlined in the introduction.

An interesting avenue of future research is to explore the connection between the learnable matrix and a coarse grained hamiltonian. While in the most general case the MFN architecture has a learnable matrix function, the functional form could be restricted to the typical Fermi-Dirac distribution. Furthermore, the MFN architecture could be combined with electrostatic models. One approach would be to implement a charge equilibration method, where the energy function is a MFN. Charge equilibration MLFFs assign a certain energy penalty ϵ to a specific atom centred charge distribution. The total energy can be obtained by minimising the total energy by varying the atom centred charges

$$\min_{\{q_i\}} \epsilon(\{x_i, q_i\}).$$

Various functional forms of ϵ have been explored. In the CENT and kQEq architectures, the energy function is quadratic in the local charge and given by

$$\sum_i \left(E_i(\{x_j\}) + \chi_i q_i + \eta_i q_i^2 + \sum_{i < j} \frac{q_i q_j}{f(r_{ij})} \right), \quad (6.34)$$

where χ are the electronegativity and η the hardness, which can both be learnable parameters of the local environment, and the second summation describes the interaction of charges with each-other. In principle this architecture with a smooth energy penalty is best suited for metallic systems, where charges are delocalised. Using the MFN architecture as the energy function may allow for the treatment of both insulating and conducting systems due to the discrete nature of the learnt matrix's spectrum.

6.9 Discussion

In this chapter we introduced Matrix Function Networks (MFNs), an architecture designed to address the limitations of existing GNNs and MPNNs in modelling non-local many-body interactions. Utilizing a resolvent expansion, MFNs achieve potentially linear scaling with respect to system size, offering a computationally

efficient solution. We demonstrate that our architecture is capable of modelling the complex non-local interactions of cumulene systems. Future work could focus on extending MFNs to other complex systems, further validating their adaptability and efficiency, and exploring their interpretability. Furthermore, MFNs may be combined with electrostatic models which explicitly model the electron density. This may allow for accurate modelling of charged molecules, charge transfer and the effect of external fields.

Chapter 7

Conclusions and Outlook

7.1 Conclusions

This thesis explores the use of machine learning force fields to accelerate atomistic simulations of reactions at interfaces. The first contribution is an automated training protocol, which simplifies the task of curating quantum mechanical datasets for computing accurate energy barriers in catalytic reactions. The use of active learning and systematic curation of a training dataset significantly reduces the barrier of entry for the general catalysis community. We demonstrate the utility of machine learning force fields not only in reducing the computational cost of routine *ab-initio* computational catalysis tasks, but also in enabling more in-depth investigations. Using the active learning workflow, we identify a lower reaction path for the highest energy barrier and find a preferred active site for formaldehyde hydrogenation to methanol over indium oxide.

In [Chapter 4](#), we explore the utility of using recently developed atomistic foundation models for simulating catalytic reactions. We find that models that are trained only on solid-state crystal structures generalise remarkably well for modelling reactions at interfaces. We show that fine-tuning allows for obtaining low-energy barriers for reactions at interfaces with minimal computational effort.

Furthermore, we demonstrate the utility of machine learning force fields in helping us to understand experimental observations. We use the previously mentioned active learning workflow and fine-tuning approach to model porous metal-organic frameworks. The machine learning force fields can explain observed NMR spectra. We find that capturing short (ns) and long-time-scale dynamics (ms) is needed for agreement between the *ab-initio* and experimental spectra.

Finally, we address a key limitation of local machine learning models, their inability to capture long-range interactions. Matrix function neural networks parametrise a matrix operator of the molecular graph, and introduce non-locality by applying a learnable matrix function to it. We show that MFNs are able to model intrinsically non-local systems, namely cumulenes transferably. Compared to models which use fully connected nodes with non-local attention, we find that MFNs correctly generalise to unseen cumulene lengths.

7.2 Outlook

The remarkable performance of foundation models has made accurate and fast atomistic modelling as accessible as ever. The work in this thesis has shown how a reduced timeline for training accurate machine learning force fields, without the need to curate extensive datasets, facilitates collaboration with experimental groups. Nevertheless, exciting challenges remain for the future, both from an application and method development standpoint. Screening large catalytic reaction networks on multiple surfaces is now technically feasible, but easy-to-use and automated pipelines are missing. The reduced computational cost enables a more exhaustive exploration of reaction pathways, avoiding the need for simplifying assumptions such as pre-selecting a “rate-determining” step.

Another significant development, especially relevant when modelling catalysis, is going beyond DFT accuracy. Currently, the datasets used for training large foundation models are dominated by relatively inexpensive and widely applicable DFT functionals. As catalytic observables depend strongly on the functional, true experimental predicability may only be obtained with higher levels of theory, such as random phase approximation (RPA) or coupled cluster methods (CCSD(T)). How to benefit from large datasets computed at lower levels of theory while also training on smaller more accurate datasets remains an interesting avenue of research.

Another interesting under-explored domain is using MLFFs for electrocatalysis. Modelling reactions under applied electric fields with explicit water requires innovation in both force field architectures and dataset generation approaches. This will allow for modelling reactions such as water splitting or carbon dioxide hydrogenation, which are particularly relevant for renewable energy storage as liquid fuels.

Overall, developing and applying machine learning force fields has transformed the potential of computational atom scale modelling. The remarkable performance and promise for use in material and molecular design have provoked increasing

commercial interest. Developments in the field may soon have a real-world impact in the form of novel catalysts with improved properties.

References

1. Behler, J. First Principles Neural Network Potentials for Reactive Simulations of Large Molecular and Condensed Systems. *Angewandte Chemie International Edition* **56**, 12828–12840. ISSN: 1521-3773 (2017).
2. Bartók, A. P., Payne, M. C., Kondor, R. & Csányi, G. Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons. *Phys. Rev. Lett.* **104**, 136403. ISSN: 0031-9007, 1079-7114 (Apr. 1, 2010).
3. Sanbonmatsu, K. Y. & Tung, C. -. High Performance Computing in Biology: Multimillion Atom Simulations of Nanoscale Systems. *Journal of Structural Biology. Advances in Molecular Dynamics Simulations* **157**, 470–480. ISSN: 1047-8477 (Mar. 1, 2007).
4. Eng, A. Y. S., Soni, C. B., Lum, Y., Khoo, E., Yao, Z., Vineeth, S. K., Kumar, V., Lu, J., Johnson, C. S., Wolverton, C. & Seh, Z. W. Theory-Guided Experimental Design in Battery Materials Research. *Sci. Adv.* **8**, eabm2422. ISSN: 2375-2548 (May 13, 2022).
5. Feynman, R. P., Leighton, R. B., Sands, M. & Hafner, E. M. The Feynman Lectures on Physics; Vol. I. *American Journal of Physics* **33**, 750–752. ISSN: 0002-9505, 1943-2909 (Sept. 1, 1965).
6. Rowe, P., Deringer, V. L., Gasparotto, P., Csányi, G. & Michaelides, A. An Accurate and Transferable Machine Learning Potential for Carbon. *The Journal of Chemical Physics* **153**, 034702. ISSN: 0021-9606 (July 15, 2020).
7. Geim, A. K. & Novoselov, K. S. The Rise of Graphene. *Nature Mater* **6**, 183–191. ISSN: 1476-4660 (Mar. 2007).
8. Chalmers, A. Retracing the Ancient Steps to Atomic Theory. *Science & Education* **7**, 69–84. ISSN: 1573-1901 (Jan. 1, 1998).
9. Thomson, J. On the Number of Corpuscles in an Atom. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **11**, 769–781. ISSN: 1941-5982, 1941-5990 (June 1906).
10. Rutherford, E. The Scattering of α and β Particles by Matter and the Structure of the Atom. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **21**, 669–688. ISSN: 1941-5982, 1941-5990 (May 1911).
11. Bohr, N. On the Constitution of Atoms and Molecules. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **26**, 1–25. ISSN: 1941-5982, 1941-5990 (July 1913).
12. Schrödinger, E. Quantisierung als Eigenwertproblem. *Annalen der Physik* **385**, 437–490. ISSN: 0003-3804, 1521-3889 (Jan. 1926).

13. Volta, A. On the Electricity Excited by the Mere Contact of Conducting Substances of Different Kinds. In a Letter from Mr. Alexander Volta, F. R. S. Professor of Natural Philosophy in the University of Pavia, to the Rt. Hon. Sir Joseph Banks, Bart. K.B. P. R. S. *Philosophical Transactions of the Royal Society of London* **90**, 403–431 (Jan. 1800).
14. Planck, M. in *Von Kirchhoff bis Planck: Theorie der Wärmestrahlung in historisch-kritischer Darstellung* (ed Schöpf, H.-G.) 178–191 (Vieweg+Teubner Verlag, Wiesbaden, 1978). ISBN: 978-3-663-13885-3.
15. Einstein, A. Über Einen Die Erzeugung Und Verwandlung Des Lichtes Betreffenden Heuristischen Gesichtspunkt. *Annalen der Physik* **322**, 132–148. ISSN: 1521-3889 (1905).
16. De Broglie, L. Recherches sur la théorie des Quanta. *Annales de Physique* (1925).
17. Heisenberg, W. Über quantentheoretische Umdeutung kinematischer und mechanischer Beziehungen. *Z. Physik* **33**, 879–893. ISSN: 0044-3328 (Dec. 1, 1925).
18. Helgaker, T., Jorgensen, P. & Olsen, J. *Molecular Electronic-Structure Theory* 949 pp. ISBN: 978-1-118-53147-1. Google Books: [APjLWFFxWkQC](#) (John Wiley & Sons, Feb. 18, 2013).
19. Atkins, P. W. & Friedman, R. S. *Molecular Quantum Mechanics* 552 pp. ISBN: 978-0-19-954142-3 (OUP Oxford, 2011).
20. Schumann, J., Stamatakis, M., Michaelides, A. & Réocreux, R. Ten-Electron Count Rule for the Binding of Adsorbates on Single-Atom Alloy Catalysts. *Nat. Chem.* **16**, 749–754. ISSN: 1755-4330, 1755-4349 (May 2024).
21. Nørskov, J. K., Bligaard, T., Rossmeisl, J. & Christensen, C. H. Towards the Computational Design of Solid Catalysts. *Nature Chem* **1**, 37–46. ISSN: 1755-4330, 1755-4349 (Apr. 2009).
22. Hannagan, R. T., Giannakakis, G., Réocreux, R., Schumann, J., Finzel, J., Wang, Y., Michaelides, A., Deshlahra, P., Christopher, P., Flytzani-Stephanopoulos, M., Stamatakis, M. & Sykes, E. C. H. First-Principles Design of a Single-Atom–Alloy Propane Dehydrogenation Catalyst. *Science* **372**, 1444–1447 (June 25, 2021).
23. Behler, J. & Parrinello, M. Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces. *Phys. Rev. Lett.* **98**, 146401 (Apr. 2, 2007).
24. Shapeev, A. V. Moment Tensor Potentials: A Class of Systematically Improvable Interatomic Potentials. *Multiscale Model. Simul.* **14**, 1153–1173. ISSN: 1540-3459 (Jan. 2016).
25. Drautz, R. Atomic Cluster Expansion for Accurate and Transferable Interatomic Potentials. *Physical Review B* **99**, 014104 (2019).
26. Batzner, S., Musaelian, A., Sun, L., Geiger, M., Mailoa, J. P., Kornbluth, M., Molinari, N., Smidt, T. E. & Kozinsky, B. E(3)-Equivariant Graph Neural Networks for Data-Efficient and Accurate Interatomic Potentials. *Nat Commun* **13**, 2453. ISSN: 2041-1723 (May 4, 2022).

27. Batatia, I., Kovacs, D. P., Simm, G., Ortner, C. & Csanyi, G. MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields. *Advances in Neural Information Processing Systems* **35**, 11423–11436 (Dec. 2022).
28. Batatia, I., Geiger, M., Munoz, J., Smidt, T., Silberman, L. & Ortner, C. *A General Framework for Equivariant Neural Networks on Reductive Lie Groups* 2023. arXiv: 2306.00091.
29. Behler, J. Four Generations of High-Dimensional Neural Network Potentials. *Chemical Reviews* **121**, 10037–10072 (2021).
30. Kovács, D. P., Moore, J. H., Browning, N. J., Batatia, I., Horton, J. T., Kapil, V., Witt, W. C., Magdău, I.-B., Cole, D. J. & Csányi, G. *MACE-OFF23: Transferable Machine Learning Force Fields for Organic Molecules* arXiv: 2312.15211. (2024). Pre-published.
31. Deringer, V. L., Caro, M. A. & Csányi, G. A General-Purpose Machine-Learning Force Field for Bulk and Nanostructured Phosphorus. *Nat Commun* **11**, 5461. ISSN: 2041-1723 (1 Oct. 29, 2020).
32. Iftimie, R., Minary, P. & Tuckerman, M. E. *Ab Initio* Molecular Dynamics: Concepts, Recent Developments, and Future Trends. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 6654–6659. ISSN: 0027-8424, 1091-6490 (May 10, 2005).
33. Vopson, M. M. Estimation of the Information Contained in the Visible Matter of the Universe. *AIP Advances* **11**, 105317. ISSN: 2158-3226 (Oct. 1, 2021).
34. Hohenberg, P. & Kohn, W. Inhomogeneous Electron Gas. *Phys. Rev.* **136**, B864–B871. ISSN: 0031-899X (3B Nov. 9, 1964).
35. Kohn, W. Density Functional and Density Matrix Method Scaling Linearly with the Number of Atoms. *Phys. Rev. Lett.* **76**, 3168–3171. ISSN: 0031-9007, 1079-7114 (Apr. 22, 1996).
36. Kohn, W. & Sham, L. J. Self-Consistent Equations Including Exchange and Correlation Effects. *Phys. Rev.* **140**, A1133–A1138. ISSN: 0031-899X (4A Nov. 15, 1965).
37. Cohen, A. J., Mori-Sánchez, P. & Yang, W. Insights into Current Limitations of Density Functional Theory. *Science* **321**, 792–794. ISSN: 0036-8075, 1095-9203 (Aug. 8, 2008).
38. Kulik, H. J. Perspective: Treating Electron over-Delocalization with the DFT+U Method. *The Journal of Chemical Physics* **142**, 240901. ISSN: 0021-9606, 1089-7690 (June 28, 2015).
39. Mori-Sánchez, P., Cohen, A. J. & Yang, W. Localization and Delocalization Errors in Density Functional Theory and Implications for Band-Gap Prediction. *Phys. Rev. Lett.* **100**, 146401. ISSN: 0031-9007, 1079-7114 (Apr. 7, 2008).
40. Christensen, R., Hansen, H. A. & Vegge, T. Identifying Systematic DFT Errors in Catalytic Reactions. *Catal. Sci. Technol.* **5**, 4946–4949. ISSN: 2044-4753, 2044-4761 (2015).
41. Goerigk, L., Hansen, A., Bauer, C., Ehrlich, S., Najibi, A. & Grimme, S. A Look at the Density Functional Theory Zoo with the Advanced GMTKN55 Database for General Main Group Thermochemistry, Kinetics and Noncovalent Interactions. *Phys. Chem. Chem. Phys.* **19**, 32184–32215. ISSN: 1463-9076, 1463-9084 (2017).

42. Grimme, S., Hansen, A., Brandenburg, J. G. & Bannwarth, C. Dispersion-Corrected Mean-Field Electronic Structure Methods. *Chem. Rev.* **116**, 5105–5154. ISSN: 0009-2665, 1520-6890 (May 11, 2016).
43. Jones, J. E. & Chapman, S. On the Determination of Molecular Fields. —II. From the Equation of State of a Gas. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* **106**, 463–477 (Oct. 1924).
44. Jorgensen, W. L., Chandrasekhar, J., Madura, J. D., Impey, R. W. & Klein, M. L. Comparison of Simple Potential Functions for Simulating Liquid Water. *The Journal of Chemical Physics* **79**, 926–935. ISSN: 0021-9606 (July 15, 1983).
45. Senftle, T. P. *et al.* The ReaxFF Reactive Force-Field: Development, Applications and Future Directions. *npj Comput Mater* **2**, 1–14. ISSN: 2057-3960 (1 Mar. 4, 2016).
46. Neese, F., Wennmohs, F., Becker, U. & Riplinger, C. The ORCA Quantum Chemistry Program Package. *The Journal of Chemical Physics* **152**, 224108. ISSN: 0021-9606 (June 12, 2020).
47. Schwerdtfeger, P. & Wales, D. J. 100 Years of the Lennard-Jones Potential. *J. Chem. Theory Comput.* **20**, 3379–3405. ISSN: 1549-9618 (May 14, 2024).
48. Georgi, H. *Lie Algebras In Particle Physics: From Isospin To Unified Theories* 340 pp. ISBN: 978-0-429-49921-0 (CRC Press, Boca Raton, May 4, 2018).
49. Zee, A. *Group Theory in a Nutshell for Physicists* 613 pp. ISBN: 978-0-691-16269-0 (Princeton University Press, Princeton, 2016).
50. Darby, J. P., Kovács, D. P., Batatia, I., Caro, M. A., Hart, G. L. W., Ortner, C. & Csányi, G. *Tensor-Reduced Atomic Density Representations* arXiv: 2210.01705. (2022). Pre-published.
51. Musil, F., Grisafi, A., Bartók, A. P., Ortner, C., Csányi, G. & Ceriotti, M. Physics-Inspired Structural Representations for Molecules and Materials. *Chem. Rev.* **121**, 9759–9815. ISSN: 0009-2665 (Aug. 25, 2021).
52. Ortner, C., Thomas, J. & Chen, H. Locality of Interatomic Forces in Tight Binding Models for Insulators. *ESAIM: M2AN* **54**, 2295–2318. ISSN: 0764-583X, 1290-3841 (6 Nov. 1, 2020).
53. Deringer, V. L., Bartók, A. P., Bernstein, N., Wilkins, D. M., Ceriotti, M. & Csányi, G. Gaussian Process Regression for Materials and Molecules. *Chem. Rev.* **121**, 10073–10141. ISSN: 0009-2665, 1520-6890 (Aug. 25, 2021).
54. Schaaf, L. L., Fako, E., De, S., Schäfer, A. & Csányi, G. Accurate Energy Barriers for Catalytic Reaction Pathways: An Automatic Training Protocol for Machine Learning Force Fields. *npj Computational Materials* **9**, 180 (2023).
55. Duval, A., Schmidt, V., Garcia, A. H., Miret, S., Malliaros, F. D., Bengio, Y. & Rolnick, D. *FAENet: Frame Averaging Equivariant GNN for Materials Modeling* arXiv: 2305.05577. (2026). Pre-published.
56. Pozdnyakov, S. N. & Ceriotti, M. *Smooth, Exact Rotational Symmetrization for Deep Learning on Point Clouds* arXiv: 2305.19302. (2026). Pre-published.

57. Deringer, V. L., Bartók, A. P., Bernstein, N., Wilkins, D. M., Ceriotti, M. & Csányi, G. Gaussian Process Regression for Materials and Molecules. *Chemical Reviews* **121**, 10073–10141 (2021).
58. Bartók, A. P., Kondor, R. & Csányi, G. On Representing Chemical Environments. *Phys. Rev. B* **87**, 184115. ISSN: 1098-0121, 1550-235X (May 28, 2013).
59. Deringer, V. L., Merlet, C., Hu, Y., Lee, T. H., Kattirtzi, J. A., Pecher, O., Csányi, G., Elliott, S. R. & Grey, C. P. Towards an Atomistic Understanding of Disordered Carbon Electrode Materials. *Chem. Commun.* **54**, 5988–5991. ISSN: 1359-7345, 1364-548X (2018).
60. Stocker, S., Jung, H., Csanyi, G., Franklin, C., Reuter, K. & Margraf, J. T. Estimating Free Energy Barriers for Heterogeneous Catalytic Reactions with Machine Learning Potentials and Umbrella Integration. *arxiv* (2023).
61. Batatia, I., Batzner, S., Kovács, D. P., Musaelian, A., Simm, G. N. C., Drautz, R., Ortner, C., Kozinsky, B. & Csányi, G. *The Design Space of $E(3)$ -Equivariant Atom-Centered Interatomic Potentials* arXiv: 2205.06643. (2022). Pre-published.
62. Pozdnyakov, S. N., Willatt, M. J., Bartók, A. P., Ortner, C., Csányi, G. & Ceriotti, M. Incompleteness of Atomic Structure Representations. *Phys. Rev. Lett.* **125**, 166001 (Oct. 12, 2020).
63. Ho, C. H., Gutleb, T. S. & Ortner, C. *Atomic Cluster Expansion without Self-Interaction* arXiv: 2401.01550. (2025). Pre-published.
64. Schütt, K. T., Sauceda, H. E., Kindermans, P.-J., Tkatchenko, A. & Müller, K.-R. SchNet – A Deep Learning Architecture for Molecules and Materials. *J. Chem. Phys.* **148**, 241722. ISSN: 0021-9606 (June 28, 2018).
65. Bronstein, M. M., Bruna, J., Cohen, T. & Veličković, P. *Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges* arXiv: 2104.13478. (2025). Pre-published.
66. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O. & Dahl, G. E. *Neural Message Passing for Quantum Chemistry* in *Proceedings of the 34th International Conference on Machine Learning-Volume 70* (JMLR. org, 2017), 1263–1272.
67. Satorras, V. G., Hoogeboom, E. & Welling, M. *$E(n)$ Equivariant Graph Neural Networks* arXiv: 2102.09844. (2025). Pre-published.
68. Norwood, S. W., Schaaf, L. L., Batatia, I., Csányi, G. & Bhowmik, A. Enhancing the Local Expressivity of Geometric Graph Neural Networks. *NeurIPS 2023 Workshop: Machine Learning and the Physical Sciences* (2023).
69. He, K., Zhang, X., Ren, S. & Sun, J. *Deep Residual Learning for Image Recognition* arXiv: 1512.03385. (2025). Pre-published.
70. Kovacs, D. P., Batatia, I., Arany, E. S. & Csanyi, G. *Evaluation of the MACE Force Field Architecture: From Medicinal Chemistry to Materials Science* arXiv, July 2023. arXiv: 2305.14247.
71. *Open Catalyst 2020 (OC20) Dataset and Community Challenges | ACS Catalysis* (2023).
72. Batatia, I. *et al.* *A Foundation Model for Atomistic Materials Chemistry* arXiv: 2401.00096. (2024). Pre-published.

73. Bartók, A. P., Kermode, J., Bernstein, N. & Csányi, G. Machine Learning a General-Purpose Interatomic Potential for Silicon. *Phys. Rev. X* **8**, 041048. ISSN: 2160-3308 (Dec. 14, 2018).
74. Deringer, V. L. & Csányi, G. Machine Learning Based Interatomic Potential for Amorphous Carbon. *Phys. Rev. B* **95**, 094203. ISSN: 2469-9950, 2469-9969 (Mar. 3, 2017).
75. Deringer, V. L., Bernstein, N., Csányi, G., Ben Mahmoud, C., Ceriotti, M., Wilson, M., Drabold, D. A. & Elliott, S. R. Origins of Structural and Electronic Transitions in Disordered Silicon. *Nature* **589**, 59–64. ISSN: 1476-4687 (7840 Jan. 2021).
76. Chmiela, S., Sauceda, H. E., Poltavsky, I., Müller, K.-R. & Tkatchenko, A. sGDML: Constructing Accurate and Data Efficient Molecular Force Fields Using Machine Learning. *Computer Physics Communications* **240**, 38–45. ISSN: 0010-4655 (July 1, 2019).
77. Bernstein, N., Csányi, G. & Deringer, V. L. De Novo Exploration and Self-Guided Learning of Potential-Energy Surfaces. *npj Comput Mater* **5**, 1–9. ISSN: 2057-3960 (1 Oct. 11, 2019).
78. Botu, V., Batra, R., Chapman, J. & Ramprasad, R. Machine Learning Force Fields: Construction, Validation, and Outlook. *J. Phys. Chem. C* **121**, 511–522. ISSN: 1932-7447, 1932-7455 (Jan. 12, 2017).
79. Podryabinkin, E. V., Tikhonov, E. V., Shapeev, A. V. & Oganov, A. R. Accelerating Crystal Structure Prediction by Machine-Learning Interatomic Potentials with Active Learning. *Phys. Rev. B* **99**, 064114. ISSN: 2469-9950, 2469-9969 (Feb. 27, 2019).
80. Jørgensen, M. S., Larsen, U. F., Jacobsen, K. W. & Hammer, B. Exploration versus Exploitation in Global Atomistic Structure Optimization. *J. Phys. Chem. A* **122**, 1504–1509. ISSN: 1089-5639 (Feb. 8, 2018).
81. Vandermause, J., Torrisi, S. B., Batzner, S., Xie, Y., Sun, L., Kolpak, A. M. & Kozinsky, B. On-the-Fly Active Learning of Interpretable Bayesian Force Fields for Atomistic Rare Events. *npj Comput Mater* **6**, 20. ISSN: 2057-3960 (Mar. 18, 2020).
82. Sivaraman, G., Krishnamoorthy, A. N., Baur, M., Holm, C., Stan, M., Csányi, G., Benmore, C. & Vázquez-Mayagoitia, A. Machine-Learned Interatomic Potentials by Active Learning: Amorphous and Liquid Hafnium Dioxide. *npj Comput Mater* **6**, 1–8. ISSN: 2057-3960 (1 July 23, 2020).
83. Young, T. A., Johnston-Wood, T., Deringer, V. L. & Duarte, F. A Transferable Active-Learning Strategy for Reactive Molecular Force Fields. *Chem. Sci.* **12**, 10944–10955. ISSN: 2041-6520, 2041-6539 (2021).
84. Vandermause, J., Xie, Y., Lim, J. S., Owen, C. J. & Kozinsky, B. Active Learning of Reactive Bayesian Force Fields Applied to Heterogeneous Catalysis Dynamics of H/Pt. *Nat Commun* **13**, 5183. ISSN: 2041-1723 (Sept. 2, 2022).
85. Van der Oord, C., Sachs, M., Kovács, D. P., Ortner, C. & Csányi, G. *Hyperactive Learning (HAL) for Data-Driven Interatomic Potentials* arXiv: 2210.04225. (2022). Pre-published.

86. Eastman, P., Behara, P. K., Dotson, D. L., Galvelis, R., Herr, J. E., Horton, J. T., Mao, Y., Chodera, J. D., Pritchard, B. P., Wang, Y., *et al.* Spice, a Dataset of Drug-like Molecules and Peptides for Training Machine Learning Potentials. *Scientific Data* **10**, 11 (2023).
87. Dang, S., Qin, B., Yang, Y., Wang, H., Cai, J., Han, Y., Li, S., Gao, P. & Sun, Y. Rationally Designed Indium Oxide Catalysts for CO₂ Hydrogenation to Methanol with High Activity and Selectivity. *Sci. Adv.* **6**, eaaz2060. ISSN: 2375-2548 (June 2020).
88. Bielz, T., Lorenz, H., Amann, P., Klötzer, B. & Penner, S. Water-Gas Shift and Formaldehyde Reforming Activity Determined by Defect Chemistry of Polycrystalline In₂O₃. *J. Phys. Chem. C* **115**, 6622–6628. ISSN: 1932-7447, 1932-7455 (Apr. 14, 2011).
89. Cai, D., Cai, Y., Tan, K. B. & Zhan, G. Recent Advances of Indium Oxide-Based Catalysts for CO₂ Hydrogenation to Methanol: Experimental and Theoretical. *Materials* **16**, 2803. ISSN: 1996-1944 (Mar. 31, 2023).
90. Cao, A., Wang, Z., Li, H. & Nørskov, J. K. Relations between Surface Oxygen Vacancies and Activity of Methanol Formation from CO₂ Hydrogenation over In₂O₃ Surfaces. *ACS Catal.* **11**, 1780–1786. ISSN: 2155-5435, 2155-5435 (Feb. 5, 2021).
91. Kauppinen, M., Posada-Borbón, A. & Grönbeck, H. Methanol Synthesis Over PdIn, In₂O₃, and CuZn From First-Principles Microkinetics: Similarities and Differences. *J. Phys. Chem. C* **126**, 15235–15246. ISSN: 1932-7447, 1932-7455 (Sept. 15, 2022).
92. Martin, O., Martín, A. J., Mondelli, C., Mitchell, S., Segawa, T. F., Hauert, R., Drouilly, C., Curulla-Ferré, D. & Pérez-Ramírez, J. Indium Oxide as a Superior Catalyst for Methanol Synthesis by CO₂ Hydrogenation. *Angewandte Chemie International Edition* **55**, 6261–6265. ISSN: 1521-3773 (2016).
93. Qin, B. & Li, S. First Principles Investigation of Dissociative Adsorption of H₂ during CO₂ Hydrogenation over Cubic and Hexagonal In₂O₃ Catalysts. *Phys. Chem. Chem. Phys.* **22**, 3390–3399. ISSN: 1463-9076, 1463-9084 (2020).
94. Qin, B., Zhou, Z., Li, S. & Gao, P. Understanding the Structure-Performance Relationship of Cubic In₂O₃ Catalysts for CO₂ Hydrogenation. *Journal of CO₂ Utilization* **49**, 101543. ISSN: 22129820 (July 2021).
95. Rui, N., Wang, Z., Sun, K., Ye, J., Ge, Q. & Liu, C.-j. CO₂ Hydrogenation to Methanol over Pd/In₂O₃: Effects of Pd and Oxygen Vacancy. *Applied Catalysis B: Environmental* **218**, 488–497. ISSN: 09263373 (Dec. 2017).
96. Tian, G., Wu, Y., Wu, S., Huang, S. & Gao, J. CO₂ Hydrogenation to Methanol over Pd/MnO/In₂O₃ Catalyst. *Journal of Environmental Chemical Engineering* **10**, 106965. ISSN: 22133437 (Feb. 2022).
97. Wang, J., Sun, K., Jia, X. & Liu, C.-j. CO₂ Hydrogenation to Methanol over Rh/In₂O₃ Catalyst. *Catalysis Today* **365**, 341–347. ISSN: 09205861 (Apr. 2021).
98. Wang, W., Chen, Y. & Zhang, M. Facet Effect of In₂O₃ for Methanol Synthesis by CO₂ Hydrogenation: A Mechanistic and Kinetic Study. *Surfaces and Interfaces* **25**, 101244. ISSN: 24680230 (Aug. 2021).

99. Ye, J., Liu, C., Mei, D. & Ge, Q. Active Oxygen Vacancy Site for Methanol Synthesis from CO₂ Hydrogenation on In₂O₃ (110): A DFT Study. *ACS Catal.* **3**, 1296–1306. ISSN: 2155-5435, 2155-5435 (June 7, 2013).
100. Taseska, T., Yu, W., Wilsey, M. K., Cox, C. P., Meng, Z., Ngarnim, S. S. & Müller, A. M. Analysis of the Scale of Global Human Needs and Opportunities for Sustainable Catalytic Technologies. *Top Catal.* **66**, 338–374. ISSN: 1572-9028 (Mar. 1, 2023).
101. Laverdura, U. P., Rossi, L., Ferella, F., Courson, C., Zarli, A., Alhajjoussef, R. & Gallucci, K. Selective Catalytic Hydrogenation of Vegetable Oils on Lindlar Catalyst. *ACS Omega* **5**, 22901–22913 (Sept. 15, 2020).
102. Humphreys, J., Lan, R. & Tao, S. Development and Recent Progress on Ammonia Synthesis Catalysts for Haber–Bosch Process. *Advanced Energy and Sustainability Research* **2**, 2000043. ISSN: 2699-9412 (2021).
103. Zhou, L., Lu, S.-Y. & Guo, S. Recent Progress on Precious Metal Single Atom Materials for Water Splitting Catalysis. *SusMat* **1**, 194–210. ISSN: 2692-4552 (2021).
104. Crabtree, G. W., Dresselhaus, M. S. & Buchanan, M. V. The Hydrogen Economy. *Physics Today* **57**, 39–44. ISSN: 0031-9228 (Dec. 1, 2004).
105. Räuchle, K., Plass, L., Wernicke, H.-J. & Bertau, M. Methanol for Renewable Energy Storage and Utilization. *Energy Technology* **4**, 193–200. ISSN: 2194-4296 (2016).
106. Palkovits, R. Sustainable Carbon Sources and Renewable Energy: Challenges and Opportunities at the Interface of Catalysis and Reaction Engineering. *Chemie Ingenieur Technik* **90**, 1699–1708. ISSN: 1522-2640 (2018).
107. Timmermann, J. *et al.* IrO₂ Surface Complexions Identified through Machine Learning and Surface Investigations. *Phys. Rev. Lett.* **125**, 206101. ISSN: 0031-9007, 1079-7114 (Nov. 10, 2020).
108. Dattila, F., Garcé-Muelas, R. & López, N. Active and Selective Ensembles in Oxide-Derived Copper Catalysts for CO₂ Reduction. *ACS Energy Lett.* **5**, 3176–3184. ISSN: 2380-8195, 2380-8195 (Oct. 9, 2020).
109. Peleg, M., Normand, M. D. & Corradini, M. G. The Arrhenius Equation Revisited. *Critical Reviews in Food Science and Nutrition* **52**, 830–851. ISSN: 1040-8398, 1549-7852 (Sept. 2012).
110. Eyring, H. The Activated Complex in Chemical Reactions. *The Journal of Chemical Physics* **3**, 107–115. ISSN: 0021-9606 (Feb. 1, 1935).
111. Sauer, J. Ab Initio Calculations for Molecule–Surface Interactions with Chemical Accuracy. *Acc. Chem. Res.* **52**, 3502–3510. ISSN: 0001-4842, 1520-4898 (Dec. 17, 2019).
112. Kozuch, S. & Shaik, S. How to Conceptualize Catalytic Cycles? The Energetic Span Model. *Acc. Chem. Res.* **44**, 101–110. ISSN: 0001-4842, 1520-4898 (Feb. 15, 2011).
113. Stocker, S., Csányi, G., Reuter, K. & Margraf, J. T. Machine Learning in Chemical Reaction Space. *Nat Commun* **11**, 5505. ISSN: 2041-1723 (Dec. 2020).

114. Marcinkowski, M. D., Jewell, A. D., Stamatakis, M., Boucher, M. B., Lewis, E. A., Murphy, C. J., Kyriakou, G. & Sykes, E. C. H. Controlling a Spillover Pathway with the Molecular Cork Effect. *Nature Mater* **12**, 523–528. ISSN: 1476-1122, 1476-4660 (June 2013).
115. Darby, M. T., Lucci, F. R., Marcinkowski, M. D., Therrien, A. J., Michaelides, A., Stamatakis, M. & Sykes, E. C. H. Carbon Monoxide Mediated Hydrogen Release from PtCu Single-Atom Alloys: The Punctured Molecular Cork Effect. *J. Phys. Chem. C* **123**, 10419–10428. ISSN: 1932-7447, 1932-7455 (Apr. 25, 2019).
116. Berger, F., Schumann, J., Réocreux, R., Stamatakis, M. & Michaelides, A. Bringing Molecules Together: Synergistic Coadsorption at Dopant Sites of Single Atom Alloys. *J. Am. Chem. Soc.* **146**, 28119–28130. ISSN: 0002-7863 (Oct. 16, 2024).
117. Weitkamp, J. Zeolites and Catalysis. *Solid State Ionics* **131**, 175–188. ISSN: 0167-2738 (June 1, 2000).
118. Csicsery, S. M. Shape-Selective Catalysis in Zeolites. *Zeolites* **4**, 202–213. ISSN: 0144-2449 (July 1, 1984).
119. Vandermause, J., Xie, Y., Lim, J. S., Owen, C. J. & Kozinsky, B. Active Learning of Reactive Bayesian Force Fields Applied to Heterogeneous Catalysis Dynamics of H/Pt. *Nat Commun* **13**, 5183. ISSN: 2041-1723. PMID: 36055982 (Sept. 2, 2022).
120. Giannozzi, P. *et al.* Advanced Capabilities for Materials Modelling with Quantum ESPRESSO. *J Phys Condens Matter* **29**, 465901. ISSN: 1361-648X. PMID: 29064822 (Nov. 22, 2017).
121. Monserrat, B., Brandenburg, J. G., Engel, E. A. & Cheng, B. Liquid Water Contains the Building Blocks of Diverse Ice Phases. *Nat Commun* **11**, 5757. ISSN: 2041-1723 (1 Nov. 13, 2020).
122. Schütt, K. T., Kindermans, P.-J., Sauceda, H. E., Chmiela, S., Tkatchenko, A. & Müller, K.-R. *SchNet: A Continuous-Filter Convolutional Neural Network for Modeling Quantum Interactions* Dec. 19, 2017. arXiv: 1706.08566.
123. Schran, C., Brezina, K. & Marsalek, O. Committee Neural Network Potentials Control Generalization Errors and Enable Active Learning. *The Journal of Chemical Physics* **153**, 104105. ISSN: 0021-9606 (Sept. 8, 2020).
124. Novikov, I. S., Gubaev, K., Podryabinkin, E. V. & Shapeev, A. V. The MLIP Package: Moment Tensor Potentials with MPI and Active Learning. *Mach. Learn.: Sci. Technol.* **2**, 025002. ISSN: 2632-2153 (Jan. 1, 2021).
125. Titsias, M. *Variational Learning of Inducing Variables in Sparse Gaussian Processes in Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics* Artificial Intelligence and Statistics (PMLR, Apr. 15, 2009), 567–574.
126. Dai, J., Adhikari, S. & Wen, M. *Uncertainty Quantification and Propagation in Atomistic Machine Learning* arXiv: 2405.02461. (2024). Pre-published.
127. Imbalzano, G., Zhuang, Y., Kapil, V., Rossi, K., Engel, E. A., Grasselli, F. & Ceriotti, M. *Uncertainty Estimation for Molecular Dynamics and Sampling* arXiv: 2011.08828. (2024). Pre-published.

128. Jinnouchi, R., Karsai, F. & Kresse, G. On-the-Fly Machine Learning Force Field Generation: Application to Melting Points. *Phys. Rev. B* **100**, 014105. ISSN: 2469-9950, 2469-9969 (July 17, 2019).
129. Smith, J. S., Nebgen, B., Lubbers, N., Isayev, O. & Roitberg, A. E. Less Is More: Sampling Chemical Space with Active Learning. *The Journal of Chemical Physics* **148**, 241733. ISSN: 0021-9606 (May 22, 2018).
130. Jain, A., Ong, S. P., Hautier, G., Chen, W., Richards, W. D., Dacek, S., Cholia, S., Gunter, D., Skinner, D., Ceder, G. & Persson, K. A. Commentary: The Materials Project: A Materials Genome Approach to Accelerating Materials Innovation. *APL Materials* **1**, 011002. ISSN: 2166-532X (July 18, 2013).
131. Boes, J. R., Mamun, O., Winther, K. & Bligaard, T. Graph Theory Approach to High-Throughput Surface Adsorption Structure Generation. *J. Phys. Chem. A* **123**, 2281–2285. ISSN: 1089-5639, 1520-5215 (Mar. 21, 2019).
132. De, S., Bartók, A. P., Csányi, G. & Ceriotti, M. Comparing Molecules and Solids across Structural and Alchemical Space. *Phys. Chem. Chem. Phys.* **18**, 13754–13769. ISSN: 1463-9076, 1463-9084 (2016).
133. Jónsson, H., Mills, G. & Jacobsen, K. W. in *Classical and Quantum Dynamics in Condensed Phase Simulations* 385–404 (WORLD SCIENTIFIC, June 1998). ISBN: 978-981-02-3498-0.
134. Henkelman, G., Uberuaga, B. P. & Jónsson, H. A Climbing Image Nudged Elastic Band Method for Finding Saddle Points and Minimum Energy Paths. *The Journal of Chemical Physics* **113**, 9901–9904. ISSN: 0021-9606, 1089-7690 (Dec. 8, 2000).
135. Henkelman, G. & Jónsson, H. A Dimer Method for Finding Saddle Points on High Dimensional Potential Surfaces Using Only First Derivatives. *The Journal of Chemical Physics* **111**, 7010–7022. ISSN: 0021-9606 (Oct. 15, 1999).
136. Ye, J., Liu, C. & Ge, Q. DFT Study of CO₂ Adsorption and Hydrogenation on the In₂O₃ Surface. *J. Phys. Chem. C* **116**, 7817–7825. ISSN: 1932-7447, 1932-7455 (Apr. 12, 2012).
137. Ye, J., Liu, C., Mei, D. & Ge, Q. Active Oxygen Vacancy Site for Methanol Synthesis from CO₂ Hydrogenation on In₂O₃(110): A DFT Study. *ACS Catal.*, **11** (2013).
138. Zhang, M., Wang, W. & Chen, Y. Insight of DFT and Ab Initio Atomistic Thermodynamics on the Surface Stability and Morphology of In₂O₃. *Applied Surface Science* **434**, 1344–1352. ISSN: 01694332 (Mar. 2018).
139. Frei, M. S., Capdevila-Cortada, M., García-Muelas, R., Mondelli, C., López, N., Stewart, J. A., Curulla Ferré, D. & Pérez-Ramírez, J. Mechanism and Microkinetics of Methanol Synthesis via CO₂ Hydrogenation on Indium Oxide. *Journal of Catalysis* **361**, 313–321. ISSN: 0021-9517 (May 1, 2018).
140. Dou, M., Zhang, M., Chen, Y. & Yu, Y. DFT Study of In₂O₃-Catalyzed Methanol Synthesis from CO₂ and CO Hydrogenation on the Defective Site. *New J. Chem.* **42**, 3293–3300. ISSN: 1144-0546, 1369-9261 (2018).
141. Nakamura, J., Choi, Y. & Fujitani, T. On the Issue of the Active Site and the Role of ZnO in Cu/ZnO Methanol Synthesis Catalysts. *Topics in Catalysis* **22**, 277–285. ISSN: 1572-9028 (Apr. 1, 2003).

142. Studt, F., Behrens, M., Kunkes, E. L., Thomas, N., Zander, S., Tarasov, A., Schumann, J., Frei, E., Varley, J. B., Abild-Pedersen, F., Nørskov, J. K. & Schlögl, R. The Mechanism of CO and CO₂ Hydrogenation to Methanol over Cu-Based Catalysts. *ChemCatChem* **7**, 1105–1111. ISSN: 1867-3899 (2015).
143. Cao, A., Wang, Z., Li, H. & Nørskov, J. K. Relations between Surface Oxygen Vacancies and Activity of Methanol Formation from CO₂ Hydrogenation over In₂O₃ Surfaces. *ACS Catal.* **11**, 1780–1786. ISSN: 2155-5435, 2155-5435 (Feb. 5, 2021).
144. Bielz, T., Lorenz, H., Jochum, W., Kaindl, R., Klauser, F., Klötzer, B. & Penner, S. Hydrogen on In₂O₃: Reducibility, Bonding, Defect Formation, and Reactivity. *J. Phys. Chem. C* **114**, 9022–9029. ISSN: 1932-7447, 1932-7455 (May 20, 2010).
145. Berne, B. J., Ciccotti, G. & Coker, D. F. *Classical And Quantum Dynamics In Condensed Phase Simulations: Proceedings Of The International School Of Physics* 881 pp. ISBN: 978-981-4496-05-6. Google Books: [UoTVCgAAQBAJ](#) (World Scientific, June 17, 1998).
146. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Phys. Rev. Lett.* **77**, 3865–3868 (Oct. 28, 1996).
147. Blöchl, P. E. Projector Augmented-Wave Method. *Phys. Rev. B* **50**, 17953–17979 (Dec. 15, 1994).
148. Kresse, G. & Joubert, D. From Ultrasoft Pseudopotentials to the Projector Augmented-Wave Method. *Phys. Rev. B* **59**, 1758–1775 (Jan. 15, 1999).
149. Dal Corso, A. Pseudopotentials Periodic Table: From H to Pu. *Computational Materials Science* **95**, 337–350. ISSN: 0927-0256 (Dec. 1, 2014).
150. Caro, M. A. Optimizing Many-Body Atomic Descriptors for Enhanced Computational Performance of Machine Learning Based Interatomic Potentials. *Phys. Rev. B* **100**, 024112. ISSN: 2469-9950, 2469-9969 (July 30, 2019).
151. Kermode, J. R. F90wrap: An Automated Tool for Constructing Deep Python Interfaces to Modern Fortran Codes. *J. Phys.: Condens. Matter* **32**, 305901. ISSN: 0953-8984, 1361-648X (July 15, 2020).
152. Csanyi, G., Winfield, S., Kermode, J., Payne, M. C., Comisso, A., Vita, A. D. & Bernstein, N. Expressive Programming for Computational Physics in Fortran 950+. *NCP*, 1–24 (2007).
153. Larsen, A. H. *et al.* The Atomic Simulation Environment—a Python Library for Working with Atoms. *Journal of Physics: Condensed Matter* **29**, 273002. ISSN: 0953-8984 (June 2017).
154. Makri, S., Ortner, C. & Kermode, J. R. A Preconditioning Scheme for Minimum Energy Path Finding Methods. *J. Chem. Phys.* **150**, 094109. ISSN: 0021-9606, 1089-7690 (Mar. 7, 2019).
155. Smidstrup, S., Pedersen, A., Stokbro, K. & Jónsson, H. Improved Initial Guess for Minimum Energy Path Calculations. *J. Chem. Phys.* **140**, 214106. ISSN: 0021-9606 (June 7, 2014).
156. Fiorin, G., Klein, M. L. & Hémin, J. Using Collective Variables to Drive Molecular Dynamics Simulations. *Molecular Physics* **111**, 3345–3362. ISSN: 0026-8976, 1362-3028 (Dec. 2013).

157. Thompson, A. P. *et al.* LAMMPS - a Flexible Simulation Tool for Particle-Based Materials Modeling at the Atomic, Meso, and Continuum Scales. *Computer Physics Communications* **271**, 108171. ISSN: 00104655 (Feb. 2022).
158. Stecher, T., Bernstein, N. & Csányi, G. Free Energy Surface Reconstruction from Umbrella Samples Using Gaussian Process Regression. *J. Chem. Theory Comput.* **10**, 4079–4097. ISSN: 1549-9618 (Sept. 9, 2014).
159. Li, X. & Rupprechter, G. Sum Frequency Generation Spectroscopy in Heterogeneous Model Catalysis: A Minireview of CO-related Processes. *Catal. Sci. Technol.* **11**, 12–26. ISSN: 2044-4753, 2044-4761 (2021).
160. Khatamirad, M., Fako, E., Boscagli, C., Müller, M., Ebert, F., d'Alnoncourt, R. N., Schaefer, A., Andreas Schunk, S., Jevtovikj, I., Rosowski, F. & De, S. A Data-Driven High-Throughput Workflow Applied to Promoted In-oxide Catalysts for CO₂ Hydrogenation to Methanol. *Catalysis Science & Technology* **13**, 2656–2661 (2023).
161. Xie, W., Xu, J., Chen, J., Wang, H. & Hu, P. Achieving Theory–Experiment Parity for Activity and Selectivity in Heterogeneous Catalysis Using Microkinetic Modeling. *Acc. Chem. Res.* **55**, 1237–1248. ISSN: 0001-4842, 1520-4898 (May 3, 2022).
162. Kozuch, S. & Shaik, S. How to Conceptualize Catalytic Cycles? The Energetic Span Model. *Acc. Chem. Res.* **44**, 101–110. ISSN: 0001-4842 (Feb. 15, 2011).
163. Goedecker, S. Minima Hopping: An Efficient Search Method for the Global Minimum of the Potential Energy Surface of Complex Molecular Systems. *The Journal of Chemical Physics* **120**, 9911–9917. ISSN: 0021-9606, 1089-7690 (June 2004).
164. Jung, H., Sauerland, L., Stocker, S., Reuter, K. & Margraf, J. T. Machine-Learning Driven Global Optimization of Surface Adsorbate Geometries. *npj Comput Mater* **9**, 1–8. ISSN: 2057-3960 (1 June 26, 2023).
165. Frei, M. S., Mondelli, C., García-Muelas, R., Morales-Vidal, J., Philipp, M., Safonova, O. V., López, N., Stewart, J. A., Ferré, D. C. & Pérez-Ramírez, J. Nanostructure of Nickel-Promoted Indium Oxide Catalysts Drives Selectivity in CO₂ Hydrogenation. *Nat Commun* **12**, 1960. ISSN: 2041-1723 (1 Mar. 30, 2021).
166. Frei, M. S., Mondelli, C., García-Muelas, R., Kley, K. S., Puértolas, B., López, N., Safonova, O. V., Stewart, J. A., Curulla Ferré, D. & Pérez-Ramírez, J. Atomic-Scale Engineering of Indium Oxide Promotion by Palladium for Methanol Production via CO₂ Hydrogenation. *Nat Commun* **10**, 3377. ISSN: 2041-1723 (1 July 29, 2019).
167. Reuter, K. Ab Initio Thermodynamics and First-Principles Microkinetics for Surface Catalysis. *Catal Lett* **146**, 541–563. ISSN: 1572-879X (Mar. 1, 2016).
168. Johnson, M. S., Gierada, M., Hermes, E. D., Bross, D. H., Sargsyan, K., Najm, H. N. & Zádor, J. Pynta-An Automated Workflow for Calculation of Surface and Gas–Surface Kinetics. *J. Chem. Inf. Model.* ISSN: 1549-9596 (Aug. 9, 2023).
169. Rhodes, B., Vandenhaute, S., Šimkus, V., Gin, J., Godwin, J., Duignan, T. & Neumann, M. *Orb-v3: Atomistic Simulation at Scale* arXiv: 2504.06231. (2026). Pre-published.

170. Rhodes, B. J., Schaaf, L. L., Zick, M. E., Pugh, S. M., Hilliard, J. S., Sharma, S., Wade, C. R., Milner, P. J., Csányi, G. & Forse, A. C. 17O NMR Spectroscopy Reveals CO₂ Speciation and Dynamics in Hydroxide-Based Carbon Capture Materials. *ChemPhysChem*, e202400941. ISSN: 1439-7641 (2024).
171. Schaaf, L. L., Rhodes, B. J., Zick, M. E., Pugh, S. M., Hilliard, J. S., Sharma, S., Wade, C. R., Milner, P. J., Csányi, G. & Forse, A. C. ML Force Fields for Computational NMR Spectra of Dynamic Materials across Time-Scales. *Neurips 2024 Workshop on AI for Accelerated Materials Design* (2024).
172. Devereux, C., Smith, J. S., Huddleston, K. K., Barros, K., Zubatyuk, R., Isayev, O. & Roitberg, A. E. Extending the Applicability of the ANI Deep Learning Molecular Potential to Sulfur and Halogens. *J. Chem. Theory Comput.* **16**, 4192–4202. ISSN: 1549-9618, 1549-9626 (July 14, 2020).
173. Smith, J. S., Isayev, O. & Roitberg, A. E. ANI-1: An Extensible Neural Network Potential with DFT Accuracy at Force Field Computational Cost. *Chem. Sci.* **8**, 3192–3203. ISSN: 2041-6520, 2041-6539 (2017).
174. Anstine, D., Zubatyuk, R. & Isayev, O. *AIMNet2: A Neural Network Potential to Meet Your Neutral, Charged, Organic, and Elemental-Organic Needs* (2024). Pre-published.
175. Zubatyuk, R., Smith, J. S., Nebgen, B. T., Tretiak, S. & Isayev, O. Teaching a Neural Network to Attach and Detach Electrons from Molecules. *Nat Commun* **12**, 4870. ISSN: 2041-1723 (Aug. 11, 2021).
176. Zubatyuk, R., Smith, J. S., Leszczynski, J. & Isayev, O. Accurate and Transferable Multitask Prediction of Chemical Properties with an Atoms-in-Molecules Neural Network. *Sci. Adv.* **5**, eaav6490. ISSN: 2375-2548 (Aug. 2, 2019).
177. Smith, J. S., Nebgen, B. T., Zubatyuk, R., Lubbers, N., Devereux, C., Barros, K., Tretiak, S., Isayev, O. & Roitberg, A. E. Approaching Coupled Cluster Accuracy with a General-Purpose Neural Network Potential through Transfer Learning. *Nat Commun* **10**, 2903. ISSN: 2041-1723 (July 1, 2019).
178. Gasteiger, J., Giri, S., Margraf, J. T. & Günnemann, S. *Fast and Uncertainty-Aware Directional Message Passing for Non-Equilibrium Molecules* arXiv: 2011.14115. (2024). Pre-published.
179. Liao, Y.-L., Wood, B., Das, A. & Smidt, T. *EquiformerV2: Improved Equivariant Transformer for Scaling to Higher-Degree Representations* arXiv: 2306.12059. (2024). Pre-published.
180. Abdelmaqsoud, K., Shuaibi, M., Kolluru, A., Cheula, R. & Kitchin, J. R. Investigating the Error Imbalance of Large-Scale Machine Learning Potentials in Catalysis. *Catal. Sci. Technol.* **14**, 5899–5908. ISSN: 2044-4761 (Oct. 14, 2024).
181. Wander, B., Shuaibi, M., Kitchin, J. R., Ulissi, Z. W. & Zitnick, C. L. *CatTSunami: Accelerating Transition State Energy Calculations with Pre-trained Graph Neural Networks* version 2. arXiv: 2405.02078. (2024). Pre-published.
182. Chen, C., Ye, W., Zuo, Y., Zheng, C. & Ong, S. P. Graph Networks as a Universal Machine Learning Framework for Molecules and Crystals. *Chem. Mater.* **31**, 3564–3572. ISSN: 0897-4756, 1520-5002 (May 14, 2019).

183. Deng, B., Zhong, P., Jun, K., Riebesell, J., Han, K., Bartel, C. J. & Ceder, G. CHGNet as a Pretrained Universal Neural Network Potential for Charge-Informed Atomistic Modelling. *Nat Mach Intell* **5**, 1031–1041. ISSN: 2522-5839 (Sept. 2023).
184. Merchant, A., Batzner, S., Schoenholz, S. S., Aykol, M., Cheon, G. & Cubuk, E. D. Scaling Deep Learning for Materials Discovery. *Nature* **624**, 80–85. ISSN: 1476-4687 (Dec. 2023).
185. Shoghi, N., Kolluru, A., Kitchin, J. R., Ulissi, Z. W., Zitnick, C. L. & Wood, B. M. *From Molecules to Materials: Pre-training Large Generalizable Models for Atomic Property Prediction* arXiv: 2310.16802. (2024). Pre-published.
186. McInnes, L., Healy, J. & Melville, J. *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction* arXiv: 1802.03426. (2024). Pre-published.
187. Kim, S. *et al.* PubChem 2025 Update. *Nucleic Acids Research* **53**, D1516–D1525. ISSN: 1362-4962 (Jan. 6, 2025).
188. Deng, B., Choi, Y., Zhong, P., Riebesell, J., Anand, S., Li, Z., Jun, K., Persson, K. A. & Ceder, G. *Overcoming Systematic Softening in Universal Machine Learning Interatomic Potentials by Fine-Tuning* arXiv: 2405.07105. (2024). Pre-published.
189. Kentgens, A. P. M. A Practical Guide to Solid-State NMR of Half-Integer Quadrupolar Nuclei with Some Applications to Disordered Systems. *Geoderma* **80**, 271–306. ISSN: 0016-7061 (1997).
190. Engel, E. A., Kapil, V. & Ceriotti, M. Importance of Nuclear Quantum Effects for NMR Crystallography. *J. Phys. Chem. Lett.* **12**, 7701–7707. ISSN: 1948-7185, 1948-7185 (Aug. 19, 2021).
191. Tsitsvero, M., Pirillo, J., Hijikata, Y. & Komatsuzaki, T. NMR Spectrum Prediction for Dynamic Molecules by Machine Learning: A Case Study of Trefoil Knot Molecule. *The Journal of Chemical Physics* **158**, 194108. ISSN: 0021-9606 (May 17, 2023).
192. Vold, R. L. & Hoatson, G. L. Effects of Jump Dynamics on Solid State Nuclear Magnetic Resonance Line Shapes and Spin Relaxation Times. *Journal of Magnetic Resonance* **198**, 57–72. ISSN: 1090-7807 (2009).
193. Günther, H. *NMR Spectroscopy: Basic Principles, Concepts and Applications in Chemistry* 739 pp. ISBN: 978-3-527-33000-3. Google Books: [9eqdAAAAQBAJ](#) (John Wiley & Sons, Nov. 4, 2013).
194. Ramsey, N. F. Magnetic Shielding of Nuclei in Molecules. *Phys. Rev.* **78**, 699–703 (June 15, 1950).
195. Pickard, C. J. & Mauri, F. All-Electron Magnetic Response with Pseudopotentials: NMR Chemical Shifts. *Physical Review B* **63**, 245101 (2001).
196. Yates, J. R., Pickard, C. J. & Mauri, F. Calculation of NMR Chemical Shifts for Extended Systems Using Ultrasoft Pseudopotentials. *Physical Review B* **76**, 024401 (2007).
197. Bonhomme, C., Gervais, C., Babonneau, F., Coelho, C., Pourpoint, F., Azais, T., Ashbrook, S. E., Griffin, J. M., Yates, J. R., Mauri, F., *et al.* First-Principles Calculation of NMR Parameters Using the Gauge Including Projector Augmented Wave Method: A Chemist’s Point of View. *Chemical Reviews* **112**, 5733–5779 (2012).

198. Valenzuela Reina, J., Civaia, F., Harper, A. F., Scheurer, C. & Köcher, S. The EFG Rosetta Stone: Translating between DFT Calculations and Solid State NMR Experiments. *Faraday Discuss.*, 10.1039.D4FD00075G. ISSN: 1359-6640, 1364-5498 (2024).
199. Helgaker, T., Jaszuński, M. & Ruud, K. Ab Initio Methods for the Calculation of NMR Shielding and Indirect Spin-Spin Coupling Constants. *Chem. Rev.* **99**, 293–352. ISSN: 0009-2665, 1520-6890 (Jan. 13, 1999).
200. Bonhomme, C., Gervais, C., Babonneau, F., Coelho, C., Pourpoint, F., Azaïs, T., Ashbrook, S. E., Griffin, J. M., Yates, J. R., Mauri, F. & Pickard, C. J. First-Principles Calculation of NMR Parameters Using the Gauge Including Projector Augmented Wave Method: A Chemist's Point of View. *Chem. Rev.* **112**, 5733–5779. ISSN: 0009-2665 (Nov. 14, 2012).
201. Goedecker, S. Minima Hopping: An Efficient Search Method for the Global Minimum of the Potential Energy Surface of Complex Molecular Systems. *The Journal of chemical physics* **120**, 9911–9917 (2004).
202. Hu, X. E. *et al.* NMR Techniques and Prediction Models for the Analysis of Species Formed in CO₂ Capture Processes with Amine-Based Sorbents: A Critical Review. *ACS Sustainable Chemistry & Engineering* **8**, 6173–6193 (2020).
203. Pugh, S. M. & Forse, A. C. Nuclear Magnetic Resonance Studies of Carbon Dioxide Capture. *Journal of Magnetic Resonance* **346**, 107343. ISSN: 1090-7807 (2023).
204. Pereira, D., Fonseca, R., Marin-Montesinos, I., Sardo, M. & Mafra, L. Understanding CO₂ Adsorption Mechanisms in Porous Adsorbents: A Solid-State NMR Survey. *Current Opinion in Colloid & Interface Science*, 101690. ISSN: 1359-0294 (2023).
205. Forse, A. C. *et al.* Elucidating CO₂ Chemisorption in Diamine-Appended Metal–Organic Frameworks. *Journal of the American Chemical Society* **140**, 18016–18031. ISSN: 0002-7863 (2018).
206. Shimon, D., Chen, C.-H., Lee, J. J., Didas, S. A., Sievers, C., Jones, C. W. & Hayes, S. E. ¹⁵N Solid State NMR Spectroscopic Study of Surface Amine Groups for Carbon Capture: 3-Aminopropylsilyl Grafted to SBA-15 Mesoporous Silica. *Environmental Science & Technology* **52**, 1488–1495. ISSN: 0013-936X (2018).
207. Yang, H., Singh, M. & Schaefer, J. Humidity-Swing Mechanism for CO₂ Capture from Ambient Air. *Chemical Communications* **54**, 4915–4918. ISSN: 1359-7345 (2018).
208. Shi, X., Xiao, H., Kanamori, K., Yonezu, A., Lackner, K. S. & Chen, X. Moisture-Driven CO₂ Sorbents. *Joule* **4**, 1823–1837. ISSN: 2542-4351 (2020).
209. Berge, A. H., Pugh, S. M., Short, M. I. M., Kaur, C., Lu, Z., Lee, J.-H., Pickard, C. J., Sayari, A. & Forse, A. C. Revealing Carbon Capture Chemistry with ¹⁷Oxygen NMR Spectroscopy. *Nature Communications* **13**, 7763. ISSN: 2041-1723 (2022).
210. Bien, C. E., Chen, K. K., Chien, S.-C., Reiner, B. R., Lin, L.-C., Wade, C. R. & Ho, W. S. W. Bioinspired Metal–Organic Framework for Trace CO₂ Capture. *Journal of the American Chemical Society* **140**, 12662–12666. ISSN: 0002-7863 (2018).

211. Bien, C. E., Liu, Q. & Wade, C. R. Assessing the Role of Metal Identity on CO₂ Adsorption in Mofs Containing m-OH Functional Groups. *Chemistry of Materials* **32**, 489–497. ISSN: 0897-4756 (2020).
212. Liao, P.-Q., Chen, H., Zhou, D.-D., Liu, S.-Y., He, C.-T., Rui, Z., Ji, H., Zhang, J.-P. & Chen, X.-M. Monodentate Hydroxide as a Super Strong yet Reversible Active Site for CO₂ Capture from High-Humidity Flue Gas. *Energy & Environmental Science* **8**, 1011–1016. ISSN: 1754-5692 (2015).
213. Wright, A. M., Wu, Z., Zhang, G., Mancuso, J. L., Comito, R. J., Day, R. W., Hendon, C. H., Miller, J. T. & Dincă, M. A Structural Mimic of Carbonic Anhydrase in a Metal-Organic Framework. *Chem* **4**, 2894–2901. ISSN: 2451-9294 (2018).
214. Cai, Z., Bien, C. E., Liu, Q. & Wade, C. R. Insights into CO₂ Adsorption in m-OH Functionalized Mofs. *Chemistry of Materials* **32**, 4257–4264. ISSN: 0897-4756 (2020).
215. Zick, M. E., Pugh, S. M., Lee, J.-H., Forse, A. C. & Milner, P. J. Carbon Dioxide Capture at Nucleophilic Hydroxide Sites in Oxidation-Resistant Cyclodextrin-Based Metal-Organic Frameworks. *Angewandte Chemie International Edition* **61**, e202206718. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/anie.202206718> (2022).
216. Dopieralski, P. D., Latajka, Z. & Olovsson, I. Proton-Transfer Dynamics in the (HCO₃⁻)₂ Dimer of KHCO₃ from Car-Parrinello and Path-Integrals Molecular Dynamics Calculations. *Acta Crystallogr B Struct Sci* **66**, 222–228. ISSN: 0108-7681 (Apr. 1, 2010).
217. Dong, S., Yamada, K. & Wu, G. Oxygen-17 Nuclear Magnetic Resonance of Organic Solids. **55**, 21–28 (2000).
218. Kong, X., Shan, M., Terskikh, V., Hung, I., Gan, Z. & Wu, G. Solid-State ¹⁷O NMR of Pharmaceutical Compounds: Salicylic Acid and Aspirin. *The Journal of Physical Chemistry B* **117**, 9643–9654. ISSN: 1520-6106 (2013).
219. Wu, G., Hung, I., Gan, Z., Terskikh, V. & Kong, X. Solid-State ¹⁷O NMR Study of Carboxylic Acid Dimers: Simultaneously Accessing Spectral Properties of Low- and High-Energy Tautomers. *The Journal of Physical Chemistry A* **123**, 8243–8253. ISSN: 1089-5639 (2019).
220. Chen, C.-H., Goldberga, I., Gaveau, P., Mitteleite, S., Špačková, J., Mullen, C., Petit, I., Métro, T.-X., Alonso, B., Gervais, C. & Laurencin, D. Looking into the Dynamics of Molecular Crystals of Ibuprofen and Terephthalic Acid Using ¹⁷O and ²H Nuclear Magnetic Resonance Analyses. *Magnetic Resonance in Chemistry* **59**, 975–990. ISSN: 0749-1581 (2021).
221. Smaldone, R. A., Forgan, R. S., Furukawa, H., Gassensmith, J. J., Slawin, A. M. Z., Yaghi, O. M. & Stoddart, J. F. Metal-Organic Frameworks from Edible Natural Products. *Angewandte Chemie International Edition* **49**, 8630–8634. ISSN: 1433-7851 (2010).
222. Pham, H. D. M. & Khaliullin, R. Z. Unraveling the Origins of Strong and Reversible Chemisorption of Carbon Dioxide in a Green Metal-Organic Framework. *The Journal of Physical Chemistry C* **125**, 24719–24727. ISSN: 1932-7447 (2021).

223. Harper, A. F., Köcher, S. S., Reuter, K. & Scheurer, C. *Performance Metrics for Tensorial Learning: Prediction of Li4Ti5O12 Nuclear Magnetic Resonance Observables at Experimental Accuracy* (2024). Pre-published.
224. Zhu, J. *et al.* Data-Driven Capacity Estimation of Commercial Lithium-Ion Batteries from Voltage Relaxation. *Nat Commun* **13**, 2261. ISSN: 2041-1723 (Apr. 27, 2022).
225. Mattoni, A. & Meloni, S. Defect Dynamics in MAPbI3 Polycrystalline Films: The Trapping Effect of Grain Boundaries. *Helvetica Chimica Acta* **103**, e2000110. ISSN: 1522-2675 (2020).
226. *Incorporating Electronic Information into Machine Learning Potential Energy Surfaces via Approaching the Ground-State Electronic Energy as a Function of Atom-Based Electronic Populations | Journal of Chemical Theory and Computation* (2025).
227. Ko, T. W., Finkler, J. A., Goedecker, S. & Behler, J. A Fourth-Generation High-Dimensional Neural Network Potential with Accurate Electrostatics Including Non-Local Charge Transfer. *Nat Commun* **12**, 398. ISSN: 2041-1723 (Jan. 15, 2021).
228. Huguenin-Dumittan, K. K., Loche, P., Haoran, N. & Ceriotti, M. *Physics-Inspired Equivariant Descriptors of Non-bonded Interactions* arXiv: 2308.13208. (2025). Pre-published.
229. Faraji, E., Kurian, P., Franzosi, R., Mancini, S., Floriani, E., Calandrini, V., Pettini, G. & Pettini, M. *Electrodynamic Forces Driving DNA-protein Interactions at Large Distances* arXiv: 2412.12127. (2025). Pre-published.
230. Unke, O. T. & Meuwly, M. PhysNet: A Neural Network for Predicting Energies, Forces, Dipole Moments, and Partial Charges. *J. Chem. Theory Comput.* **15**, 3678–3693. ISSN: 1549-9618 (June 11, 2019).
231. Darley, M. G., Handley, C. M. & Popelier, P. L. A. Beyond Point Charges: Dynamic Polarization from Neural Net Predicted Multipole Moments. *J. Chem. Theory Comput.* **4**, 1435–1448. ISSN: 1549-9618, 1549-9626 (Sept. 9, 2008).
232. Behler, J. Four Generations of High-Dimensional Neural Network Potentials. *Chem. Rev.* **121**, 10037–10072. ISSN: 0009-2665 (Aug. 25, 2021).
233. Zhou, H.-X. & Pang, X. Electrostatic Interactions in Protein Structure, Folding, Binding, and Condensation. *Chem. Rev.* **118**, 1691–1741. ISSN: 0009-2665 (Feb. 28, 2018).
234. Zhou, H.-X. Interactions of macromolecules with salt ions: An electrostatic theory for the Hofmeister effect. *Proteins: Structure, Function, and Bioinformatics* **61**, 69–78. ISSN: 1097-0134 (2005).
235. Batatia, I., Schaaf, L. L., Chen, H., Csányi, G., Ortner, C. & Faber, F. A. Equivariant Matrix Function Neural Networks. *International Conference on Learning Representations 2024*. arXiv: 2310.10434 (Oct. 16, 2023).
236. Friede, M., Hölzer, C., Ehlert, S. & Grimme, S. Dxtb—An Efficient and Fully Differentiable Framework for Extended Tight-Binding. *The Journal of Chemical Physics* **161**, 062501. ISSN: 0021-9606 (Aug. 9, 2024).

237. Bannwarth, C., Ehlert, S. & Grimme, S. GFN2-xTB—An Accurate and Broadly Parametrized Self-Consistent Tight-Binding Quantum Chemical Method with Multipole Electrostatics and Density-Dependent Dispersion Contributions. *J. Chem. Theory Comput.* **15**, 1652–1671. ISSN: 1549-9618, 1549-9626 (Mar. 12, 2019).
238. Froitzheim, T., Müller, M., Hansen, A. & Grimme, S. *G-xTB: A General-Purpose Extended Tight-Binding Electronic Structure Method For the Elements H to Lr (Z=1–103)* (2026). Pre-published.
239. Westermayr, J. & Maurer, R. J. Physically Inspired Deep Learning of Molecular Excitations and Photoemission Spectra. *Chem. Sci.* **12**, 10755–10764. ISSN: 2041-6520, 2041-6539 (2021).
240. Nigam, J., Willatt, M. J. & Ceriotti, M. Equivariant Representations for Molecular Hamiltonians and N-center Atomic-Scale Properties. *The Journal of Chemical Physics* **156**, 014115. ISSN: 0021-9606 (Jan. 4, 2022).
241. Cignoni, E., Suman, D., Nigam, J., Cupellini, L., Mennucci, B. & Ceriotti, M. Electronic Excited States from Physically Constrained Machine Learning. *ACS Cent. Sci.* **10**, 637–648. ISSN: 2374-7943 (Mar. 27, 2024).
242. Stöhr, M., Sandonas, L. M. & Tkatchenko, A. *Accurate Many-Body Repulsive Potentials for Density-Functional Tight-Binding from Deep Tensor Neural Networks* arXiv: 2006.10429. (2024). Pre-published.
243. Zhang, L., Onat, B., Dusson, G., Anand, G., Maurer, R. J., Ortner, C. & Kermode, J. R. Equivariant Analytical Mapping of First Principles Hamiltonians to Accurate and Transferable Materials Models. *npj Computational Materials* **8** (2022).
244. Zhang, L., Onat, B., Dusson, G., McSloy, A., Anand, G., Maurer, R. J., Ortner, C. & Kermode, J. R. Equivariant Analytical Mapping of First Principles Hamiltonians to Accurate and Transferable Materials Models. *npj Comput Mater* **8**, 1–14. ISSN: 2057-3960 (July 22, 2022).
245. Ioffe, S. & Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift.
246. Ba, J. L., Kiros, J. R. & Hinton, G. E. *Layer Normalization* arXiv: 1607.06450. (2025). Pre-published.
247. Lin, L., Lu, J., Ying, L., Car, R. & E, W. Fast Algorithm for Extracting the Diagonal of the Inverse Matrix with Application to the Electronic Structure Analysis of Metallic Systems. *Comm. Math. Sci.* **7**, 755 (2009).
248. Goedecker, S. Linear Scaling Electronic Structure Methods. *Rev. Mod. Phys.* **71**, 1085–1123 (July 1999).
249. Lin, L., Chen, M., Yang, C. & He, L. Accelerating Atomic Orbital-Based Electronic Structure Calculation via Pole Expansion and Selected Inversion. *J. Phys. Condens. Matter* **25**, 295501 (2013).
250. Higham, N. J. *Functions of Matrices: Theory and Computation* 432 pp. ISBN: 978-0-89871-777-8. Google Books: [S6gpNn1JmbgC](#) (SIAM, Jan. 1, 2008).
251. Combes, J. M. & Thomas, L. Asymptotic Behaviour of Eigenfunctions for Multiparticle Schrödinger Operators. *Communications in Mathematical Physics* **34**, 251–270 (Dec. 1973).

252. Chung, F. R. *Spectral Graph Theory* (American Mathematical Soc., 1997).
253. Unke, O. T., Chmiela, S., Sauceda, H. E., Gastegger, M., Poltavsky, I., Schütt, K. T., Tkatchenko, A. & Müller, K.-R. Machine Learning Force Fields. *Chem. Rev.* **121**, 10142–10186. ISSN: 0009-2665 (Aug. 25, 2021).
254. Joshi, C. K., Bodnar, C., Mathis, S. V., Cohen, T. & Liò, P. *On the Expressive Power of Geometric Graph Neural Networks in International Conference on Machine Learning* (2023).
255. Frank, T., Unke, O. & Müller, K.-R. So3krates: Equivariant Attention for Interactions on Arbitrary Length-Scales in Molecular Systems. *Advances in Neural Information Processing Systems* **35**, 29400–29413 (2022).
256. Chmiela, S., Vassilev-Galindo, V., Unke, O. T., Kabylda, A., Sauceda, H. E., Tkatchenko, A. & Müller, K.-R. Accurate Global Machine Learning Force Fields for Molecules with Hundreds of Atoms. *Sci. Adv.* **9**, eadf0873. ISSN: 2375-2548 (Jan. 13, 2023).
257. Batatia, I., Kovacs, D. P., Simm, G., Ortner, C. & Csanyi, G. MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields. *Advances in Neural Information Processing Systems* **35**, 11423–11436 (Dec. 6, 2022).

List of figures

2.1	Atomistic simulations and the trade-off between accuracy and computational cost	16
2.2	Machine learning force fields and relevant symmetries	29
3.1	Challenges in computational catalysis	47
3.2	Predicted energy uncertainty during molecular dynamics (MD) . . .	54
3.3	Automated training protocol	56
3.4	Energy barrier accuracy throughout active learning	61
3.5	Lower rate-limiting step	64
3.6	Free energy barriers for formaldehyde production	65
3.7	Adsorbate configuration exploration	67
3.8	Reduced surface reaction pathway	68
3.9	Training set transferability to other MLFF frameworks	69
4.1	Transition state searching using foundation model	77
4.2	Molecular dynamics snapshot of an indium oxide slab covered with diverse molecules	78
4.3	Fine-tuning foundation models	79
5.1	Capturing dynamic effects of NMR spectra across timescales	90
5.2	Dynamics of adsorbed CO ₂ in MFU-4l MOF across chemical exchange regimes	93
5.3	Fast regime dynamic modelling of KHCO ₃ (s)	94
5.4	The NMR spectra and assignments of CO ₂ dosed KHCO ₃ CD-MOF .	95
5.5	Proposed CO ₂ binding mechanism for KHCO ₃ -CD-MOF	96
6.1	Matrix function network architecture	101
6.2	Edge feature construction	103
6.3	Block structure of a Euclidean MFN operator, \mathbf{H}	105
6.4	Matrix block decomposition for SO(3)	106

6.5	Matrix powers series and non-locality	111
6.6	Matrix elements of the resolvent around $z_l = x + iy$	113
6.7	Visualizing MFN expressivity on cumulene chains	116
6.8	Angles and distances that define a cumulene graph	119

List of tables

2.1	Multiplication table of the rank 3 permutation group, S_3	20
6.1	Guaranteed non-local cumulene dataset	118
6.2	Model training parameters	120
A.1	Overview of training data-set size throughout the active learning blocks	152
A.2	DFT setting for approximate intermediates	152

Appendix A

Computational details

CO₂ hydrogenation to methanol on In₂O₃ surface

Please see further details about the computational setup for [Chapter 3](#). For visualisations of the reactions paths, please see the publication in reference [\[54\]](#) directly.

Supplementary Note 1: Dataset overview

The dataset accumulated throughout active learning can be obtained from the Zenodo repository, <https://doi.org/10.5281/zenodo.8268725>, as detailed in the data availability section. [Table A.1](#) outlines the number of configurations in the training set throughout the training protocol for the training dataset curated for validation against Dang et al [\[87\]](#). Each configuration in the online repository is labelled with details on which active learning (AL) block and which iteration the configurations originate from.

Supplementary Note 2: Obtaining approximate intermediates

A centerpiece of our investigation is that we can benchmark our protocol against a DFT study. We hence focus on the exact reaction pathway explored by *Ref* [\[87\]](#). To ensure we investigate similar local minima, we need to choose an adsorbate configuration close to the one found in the paper. To emulate the results from an early path search, we take the minima from the paper and geometry optimize the molecular geometries with less accurate and therefore less computationally demanding DFT settings. The details of the changed DFT settings are listed in [Table A.2](#). In comparison to the convergence parameters used for creating the training set, the CPU cost is a factor of 27 lower on a single core (as tested on a slab configuration). Note also, that the Dang investigation used VASP rather than Quantum Espresso (QE). Ultimately

Table A.1 Overview of training data-set size throughout the active learning blocks. The majority of the training configurations contain molecules and surfaces.

AL block	block name	iterations	counts
0	initial dataset	0: isolated	4
		0: material project	1
1	bulk	0: rattled + deformed	10
		1	7
2	surfaces	0	2
		1	20
3	molecules-md	0: approximate intermediates	11
		0: CatKit	36
		0: dimers	24
		1-6	66
4	molecules-geom-opt	1-12	144
5	molecules-neb	0: interpolated	5
		1-6	292
Sum			622

Table A.2 DFT setting for approximate intermediates. A comparison with the default setting used to curate the training data.

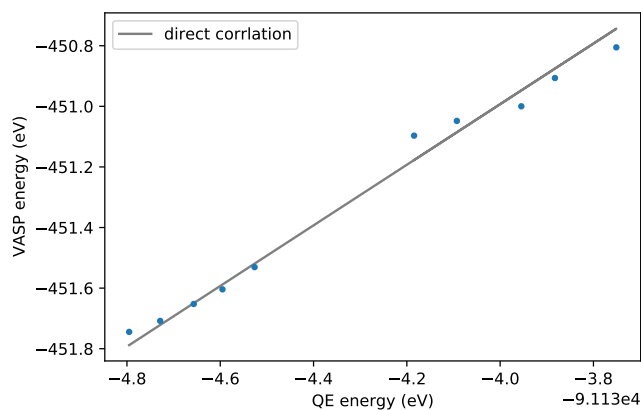
DFT setting	Training set	Approximate Intermediates
SCF convergence criteria	10^{-7}	10^{-4}
plane-wave cut-off	65 Ry	25 Ry
k-point spacing	0.25 \AA^{-1}	gamma

we find minima, that when evaluated with the tight QE settings, the intermediates have forces as large as 1 eV/\AA . The large forces on the atoms imply that the initial configurations are sufficiently far away from the true minima, that the convergence of the adsorption energies is still a meaningful task.

Supplementary Methods

Supplementary Method 1: Vasp vs QE

Throughout the investigation, we use the open-source electronic structure code Quantum Espresso (QE). Previous studies in the literature have used VASP. Here we compare the total energies of the path illustrated in Figure 4 of the main text ($\text{HCOO} + \text{H} \rightarrow \text{H}_2\text{COO}$) between the two DFT methods. The QE settings can be found in the main text, while for the VASP calculations, we use the same settings as Ref [87].



Quantum Espresso total energies vs VASP total energies. A DFT comparison for the images along the path illustrated in Figure 4 of the main text ($\text{HCOO} + \text{H} \rightarrow \text{H}_2\text{COO}$).

Supplementary Method 2: GAP Parameter list

MLFF descriptor settings

Throughout the investigation, we use the Gaussian Approximation Potential (GAP) [2, 53] in combination with a double turbo SOAPs [150] as our many body descriptor, with 4 Å and 6 Å cut-offs. The exact parameters are found below.

```

l_max: 4,
n_species: 3,
rcut_hard: 4.0/6.0,
rcut_soft: 3.5/5.5,
basis: poly3,
scaling_mode: polynomial,
radial_enhancement: 1,
covariance_type: dot_product,
add_species: F,
delta: 0.1,
n_sparse: 2500,
sparse_method: cur_points,
zeta: 4,
alpha_max: 8,
atom_sigma_r: 0.3,
atom_sigma_t: 0.3,
atom_sigma_r_scaling: 0.0,

```

```
atom_sigma_t_scaling: 0.0,  
amplitude_scaling: 1.0,  
central_weight: 1.0  
  
sparse_jitter: 1e-8,  
default_sigma: {0.001 0.01 0.01 0.0},  
force_parameter_name: qe4_forces,  
virial_parameter_name: qe4_virial,  
energy_parameter_name: qe4_fenergy
```

Filtering unique CO₂ adsorbate structures

To select unique carbon dioxide adsorbate configurations, we compare the atomic environment of the carbon atom. Here we use a standard soap vector:

```
soap add_species=T l_max=4 n_max=8 atom_sigma=0.4 cutoff=6.0  
cutoff_transition_width=0.5 central_weight=1.0
```

Supplementary Methods 3: MACE Parameter list

For the dopant task, the following settings were used with the MACE [257] framework.

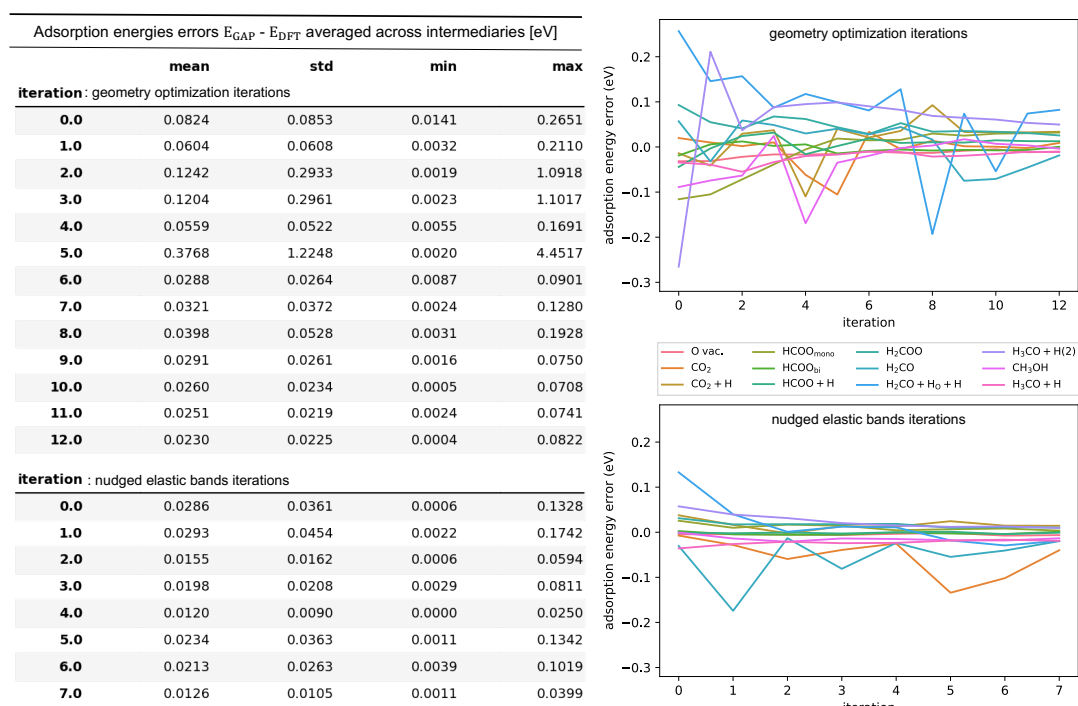
```
device: cuda  
valid_fraction: 0.1  
energy_weight: 1.0  
forces_weight: 10.0  
start_swa: 1200  
swa_energy_weight: 1000.0  
swa_forces_weight: 1.0  
swa_lr: 1e-3  
lr: 1e-2  
E0s: '{1: -12.54, 6: -241.86, 8: -560.11,  
49: -1968.24, 78: -2860.81}'  
model: "MACE"  
interaction_first: "RealAgnosticResidualInteractionBlock"  
interaction: "RealAgnosticResidualInteractionBlock"  
num_interactions: 2  
max_ell: 3
```

```


hidden_irreps: '256x0e + 256x1o'
num_cutoff_basis: 5
correlation: 3
r_max: 5.0
scaling: 'rms_forces_scaling'
batch_size: 8
max_num_epochs: 2000
patience: 800
weight_decay: 5e-7
ema_decay: 0.99

```

Supplementary Figures




Supplementary Figure 1: Adsorption energy error across iterations. For validating our method against the Ref [87] path, we show how the GAP adsorption energies converge to the true DFT ones. Above we see the adsorption energy error across geometry optimization and nudged elastic band iterations. The tables show the mean, standard deviation, minimum, and maximum energy errors across all 12 intermediates for each iteration. The graphs give more detailed insight into each individual intermediate.




Supplementary Figure 2: GAP minima characterisation. At each iteration the minima found by GAP are evaluated using single-point DFT. The DFT force quantifies how close the GAP minima are to true DFT minima. After the 11th iteration, all minima are below the 0.2 eV/Å cutoff, and hence the termination criteria for the geometry optimization active-learning block is reached.


A.0.1 Visualisation of paths




Supplementary Figure 3: Ref [87] path for $\text{HCOO} + \text{H} \rightarrow \text{H}_2\text{COO}$ Note that this is simply a trajectory, that follows a similar path to Ref [87]. It is not the exact images of the Ref [87] NEB, which include four images. Here we show all 10 images of the NEBs found during this investigation. Note on all visualisations: the surface atoms are scaled down to 70% of their default size and connected by bonds. The molecular atoms are kept at 100% size. The difference in the oxygen coordination of the two indium atoms of this oxygen vacancy can be easily distinguished in this view. Note that in the paper, surface atoms are colored in grey for clarity. In the main text, Indium and Oxygen atoms can be distinguished by their size.




Supplementary Figure 4: Optimised path for $\text{HCOO} + \text{H} \rightarrow \text{H}_2\text{COO}$ This is the newly found minimum energy path for the reaction step that was thought to be rate limiting.



Supplementary Figure 5: Ref [87] path for $\text{H}_2\text{COO} \rightarrow \text{H}_2\text{CO} + \text{O}$ Note that this is simply a trajectory, that follows a similar path to Ref [87] It is not the exact images of the Ref [87] NEB, which included four images. Here we show all 10 images of the NEBs found during this investigation.



Supplementary Figure 6: Thermodynamically relevant minimum energy path for $\text{H}_2\text{COO} \rightarrow \text{H}_2\text{CO} + \text{O}$. This is the newly found minimum energy path for the third reaction step.



Supplementary Figure 7: Visualisation of the top-layer reduced reaction pathway

