

Accelerating Molecular Simulation Using Machine Learning: From Wave Functions to Thermodynamics

Sander Vandenhaute

Doctoral dissertation submitted to obtain the academic degree of
Doctor of Engineering Physics

Supervisor

Prof. Veronique Van Speybroeck, PhD
Department of Applied Physics
Faculty of Engineering and Architecture, Ghent University

January 2025



ISBN 978-94-6355-947-8

NUR 928

Wettelijk depot: D/2025/10.500/7

Members of the Examination Board

Chair

Prof. Hennie De Schepper, PhD, Ghent University

Other members entitled to vote

Prof. Jutho Haegeman, PhD, Ghent University

Prof. Venkat Kapil, PhD, University College London, United Kingdom

Prof. Alexandre Tkatchenko, PhD, Université du Luxembourg, Luxembourg

Prof. Louis Vanduyfhuys, PhD, Ghent University

Supervisor

Prof. Veronique Van Speybroeck, PhD, Ghent University

JM: *“If this was going to be viewed forever by young high school kids and college kids — young entrepreneurs who want to go out and do something while they’re still young. [...] What advice would you give them?”*

SJ: *“The thing I would say is, when you grow up, you tend to get told that the world is the way it is, and your life is just to live your life inside the world. Try not to bash into the walls too much. Try to have a nice family life, have fun, save a little money. But that’s a very limited life. Life can be much broader once you discover one simple fact, and that is: everything around you that you call life was made up by people that were no smarter than you. And you can change it. You can influence it. You can build your own things that other people can use. And the minute that you understand that you can poke life, and actually something will, you know, if you push in, something will pop out the other side, that you can change it. You can mold it. That’s maybe the most important thing, to shake off this erroneous notion that life is there and you’re just going to live in it, versus embrace it. Change it. Improve it. Make your mark upon it. I think that’s very important. And however you learn that, once you learn it, you’ll want to change life and make it better, because it’s kind of messed up in a lot of ways. Once you learn that, you’ll never be the same again”*

Transcript from an interview between John McLaughlin from the Santa Clara Valley Historical Association, and Steve Jobs (November 11, 1994).

Table of Contents

Acknowledgements	i
Summary	iii
Samenvatting	vii
List of Symbols	xiii
List of Abbreviations	15
Part I: Accelerating Molecular Simulation using Machine Learning: From Wave Functions to Thermodynamics	
1. Introduction	3
Electrons and Nuclei	3
Quantum mechanics	4
Statistical mechanics	9
Molecular simulation	12
Goal and outline	14
2. Computational Challenges	17

The interaction energy	18
The time scale problem	23
Classical interaction potentials	26
Summary	30
3. Learning the Interaction Energy	33
Introduction	34
The manifold hypothesis and neural networks	35
Graph neural networks	36
On-the-fly learning	44
Case Study: Size Effects in Phase Transitions	49
Transfer learning	50
Conclusions	55
4. Learned Rare Event Sampling	57
Learned chemical representations	57
Basins of attraction	60
Architecture & training	62
Conclusion	64
5. psiflow: a Modular and Scalable Simulation Framework	67
Overview	68
Asynchronous execution	68
Hamiltonians	71
Sampling	72
Free energy calculations	73
Conclusions	74
6. Conclusions and Perspectives	75

Part II: Peer-Reviewed Publications, Preprints, and Conference Contributions

A. Selected Publications and Preprints	83
Paper I	84
Paper II	109
Paper III	118
Paper IV	171
B. Appendices	185
Empirical Scaling of Quantum Chemistry Methods	185
Extended Hessian calculations	187
C. List of Publications and Conference Contributions	191
Publications and Preprints	191
Presentations at Conferences & Workshops	192
Bibliography	195

Acknowledgements

This work stands on the shoulders of many: my family, my advisor, the jury members, colleagues, friends, but also external collaborators, other researchers in the community, contributors to open-source software packages, funding agencies, and even my high school mathematics teacher; all of them have played a significant role in one way or another. Instead of a complete enumeration, I will highlight a few of the people whose contribution to this work has been most important.

First, I want to thank my advisor, Veronique. From day one, you have given me *carte blanche* and encouraged me to pursue my own research ideas; you created precisely the kind of environment I needed to perform at my very best. Beyond your academic guidance, you have always been very patient towards my somewhat chaotic work habits, and only now do I fully appreciate that. I hope that my work will, in some way, prove useful to your future research purposes.

Second, I want to thank my jury members. I am really happy with the rigorous discussion we had during my internal defense. Honestly, I wish I could have convened this committee of researchers regularly throughout my PhD as to benefit even more from your expertise and critical review. Thank you for your careful reading of an earlier version of this thesis; I have tried to incorporate your helpful comments and suggestions to the best of my ability.

I used to believe that the universe was made from elementary particles which follow the laws of physics. But then I look at my parents, and I realize there must be something more. Thank you so, so much.

– Sander

Summary

Historically, the most significant technological breakthroughs were often preceded by the discovery of a particular material or molecule. For example, *semiconductor* materials induced the transformation from physically enormous mechanical computers in the 1950s into today's powerful handheld devices, and the *penicillin* molecule revolutionized health care as the first effective treatment against deadly bacteria. Today's society is faced with monumental challenges, and we will likely require specialized materials and molecules to tackle them. Fortunately, the properties of materials and molecules can be understood based on the laws of quantum and statistical mechanics, and it is therefore theoretically possible to design them for specific applications. Unfortunately, their design space is incredibly large – with approximately 94 naturally occurring elements that can be combined in virtually infinite ways – and the equations for predicting their properties are far too complicated to be solved by hand.

Computer simulations may provide a viable alternative. Computing power has been growing almost exponentially since the 1970s, and today's exascale systems may provide us with enough capacity to predict the physical and chemical properties of many molecules and materials. As such, it would be possible to perform a *screening* across all possible combinations of atomic arrangements and find, for example, new materials with superconductive behavior at ambient conditions, or new molecules which bind to specific target proteins. One of the key challenges in this approach is the enormous computational cost that is associated with explicit simulation of the equations of quantum and statistical mechanics. Over the past few decades, researchers have developed a large variety of approximate methods which can be used to alleviate a large portion of the total cost. Despite significant advances in approximation techniques, current methods present an unsatisfactory trade-off between accuracy and computational efficiency: methods that achieve

chemical accuracy remain prohibitively expensive for large-scale screening applications, while faster approximations introduce errors that render their predictions unreliable. Even with projected increases in computing power over the next decade, this fundamental limitation would persist.

Machine learning has recently established itself as an extraordinary tool for scientific discovery; the 2024 Nobel prizes in both physics and chemistry were awarded to *machine learning* scientists. For molecular simulation in particular, machine learning has the ability to improve the accuracy-cost tradeoff by orders of magnitude. However, its application to complex molecular systems remains very challenging. As such, the main goal of this work is to develop scalable machine learning algorithms to advance and accelerate computational research into chemical and physical processes in nanostructured materials, with a particular focus on metal-organic frameworks and zeolites because of their extraordinary properties in the context of gas storage and catalysis applications.

To understand how machine learning can improve the accuracy-cost tradeoff, we must first examine what makes molecular simulations computationally intensive. On an abstract level, dynamic simulations can be viewed as a long, repeating sequence of steps. Each step represents a single application of Newton's equations of motion: we evaluate the interatomic forces with quantum mechanics and use them to advance the atomic positions over a very small time interval. By repeating this over and over, we simulate the dynamics of individual atoms and are able to predict their physical properties. The computational challenges in such simulations are therefore twofold: how do we reduce the cost of a single such step (i.e. the quantum mechanical energy and force evaluation), and how do we reduce the total number of steps (i.e. the total simulation time). The main contribution of the research in this thesis is the development of new machine learning techniques which advance the state of the art for both these challenges.

For the first challenge, we build upon a large body of prior research on the application of machine learning to materials and molecules. A machine learning model considers a material or molecule as a three-dimensional arrangements of atoms (with or without periodicity), and its task is to predict the quantum mechanical interaction energy and the interatomic forces (as required for dynamic simulations). Such models have demonstrated excellent

performance on benchmark datasets, but their application to nanoporous materials had not yet been explored. The key difficulty relates to the generation of high-quality training data which contain the target energy and forces that the model needs to learn to predict. On one hand, the generation of such training data can be very challenging for metal-organic frameworks and zeolites due to their chemical variety and potential metastable behavior. On the other hand, the most accurate quantum mechanical (wave-function-based) methods can sometimes only compute the interaction energy and not the interatomic forces even though they are required during training. We tackle these problems using a combination of on-the-fly learning for data generation and transfer learning to bypass the need for high fidelity interatomic forces.

For the second challenge – reducing the total simulation time – we are inspired by existing *rare event* sampling algorithms which rely on system-specific knowledge to accelerate the effective dynamics such that slow or rare events happen much more quickly or frequently. The main bottleneck in these approaches is the requirement of a system-specific reaction coordinate whose definition usually involves specialized knowledge and at least some level of human trial and error. We developed a universal machine learning approach which augments the energy prediction models with knowledge regarding the rare event in a way that allows us to generate reaction coordinates for chemical and physical transformations in a generic manner. Moreover, by careful consideration of the required symmetries, we demonstrate state of the art performance on a variety of applications.

We conclude this summary with a brief outline of the thesis. Chapters 1 & 2 introduce the fundamentals of molecular simulation, and elaborate in more detail on the twofold computational challenge as well as traditional approximations and techniques to overcome them. Chapters 3 & 4 present the main research results, with Chapter 3 focusing on the first challenge using on-the-fly learning and transfer learning, and Chapter 4 focusing on learned rare event sampling. Chapter 5 integrates the research results into *psiflow*: a molecular simulation framework that facilitates the deployment of these techniques on large amounts of compute resources with minimal human effort. Finally, we formulate our overall conclusions and perspectives for future research in Chapter 6.

Samenvatting

Doorheen de geschiedenis werden de belangrijkste technologische doorbraken vaak voorafgegaan door de ontdekking van een specifiek materiaal of molecule. Bijvoorbeeld, *halfgeleider* materialen leidden tot de transformatie van de gigantische mechanische computers in de jaren 50 naar de krachtige draagbare apparaten van vandaag, en de ontdekking van *penicilline* zorgde voor een revolutie in de geneeskunde als de eerste effectieve behandeling tegen (dodelijke) bacteriën. De huidige samenleving staat voor enorme uitdagingen, en we gaan waarschijnlijk gespecialiseerde materialen en moleculen nodig hebben om deze aan te pakken. Gelukkig kunnen de eigenschappen van materialen en moleculen worden begrepen op basis van de wetten van de kwantum- en statistische fysica, en dat maakt het theoretisch mogelijk om ze te ontwerpen voor specifieke toepassingen. Helaas is het aantal mogelijkheden ongelooflijk groot – er bestaan ongeveer 94 natuurlijk voorkomende chemische elementen die op oneindig veel manieren kunnen worden gecombineerd – en de vergelijkingen voor het voorspellen van hun eigenschappen zijn veel te ingewikkeld om met de hand op te lossen.

Computersimulaties kunnen een haalbaar alternatief bieden. Rekenkracht is sinds de jaren 70 bijna exponentieel gegroeid, en de huidige supercomputers zouden ons in principe genoeg capaciteit kunnen bieden om de fysische en chemische eigenschappen van moleculen en materialen te voorspellen. Aldus moet het mogelijk zijn om een *screening* uit te voeren over alle mogelijke ruimtelijke composities van atomen, en bijvoorbeeld nieuwe materialen te vinden met supergeleidend gedrag bij kamertemperatuur of nieuwe moleculen die zich specifiek binden aan bepaalde eiwitten. Eén van de belangrijkste uitdagingen in deze aanpak is de enorme computationele kost die gepaard gaat met expliciete simulatie van de vergelijkingen van quantum- en statistische mechanica. In de afgelopen decennia hebben onderzoekers een zeer grote verscheidenheid aan benaderende methoden ontwikkeld die

kunnen worden gebruikt om een groot deel van de totale kosten te verlichten. Desondanks bieden de huidige methoden nog niet de juiste balans tussen nauwkeurigheid en computationele efficiëntie: methoden die chemische nauwkeurigheid bereiken blijven te duur voor grootschalige screening-toepassingen, terwijl snellere benaderingen fouten introduceren die hun voorspellingen onbetrouwbaar maken. Zelfs met de verwachte toename in rekenkracht in het komende decennium, zou deze fundamentele beperking blijven bestaan.

Machine learning heeft zich recent gevestigd als een buitengewoon hulpmiddel voor wetenschappelijke vooruitgang; de Nobelprijzen voor zowel natuurkunde als scheikunde werden in 2024 toegekend aan *machine learning* wetenschappers. Voor moleculaire simulatie in het bijzonder kan machine learning simultaan de nauwkeurigheid *en* de computationele efficiëntie met meerdere grootteorden verbeteren. Desalniettemin blijft het toepassen van machine learning methoden op complexe moleculaire systemen bijzonder uitdagend. Het hoofddoel van dit werk situeert zich in deze context, en beoogt om schaalbare machine learning methoden te ontwikkelen die het computationeel onderzoek naar chemische en fysische processen in nanogestructureerde materialen kunnen versnellen en verbeteren, met een specifieke focus op metaal-organische roosters en zeolieten vanwege hun buitengewone eigenschappen in de context van gas adsorptie en katalyse.

Om dit te begrijpen, moeten we eerst onderzoeken wat moleculaire simulaties computationeel intensief maakt. Op een abstract niveau kunnen dynamische simulaties worden gezien als een lange, herhalende reeks stappen. Elke stap vertegenwoordigt een enkele toepassing van de bewegingsvergelijkingen van Newton: we evalueren de interatomaire krachten via kwantummechanische berekeningen en gebruiken deze om de atomaire posities over een zeer klein tijdsinterval te propageren. Door dit steeds te herhalen, simuleren we de dynamica van individuele atomen en kunnen we hun fysische eigenschappen voorspellen. De computationele uitdagingen in dergelijke simulaties zijn daarom tweeledig: hoe verminderen we de kosten van een enkele dergelijke stap (d.w.z. de kwantummechanische evaluatie van de interatomaire krachten), en hoe verminderen we het totale aantal stappen (d.w.z. de totale simulatietijd). De belangrijkste bijdrage van het onderzoek

in dit werk is de ontwikkeling van nieuwe machine learning technieken die de state of the art voor beide uitdagingen vooruit helpen.

Voor de eerste uitdaging bouwen we voort op een groot aantal eerdere onderzoeken naar de toepassing van machine learning op materialen en moleculen. Een machine learning model beschouwt een materiaal of molecule als een driedimensionale rangschikking van atomen (met of zonder periodiciteit), en zijn taak is om de kwantummechanische interactie-energie en de interatomaire krachten te voorspellen (zoals vereist voor dynamische simulaties). Zulke modellen hebben uitstekende prestaties laten zien op benchmark datasets, maar hun toepassing op nanoporeuze materialen was nog niet onderzocht. De belangrijkste moeilijkheid heeft betrekking op het genereren van hoogwaardige trainingsdata die de kwantummechanische energie en krachten bevatten die het model moet leren voorspellen. Aan de ene kant kan het genereren van die trainingsdata zeer uitdagend zijn voor metaal-organische frameworks en zeolieten vanwege hun chemische verscheidenheid en potentieel metastabiel gedrag. Aan de andere kant kunnen de meest accurate kwantummechanische methoden soms enkel de energie berekenen en niet de interatomaire krachten, hoewel deze vereist zijn tijdens de training. We pakken deze problemen aan met een combinatie van on-the-fly learning voor data generatie en transfer learning om de noodzaak voor zeer nauwkeurige interatomaire krachten te omzeilen.

Voor de tweede uitdaging – het verminderen van de totale simulatietijd – laten we ons inspireren door bestaande *rare event* sampling algoritmen die vertrouwen op systeem-specifieke kennis om de effectieve dynamica te versnellen zodat trage of zeldzame gebeurtenissen veel sneller of frequenter plaatsvinden. De belangrijkste bottleneck in deze benaderingen is de vereiste van een systeem-specifieke reactiecoördinaat waarvan de definitie meestal gespecialiseerde kennis vereist en/of een minimum aan menselijke trial and error. We ontwikkelden een universele machine learning aanpak die de energievoorspellingsmodellen aanvult met kennis over de zeldzame gebeurtenis op een manier die ons in staat stelt om reactiecoördinaten voor chemische en fysische transformaties op een generieke manier te berekenen. Bovendien demonstreren we door zorgvuldige overweging van de vereiste symmetrieën dat onze methode het qua prestaties significant beter doet ten opzichte van alternatieven in de literatuur.

We sluiten deze samenvatting af met een kort overzicht van het proefschrift. Hoofdstukken 1 & 2 introduceren de fundamentele van moleculaire simulatie, en gaan meer in detail over de tweeledige computationele uitdaging en traditionele benaderingen en technieken om deze te overwinnen. Hoofdstukken 3 & 4 bespreken de belangrijkste onderzoeksresultaten, waarbij Hoofdstuk 3 zich richt op de eerste uitdaging met behulp van on-the-fly learning en transfer learning, en Hoofdstuk 4 zich richt op het gericht versnellen van de dynamica. Hoofdstuk 5 integreert de onderzoeksresultaten in 'psiflow': een softwarepakket voor moleculaire simulaties dat het toelaat om deze technieken op grote supercomputers toe te passen met minimale menselijke interventie. Ten slotte formuleren we onze algemene conclusies en perspectieven voor toekomstig onderzoek in Hoofdstuk 6.

List of Symbols

\mathbf{x}	three-dimensional position vector ($\mathbf{x} \in \mathbb{R}^3$)
$\psi(\mathbf{x})$	time-independent wave function
\hat{H}	Hamiltonian operator
$\Psi(\mathbf{x}, t)$	time-dependent wave function
\mathbf{r}	cartesian coordinates of a collection of atomic nuclei
\mathbf{r}_e	cartesian coordinates of a collection of electrons
\mathbf{r}_i	cartesian coordinates of atom i
\mathbf{r}_{ij}	distance vector $\mathbf{r}_i - \mathbf{r}_j$
$E(\mathbf{r})$	ground-state electronic energy at fixed nuclear coordinates \mathbf{r}
\hbar	reduced Planck's constant
k	Boltzmann constant
T	temperature (in thermodynamic equilibrium)
β	inverse temperature, defined as $(kT)^{-1}$
P	pressure (in thermodynamic equilibrium)
$p(\mathbf{r})$	Boltzmann distribution in various ensembles
F	Helmholtz free energy
Δt	discretization step in dynamic simulations
cp	closed-pore phase
lp	closed-pore phase
\mathbf{h}_i	node embedding across layers in a graph neural network
E_i	atomic energy of atom i
\mathbf{q}	log probability vector; the number of components is equal to the number of phases which need to be discriminated against

List of Abbreviations

QM	quantum mechanical
CPU	central processing unit
GPU	graphics processing unit
Flop	floating point operation
HF	Hartree-Fock
DFT	density functional theory
MP	Møller-Plesset perturbation theory
RPA	random phase approximation
CCSD(T)	coupled cluster with singles, doubles, and perturbative triples excitation
MD	molecular dynamics
PBE	Perdew-Berke-Ernzerhof functional
GGA	generalized gradient approximation
D3(BJ)	Becke-Jonhson dispersion correction
D3(TS)	Tkatchenko-Scheffler dispersion correction
DLPNO	domain local pair natural orbital localization method
COSX	chain of spheres approximation
RI	resolution of identity approximation
PBCs	periodic boundary conditions
MOF	metal-organic framework
ML	machine learning
GNN	graph neural networks
MLP	multilayer perceptron
SBC	smooth basin classification
SOTA	state of the art

Part I

Accelerating Molecular Simulation using
Machine Learning: From Wave Functions
to Thermodynamics

1

Introduction

Materials and molecules, and collections of atoms in general, consist of electrons and nuclei that interact electrostatically with each other. Simulations of such systems may at first sight seem similar to large cosmological simulations with stars and planets which interact through gravity instead of electrostatics. The crucial distinction is that electrons and nuclei are both incredibly small (less than one nanometer in size) and many orders of magnitude lighter than stars and planets. These fundamental differences lead to entirely different laws governing their behavior. In this Chapter, we will briefly revise the basic physical concepts that are necessary to simulate matter at the atomic scale, and conclude with an outlook on why computer simulations at this scale are so challenging. The goal is to give the reader an intuitive understanding of how the laws of physics can be converted into practical simulation algorithms. For a more rigorous and complete treatment of the material in this Chapter, we refer the reader to specialized textbooks [1–3].

1.1 Electrons and Nuclei

The vast majority of empirical observations ever recorded by humankind can be understood based on two key premises. First, materials and molecules are built from *electrons* and *nuclei*. Both the electron and the nucleus possess a few intrinsic properties such as *charge*, *mass*, and *spin*. In particular, nuclei are relatively heavy particles which possess a variable but positive charge. Electrons are more than thousand times lighter than even the lightest nucleus, and they carry a fixed, negative charge. Interestingly, electrons and

nuclei can form groups with well-defined compositions; these are referred to as *atoms*. The periodic table summarizes the frequently observed atoms (Figure 1.1).

H ¹																	He ²
Li ³	Be ⁴											B ⁵	C ⁶	N ⁷	O ⁸	F ⁹	Ne ¹⁰
Na ¹¹	Mg ¹²											Al ¹³	Si ¹⁴	P ¹⁵	S ¹⁶	Cl ¹⁷	Ar ¹⁸
K ¹⁹	Ca ²⁰	Sc ²¹	Ti ²²	V ²³	Cr ²⁴	Mn ²⁵	Fe ²⁶	Co ²⁷	Ni ²⁸	Cu ²⁹	Zn ³⁰	Ga ³¹	Ge ³²	As ³³	Se ³⁴	Br ³⁵	Kr ³⁶
Rb ³⁷	Sr ³⁸	Y ³⁹	Zr ⁴⁰	Nb ⁴¹	Mo ⁴²	Tc ⁴³	Ru ⁴⁴	Rh ⁴⁵	Pd ⁴⁶	Ag ⁴⁷	Cd ⁴⁸	In ⁴⁹	Sn ⁵⁰	Sb ⁵¹	Te ⁵²	I ⁵³	Xe ⁵⁴
Cs ⁵⁵	Ba ⁵⁶	Lu ⁷¹	Hf ⁷²	Ta ⁷³	W ⁷⁴	Re ⁷⁵	Os ⁷⁶	Ir ⁷⁷	Pt ⁷⁸	Au ⁷⁹	Hg ⁸⁰	Tl ⁸¹	Pb ⁸²	Bi ⁸³	Po ⁸⁴	At ⁸⁵	Rn ⁸⁶
Fr ⁸⁷	Ra ⁸⁸	Lr ¹⁰³	Rf ¹⁰⁴	Db ¹⁰⁵	Sg ¹⁰⁶	Bh ¹⁰⁷	Hs ¹⁰⁸	Mt ¹⁰⁹	Ds ¹¹⁰	Rg ¹¹¹	Cn ¹¹²	Nh ¹¹³	Fl ¹¹⁴	Mc ¹¹⁵	Lv ¹¹⁶	Ts ¹¹⁷	Og ¹¹⁸
		La ⁵⁷	Ce ⁵⁸	Pr ⁵⁹	Nd ⁶⁰	Pm ⁶¹	Sm ⁶²	Eu ⁶³	Gd ⁶⁴	Tb ⁶⁵	Dy ⁶⁶	Ho ⁶⁷	Er ⁶⁸	Tm ⁶⁹	Yb ⁷⁰		
		Ac ⁸⁹	Th ⁹⁰	Pa ⁹¹	U ⁹²	Np ⁹³	Pu ⁹⁴	Am ⁹⁵	Cm ⁹⁶	Bk ⁹⁷	Cf ⁹⁸	Es ⁹⁹	Fm ¹⁰⁰	Md ¹⁰¹	No ¹⁰²		

Figure 1.1: The periodic table of elements arranges all chemical species sequentially based on the number of electrons. It was at least partially developed by Dmitri Ivanovich Mendeleev (1834 - 1907).

The second premise is that the behavior of electrons and nuclei through time is not just completely random, but it obeys certain physical laws. The ones that are relevant in this thesis are the laws of *quantum mechanics* [1] and *statistical mechanics* [2].

1.2 Quantum mechanics

Quantum mechanics dictates how electrons and nuclei interact with each other, and how those interactions determine their behavior through time. The cornerstone is that electrons and nuclei cannot be considered as infinitely small, point-like particles with a well-defined position and velocity. Instead, quantum mechanics characterizes particles using their *wave function*. For all practical purposes, and for the remainder of this thesis, the wave function of a single particle is a scalar, dimensionless, and complex-valued function which maps any (three-dimensional) spatial coordinate $\boldsymbol{x} \in \mathbb{R}^3$ into a complex number $\psi(\boldsymbol{x}) \in \mathbb{C}$. The physical interpretation associated with ψ is that it determines a probability density of observing the particle at any given

location in space. More specific, the probability of observing the particle at location \mathbf{x} in space is proportional¹ to the modulus squared of the wave function, i.e. $|\psi(\mathbf{x})|^2$. In particular, this implies that electrons (and to a lesser extent nuclei) can no longer be localized into a particular point in space.

The time behavior of a particle is formalized by using a time-dependent wave function, which is by convention denoted with a capital greek letter: $\Psi(\mathbf{x}, t)$. The well-known *Schrödinger* equation establishes a relationship between the time evolution of the wave function $\Psi(\mathbf{x}, t)$ and the particle's *hamiltonian* operator \hat{H} (the mathematical entity which encodes the kinetic and potential energy of the particle) [1]:

$$i\hbar \frac{\partial \Psi}{\partial t}(\mathbf{x}, t) = \hat{H}\Psi(\mathbf{x}, t) \quad (1)$$

In theory, we can solve Equation 1 to calculate the wave function $\Psi(\mathbf{x}, t)$ of the particle as a function of space and time, and compute the probability distribution $|\Psi(\mathbf{x}, t)|^2$ of actually observing the particle at any given position \mathbf{x} and time t .

As Figure 1.1 demonstrates, electrons and nuclei agglomerate into atoms, and atoms are known to further organize themselves into molecules and materials. Therefore, if we aim to predict the time behavior of molecules and materials, we should be able to extend Equation 1 such that it can deal with multiple electrons and multiple nuclei. Unsurprisingly, the wave function for a system of N particles depends on $3N$ spatial coordinates $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$ as well as time: $\Psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N, t)$. The hamiltonian \hat{H} of the system now contains the kinetic energy of all N particles as well as $O(N^2)$ interaction terms because of the charge each particle carries. While these electrostatic interactions between the positively charged nuclei and negatively charged electrons are essential to the stability of the materials and molecules around us, they also introduce a lot of complexity into Equation 1. In fact, Equation 1 is not analytically solvable for any system involving more than two particles. With the exception of an isolated hydrogen atom – which is a two-particle

¹The typical wording in physics textbooks used in this context is that $|\psi(\mathbf{x})|^2 d\mathbf{x}$ encodes the probability of observing the particle at an infinitesimal volume $[\mathbf{x}, \mathbf{x} + d\mathbf{x}]$. Nevertheless, $|\psi(\mathbf{x})|^2$ is formally a continuous probability density function, and its rigorous mathematical definition involves measure theory, which is rather complicated and fully outside the scope of this thesis.

system – it is necessary to develop and apply systematic approximations in order to calculate the wave function of the system.

The physical *ansatz* for the most important approximations in this thesis start from the observation that electrons and nuclei are not just somewhat similar but oppositely charged particles; their masses differ by over three orders of magnitude! Inspired by this, Max Born and Robert Oppenheimer demonstrated that this large difference in mass induces a time scale separation in Equation 1 [4]. The most important consequence is that the many-particle wave function Ψ can actually be factorized into two contributions: one coming from the lightweight, fast-moving electrons, and one from the heavy, slowly-moving nuclei. Before we proceed, let us redefine our notation of the spatial coordinates of the particles to emphasize the difference between electronic and nuclear coordinates²:

$$\begin{aligned} (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) &= (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_m, \mathbf{r}_{e_1}, \mathbf{r}_{e_2}, \dots, \mathbf{r}_{e_n}) \\ &= (\mathbf{r}, \mathbf{r}_e) \end{aligned} \quad (2)$$

for a system with m nuclei and n electrons (and naturally $m + n = N$). With this notation, we can express³ the so-called Born-Oppenheimer approximation in the context of this thesis as follows: the full N -particle wave function $\Psi(\mathbf{r}, \mathbf{r}_e, t)$ can be written as a product of a nuclear wave function χ which depends on the spatial coordinates of the nuclei \mathbf{r} and on the time t , and an electronic wave function ψ_e which depends explicitly on the electron coordinates \mathbf{r}_e and implicitly on the nuclear coordinates \mathbf{r} [5,6]:

$$\Psi(\mathbf{r}, \mathbf{r}_e, t) = \psi_e(\mathbf{r}_e | \mathbf{r})\chi(\mathbf{r}, t) \quad (3)$$

The dynamic behavior of a system in the Born-Oppenheimer approximation is therefore fully contained in the (time-dependent) nuclear wave function χ . From the perspective of the electrons, the nuclei represent positive point

²In theory, the wave function Ψ also exhibits a dependence on the intrinsic spin variables σ_i of each particle. In addition, because electrons are spin- $\frac{1}{2}$ particles, the many-body wave function Ψ should satisfy the *Pauli exclusion principle*, which states that two electrons can never occupy the same quantum state – or alternatively that the wave function Ψ is required to switch sign whenever two electrons are interchanged. For the qualitative discussion in the present Chapter, we choose to simplify the notation and omit the spin dependence of the wave function.

³We implicitly assume that the electronic eigenvalues in Equation 4 are sufficiently separated, and that the electrons occupy their ground-state energy level. This will always be the case in this thesis.

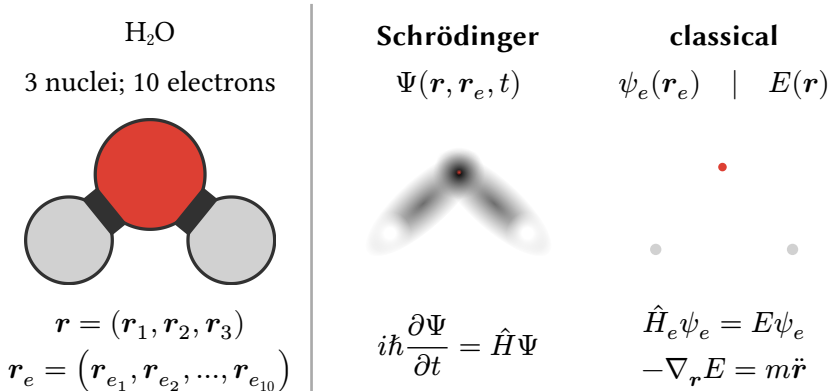


Figure 1.2: Illustration of the wave function formalism in quantum mechanics and the classical approximation that is assumed for the remainder of this thesis, for a single water molecule. Pure quantum mechanics considers all particles as wave-like, with their time evolution given by Equation 1. In the classical picture, the nuclei are considered as Newtonian point-like particles which are subject to effective quantum mechanical forces determined by the electronic wave function ψ_e .

charges which are fixed or “clamped” in space at positions \mathbf{r} . The electrostatic interaction with the clamped nuclei is what causes the electronic wave function to depend parametrically on the nuclear positions \mathbf{r} . The electronic Hamiltonian \hat{H}_e encodes the kinetic energy of all electrons as well as the electron-electron and electron-nucleus interactions, and it defines the eigenvalue equation which is used to calculate⁴ ψ_e [7]:

$$\hat{H}_e \psi_e(\mathbf{r}_e \mid \mathbf{r}) = E(\mathbf{r}) \psi_e(\mathbf{r}_e \mid \mathbf{r}) \quad (4)$$

The eigenvalue $E(\mathbf{r})$ represents the electronic energy, and it naturally depends on the nuclear coordinates as well. Chapter 2 will go into more detail regarding the practical solution of Equation 4.

Aside from the separation of the electronic and nuclear degrees of freedom, there is a second important approximation that can be made here. In general, experiments have shown that the degree to which particles exhibit quantum mechanical behavior decreases with the mass of a particle.⁵ Fortu-

⁴In many cases, there are infinitely many solutions, each of which has a particular value for the eigenvalue E . In the context of the Born-Oppenheimer approximation and this thesis, the relevant solution is the *ground state*, i.e. the eigenstate with the lowest value for the energy E

⁵Fortunately, the theory of quantum mechanics supports these empirical observations; Equation 1 and the associated wave function description approaches the theory of classical

nately, it has been shown that for a large variety of physical properties and for all but the lightest nuclei, we may actually treat the dynamic behavior of nuclei in a classical manner, i.e. using the familiar equations of classical mechanics⁶. In this approach, the electrons are still treated quantum mechanically based on their time-independent wave function ψ_e as it arises from the Born-Oppenheimer approximation, but the nuclei are assumed to behave as classical point-like particles with positions \mathbf{r} , velocities $\dot{\mathbf{r}}$, and accelerations $\ddot{\mathbf{r}}$. The movement of the nuclei is determined by the quantum mechanical interaction energy $E(\mathbf{r})$ as defined by the electronic eigenvalue equation (Equation 4). Effectively, $E(\mathbf{r})$ represents a so-called *Born-Oppenheimer* surface, which functions as a potential energy surface that defines the time evolution of the nuclei according to classical mechanics.

Figure 1.2 summarizes our discussion thus far based on the most common substance in the universe: water. It consists of three nuclei (with charges⁷ +8, +1, and +1) as well as 10 electrons (all of which have charge -1), or in terms of atoms: two hydrogens and one oxygen. According to Equation 1, the dynamics of this system requires knowledge of the full many-particle wavefunction $\Psi(\mathbf{r}, \mathbf{r}_e, t)$ as a function of the coordinates of all nuclei $\mathbf{r} = (\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3)$ (for oxygen and the two hydrogens), the electrons $\mathbf{r}_e = (\mathbf{r}_{e_1}, \mathbf{r}_{e_2}, \dots, \mathbf{r}_{e_{10}})$, and time – which sums up to $13 \times 3 + 1 = 40$ dimensions. Combining the Born-Oppenheimer time scale separation in Equation 3 and a classical approximation of the atomic nuclei, the dynamic behavior of the entire system can be simplified into an equivalent *classical* system in which the atomic nuclei are subject to an effective interaction potential $E(\mathbf{r})$ which contains the quantum mechanical effects of the electrons [4]. As stated in Equation 4, this interaction potential $E(\mathbf{r})$ is defined as the smallest eigenvalue of Equation 4, and it is a smooth function of the nuclear coordinates \mathbf{r} . In particular, the gradient of the interaction energy defines a set of forces on the atomic nuclei, which can be inserted into Newton’s second law in order to obtain the resulting acceleration:

mechanics in the limit of large masses – this is an informal description of what is known as the Ehrenfest theorem [1].

⁶For systems involving light nuclei such as hydrogen, the delocalized nature of the nuclei can in fact influence their physical properties. This influence becomes more pronounced at low temperatures.

⁷In units of e , the elementary unit of charge; $e = 1.602176634 \cdot 10^{-19}$ Coulomb

$$-\nabla_r E = m\ddot{\mathbf{r}} \quad (5)$$

in which $\ddot{\mathbf{r}}$ represents the acceleration as the second derivative of the nuclei positions with respect to time. In this sense, the time behavior of a system of electrons and nuclei is not terribly different from, say, a system of stars in the galaxy.

While the Born-Oppenheimer picture in Figure 1.2 is a useful and intuitive point of view, it is by no means sufficient to help us understand why ice melts and why steel becomes more brittle with increasing fractions of carbon to iron. Typical experiments in physics and chemistry do not deal with just one atom, but with an enormously large number of atoms – on the order of Avogadro’s number, i.e. 10^{23} . Brute force solution of either Equation 1 or the Born-Oppenheimer equivalent is and always will be completely intractable due to the inconceivably large number of atoms in any experimental setup. Even though the microscopic laws which govern the behavior of electrons and nuclei are now known, it is still completely unclear how we can ever apply them to understand empirical observations at the macroscopic level.

Let us first reflect on how typical scientific experiments are performed [8]. Samples are prepared according to a well-defined procedure; its environment is carefully controlled; and the property of interest is measured over some finite interval of time. Typically, the experiment is repeated multiple times, and the final value is taken as the mean of all individual measurements, with a standard error defined by the variation between each measurement. Even if we could somehow predict the behavior of 10^{23} particles by solving Equation 4 and mapping each of their trajectories in space and time, how would we compare this theoretical prediction with the outcome of such an experiment? There is no experimental technique which allows us to measure the simultaneous position and velocity of every atom, so the outcome of our incredibly expensive simulations can not even be validated even if we could perform them.

This is where *statistical mechanics* comes into play [2,8].

1.3 Statistical mechanics

Statistical mechanics and thermodynamics are based on real experiments. As mentioned before, real experiments require the preparation of some sample,




Figure 1.3: Schematic illustration of a statistical mechanical system and its surroundings. Systems can optionally interact with their surroundings by exchanging energy, volume, and/or particles.

immerse it in a controlled environment, and measure some property of interest over a sufficiently long time interval. In an analogous but more abstract fashion, statistical mechanics partitions the universe into a *system* and its *surroundings* [2]. In line with a classical approximation of the dynamics of atomic nuclei, the state of the system is defined by the positions and velocities of the nuclei at any given instant. In line with real experiments, the state of the surroundings is only known in a macroscopic sense, i.e. based on macroscopic variables such as temperature or pressure. The system and its surroundings may interact with each other; the surroundings might transfer heat into the system, or impose a mechanical force (and vice-versa). In this setup, neither quantum nor classical mechanics can help us in deducing properties regarding the system, because the precise behavior of the atoms in the surroundings is not really known. The goal of statistical mechanics is to determine the *statistical* behavior of the atoms in the system as a function of the macroscopic variables which characterize the environment. The adjective “statistical” refers to the fact that we do not aim to predict the precise movement of each and every atom. Instead, we aim to understand and predict the *probability* of observing the system in a particular *configuration*, which is defined by the set of atomic positions \mathbf{r} and velocities \mathbf{v} ⁸. Formally, this distribution is characterized by a probability density $p(\mathbf{r}, \mathbf{v})$, the specifics of which depend on how the system and the environment interact with each other as in Figure 1.3. Let us briefly discuss three important types of ensembles:

⁸We employ \mathbf{v} and no longer $\dot{\mathbf{r}}$ because in statistical mechanics, positions and velocities are treated as independent variables

- **isolated:** the system is entirely isolated from its environment. With respect to Figure 1.3, we can say that there is no exchange of atoms, no exchange of energy, and no exchange of system volume between the system and its environment. In that case, the number of atoms \mathcal{N} , the total volume \mathcal{V} , and the total (potential and kinetic) energy \mathcal{E} of the system remain constant. Because there is no interaction between the system and its environment, we can use the familiar equations of classical mechanics to fully describe its time behavior. According to Liouville's theorem, for such a system, every state (\mathbf{r}, \mathbf{v}) which is energetically accessible is equally likely [9]:

$$p(\mathbf{r}, \mathbf{v}) \propto \delta(\mathcal{E} - E(\mathbf{r}) - K(\mathbf{v})) \quad (6)$$

with $E(\mathbf{r})$ the quantum mechanical interaction energy as resulting from Equation 4, $K(\mathbf{v})$ the kinetic energy of all atoms, and δ the Dirac distribution. This ensemble is commonly referred to as the microcanonical ensemble, and can be denoted as (N, V, \mathcal{E}) .

- **thermal equilibrium:** the most common scenario; the system is in thermal equilibrium with its environment, at some temperature T . This implies that there is energy flowing between the system and its environment. In this scenario, it is possible to simplify our discussion and consider the probability distribution as a function of the positions \mathbf{r} only, not the velocities⁹ \mathbf{v} :

$$p(\mathbf{r}) \propto e^{-\beta E(\mathbf{r})} \quad (7)$$

with $\beta = (kT)^{-1}$ and k the familiar Boltzmann's constant. The interpretation of Equation 7 is that states \mathbf{r} with lower energy are exponentially more likely than states with higher energy, but no state is fundamentally inaccessible. This ensemble is commonly denoted by (N, V, T) .

- **thermal and mechanical equilibrium:** a common scenario in materials science; the system is in thermal and mechanical equilibrium with its environment, as characterized by a temperature T and a (hydrostatic) pressure P . As for the isothermal case, we can ignore the velocity dependence and simply consider the configurational probability distribution:

$$p(\mathbf{r}) \propto e^{-\beta E(\mathbf{r})} e^{-\beta PV} \quad (8)$$

⁹In this case, the velocity dependence can be shown to result in an overall factor which is otherwise irrelevant to the discussion in the present Chapter.

This ensemble is commonly denoted by (N, P, T) .

For rigorous proofs of these expressions, we refer to reader to e.g. Ref [3]. The research in this thesis only ever employs the constant temperature ensembles (Equation 7 and Equation 8), and we use a single notation $p(\mathbf{r})$ to denote the (configurational part of the) ensemble probability density in both cases. Note that the probability density always contains a dependence on the interaction potential $E(\mathbf{r})$ which, in the Born-Oppenheimer approximation, represents the electronic energy as computed from Equation 4.

The foundation of statistical mechanics is that – similar to real experiments which perform n measurements of a property of interest and compute its average – the value of a particular observable A is defined as an average of its local value $A(\mathbf{r})$ over the ensemble probability density:

$$\langle A \rangle = \int p(\mathbf{r})A(\mathbf{r})d\mathbf{r} \quad (9)$$

Let us summarize our discussion thus far. Quantum and classical mechanics define the microscopic laws which govern the behavior of an isolated system of atoms. Because experimental measurements rarely ever involve a fully isolated system, it is necessary to extend the microscopic laws with a formalism that allows us to treat systems in equilibrium with their surroundings (Figure 1.3). The value of a physical observable is then defined as an average of its instantaneous value over the statistical distribution function of the system. This combination of quantum mechanics and statistical mechanics is incredibly powerful, and it provides the foundation for the remainder of this thesis. It allows us to write down specific equations for pretty much any experimentally measurable quantity. Unfortunately, that’s about it; it allows us to write these equations down, not solve them.

1.4 Molecular simulation

The equations of quantum and statistical mechanics can almost never be solved using analytic methods. It is only really possible to employ numerical algorithms on a computer, i.e. to perform computer *simulations* [8]. The integral in Equation 9 can be estimated numerically by generating atomic configurations which are distributed according to $p(\mathbf{r})$, and computing the average value of A over these configurations:

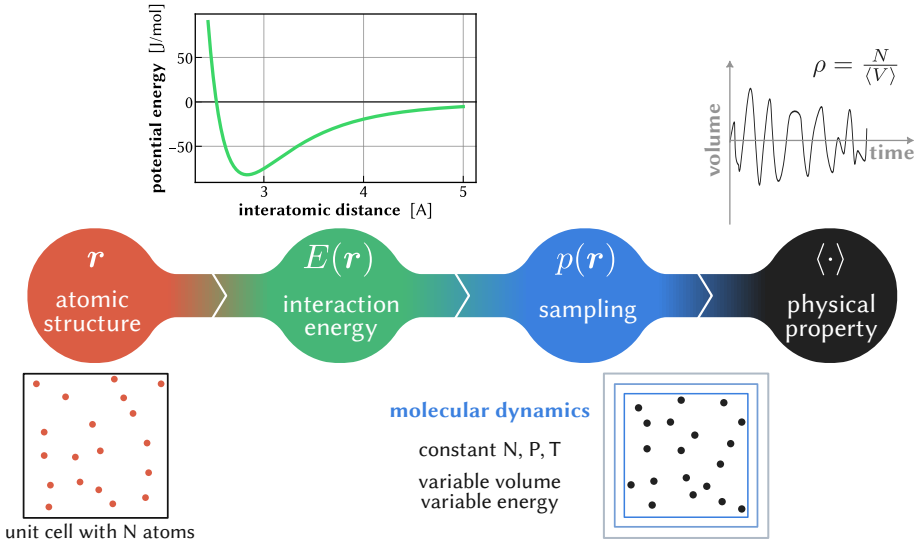


Figure 1.4: Integration of quantum and statistical mechanics into a molecular simulation pipeline, for the specific example of the density of helium gas. The system consists of N atoms in a periodic simulation cell, whose Born-Oppenheimer interaction energy is approximated with a pairwise Lennard-Jones potential. The resulting Boltzmann distribution can be sampled using e.g. molecular dynamics at constant pressure, and the average of the unit cell volume can be used to estimate the density ρ .

$$A \approx \frac{1}{n} \sum_{i=1}^n A(\mathbf{r}_i) \quad (10)$$

Over the past decades, various *sampling* algorithms have been developed which can generate such samples relatively efficiently, irrespective of the dimensionality of $p(\mathbf{r})$ or the complexity of the atomic interactions $E(\mathbf{r})$. Key to this thesis are molecular dynamics algorithms [3]. Conceptually, they represent nothing more than a discretization of the movement of atoms into small steps – typically on the order of one femtosecond, i.e. 10^{-15} s. In each time step, Newton’s equations of motion are applied to deduce the direction of the acceleration based on the instantaneous force. The instantaneous force is obtained as the gradient of the electronic energy $E(\mathbf{r})$, which itself is obtained through approximate solution of Equation 4. These are then combined with the positions and velocities of the previous step to generate new positions and velocities, after which the process is repeated. After a sufficiently large number of steps, the obtained trajectory can be considered as converged, and Equation 10 can be used to estimate physical properties




Figure 1.5: Schematic illustration of the materials investigated in this work. **(left)** metal-organic frameworks are composed of inorganic and organic building blocks joined together with coordination bonds to form porous networks. **(middle)** Zeolites are porous aluminosilicate minerals. **(right)** Metal halide perovskites are a class of semiconducting materials with the chemical formula ABX_3 , with A an organic or inorganic cation, B usually Pb or Sn, and X a halide.

based on the configurations in the trajectory. Molecular dynamics algorithms are incredibly versatile and allow to compute an impressive number of properties of the system. At any given instant, there are millions of CPUs and GPUs around the world which are executing (variations on) this basic algorithm. Lastly, we mention that there exist alternative methods to generate samples from the Boltzmann distribution (e.g. Monte Carlo methods [8]).

1.5 Goal and outline

Figure 1.4 serves as the cornerstone of the research presented in this work. While the theory behind it goes back to the 1950s and is generally well-known, its computational execution is still a subject of highly active research. The most promising materials and molecules for tackling today's technological challenges are inherently complex in nature, and application of the workflow in Figure 1.4 to these systems is incredibly challenging. To a large extent, the difficulty lies in the computational cost that is associated with evaluations of the interaction energy $E(\mathbf{r})$ and converged sampling of $p(\mathbf{r})$, as will be discussed in more detail in the next Chapter. These challenges are in essence a result of the curse of dimensionality: realistic systems contain thousands of electrons and nuclei, and the electronic energy $E(\mathbf{r})$ as well as the Boltzmann distribution $p(\mathbf{r})$ are therefore intrinsically high-dimensional quantities.

Fortunately, computer scientists and engineers have developed extraordinarily effective tools which are able to tackle computational challenges in high-dimensional spaces – machine learning, and *deep learning* in particular. Along with quantum mechanics and statistical mechanics, machine learning constitutes the foundation of the research in this work. While existing machine learning techniques show great promise in accelerating molecular simulation pipelines as in Figure 1.4, their application to complex molecular systems is rarely straightforward. The goal of this thesis is to develop and extend machine learning techniques which are effective for materials with a large chemical and structural variety. The most important areas of application in this thesis are shown in Figure 1.5; nanoporous materials (such as metal-organic frameworks and zeolites) and metal-halide perovskites (such as CsPbI₃).

We conclude with an outline of the thesis. Chapter 2 defines the two main challenges associated with Figure 1.4, namely the calculation of the quantum mechanical interaction energy $E(\mathbf{r})$ and the generation of decorrelated samples from the Boltzmann distribution $p(\mathbf{r})$. Chapters 3 and 4 present the main research results, with Chapter 3 targeting the efficient approximation of $E(\mathbf{r})$ and Chapter 4 the efficient sampling of $p(\mathbf{r})$. Finally, in Chapter 5, we present a computational framework which integrates the developments of Chapters 3 and 4 into a modular Python framework which largely removes the research overhead associated with the technical execution of molecular simulation pipelines. The framework is termed `psiflow`, and allows for a concise definition of arbitrarily complex simulation workflows that can be executed autonomously on modern computing infrastructure. Finally, Chapter 6 presents our overall conclusions and perspectives for future research.

2

Computational Challenges

The equations from quantum and statistical mechanics cannot be solved analytically for any practical system of interest. The only possible solution strategy is based on *computer simulations*. Unsurprisingly, such simulations require large amounts of computing power, and they are therefore executed on supercomputers. Figure 2.1 depicts the capacity of the largest general-purpose supercomputers in the world as a function of time over the past decades. Compute capacity is measured in the number of floating point operations per second, or Flop/s. While the growth has slowed down in recent years – Moore’s law no longer holds – it is still approximately exponential, and the amount of compute that is currently available specifically for academic




Figure 2.1: Compute capacity of the largest supercomputer in the world at any given moment during 1993 – 2024. The y-axis represents R_{\max} in units of teraflops, i.e. 10^{12} floating point operations per second. (Un)surprisingly, the vast majority of these supercomputers were operational in either Japan or the United States. Data was scraped from top500.org .

purposes is by all measures incredible. Petascale systems started appearing around 2010 (with 10^3 TFlop/s), and this decade has seen the emergence of exascale systems (with 10^6 TFlop/s). The National Energy Research Scientific Computing Center (NERSC) from the United States Department of Energy concluded that almost half of the total compute capacity is targeted towards molecular simulation¹⁰. The harsh reality is that the equations in Section 1.4 contain two incredibly hard problems, each of which can only be tackled with large amounts of compute power. The first problem is the determination of the electronic ground-state energy, which defines the effective interactions between the nuclei (Equation 4, but without arguments for simplicity):

$$\hat{H}_e \psi_e = E \psi_e \quad (1)$$

The second problem is to compute converged averages over $p(\mathbf{r})$; an incredibly complex and high-dimensional Boltzmann probability distribution:

$$p(\mathbf{r}) \propto e^{-\beta E(\mathbf{r})} \quad (2)$$

The majority of the research presented in this work is focused on developing effective tools to tackle both problems while maintaining a high level of accuracy. First, it is necessary to understand in more detail why these are so challenging.

2.1 The interaction energy

Recall from the previous Chapter that the Born-Oppenheimer approximation defines the electronic ground-state energy $E(\mathbf{r})$ as a smooth function of the nuclear coordinates \mathbf{r} based on the electronic eigenvalue equation (Equation 4):

$$\hat{H}_e \psi_e(\mathbf{r}_e | \mathbf{r}) = E(\mathbf{r}) \psi_e(\mathbf{r}_e | \mathbf{r}) \quad (3)$$

where we introduce the notation $\psi_e(\mathbf{r}_e | \mathbf{r})$ to indicate that the electronic wave function is explicitly dependent on the dynamic variables of the electrons \mathbf{r}_e while also parametrically depending on the clamped nuclei positions \mathbf{r} because they act as Coulomb interaction centers for the electrons. To

¹⁰see <https://www.nersc.gov/research-and-development/workload-analysis/> – combining the density functional theory and molecular dynamics fractions yields almost half of the pie chart.

understand why it is so difficult to solve this equation, let us first expand the electronic Hamiltonian operator into individual kinetic \hat{T} and potential \hat{V} contributions:

$$\hat{H}_e = \hat{T}_e + \hat{V}_{ee} + \hat{V}_{en} + \hat{V}_{nn} \quad (4)$$

The subscripts indicate the origin of each contribution – e for electrons and n for nuclei. Assume the system consists of n electrons and m nuclei. Then we have:

$$\hat{T}_e = \sum_i^n \frac{-\hbar^2}{2m_e} \nabla_i^2 \quad (5)$$

$$\hat{V}_{ee} = \sum_{i,j>i}^n \frac{e^2}{4\pi\epsilon_0} \frac{1}{|\mathbf{r}_{e_i} - \mathbf{r}_{e_j}|} \quad (6)$$

$$\hat{V}_{en} = - \sum_{i,j}^{n,m} \frac{e^2}{4\pi\epsilon_0} \frac{Z_n}{|\mathbf{r}_{e_i} - \mathbf{r}_j|} \quad (7)$$

$$\hat{V}_{nn} = \sum_{i,j>i}^m \frac{e^2}{4\pi\epsilon_0} \frac{Z_i Z_j}{|\mathbf{r}_i - \mathbf{r}_j|} \quad (8)$$

in which ϵ_0 denotes the dielectric constant in vacuum, e the elementary charge of an electron, and Z_i the atomic number of atom i . Note that the nuclei kinetic energies are omitted in the electronic Hamiltonian \hat{H}_e because of the assumed time scale separation between electrons and nuclei in the Born-Oppenheimer picture. The nucleus-nucleus Coulomb interaction does not depend on any of the electronic degrees of freedom \mathbf{r}_e and hence is irrelevant when determining the ground-state solution ψ_e . It does however contribute to the interaction energy $E(\mathbf{r})$.

Equation 3 is not analytically solvable for systems involving more than one electron. The culprit is the Coulomb repulsion between electrons, i.e. Equation 6, which effectively couples or correlates the behavior of the electrons in a very complicated manner. To overcome this, a plethora of approximations and numerical methods have been developed since the late 1920s; all of them aim to determine the ground-state energy $E(\mathbf{r})$ and its gradients as accurate as possible, without excessive computational cost. In what follows, we will briefly discuss the key characteristics of some of the

most commonly used methods. For theoretical derivations, we refer the reader to the original research papers and/or specialized textbooks [7,10–15]. Hartree-Fock (HF) is one of the oldest approximate solution methods and can be considered as a mean-field approximation to Equation 3, i.e. each electron experiences the average electrostatic potential that arises from all the other electrons [16,17]. HF itself is often a very crude approximation because the direct electron-electron repulsion (Equation 6) becomes non-negligible when any two electrons are very close to each other. In the literature, the *electron correlation energy* is defined as the energy difference between the exact ground-state energy and the HF (i.e mean-field) energy. There exist a variety of post-HF methods which aim to determine the correlation energy starting from the HF solution. Important examples are Møller-Plesset perturbation theory (MP) [18], the random phase approximation (RPA) [19], and coupled-cluster methods [20]. In contrast to the aforementioned techniques which can be regarded as wave-function-based, there exist alternative approaches which rely on the electron *density* as fundamental variable; these are referred to as density functional theory (DFT) methods. DFT methods are founded on the Hohenberg-Kohn theorems, which demonstrate a unique relationship between the ground-state electronic wave function, the ground-state electron density, and the external potential that arises from the nuclei [21]. In particular, these theorems prove that the ground-state energy is a functional of the electron density – although the exact expression for this density functional is unknown. In practice, DFT methods therefore aim to approximate the exact density functional in various ways, and they will be used extensively throughout this thesis [22].

By way of example, let us analyze the computational efficiency of a few of these methods. To keep the discussion tractable, we will limit our analysis in this paragraph to a measurement of the calculation runtime as observed with one particular software package (ORCA 5.0.4 [23]) and one particular class of system (a cluster of water molecules of varying size). We will consider four topical methods:

- **DFT** | Perdew-Berke-Ernzerhof (PBE) functional: this is perhaps the most popular DFT functional in the generalized gradient approximation (GGA) [24]; it is ubiquitous in both materials science and chemistry applications.

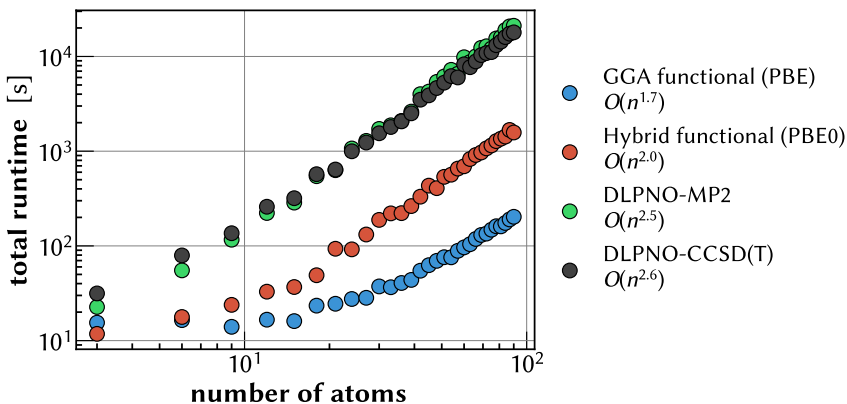


Figure 2.2: Empirical scaling behavior of the most commonly employed quantum chemistry methods, as obtained using ORCA 5.0.4 [23], on clusters of water molecules of varying size. The slope on the log-log representation provides an estimate for the exponent in the computational complexity, and is reported for each method in the legend on the right. All calculations perform energy-only evaluations at production settings for each of the methods, with sufficiently large basis sets and including common algorithmic optimizations (see Appendix B.1).

GGA functionals are commonly augmented with empirical dispersion corrections (e.g. D3 [25] or Tkatchenko-Scheffler [26]).

- **DFT** | PBE0 functional: adaptation of PBE in which a fraction of the exchange energy is computed explicitly using Hartree-Fock integrals [27]. Such functionals are also referred to as hybrid functionals, and are more expensive than pure GGA functionals though slightly more accurate.
- **post-HF** | MP2: second-order perturbative correction to HF [18]. This reintroduces most of the electron correlation which is lost by (hybrid) GGA functionals.
- **post-HF** | CCSD(T): coupled-cluster method using singles, doubles, and perturbative triples excitations [20]. This method is often considered as the gold standard in quantum chemistry for computing the interaction energy of molecular clusters, and can usually retrieve the full correlation energy.

Figure 2.2 summarizes the results for each of the four methods. We begin by emphasizing that the actual implementations of these methods include decades of algorithmic developments and low-level optimizations, and a full understanding of their computational characteristics is therefore absolutely nontrivial. Appendix B.1 discusses some of the algorithmic optimizations

that were employed in the calculations in Figure 2.2, and provides a full specification of the input as necessary to reproduce these results. Nevertheless, Figure 2.2 allows us to make three important observations: (1) more accurate methods tend to be more computationally expensive; (2) the time complexity becomes worse for higher-level methods, and (3) typical energy evaluations for systems with hundreds of atoms require minutes to hours.

We conclude this section with a brief discussion on accuracy. In general, hybrid DFT is almost always an improvement with respect to GGA DFT, while MP2 is not always more accurate than (hybrid) DFT. There exist various empirical variations of MP2 (spin-component scaled MP2, SCS-MP2) which are observed to improve its accuracy, and at the same time there exist so-called double-hybrid DFT functionals which incorporate exact exchange as well as a second-order perturbative treatment of electron correlation such as MP2. In general, the accuracy of MP2-based methods (including double hybrid functionals) is highly system-dependent. On the other hand, CCSD(T) is overall the most accurate method out of the four, and is generally regarded as the best single-reference quantum chemistry method. In Figure 2.3, we demonstrate how very large energy differences (>10 kJ/mol) between these methods can exist even for a seemingly simple case such as the proton exchange between two formic acid molecules.




Figure 2.3: Potential energy curves for a proton exchange reaction in a formic acid dimer for various levels of theory. We employ PBE-D3 and PBE0-D3 to represent the GGA and hybrid GGA levels of theory using a cc-pVTZ basis; and MP2 and CCSD(T) in the converged basis set limit (using the explicitly correlated cc-pVTZ-F12).

2.2 The time scale problem

The second computational challenge consists of computing averages over the Boltzmann distribution $p(\mathbf{r})$ in a comprehensive manner. As explained in Section 1.4, the most efficient and well-known method to do this is molecular dynamics simulation using the Born–Oppenheimer interaction energy $E(\mathbf{r})$ and its gradients. Newton’s equations of motion are differential equations, and their numerical solution requires choosing a discretization or time step Δt . Smaller time steps lead to lower numerical errors but increase the total number of propagation steps at fixed simulation time $\tau = N\Delta t$. Because the propagation at every step requires knowledge of the interatomic forces – $E(\mathbf{r})$ and its gradients – the total cost of a simulation increases linearly with the number of steps N , and it is in our best interest to choose Δt as large as possible. Large time steps imply large displacements, and if the potential energy surface $E(\mathbf{r})$ contains directions along which it is rapidly changing, such large displacements can (and generally will) cause diverging and unphysical behavior. An equivalent viewpoint relies on the *Nyquist–Shannon*¹¹ sampling theorem and states that the time step needs to be sufficiently small as to resolve the fastest motions in the system. In practice, this means that organic compounds require $\Delta t \sim 1$ fs whereas some inorganic solid state systems tolerate $\Delta t \sim 5$ fs.

For example, Figure 2.4 depicts the hierarchical structure of complex proteins starting from the atomic composition of the constituting amino acids. Since these are organic compounds with fast O–H vibrations, the largest possible timestep for proteins is typically around 0.5 fs. On the other hand, the physical properties of such systems – the shape of their folded structure and transition rates to misfolded structures – are determined on timescales at least ten (!) orders of magnitude larger [29]. This means that, if a simulation starts in one conformation, it will take on the order of 10^{10} steps (and an equal number of force evaluations) for a protein to adopt a different conformation. Especially for systems with thousands of atoms, such an enormous number of steps is absolutely infeasible with the computing power available now or even within the next few decades. Essentially the same conclusion holds for chemical reactions and/or phase transitions in solid-state systems; the number of steps required in dynamic simulations in order to make the

¹¹although originally discovered by E. T. Whittaker [28]




Figure 2.4: Hierarchical overview of the atomic composition of large protein complexes, starting from the amino acid sequence. Molecular vibrations in the amino acids occur on a timescale from femto- to picoseconds, whereas large conformational changes or protein assemblies happen on a timescale of milliseconds or larger.




Figure 2.5: Abstract representation of two free energy minima in high-dimensional phase space. In a dynamic simulation, the system can rarely ever escape the stable regions A or B because it needs to cross an energetically unfavorable region, for which the associated Boltzmann probability becomes exponentially small.

targeted “event” happen is usually orders of magnitude larger than what is computationally feasible. In addition, if we take into account the computational cost per step, i.e. the computational cost of evaluating the interatomic forces using any of the methods in Figure 2.2, we must conclude that brute force molecular dynamics simulation using quantum chemistry methods is almost never a computationally feasible option. This problem can be broadly referred to as the time scale problem.

The origin of the incredibly large range of time scales in many physical systems is the exponential dependence of the Boltzmann probability density $p(\mathbf{r})$ on the dimensionless interaction energy $\beta E(\mathbf{r})$. At room temperature, the thermal energy β^{-1} amounts to about 2.5 kJ/mol, which implies a very rapid decrease in the relative likelihood of a given configuration \mathbf{r} with increasing energy $E(\mathbf{r})$. The consequence of this is that states which possess an energy that is significantly higher than the thermal energy become exponentially less probable to be sampled. For example, the transition barrier for the proton transfer reaction in Figure 2.3 is about 35 kJ/mol higher in energy than the stable minimum, which makes those states about e^{-15} or 1,000,000 times less likely to be sampled. The proton transfer transition event can therefore be referred to as a *rare event*, because the associated time scale will be roughly a million times higher than e.g. the intramolecular vibrations. In practice, this means that the *ergodic hypothesis* is almost always violated; simulations are unable to escape regions of high probability because the energy barriers separating the minima exceed the thermal energy kT by a large factor. This makes it very hard to generate samples across multiple ‘disconnected’ regions in phase space and produce unbiased property predictions.

One of the most common physical properties to predict is the relative free energy difference between two regions in phase space [8]. Such regions may be explicitly defined in terms of the atomic coordinates (e.g. a liquid versus a solid phase), or implicitly defined based on the thermodynamic control variables (temperature, pressure, ...). For explicitly defined regions, we define the free energy difference ΔF as follows:

$$\begin{aligned} \beta\Delta F &= \beta(F_B - F_A) \\ &= \log \frac{\int_A p(\mathbf{r}) d\mathbf{r}}{\int_B p(\mathbf{r}) d\mathbf{r}} \end{aligned} \quad (9)$$

where $\beta = (kT)^{-1}$ represents the inverse temperature. Intuitively, ΔF quantifies the ratio of energetically accessible volumes of phase space in A versus B. There exists a large variety of methods to compute Equation 9 in an efficient manner [30,31]. A common denominator in nearly all of these methods is that they compute ΔF between disjoint regions A and B by cre-

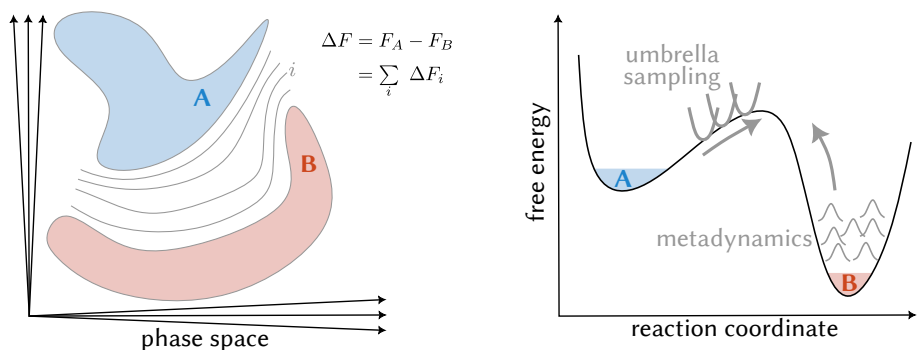


Figure 2.6: Abstract illustration of the calculation of ΔF . Typically, it is necessary to create intermediate sampling distributions which bridge the thermodynamically stable regions A and B with each other. In many cases, this is achieved by using a reaction coordinate in combination with e.g. biasing potentials.

ating intermediate distributions p_i which bridge regions A and B gradually and with sufficient overlap [32]. The total free energy difference ΔF is then computed as a sum over smaller differences ΔF_i . A topical example is (al-chemical) thermodynamic integration [33–37], but also reaction-coordinate-based methods such as umbrella sampling or on-the-fly probability enhanced sampling (OPES) can be framed in this sense [38–40]. Implicitly, such methods modify the standard Boltzmann probability $e^{-\beta E}$ with either additional energy contributions or using modified thermodynamic conditions such that the sampling in each subregion becomes approximately ergodic again. While such methods have demonstrated their enormous power, they very often require specialized prior knowledge about the system of interest and the particular energy barriers that are present.

2.3 Classical interaction potentials

Both the Born-Oppenheimer approximation and most of today’s quantum chemistry methods were already developed by the second half of the twentieth century. At the time, it was very clear that dynamic simulations of atoms and materials could never rely on explicit solution of $E(\mathbf{r})$ and its gradients because computing technology was still at a very early stage. As such, the first applications of molecular dynamics simulations were actually employing classical force fields, i.e. simple analytical functions which were designed to reproduce certain qualitative characteristics of the system under

bonded		nonbonded	
bond	$\frac{K}{2}(r - r_0)^2$	Lennard-Jones	$4\epsilon \left(\frac{\sigma}{r_{ij}}\right)^{12} - 4\epsilon \left(\frac{\sigma}{r_{ij}}\right)^6$
angle	$\frac{K}{2}(\theta - \theta_0)^2$	Coulomb (point charge)	$\frac{q}{4\pi\epsilon_0 r_{ij}}$
torsion	$K \cos(m\psi - \psi_0)$	Coulomb (Gaussian [41,42])	$\frac{q}{4\pi\epsilon_0 r_{ij}} \operatorname{erf}\left(r_{ij} \frac{\pi}{\sqrt{2}a}\right)$

Table 2.1: Overview of the most common bonded and nonbonded interaction functions in classical force fields. Bonded terms require manual specification of the particles which define the *internal coordinates* r , θ , or ψ . The constants K represent interaction strengths whereas θ_0 , ψ_0 , and m define the rest values (values at which the interaction energy is zero). Nonbonded interactions are typically considered between any pair of particles whose distance r_{ij} is smaller than a predefined cutoff (although additional exclusion rules may apply). Various functional forms exist which aim to describe either dispersion or electrostatic interactions. Each atom carries its own parameters q_i , a_i , σ_i , ϵ_i with the appropriate mixing rules (e.g $q = q_i q_j$, $\sigma = \sqrt{\sigma_i^2 + \sigma_j^2}$).

study [43,44]. Such potentials separate the total interaction energy into *bonded* and *nonbonded* contributions:

$$E_{\text{classical}} = E_{\text{bonded}} + E_{\text{nonbonded}} \quad (10)$$

An overview of common interaction terms is given in Table 2.1. The bonded contribution aims to describe the “covalent” part of the interaction energy, and consists mostly of harmonic functions of various internal coordinates such as bond distances, (torsion) angles, and others. Their key feature is that they enforce strict chemical connectivity at a local level; atoms which are bonded at the start of the simulation will remain bonded. This is enforced by hardcoding the indices of the atoms that constitute a particular bond distance, angle, or other internal coordinate in the system. The nonbonded energy in Equation 10 represents the longer-ranged dispersion and electrostatic interactions. Most often, they are expressed as pairwise van der Waals and Coulomb contributions. While they are in theory active between any two atoms in the system, they are in practice only evaluated for atoms within a specified cutoff radius (e.g. 15 Å), or employ more advanced methods such as

the (particle mesh) Ewald summation [45] or the fast multipole method [46]. Lastly, we mention that all parameters in Equation 10 – equilibrium bond lengths, partial charges, van der Waals radii - can be determined based on both experimental and theoretical results [42,47–50].

Classical force fields enforce the exact same chemical connectivity at any point throughout the simulation. More generally, their functional form assumes a stationary electron density such that the covalent, dispersion, and electrostatic parameters of the force field continue to describe the pertinent features of $E(\mathbf{r})$ in a qualitatively correct manner. For situations in which these assumptions are valid, they constitute an incredibly efficient tool to simulate the behavior of extended systems over large time scales. From a computational point of view, evaluation of Equation 10 can be performed in a massively parallel manner, and force field simulations are therefore highly suited for execution on massively parallel hardware such as GPUs [53,54].

In **Paper I**, we exploit this efficiency and investigate the dynamics of a structural phase transition in a flexible metal-organic framework (MOF) on an unprecedented scale (Figure 2.7). Since this transition is associated with large and anisotropic unit cell changes, it requires constant pressure molecular dynamics simulations in which the unit cell shape is in no way constrained. Conventional constant pressure simulation algorithms (see MTK [55], Langevin [56], Parrinello-Rahman [57], or Berendsen [58]) require analytical evaluation of the gradients of Equation 10 with respect to the unit cell parameters, but these are cumbersome to implement and represent an additional computational cost. Importantly, because most MD engines are targeted towards biomolecular force fields, they do not support all features of more specialized force fields used in (soft) condensed matter. Examples are the use of delocalized atomic charges (in combination with particle mesh Ewald), custom dispersion interactions and tail corrections, and alternative three- and four-body covalent interactions [42,50]. To circumvent the need for implementing such gradients manually, we developed a new pressure control algorithm which is fully anisotropic and which does not require such gradients during the simulation. Effectively, the algorithm uses molecular dynamics to dynamically evolve the atomic positions in time, whereas the sampling of various unit cells is achieved using Monte Carlo sampling. After every k steps of position updates, a trial move in the unit cell components

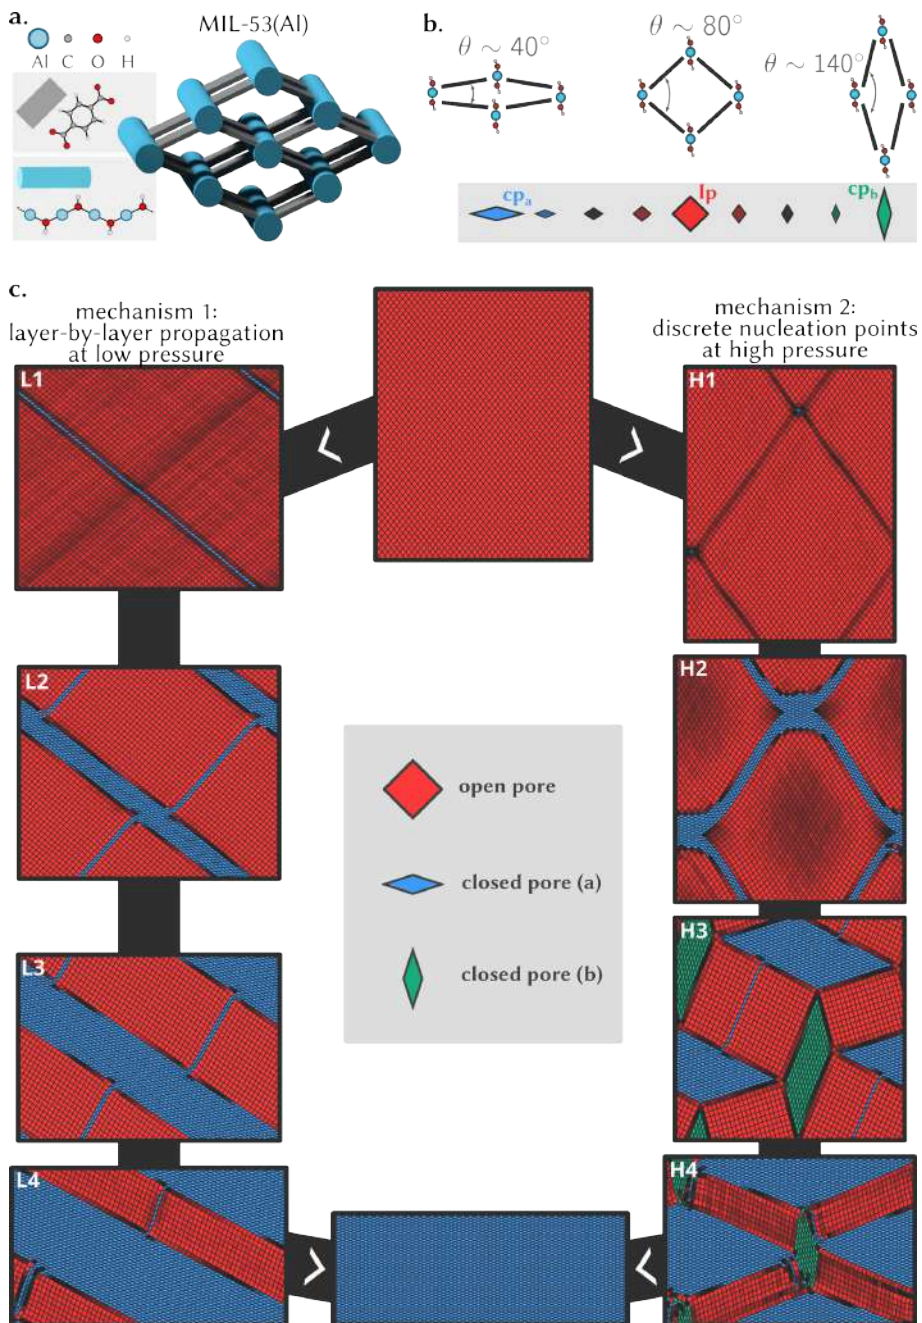


Figure 2.7: **a** Winerack topology of MIL-53(Al) [51]. **b** Overview of the closed pore (*cp*) and open pore (*op*) phases as observed in experiment and in small-scale atomic simulations. **c** Visualization of large-scale molecular dynamics simulations of the phase transition. The nucleation of the transition depends on the applied pressure, and is one-dimensional for low pressures and zero-dimensional at high pressures. Figure adopted from Ref. [52]

is attempted, whereby usually $10 < k < 100$. To guarantee detailed balance and ensure that the statistical fluctuations in the cell parameters are correctly reproduced, it is necessary to start from the exact isothermal-isobaric partition function [59] which takes into account all degrees of freedom of the unit cell $\mathbf{g} \in \mathbb{R}^{3 \times 3}$:

$$\mathcal{Z}(N, P, T) \propto \int e^{-\beta E(\mathbf{r})} e^{-\beta P \det(\mathbf{g})} \det(\mathbf{g})^{-2} d\mathbf{r} d\mathbf{g} \quad (11)$$

based on which it is possible to derive a Metropolis-like acceptance criterion for trial moves in any of the unit cell components. The details of this approach as well as extensive validation experiments are presented in the Supporting Material of **Paper I**.

We implemented the new pressure control algorithm in OpenMM and used it to simulate the phase transition in MIL-53(Al) at an unprecedented scale for these systems – our largest supercell contained over 1 million atoms. We revealed two qualitatively different transition mechanisms depending on the applied pressure. At lower pressures (100 MPa), we observed transitions occurring through a layer-by-layer mechanism where the framework maintains its rhombus-like shape and changes happen cooperatively within individual layers, propagating gradually across the material. In contrast, at higher pressures (300 MPa), we found transitions initiating at discrete nucleation points throughout the framework, leading to multiple domain formation and significant structural deformation.

2.4 Summary

This Chapter outlines the key challenges that are associated with molecular simulations of materials and molecules, namely (1) the evaluation of the Born-Oppenheimer interaction energy $E(\mathbf{r})$ and (2) the calculation of unbiased and converged averages over the Boltzmann distribution $p(\mathbf{r})$. The two challenges are not independent from each other, given that sampling algorithms for $p(\mathbf{r})$ – molecular dynamics and Monte Carlo – internally require evaluation(s) of the energy and/or its gradients. However, even if we could determine $E(\mathbf{r})$ in only a fraction of a second, Figure 2.4 still suggests that it takes billions of timesteps to simulate a single protein folding event, let alone characterize the relative stability between two conformations. While

fast evaluation of $E(\mathbf{r})$ is a necessary condition for being able to sample $p(\mathbf{r})$, it is by no means a sufficient condition. Overcoming free energy barriers requires dedicated sampling algorithms and often becomes highly system-specific – as shown in Figure 2.5 and Figure 2.6. For example, in **Paper I**, we incorporated a priori knowledge on the phase transition of the framework by selecting the unit cell volume as biasing coordinate. Such knowledge is not always available, especially for realistic systems with many components, impurities, and/or defects. In addition, the use of classical force fields assumes fixed chemical connectivity, fixed local charge density, and overall a simple parameterization for the covalent interaction energy. Such assumptions break down in the majority of practical scenarios, and more general solution strategies are necessary.

The following two Chapters present the most important research milestones in this work; Chapter 3 demonstrates a general approach to reduce the computational cost of evaluating $E(\mathbf{r})$ by orders of magnitude, and Chapter 4 presents a new framework for characterizing transition paths and relative free energies.

3

Learning the Interaction Energy

Computational physicists and chemists have worked for decades trying to come up with clever algorithms and approximations which allow to overcome the challenges posed in Chapter 2. Some of these have proven extraordinarily effective, such as the myriad of enhanced sampling algorithms, or local correlation methods such as DLPNO. In some cases, the exponential increase in computing power as shown in Figure 2.1 allows us – now or in the near future – to brute force a solution at the cost of compute power and a large time-to-solution. Nevertheless, many research questions in materials science and chemistry remain unanswered due to the lack of a computationally feasible solution. Instead of looking for approximations *within* a physics context, it could be refreshing to look for approximations *outside* of it, by regarding the interaction energy $E(\mathbf{r})$ as just another mathematical function to be approximated, and $p(\mathbf{r})$ as just another probability distribution to be sampled. In that sense, “finding approximations” can be rephrased as “learning approximations”, and that branch of science is no longer called physics, mathematics, or chemistry – it is called computer science, and more specifically *machine learning*.

This Chapter will introduce the essentials of the machine learning techniques which have been employed and exploited in this thesis. For a general introduction to machine learning, we refer the reader to excellent textbooks by e.g. Bishop [60] or Goodfellow [61]. For a general introduction to deep learning in chemistry and materials science, we highly recommend White [62].

3.1 Introduction

In machine learning, we aim to approximate the interaction energy $E(\mathbf{r})$ by a function $f : \mathbf{r} \rightarrow \mathbb{R}$ which has parameters or weights \mathbf{w} : $f \triangleq f_{\mathbf{w}}$. The weights should be *learned* based on some amount of labeled reference data, that is, atomic geometries \mathbf{r} and their associated interaction energy E as obtained using the quantum chemistry methods outlined in the previous Chapter. Which considerations go into the design of $f_{\mathbf{w}}$ and the generation of the training set, and how do we find the optimal weights?

From a computer scientist's point of view, the interaction energy $E(\mathbf{r})$ represents a mapping from a *point cloud* to a real number. Each point has cartesian coordinates and an attribute which defines its chemical identity. Importantly, point clouds are invariant with respect to any permutation of points which have identical attributes (i.e. which are of the same chemical species). This coincides with our intuition that atoms of the same chemical species are fundamentally indistinguishable. Because space is also fundamentally homogeneous and isotropic [63], any physical property of the point cloud should also remain invariant with respect to global translations and global rotations. Thirdly, and somewhat more subtly, we also know that the interaction strength between any two atoms decays with increasing distance as a consequence of the $\frac{1}{r}$ dependence in Coulomb's law. These considerations represent our prior knowledge about $E(\mathbf{r})$, and it is beneficial to encode this prior knowledge into the design of $f_{\mathbf{w}}$. In machine learning, design considerations for $f_{\mathbf{w}}$ based on prior knowledge are referred to as *inductive biases*. Here, the inductive biases for learning the interactions are *symmetry* (rotational, translation, permutational), *locality* (interaction strengths decrease with increasing distance), and *continuity* (the energy is a smoothly varying and differentiable quantity).

Secondly, the point clouds that are to be dealt with are extremely high-dimensional. The range of system sizes that are encountered across physics, chemistry, and biochemistry is enormous, starting at 10-100 atoms for small unit cells and molecules to 10^3 - 10^6 atoms for biomolecular systems. To accommodate such system sizes, $f_{\mathbf{w}}$ should scale quasilinearly with system size.

3.2 The manifold hypothesis and neural networks

One of the most common assumptions or priors in machine learning is the smoothness prior; it states that the function to be learned should not change very much within a small region [61]. The k -nearest neighbors algorithm is an extreme example of this prior. It constructs a neighborhood around the input \mathbf{r} , looks at training points within a neighborhood ε , and assigns as output the average value of the neighboring training points:

$$\mathbf{r} \longrightarrow \text{mean} \{E(\mathbf{r}') \mid \|\mathbf{r}' - \mathbf{r}\| < \varepsilon\} \quad (1)$$

Instead of outputting the mean, it is possible to output some continuous weighted average which takes the relative distances into account. In addition, instead of the regular Euclidean distance $\|\cdot\|$, they are frequently employed in combination with more complex similarity measures, or *kernels*. Such methods work very well in low-dimensional input spaces, but quickly become intractable in high-dimensional settings. Consider a system of just 20 atoms; let us represent the interaction energy $E(\mathbf{r})$ on a grid with two cells per dimension. If we are to characterize the interaction energy across this grid, we would require at least one training sample (\mathbf{r}, E) per grid cell and therefore about $2^{60} \sim 10^{18}$ training samples for the entire grid. This is an incredibly large number, even though this is an extremely coarse grid and a very small system. How can machine learning overcome this curse of dimensionality?

Even though the input space is high-dimensional, only a fraction of that space will correspond to physically meaningful configurations with a nonnegligible Boltzmann probability, and it is precisely these states which contribute significantly to ensemble averages. In statistics and machine learning, the *manifold hypothesis* states that the relevant fraction of a high-dimensional input space very often consists of lower-dimensional regions or *manifolds*, and our goal should be to generate functions $f_{\mathbf{w}}$ which can learn this manifold structure based on their weights \mathbf{w} . This hypothesis is valid not just for atomic geometries but also for image recognition (a random collection of pixels will rarely represent an actual picture), speech recognition (a random waveform rarely represents correctly spoken language), and language models (a random sequence of letters will rarely represent intelligible text) [64–66]. In each of these situations, *neural networks* have demonstrated

to be extraordinarily effective as parameterizations f_w . Unsurprisingly, this has also been the case for learning general properties on point clouds, and in particular learning physical properties of atomic systems. In the next sections, we will elaborate on arguably the most powerful neural network architecture for predicting properties of atomic systems: geometric graph neural networks.

3.3 Graph neural networks

Graph neural networks (GNNs) make predictions on input data that is structured as a mathematical graph. Informally, a graph can be considered as a set of nodes or vertices \mathcal{V} between which there exist zero or more edges \mathcal{E} . Predictions on graphs can concern node-level properties, or global properties. Graphs provide an ideal framework to incorporate the inductive biases we outlined in Section 3.1, such as permutation invariance and locality. Consider a system of n atoms with cartesian coordinates \mathbf{r} . Each atom i possesses a chemical identity defined by its atomic number Z_i . To convert an atomic system into a graph \mathcal{G} , the following set of rules makes intuitive sense [67]:

- each atom corresponds to precisely one node in the graph, i.e. \mathcal{G} has n nodes. In addition, nodes are annotated with a single attribute that denotes their chemical identity; the atomic number Z_i .
- there exists an edge between nodes v_i and v_j whenever the interatomic distance $\|\mathbf{r}_j - \mathbf{r}_i\|$ is smaller than some cutoff distance r_{cut} . As such, edges encode the locality of the nodes in the graph. In addition, this approach guarantees that the size of \mathcal{G} scales linearly with the number of atoms (both in terms of nodes and edges). This is important to guarantee good scaling behavior of the models with system size.

The above construction is robust towards permutation, translation, and symmetry operations. The addition of edges between nearby atoms encodes our physical intuition that atoms which are in each other’s proximity should interact more strongly. However, the exact geometric structure of the atomic system has not yet been introduced in the graph. In geometric GNNs, the edges are augmented with specific features which are computed based on the length and orientation of relative distance vectors \mathbf{r}_{ij} in the system. Traditionally, such edge features were only based on the norm of the distance




Figure 3.1: **a** arbitrary atomic geometry to be embedded in a set of (equivariant) features. **b** graph-level representation of the same geometry, with atoms becoming nodes and edges being drawn between nearby atoms irrespective of the chemical connectivity. **c** Embedding of the graph edges using the interatomic distance vectors. The length is encoded using radial basis functions (e.g. Gaussian or Bessel), whereas the directionality is encoded in an equivariant manner using spherical harmonics of arbitrary order.

vectors because it is the most straightforward way to guarantee rotational invariance [68,69]. Early approaches to augment purely distance-based edge features with orientational awareness include e.g. DimeNet [70], in which the angles between atoms are included. However, such approaches are rather ad hoc, and never guarantee that the resulting atomic environment representation is complete [71].

While rotational invariance of the final energy prediction is absolutely essential, invariance of the internal feature representations is not. Conceptually, it is possible to use noninvariant features provided that we keep track of the rotational behavior. For example, we may perform vectorial operations on the distance vectors as long as we make sure that the final energy only ever depends on dot product combinations which are again invariant to global rotations. The idea of “keeping track of the transformation behavior of mathematical objects” has been formalized in the mathematical theory of symmetry and representation [63], and has been introduced into deep learning less than a decade ago by Max Welling and others [72–76]. *Group-equivariant* GNNs provide the formal framework which is necessary to incorporate noninvariant features into GNNs, for any symmetry group G . In essence, features are no longer constrained to invariants, but are allowed to be noninvariant. Based on symmetry theory, any noninvariant feature can be decomposed into irreducible representations of the symmetry group, each of which transforms in a well-defined manner upon application of group ele-

ments – this is the mathematical equivalent to the intuitive notion of keeping track of the transformation behavior of a quantity as mentioned earlier. In the particular case of point clouds and atomic systems, the mathematical symmetry group is the *Euclidean* group, which contains all translations, rotations, and reflections; it is mathematically denoted as $E(3)$. The subgroup of translations and rotations is referred to as the *special* Euclidean group, and denoted with $SE(3)$.

Figure 3.1 explains how the geometric graph is constructed from a collection of atoms, and how the nodes and edges are embedded into features that are used as inputs for the neural network. The node embedding consists of a (trainable) embedding vector for each chemical element in the system; the embeddings are referred to as $\mathbf{h}_i^{(0)}$ for each node i , with $\mathbf{h}_i^{(0)} = \mathbf{h}_j^{(0)}$ if and only if atom i and atom j represent the same chemical element. The edge embeddings encode the actual atomic geometry of the atoms, and can be thought of as separable into a radial and an optional angular part. The radial part of the embedding is constructed based on the distance of an edge between atoms i and j :

$$\begin{aligned} \mathbf{r}_{ij} &= \mathbf{r}_j - \mathbf{r}_i \\ r_{ij} &= \|\mathbf{r}_{ij}\| \end{aligned} \tag{2}$$

While the actual design choices may vary between different models, the common denominator for all radial embeddings is that each (scalar) distance r_{ij} is expanded using a set of basis functions R_m as it is well-established that such expansions greatly facilitate model training [61]. To guarantee continuity of $E(\mathbf{r})$ whenever atoms are moving in and out of each other’s radius of interaction, the radial embedding is constructed such that its values decay gradually to zero whenever the distance r_{ij} approaches the cutoff r_{cut} . This is achieved by means of a separate envelope function f_{env} which decays to zero at the cutoff distance in a continuously differentiable manner. Combining the radial basis function expansion with the envelope function, we can write down a generic radial embedding as follows:

$$R(r_{ij}) = \{f_{\text{env}}(r_{ij})f_m(r_{ij}) \mid m \in \{0, 1, \dots, k\}\} \tag{3}$$

Common choices of basis functions are Gaussian hills that are centered uniformly in the range $[0, r_{\text{cut}}]$, or Bessel functions of varying order. The

number of basis functions is somewhat related to the size of the interaction radius, but for most models usually ranges between 6 and 10. The envelope function is typically constructed using a polynomial function¹².

Earlier GNN architectures such as SchNet [68] restrict themselves to invariant radial embeddings as in Equation 3, mainly because it is the most straightforward way to guarantee invariance of the final model. More recent (equivariant) GNN models complement the radial embedding with noninvariant components in order to better characterize the orientation of an edge in 3D space. As stated in the previous paragraph, doing so requires careful consideration of the rotational character of the noninvariant components such that the final output prediction remains invariant. Many equivariant models therefore choose an expansion based on spherical harmonics, as these constitute irreducible representations of SE(3). More precisely, they evaluate the (real) spherical harmonics in the point on the unit sphere that corresponds to the direction $\hat{\mathbf{r}}_{ij} = \frac{\mathbf{r}_{ij}}{r_{ij}}$ of the edge:

$$\left[Y_0^0(\hat{\mathbf{r}}_{ij}) \right], \left[\begin{array}{c} Y_1^1(\hat{\mathbf{r}}_{ij}) \\ Y_1^0(\hat{\mathbf{r}}_{ij}) \\ Y_1^{-1}(\hat{\mathbf{r}}_{ij}) \end{array} \right], \left[\begin{array}{c} Y_2^2(\hat{\mathbf{r}}_{ij}) \\ Y_2^1(\hat{\mathbf{r}}_{ij}) \\ Y_2^0(\hat{\mathbf{r}}_{ij}) \\ Y_2^{-1}(\hat{\mathbf{r}}_{ij}) \\ Y_2^{-2}(\hat{\mathbf{r}}_{ij}) \end{array} \right], \dots \quad (5)$$

up to and including a given maximum rotation order l_{\max} .

After assigning appropriate feature representations to both nodes and edges, GNNs perform (learned) operations using those features in order to predict node- and graph-level properties. Meaningful predictions need to characterize the structure of the graph beyond the simple embedding depicted in Figure 3.1. In GNNs, this is achieved by means of a series of

¹²Many of the most well known GNN architectures use the following polynomial [77]:

$$u(d) = 1 - \frac{(p+1)(p+2)}{2}d^p + p(p+2)d^{p+1} - \frac{p(p+1)}{2}d^{p+2} \quad (4)$$

with $p = 6$ and $d = \frac{r_{ij}}{r_{\text{cut}}}$. The polynomial has a zero with multiplicity three at $d = 1$.

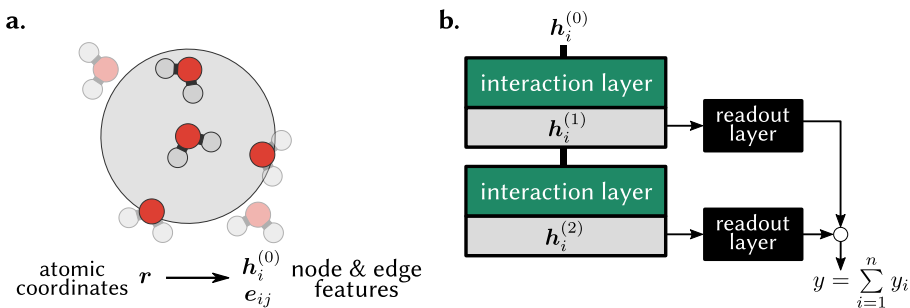


Figure 3.2: **a** atomic environments are converted into node features $\mathbf{h}_i^{(0)}$ and edge embeddings e_{ij} according to Figure 3.1. **b** A series of interaction layers refines the node features based on the local environment of each of the nodes, after which a readout function computes node-level predictions y_i based on the total learned representation \mathbf{h}_i .

layers which are referred to as interaction layers. These layers will update node representations by exchanging “messages” with neighboring nodes regarding their local environments, and it effectively allows local information to “flow” through the graph (Figure 3.2). This process is known as *message passing*, and it provides the foundation of modern graph neural networks [67]. There are typically three main steps in each interaction layer: message construction, message aggregation, and node update, although the specific implementation of these steps can vary quite significantly across different architectures [68,69,78,79]. In what follows, we denote the node features of atom i in layer l by $\mathbf{h}_i^{(l)}$, and the edge features between atoms i and j by e_{ij} . Note that the topology of the graph as well as the associated edge embedding e_{ij} remains fixed throughout the interaction layers.

1. **message construction:** for each edge in the graph, a message \mathbf{m}_{ij} is computed based on the node features of atoms i and j , as well as the edge features e_{ij} :

$$\mathbf{m}_{ij}^{(l+1)} = M \left(\mathbf{h}_j^{(l)}, \mathbf{h}_i^{(l)}, e_{ij} \right) \quad (6)$$

where M represents the message function. Importantly, the message function needs to take into account the irreducible representations for each of its inputs. For the Euclidean group in particular, features from representations l_1 and l_2 can give rise to representation orders l_3 :

$$|l_1 - l_2| \leq l_3 \leq |l_1 + l_2| \quad (7)$$

using Clebsch-Gordan coefficients [80]. Typically, the feature representations $\mathbf{h}_i^{(l)}$ are combination of multiple channels or dimensions over various rotation orders. The execution and subsequent linear mixing of all possible tensor product “paths” includes trainable weights that are obtained from the radial embedding $R(r_{ij})$ using a multilayer perceptron (MLP). This operation is known as a FullyConnectedTensorProduct in the e3nn library [81], which is used by most of the equivariant GNN architectures. For more information, we refer the interested reader to the documentation of e3nn or the accompanying papers [74,76].

2. **message aggregation:** edge messages are aggregated for all neighbors j of node i :

$$\mathbf{m}_i^{(l+1)} = \sum_j \mathbf{m}_{ij}^{(l+1)} \quad (8)$$

3. **node update:** the new node features $\mathbf{h}_i^{(l+1)}$ are obtained based on the previous node features $\mathbf{h}_i^{(l)}$ and the total message $\mathbf{m}_i^{(l+1)}$ using some update function U :

$$\mathbf{h}_i^{(l+1)} = U\left(\mathbf{h}_i^{(l)}, \mathbf{m}_i^{(l+1)}\right) \quad (9)$$

Finally, target properties are predicted as sums of per-node predictions using relatively simple readout functions \mathcal{R} . In many cases, these are simple MLPs which contain a much smaller number of trainable weights in comparison to the interaction layers. They compute a node-level prediction y_i for atom i based only on the learned feature representation for that node:

$$y_i = \mathcal{R}(\mathbf{h}_i) \quad (10)$$

Graph-level properties are then computed by aggregating the node-level predictions across the graph:

$$y = \sum_i y_i \quad (11)$$

Graph neural networks can be employed as machine-learned (ML) potentials if the predicted quantity – y from Equation 10 and Equation 11 – is trained to reproduce the interaction energy $E(\mathbf{r})$ as defined in the Born-Oppenheimer approximation. Since the energy is a global property of the atomic system,




Figure 3.3: atomic energies from a MACE-type GNN [82], for the formic acid proton exchange reaction as depicted in Figure 2.3. Atoms are colored according to the variance of the corresponding atomic energy (darker means larger variance). As such, we see that it is mostly the oxygens and the exchanging hydrogens which contribute to the total energy barrier, and not so much from the central carbon or the outside hydrogen.

the predicted energy E will be constructed as a sum over per-atom predictions E_i :

$$E = \sum_i E_i \quad (12)$$

In addition to the discussion in the previous Section, there are a few additional considerations when GNNs are employed to predict E .

First, there exists no well-defined physical observable in quantum mechanics whose eigenvalues correspond to a per-atom energy E_i . Only the total energy is quantum mechanically well-defined, as it represents the eigenvalue of the hamiltonian operator \hat{H} . In this sense, it is meaningless to attribute any physical significance to the per-atom partitioning of the energy $E = \sum_i E_i$. Nevertheless, the decomposition in Equation 12 is very useful since it preserves the permutation invariance of the energy prediction and ensures it represents an extensive quantity. In addition, we demonstrate in Figure 3.3 how trained GNNs are surprisingly good at producing intuitive partitionings of the total energy into atomic contributions. This observation inspired further research regarding the learned feature representations in these GNNs, and they will be discussed in the next Chapter. The total energy E is a global property which does not contain information on the local

interatomic interactions. Empirically, it has been observed that it is much more efficient to include the interatomic forces in the loss function as well:

$$\mathcal{L} = \lambda_e \text{RMSE}(\text{energy}) + \lambda_f \text{RMSE}(\text{forces}) + \lambda_\sigma \text{RMSE}(\text{stress}) \quad (13)$$

in which $\lambda_{e,f,\sigma}$ represent weighting factors which determine the tradeoff between the various error contributions.

Second, the continuity of E imposes constraints on the model architecture. When neighboring atoms are located just within the cutoff radius, a small displacement can cause the corresponding edge to disappear, which, due to disappearing messages, can cause discrete changes in the predicted atomic energy. To avoid discontinuities in the energy, radial embeddings are carefully modulated by a smooth cutoff function which tends to zero as the interatomic distance approaches the cutoff. Activation functions in the network should be continuous – e.g. a sigmoid linear unit (SiLU) is fine while a rectified linear unit (ReLU) is not since its first derivative is discontinuous at the origin.

Third, dynamic simulations require not only the energy but also its *gradients* for integrating the equations of motion, i.e. interatomic forces and optionally the virial stress in case of constant pressure simulations. While gradients are usually cumbersome to implement in quantum chemistry packages, they can be automatically obtained for the output energy E using automatic differentiation. Importantly, this procedure is numerically stable and works for second-order gradients as well. In Section 3.4, we will showcase how numerically stable higher-order gradients unlock fast and convenient methods to evaluate the mechanical stability of materials.

Fourth, it is important to judiciously benchmark the network hyperparameters in order to determine a suitable tradeoff between accuracy and computational efficiency. Aside from the size of the node and edge embeddings, the cutoff radius which determines the size of the local environments is a critical parameter. For all systems studied in this work, we employ a relatively large cutoff of 6.5 Å as this enabled us to obtain excellent agreement between DFT and ML potential predictions in physical properties (as will be discussed in more detail in Sections 3.4 and 3.5).

In the literature, GNN architectures are benchmarked on synthetic datasets in order to showcase their accuracy in terms of total energy and

per-atom force predictions at fixed training set size – common examples are QM9, MD17, and SPICE [83–85]. In practical scenarios, there are additional design considerations and performance indicators which should be taken into account and which are complementary to any error metrics. One of the most important ones is computational efficiency in terms of training (the time it takes to fully train a model from scratch) and inference (the time it takes to evaluate the energy and forces for a single atomic geometry). In addition, model stability during MD is important; models can perform very well on (large) datasets but still fail to produce physically relevant trajectories in phase space because the learned interaction energy contains singularities [86].

3.4 On-the-fly learning

The predictive power of ML potentials is critically dependent on their training data. In general, their predictions are reliable for atomic configurations which can be considered as *in-distribution*, i.e. which contain local environments which are sufficiently represented in the training and validation set of the model. Conversely, atomic configurations that are said to be *out-of-distribution* when they contain unseen atomic environments, and the model’s predictions for such samples will generally be quantitatively and qualitatively incorrect. As such, application of ML potentials in molecular simulations is not always straightforward since it needs to be guaranteed that all atomic environments encountered *during* the simulation should have already been included in the model’s training data *prior* to the simulation. This seems to give rise to a Catch-22 situation, since it is not really possible to know in advance the precise atomic structures which are going to be encountered in a given simulation. The most straightforward solution to this problem is to perform a relatively small amount of *first principles* molecular dynamics simulations using one of the quantum chemistry methods from Chapter 2 to evaluate the interatomic forces at every time step. The obtained trajectories are then used as training data for the ML potential, after which much longer “production-level” simulations can be performed using the trained potential. Even though this is by far the most common approach to training data generation, it suffers from two major drawbacks. First, it is demonstrably inefficient: most of the computational effort is spent

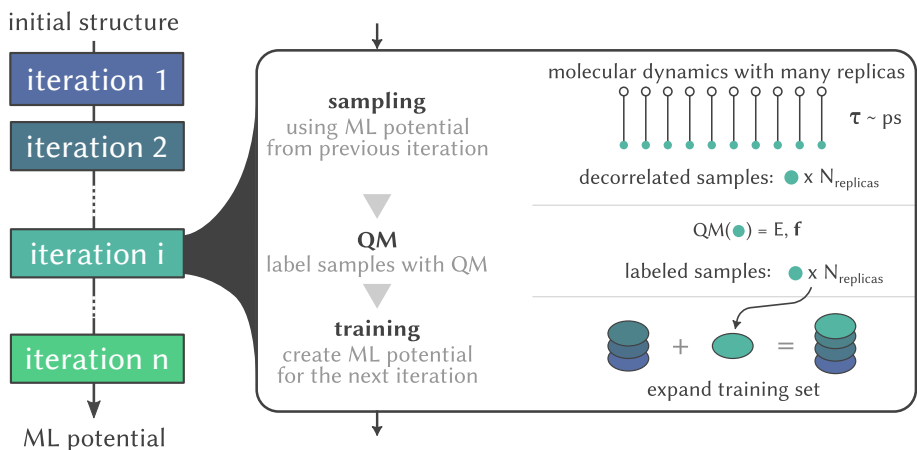


Figure 3.4: Overview of on-the-fly learning as proposed in **Paper II**. The ML potential from the previous iteration is used to perform a large number of short molecular dynamics simulations – optionally at different thermodynamic conditions and/or with biasing forces. Samples are evaluated using QM and appended to the training data. The algorithm can be bootstrapped by training on random perturbations of the initial configurations or using a pretrained ML potential from literature (e.g. MACE-MP0) [87]. Figure inspired by Ref. [88].

on computing QM energy and forces on highly correlated structures, and only a fraction of the generated trajectories will actually contain relevant information regarding the atomic interactions. Second, systems with large length and/or time scales cannot be simulated using quantum chemical approximations, and training data generation is therefore also very hard if not impossible.

In **Paper II**, we present an iterative learning algorithm which overcomes these limitations. In each iteration, atomic configurations are generated in a massively parallel fashion based on short dynamic simulations using the ML potential from the previous iteration. The sampled configurations are evaluated with QM and appended to training set, after which the ML potential is retrained. Because the ML potential itself is used to explore the phase space of the system(s) under study, we refer to this as an *on-the-fly* learning algorithm. It is computationally efficient in the sense that it only ever performs QM energy and force evaluations on decorrelated structures. In addition, on-the-fly learning guarantees stability of the ML potential over long time scales because it is trained in a self-consistent manner. Finally, we emphasize that




Figure 3.5: Building blocks and framework topology for UiO-66(Zr) and MIL-53(Al). Figure adopted from Ref. [88].

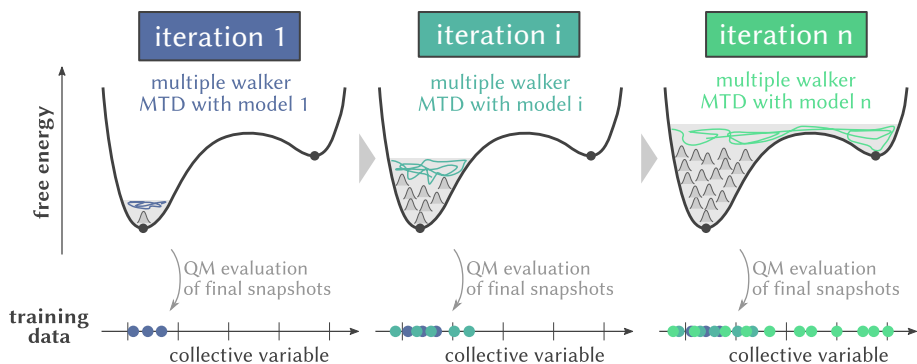


Figure 3.6: Sampled configurations during metadynamics throughout the active learning iterations from Figure 3.4. Suitable biasing forces are able to overcome the free energy barriers and yield a much broader distribution of atomic configurations in the training set. Figure adopted from Ref. [88].

phase space sampling and QM evaluation can be performed in a massively parallel fashion; a schematic overview is given in Figure 3.4.

Paper II demonstrates the efficacy of this algorithm on topical metal-organic frameworks (MOFs) a large class of hybrid organic-inorganic materials which have attracted a lot of attention over the past three decades. Figure 3.5 a, b visualize the mechanically rigid UiO-66(Zr) and the flexible MIL-53(Al) frameworks [51,89]. The unit cells of these materials contain on the order of 200-400 atoms, and direct first principles simulations are therefore highly expensive. To make matters worse, MIL-53(Al) exhibits a reversible phase transition between an open and closed pore phase, and the transition is a rare event during equilibrium simulations. These systems provide an ideal use case for the on-the-fly learning algorithm presented in Figure 3.4 because enhanced sampling algorithms can be readily incorporated in the sampling step. In this particular case, we employ metadynamics using the unit cell

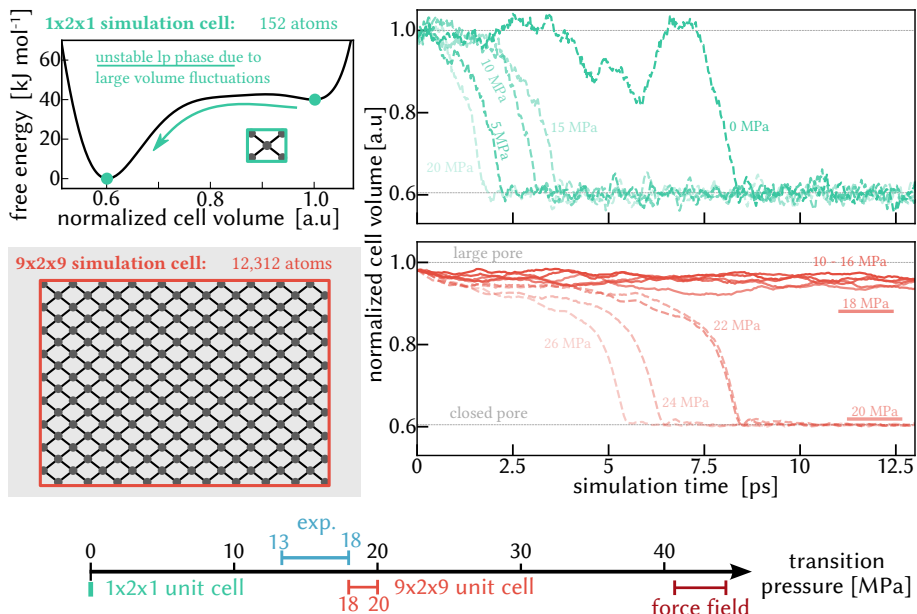


Figure 3.8: We compare estimates using a regular $1 \times 2 \times 1$ unit cell versus a large-scale $9 \times 2 \times 9$ unit cell, as obtained based on NPT dynamics. All unit cell volumes were normalized with respect to the lp minimum of the Helmholtz free energy surface at 300 K, which occurs at 2882 Å. Nevertheless, the lp phase is not stable for the $1 \times 2 \times 1$ cell due to the larger volume fluctuations. In contrast, the $9 \times 2 \times 9$ cell allows to determine a transition pressure in the range of 18–20 MPa, which is in good agreement with experiment [91]. Figure adopted from Ref. [88].

expression for the elasticity tensor is instead given by the so-called Schur complement of the full Hessian matrix [88]:

$$\mathbf{C} = \partial_{\epsilon}^2 E - (\partial_{\epsilon} \partial_r E)^T (\partial_r \partial_r E)^{-1} \partial_r \partial_{\epsilon} E \quad (16)$$

The second term in Equation 16 is ill-conditioned and can only be evaluated properly if the hessian is available with sufficient numerical precision. While traditional quantum chemistry methods are unable to provide this precision, automatic differentiation of the ML potential does allow us to compute \mathbf{C} based on a single Hessian evaluation. Figure 3.7 presents the results for both materials and shows the absolute deviation with respect to the DFT estimate. Note that the DFT results were computed using a finite difference approximation to Equation 14, and they were therefore much more expensive to compute.

3.5 Case Study: Size Effects in Phase Transitions

The ML potentials in **Paper II** were fitted to energy and forces obtained at the PBE-D3(BJ) level of theory. They represent a substantial improvement over classical force fields in terms of accuracy, while being much faster than DFT. An interesting case study which demonstrates how this unique combination of speed and accuracy improves the predictive power of simulations is the structural phase transition of MIL-53(Al), and in particular the estimation of the threshold pressure at which the material transitions from the *lp* to the *cp* phase. Before ML potentials, the mechanical behavior of these materials is modeled via either classical force fields [92] or static DFT approaches [93], and their accuracy is generally limited because of an incomplete description of the atomic interactions and anharmonic effects, respectively. The phase transition in MIL-53(Al) is accompanied by a large change in unit cell volume, and the hydrostatic pressure can be used as an external trigger to simulate the transition in a molecular dynamics simulation. The top panels in Figure 3.8 analyze the behavior of the material when simulated using a $1 \times 2 \times 1$ unit cell. While the required transition pressure is experimentally estimated at 13 - 18 MPa [91], the simulation is observed to transition to the *cp* phase for any nonzero pressure. The reason for this is that for small systems, the statistical fluctuations in the unit cell volume become relatively large with respect to the total change in unit cell volume between both phases. In that case, random fluctuations can trigger a phase transition prematurely, i.e. at a pressure for which it would not be thermodynamically favorable to proceed. As such, a reliable characterization of such transitions requires large-scale dynamics. With the ML potential, we were able to perform much larger simulations of the transition using a $9 \times 2 \times 9$ unit cell, as shown in Figure 3.8. Clearly, the increase in system size effectively dampens the magnitude of the statistical fluctuations by a factor of 9^{13} . As such, we estimated a transition pressure between 18 and 20 MPa, which is in line with earlier experimental estimates. In contrast, earlier work based on classical force fields estimated the transition to occur at >40 MPa [92].

¹³the large cell is 81 times larger than the smaller cell, so the square root of the ratio of the number of particles is 9




Figure 3.9: Translational (in)variance of CP2K output forces for a simple ice crystal. The crystal was globally translated with small increments, and the instantaneous forces were compared with the overall mean force per atom. Computational settings were chosen as to maximize numerical precision; tight SCF convergence (10^{-8} Ha), 6 grids, high plane-wave cutoff (1000 Ry) and relative cutoff (65), and sufficiently large basis set (TZV2P).

3.6 Transfer learning

The predictive power of ML potentials is critically dependent on the accuracy of the reference level of theory that is used to evaluate the target energy and forces. Typical datasets contain a few thousand atomic configurations, each of which is labeled with the total potential energy and the atomic forces. With the proposed on-the-fly learning approach, the total quantum mechanical cost of creating ML potentials therefore amounts to a few thousand energy and force evaluations. If we compare this to the cost of first principles molecular dynamics simulations, it is clear that this already represents an enormous reduction in computational effort; the cost of an entire dataset is equivalent to a single first principles simulation of just a few picoseconds. Nevertheless, QM calculations of thousands of structures remain a challenging task, especially when requiring numerically precise results (to facilitate training) and an accurate level of theory (to improve predictive power). While the energy is in most cases relatively easy to converge with respect to various computational settings (basis set and grid sizes, energy cutoffs, ...), the nuclear gradients are much more troublesome. Analytical expressions for such gradients can be obtained using the Hellman-Feynman theorem [1], but they often tend to be much more numerically unstable than the energy

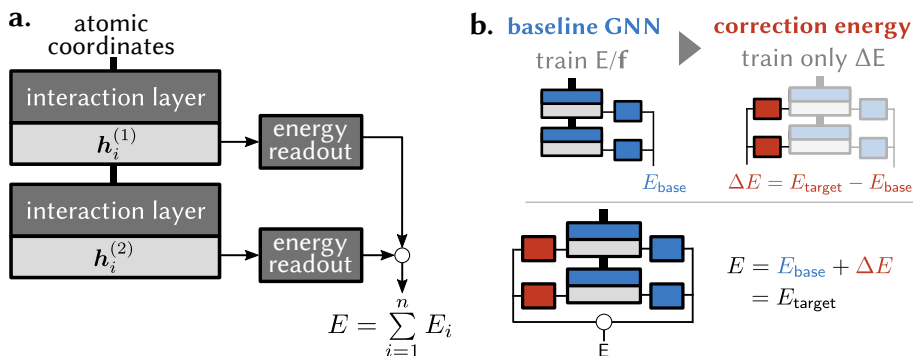


Figure 3.10: **a** overview of the GNN architecture. Most of the computational effort and trainable weights are situated in the interaction layers. **b** Two-step transfer learning approach to learn a simple yet effective $\Delta E(r)$ correction as presented in **Paper III**.

and require significantly more compute resources (runtime and memory). To make matters worse, many software packages do not even implement such analytical gradients for most higher-level methods since they were not really useful before the breakthrough of ML potentials anyway. Instead, many codes adopt a finite-difference approximation; this increases the computational complexity by an additional factor N and generally yields imprecise results. In conclusion, we can state that even with state of the art ML potentials and efficient learning schemes, we can not yet reach the required quantum chemical accuracy in terms of both energy and forces.

Inspired by empirical dispersion corrections to exchange-correlation functionals, it is natural to search for an equivalent “correction” procedure for ML potentials. The differences between lower-level methods such as (GGA) DFT and higher-level methods such as MP2 or CCSD(T) are highly systematic [94], so it is reasonable to assume such corrections can be learned on a system-specific basis with a relatively limited amount of extra data. Because the previous paragraph established that computing precise nuclear gradients represents a major computational bottleneck for high-level methods, it would be beneficial if such corrections can be learned based on energy labels *only*, i.e. based on a small dataset of atomic configurations for which only the potential energy is evaluated – not the interatomic forces. To achieve this, we revisit the high-level architecture of GNNs in Figure 3.10 (a). Because the energy readout functions contain only a small fraction of the total number of weights in the model, it is reasonable to assume that most of the heavy lifting




Figure 3.11: Reaction intermediates during a proton transfer between isobutene and an aluminum Brønsted acid site in chabazite.

is done by the interaction layers. They convert the atomic coordinates into abstract representations for the chemical environment of each of the atoms. Because these representations are by construction sufficient for computing the total interaction energy in the system, we may reasonably assume that they are sufficiently representative to learn a correction energy. Figure 3.10 (b) illustrates this in more detail; the GNN is augmented with a second readout layer that is trained to predict the correction energy $\Delta E(\mathbf{r})$ between the relatively inexpensive *baseline* level of theory and the *target* level. During training, we first optimize the interaction layers and the base readout layer to energy and force labels for the baseline level of theory, since these are by assumption relatively easy to compute. In the second step, we fix the model weights in the interaction layers and optimize only the delta readout layer using a small number of high-level energy-only labels as obtained with the target level of theory. Because the node representations \mathbf{h}_i are kept fixed and the readouts are rather small, these corrections can be learned based on a fraction ($\lesssim 10\%$) of the original dataset size.

Paper III applies this methodology to a relatively complex chemical reaction within an acidic zeolite (Figure 3.11). An isobutene molecule is adsorbed

	PBE-D3	PBE0-D3	MP2	RPA	CCSD(T) [†]
	[kJ/mol]	[kJ/mol]	[kJ/mol]	[kJ/mol]	[kJ/mol]
isobutene → cation	24.3	31.7	56.1	62.3	61.6
isobutene → alkoxide	50.9	46.4	27.2	36.5	46.4

Table 3.1: Relative energy difference of each intermediate with respect to the neutral isobutene reactant. Values were obtained from Ref. [95]. The CCSD(T) entry was obtained with an ad hoc cluster correction approach as described in Ref. [96].

into the pores of a zeolite which has an aluminum Brønsted acid site. The charge-compensating proton can transfer towards the molecule to form a charged molecular cation, leaving the framework oxygen with a negative charge. Due to electrostatic attraction between both, the initially mobile cation can again adsorb via its central carbon atom to form a surface alkoxide species. Table 3.1 illustrates how (hybrid) DFT overstabilizes the delocalized charge density of the cation, whereas post-HF methods such as MP2 or RPA can largely fix this deficiency. In the particular case of **Paper III**, we use PBE-D3(BJ) as a baseline and RPA as high-level target (see Table 3.1). The ML potential is constructed in two major steps. First, the baseline GNN is trained to a dataset of atomic configurations with both energy and force labels at the PBE-D3(BJ) level of theory. These are obtained using the on-the-fly learning approach as presented in Section 3.4. Because there exist large free energy barriers between the reaction intermediates, it is necessary to complement the phase space sampling step in the on-the-fly learning loop with biasing forces defined along the appropriate collective variables. Figure 3.12 shows how this can be achieved for the two stages of the reaction separately. A linear combination of both CVs can be used to define a global reaction coordinate for the entire transition (from isobutene to the alkoxide):

$$CV = \cos(0.148\pi)CV_1 + \sin(0.148\pi)CV_2 \quad (17)$$

To guarantee high fidelity of the ML potential across both reactions, it is necessary that the training data contains both the low and high energy states of the transition. Because transition regions are sometimes accompanied with relatively sharp changes in free energy, we choose to perform on-the-fly learning with *umbrella sampling* because it allows us to enforce an approximately uniform sampling in the CV defined in Equation 17. For this reaction, we observed that to achieve sub-meV/atom energy RMSEs, it is

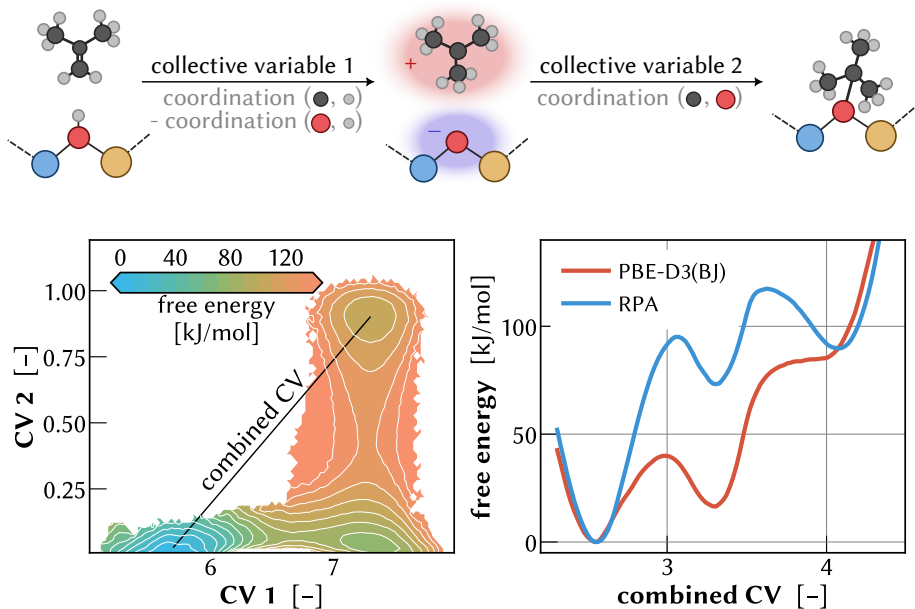


Figure 3.12: Collective variable (CV) definition for the two stages in the reaction from Figure 3.11. The reaction free energy surface $F(CV_1, CV_2)$ at the RPA level of theory is computed using two-dimensional umbrella sampling. A one-dimensional projection is computed along a linear combination of CV 1 and CV 2, for both PBE-D3(BJ) and RPA. Clearly, the inclusion of dynamic electron correlation effects is necessary to correctly predict the relative stability of the cation intermediate.

necessary to sample and evaluate a few thousand atomic geometries at the PBE-D3(BJ) level.

Next, we follow the approach presented in Figure 3.10 and add an additional energy readout to the model which is trained to parameterize the energy difference ΔE between RPA and PBE-D3(BJ). Importantly, this second readout layer is strongly limited in capacity and contains roughly 1000 weights, or about 1% relative to the total number of weights in the model. As such, this allows us to train these weights using only the RPA energy as targets – the forces are not required. In addition, we observe that we require only a few hundred RPA energy evaluations in order to achieve energy RMSEs on a 623 K test set below 1 meV/atom. This highlights the exceptional data efficiency of the approach, and demonstrates its viability on complex systems. Figure 3.12 presents the obtained two-dimensional free energy surface at the RPA level of theory (left) and contrasts both levels




Figure 3.13: Transfer-learned atomic correction energies between PBE-D3(BJ) and RPA, for the three intermediates. The model implicitly learned that both levels of theory only really differ in their description of the active site.

of theory in terms of their one-dimensional free energy profile (right). The later profile confirms the discrepancy in zero-kelvin stabilities as presented in Table 3.1.

3.7 Conclusions

Over the past decade, a wide range of powerful ML models have emerged which can be used to predict the interaction energy $E(\mathbf{r})$ of arbitrarily complex materials and geometries. While these models are orders of magnitude faster than any quantum chemical approximation, they require sufficient training data to be able to give accurate predictions on any given system. Generation of such training data can be difficult because for systems with large configurational flexibility or which exhibit multistable behavior. **Paper II** presents an on-the-fly learning approach which relies on enhanced sampling and massive parallelization to generate the training data in a computationally efficient manner. In systems for which post-HF-level accuracy is needed, such approaches can fall short because training data requires QM forces and those are not always available in a numerically stable manner. For this reason, we developed a transfer learning approach which allows for efficient ML fine-tuning without requiring numerically accurate forces at the target level of theory – see **Paper III**.

The success of transfer learning in these systems suggests that GNNs implicitly learn powerful representations for the chemical environments that are present in the training data. Figure 3.13 shows how the correction ΔE

is localized precisely at the active site where the stability of the delocalized charge density causes GGA DFT and RPA to diverge. Indeed, Figure 3.3 also demonstrates how the atomic energies E_i represented a physically reasonable partition of the interaction energy. These observations provide empirical evidence for the quality of the learned GNN representations, and we will continue to investigate and exploit this fact in the next Chapter.

4

Learned Rare Event Sampling

The combination of ML potentials and powerful learning schemes have induced a revolution in terms of the attainable system sizes, the accuracy of the interaction energy, and the time scale of atomic simulations. As such, they are rapidly being adopted throughout materials science and chemistry research communities. However, physical and chemical processes such as crystal nucleation, protein folding, large chemical reaction networks, and solid-state transformations in complex solids remain incredibly hard to characterize with dynamic simulations. The key challenge is to overcome the disparity between very fast local atomic motions (which determine the time step) and the exponentially low probability of observing a transition. Even if the computational cost of ML potentials could be further reduced by one or even two orders of magnitude, direct equilibrium simulation of such rare events remains infeasible and specialized enhanced sampling methods are necessary [30,31]. Such methods typically construct a series of overlapping intermediate distributions p_i which bridge the disjoint regions of the stable minima, but their application usually requires specialized knowledge on the system under study. Having observed the power of GNNs in predicting $E(\mathbf{r})$, it is natural to investigate how GNNs can be exploited to overcome free energy barriers and accelerate the simulation of rare events.

4.1 Learned chemical representations

GNNs learn atomic representations \mathbf{h}_i which are used to predict node- and graph-level physical properties. For example, the total interaction energy $E(\mathbf{r})$ is predicted as a sum over per-atom contributions $E_i = \mathcal{R}(\mathbf{h}_i)$. The

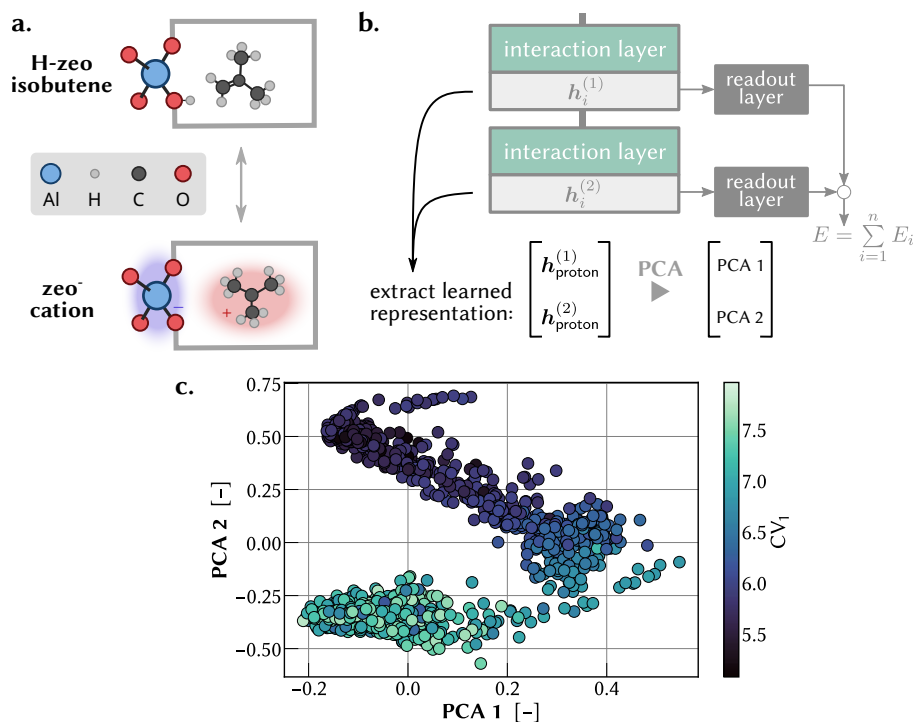


Figure 4.1: **a** schematic illustration of the proton transfer reaction from Figure 3.12. **b** GNNs produce (invariant) learned representations h_i which can be extracted for each layer. **c** visualization of the learned representation for the transferring proton over a dataset obtained by umbrella sampling. The colorcoding is achieved using CV_1 as defined in Figure 3.12 (dark green corresponds to isobutene, light green corresponds to cation).

interaction energy $E(\mathbf{r})$ and its gradients define the dynamics and equilibrium statistics of the system. As such, we can infer that these learned representations h_i contain everything there is to know about the interactions surrounding atom i , since they are by construction sufficient to predict the interaction energy. In other words, h_i are *informative* representations of the chemical environment of each atom.

To demonstrate this explicitly, we consider the learned representations of the transferring proton from the isobutene reaction considered in Figure 3.11. From the perspective of the transferring proton, the isobutene and cation state represent radically different situations (bonded to oxygen and neutral versus bonded to carbon and positive). As such, we can anticipate that they are embedded into different clusters in feature space, and that they will

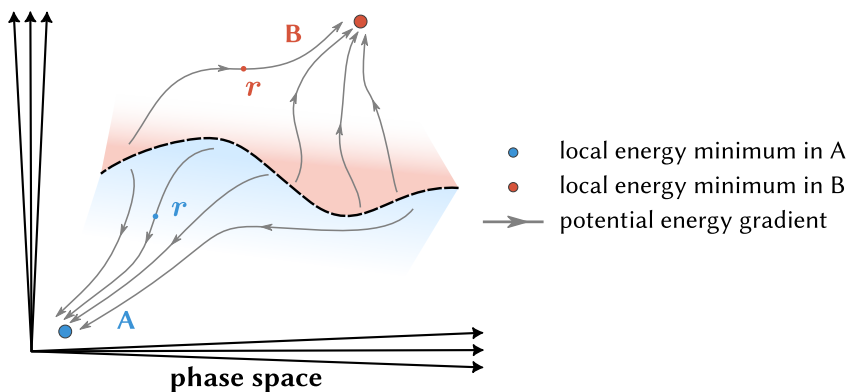


Figure 4.2: Schematic illustration of two basins of attraction A and B, in line with the definition in Equation 1.

generally also yield different atomic energies. By construction, this needs to be reflected in the atomic representation h_{proton} . This is effectively a form of *unsupervised* learning since the model was never explicitly taught which atomic configurations represent which chemical environment for the proton – the model learns this information implicitly from the data as a byproduct of minimizing the energy and forces RMSEs. In Figure 4.1, we investigate how the chemical environment of the transferring proton is embedded in this feature space, using principal component analysis (PCA) for visualization purposes and CV_1 as defined in Figure 3.12 to distinguish between the neutral and cation states. Clearly, atomic configurations which correspond to a similar CV value are also in each other’s proximity in the feature space, indicating that this is very much a physically-informed representation.

Chapter 2 discussed how the existence of large (free) energy barriers between disjoint stable regions in phase space are hard to traverse during dynamic simulations. In Chapter 3, we overcame such limitations on a system-specific basis by including biasing forces along manually defined *collective variables*. As alluded to earlier, the definition of suitable variables for any given rare event requires specialized knowledge about the specific rare event, and is commonly based on human intuition and some amount of trial and error [97]. Figure 4.1 suggests that the learned representations of the GNN can be used to discriminate between different classes of atomic environments in a continuously differentiable manner, thus being able to function as effective reaction coordinates. **Paper IV** presents smooth basin classification (SBC); a new framework which employs the feature represen-

tation of GNN-based interaction potentials to learn reaction coordinates [98]. In what follows, we present the formal framework on which SBC is built, and demonstrate its extraordinary data efficiency based on a number of examples.

4.2 Basins of attraction

The *phase* of an atomic geometry \mathbf{r} is defined as a discrete, categorical property. For any geometry \mathbf{r} , we say that \mathbf{r} belongs to phase A whenever a gradient-descent optimization starting at \mathbf{r} ends in any of the local minima that are associated with phase A¹⁴:

$$\mathbf{r} \in \text{phase A} \iff \mathbf{r} \text{ optimizes to basin A} \quad (1)$$

Note that this definition assigns each point \mathbf{r} to precisely one phase. Furthermore, it is purely¹⁵ a function of the interaction energy $E(\mathbf{r})$, which means that it is universally applicable and requires no system-specific hyperparameters to be set, in contrast to committor functions. The downside is that our definition is discrete: a geometry \mathbf{r} belongs entirely to precisely one phase, even if it is located at or near a transition state. Such discrete behavior seems undesirable given that our goal is to learn smooth, continuously differentiable functions which can be used to bias the dynamics. Fortunately, we can use ML to convert Equation 1 into a smoothly differentiable but discriminating quantity. The prediction of categorical variables is referred to as a *classification* task [60]. In our case, the task consists of predicting the phase label \mathcal{C} which is associated to a given geometry \mathbf{r} . Even though the targets are discrete class labels, classification models achieve this by outputting continuous real numbers, which are to be interpreted as *log probabilities* or logits $\mathbf{q} \in \mathbb{R}^{N_p}$ with N_p the number of classes/phases in the system. The logits \mathbf{q} are interpreted as the log likelihood of each class given an input \mathbf{r} as assigned by the model. The actual predicted class \mathcal{C} is, naturally, the one which receives the highest (log) likelihood:

$$\mathcal{C} = \text{argmax}(\mathbf{q}) \quad (2)$$

¹⁴In this definition, we tacitly assume that each free energy minimum contains at least one local potential energy minimum. Although this is not true in general, it is true in virtually all practical cases of interest.

¹⁵In theory, it is possible that different (gradient-based) optimizers assign different phases to a given geometry because the optimizations trace different paths. In practice, this has never been observed.

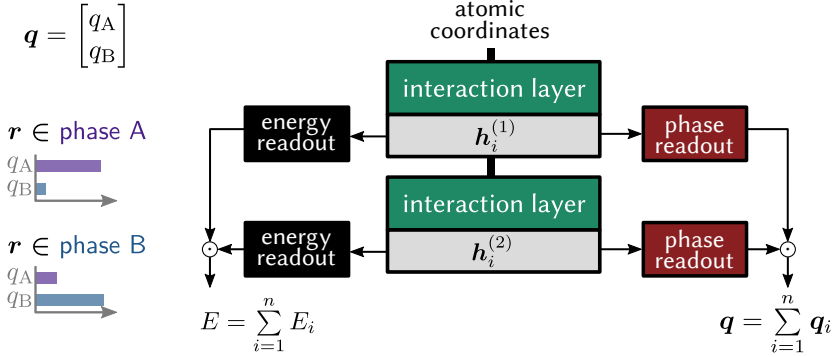


Figure 4.3: Definition of q as the log likelihood of an input geometry r belongs to a particular phase. The logits are computed as a sum over per-atom contributions as to preserve permutation invariance. Figure adopted with permission from [98].

In effect, each logit encodes how likely the model thinks a given sample r belongs to each class. As suggested by Figure 4.1, we will augment the GNN used for predicting $E(r)$ with an additional readout function for q ; see Figure 4.3. To guarantee q to be predicted as a global property that is permutation-invariant, we deconstruct q as a sum over per-atom logits q_i :

$$q = \sum_i q_i \quad (3)$$

Note that in this architecture, both the interaction energy $E(r)$ as well as the phase log likelihoods q are predicted based on the *same* learned feature representations h_i . This makes sense given that the two ‘classes’ discussed in Figure 4.1 (carbonyl and hydroxyl oxygen) are in fact linearly separable in this feature space. In addition, phase predictions are subject to the same physical symmetries as the prediction of the total energy, and by embedding in the learned feature space we are effectively guaranteed that these symmetries are rigorously satisfied (by construction of the GNN, see Chapter 3).

The logits q are continuously differentiable with respect to the GNN inputs – i.e. the atomic coordinates and optionally the unit cell components – even though they are used to predict a discrete class label as in Equation 2. However, they are not immediately usable as collective variables for driving transitions between phases or characterizing the relative stability ΔF . To this end, consider the difference of two logits:

$$\Delta q(r) = q_B(r) - q_A(r) \quad (4)$$

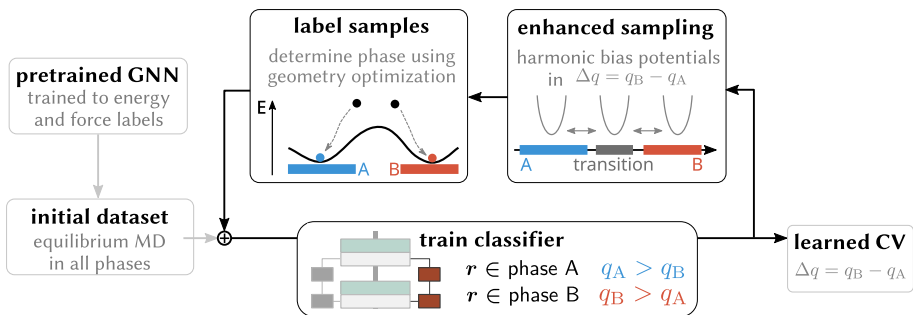


Figure 4.4: Active learning methodology for generating out-of-equilibrium samples along the transition path(s). Figure adopted with permission from [98].

for a system with two phases A and B. Effectively, Δq measures how much more likely the model thinks a given structure belongs to phase B rather than phase A. It represents a one-dimensional variable which is (trained to be) strictly negative whenever \mathbf{r} belongs to phase A, and strictly positive whenever it belongs to phase B, which means it can effectively function as a discriminator. At the same time, q_A and q_B are smooth, differentiable quantities with respect to the atomic coordinates. As such, we obtain that Δq is both smoothly differentiable and discriminating for phases A and B, which makes it a collective variable for transitions between phases A and B.

4.3 Architecture & training

The learned feature representations of a GNN-based interaction potential are by construction sufficiently expressive as to predict the total interaction energy and the interatomic forces. Because Equation 1 defines the phase of the system as a discrete quantity that depends directly on interaction energy $E(\mathbf{r})$, it makes sense to consider a joint prediction of both the energy and the phase using a model similar to Figure 4.3. As such, it is reasonable to assume that the phase readout layers can be constructed in a similar manner to the energy readouts – a small and shallow dense network with one or two nonlinearities.

In Figure 4.4, we demonstrate an efficient approach to training the augmented GNN from Figure 4.3. It starts from a pretrained ML potential (e.g. MACE-MP0) and a small initial dataset of equilibrium configurations in the stable minima of interest – i.e. no initial transition paths are required. In each

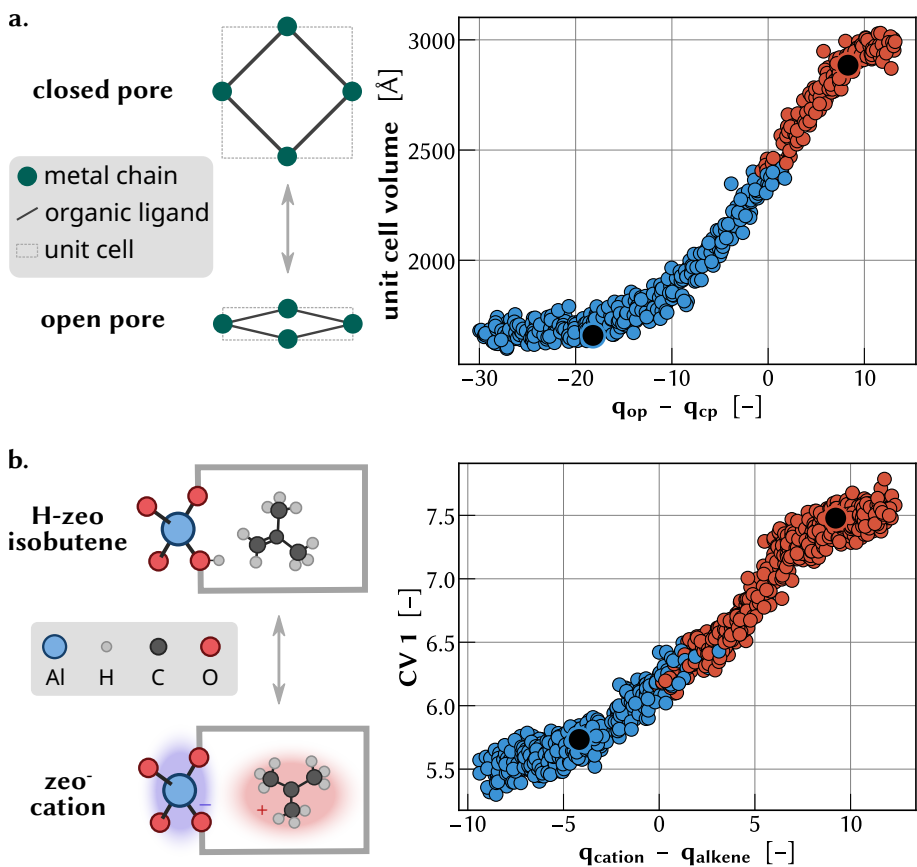


Figure 4.5: **a** application of SBC on the structural phase transition in MIL-53(Al). The transition is anisotropic and induces large volume changes in the unit cell. The scatter plot shows the value of the learned CV Δq versus a reference coordinate (the unit cell volume). Points are colored according to their phase; the black dots indicate the optimized geometries. **b** application of SBC for the protonation of isobutene in chabazite. The reference CV is the coordination number as defined in Figure 3.12. Figure adopted with permission from Ref. [98].

iteration, we use Δq (Equation 4) as biasing coordinate to perform enhanced sampling and discover new states which lie in between both phases. The sampled states are labeled with the correct phase as determined from an energy minimization, and are appended to the dataset. Training of the phase readouts proceeds along the ML paradigm of classification learning, with a cross-entropy loss term on the logits \mathbf{q} as well as a regularization term on the per-atom logits \mathbf{q}_i to prevent the model from making overconfident predictions from single atomic environments (see **Paper IV** for more details)

	training cost	
	snapshots	time [ns]
time-lagged AE [108]		750 ns
FEBILAE [112]		1500 ns
VAMPnets [109]	250,000	250 ns
LINES [106]		100 ns
LED [113])		100 ns
Ref. [114]	>300,000	
Deep-LDA [100]	10,000	
Deep-TDA [115]		40 ns
Ref. [116]	15,000	15 ns
Ref. [117]	8,000	512 ns
SBC	3,694	11 ns

Table 4.1: Performance summary of learned collective variable methods for alanine dipeptide. SBC achieves SOTA in terms of both the required number of snapshots as well as the total simulation time. See Table 1 in **Paper IV** for a more extensive comparison between SBC and existing approaches in literature.

[60,98]. As shown in **Paper IV**, this approach allows the model to discover and learn the transition region in between both phases in a highly efficient manner. In Figure 4.5, we present the application of SBC to model the structural flexibility of MIL-53(Al) as well as the proton transfer reaction as discussed in the previous Chapter. In both cases, we compare the learned logit difference Δq with a reference collective variable which is known to be a continuously differentiable discriminator for the relevant transition. Importantly, these results were obtained *without* any system-specific knowledge; the model learned correct reaction coordinates without predefined input features or transition paths, unlike most approaches in literature do [99–111].

4.4 Conclusion

We present a general framework for augmenting GNNs with knowledge about basins of attraction. By training the model to associate the correct basin of attraction with each atomic geometry, the model learns a continuously differentiable mapping from the feature representation that emerges from the interaction layers to a set of log likelihoods for each basin. The

difference between two log likelihoods quantifies how much more likely the model believes that a given atomic geometry belongs to one basin rather than another. This quantity can be used as collective variable by definition since it is discriminating and continuously differentiable. In this way, SBC is one of the first methods to achieve exact translation, rotation, and permutation invariance by construction [118], and it is *the* first method to demonstrate the efficacy of permutation invariance on reactive and conformational rare events. In addition, it benefits from years of development of GNN architectures as efficient approximators for the interaction energy $E(\mathbf{r})$, including the use of directional message passing and equivariant feature representations. In **Paper IV**, we demonstrate explicitly how these properties make SBC the most data-efficient method that is currently available in literature. We do this by learning and characterizing the conformational flexibility in alanine dipeptide, a common benchmark system in this field. Table 4.1 shows how SBC outperforms any other method in terms of the total required simulation time needed to train the model. In addition, many of the methods in Table 4.1 require the input of manually selected features which are often not permutation-invariant, and which implicitly require human intuition much like the manual definition of a reaction coordinate.

5

psiflow: a Modular and Scalable Simulation Framework

Chapters 3 & 4 demonstrate that clever usage of machine learning reduces the computational cost of molecular simulations by many orders of magnitude. The biggest downside is that the complexity of common simulation pipelines increases dramatically when employing more advanced methodologies such as the ones presented in this work. For example, on-the-fly learning usually requires 5–20 iterations, each of which constitutes a mixed workload of QM evaluation (CPU), model training (GPU), and sampling (CPU/GPU). Transfer learning requires the model and dataset to be constructed in multiple stages, with high-level QM calculations sometimes requiring specialized high-memory nodes. While the power of these methods is undisputable, it has become apparent that the human effort required to execute such workflows must not be underestimated and in fact represents a significant barrier towards widespread adoption. It is important to note that the human effort consists mostly out of administrative, dull tasks; creating submission scripts and working directories, waiting for resource allocations to start and calculations to complete before starting new ones, grouping tasks into larger allocations, processing generated data, resubmitting calculations in case of failures, ... These tasks do not require the full attention of a researcher and are – at least in theory – automatable. In other words, the research from Chapters 3 & 4 has created an engineering problem: how can we execute a large variety of complex molecular simulation pipelines with minimal user intervention? To solve this problem, we built *psiflow*.



Figure 5.1: a scalable molecular simulation framework; MIT licensed. Visit github.com/molmod/psiflow.

5.1 Overview

Psiflow is a scalable molecular simulation engine for chemistry and materials science applications. It supports:

- QM calculations at various levels of theory, including GGA and hybrid DFT as well as post-HF methods such as MP2, RPA, and CCSD(T). Calculations are executed using CP2K [119], GPAW [120], or ORCA [23].
- trainable interaction potentials as well as easy-to-use universal potentials, e.g. MACE-MP0 [121].
- a wide range of sampling algorithms: constant energy/temperature/pressure, path-integral molecular dynamics, alchemical replica exchange, metadynamics, phonon-based sampling, thermodynamic integration; using i-PI [122], PLUMED [123], ...

Users may define arbitrarily complex workflows and execute them automatically on local, HPC, and/or cloud infrastructure. To achieve this, psiflow is built using Parsl [124]: a parallel execution library which manages job submission and workload distribution. As such, psiflow can orchestrate large molecular simulation pipelines on hundreds or even thousands of nodes. In the remainder of this Chapter, we will outline the architecture and building blocks of psiflow. More information can be found in the documentation¹⁶.

5.2 Asynchronous execution

Consider the following hypothetical scenario: imagine that we have a trajectory of atomic geometries, and we wish to minimize each of their potential

¹⁶see <https://molmod.github.io/psiflow>

energies and inspect the result (for example to find the global minimum). In pseudo-code, this would look something like this:

```
minima = []
for state in trajectory:
    local_minimum = geometry_optimization(state)
    minima.append(local_minimum)
global_minimum(*minima)
```

In regular, *synchronous* execution, when entering the first iteration of the loop, Python would start executing the first geometry optimization right away and *wait* for it complete, before moving on to the next iteration. This makes little sense, since we know in advance that each of the optimizations is independent. As such, the loop can in fact be completed much quicker if we would simply execute each optimization in parallel. This is referred to as *asynchronous* execution because the execution of any optimization will now be performed independent from another. The intended way to achieve this in Python is by using the built-in `concurrent`¹⁷ library, and it provides the foundation of psiflow's efficiency and scalability. Aside from *asynchronous* execution, we also want *remote* execution. Geometry optimizations, molecular dynamics simulations, ML training, quantum chemistry calculations, ... are rarely ever performed on a local workstation. Instead, they should ideally be forwarded to e.g. a SLURM/PBS(Pro)/SGE scheduler or an Amazon/Google cloud instance. To achieve this, psiflow is built with Parsl: a scientific Python library supported by the US Department of Energy which extends the native `python.concurrent` library with the ability to offload execution towards remote compute resources [124].

Parsl (and `python.concurrent`) is founded on two concepts: apps and futures. In their simplest form, apps are just functions, and futures are the result of an app given a set of inputs. Importantly, a Future already exists before the actual calculation is performed. In essence, a Future *promises* that, at some time in the future, it will contain the actual result of the function evaluation. Take a look at the following example:

¹⁷<https://docs.python.org/3/library/concurrent.futures.html>

```

from parsl.app.app import python_app

@python_app
def sum_integers(a, b):
    return a + b

sum_future = sum_integers(3, 4) # AppFuture, not an int!
sum_future.result()             # wait for the actual result (7)!

```

The return value of Parsl apps is not the actual result (in this case, an integer), but an AppFuture that will store the result of the function evaluation after it has completed. By explicitly calling `.result()` on the future, we block the main code execution until the sum is actually computed, and literally ask for the result. For more information, check out the Parsl documentation¹⁸.

In our geometry optimization example from before, we would implement the function `geometry_optimization` as a Parsl app, and its return value `final` would no longer represent the actual optimized geometry; it would be a future of the optimized geometry. Importantly, when organized in this way, Python will go through the loop almost instantaneously, observe that we want to perform a number of `geometry_optimization` calculations, and offload those calculations separately, independently, and immediately to whatever compute resource it has available. As such, all optimizations will effectively run in parallel:

```

minima = [] # list which tracks futures of the optimizations
for state in trajectory:
    local_minimum = geometry_optimization(state) # future!
    minima.append(local_minimum)
global_minimum = find_lowest(*minima) # again a future
# all optimizations are running at this point in the code.

global_minimum.result() # waits until all optimizations completed

```

¹⁸<https://parsl.readthedocs.io/en/stable/>

5.3 Hamiltonians

In Born-Oppenheimer-based molecular simulation, atomic nuclei are treated as classical particles that are subject to effective interactions $E(\mathbf{r})$. These interactions determine the interatomic forces which are used in a dynamic simulation to propagate the atomic positions from one timestep to the next. The `Hamiltonian` class is used to represent any type of (conservative) energy contribution to the hamiltonian of the system. The most prominent example are ML potentials; either pre-trained, ‘universal’ models such as MACE-MP0 [121] or system-specific ML potentials trained within psiflow. Other examples include harmonic approximations to $E(\mathbf{r})$ in order to describe the interaction energy close to a local energy minimum. In addition, various sampling schemes employ bias potentials which are superimposed on $E(\mathbf{r})$ in order to drive the system along specific reaction coordinates (e.g. metadynamics, umbrella sampling). Such bias potentials are also constructed as `Hamiltonian` instances.

As mentioned earlier, the key feature of hamiltonians is that they represent an interaction energy between atoms, i.e. they output an energy (and its gradients) when given a geometry as input. Because hamiltonians might require specialized resources for their evaluation (e.g. an ML potential which gets executed on a GPU), evaluation of a hamiltonian does not necessarily happen instantly (e.g. if a GPU node is not immediately available). Similar to how geometry optimizations returned a future of a `Geometry` which represents the final result, hamiltonians return a future when asked to evaluate the energy/forces/stress of a particular `Geometry`.

A unique feature of psiflow `Hamiltonian` instances is the ability to create a new hamiltonian from a linear combination of two or more existing hamiltonians:

$$H = \sum_i \alpha_i H_i \quad (1)$$

In psiflow, it is possible to define such ‘combined’ hamiltonians very easily:

```
# assume H_1, H_2, H_3 are defined
H = 0.1 * H_1 + 1.6 * H_2 - 0.7 * H_3
```

This is possible irrespective of the specific nature of each H_i and its execution requirements. Large MACE models which require GPUs for evaluation can be semantically added to simple harmonic bias potentials which are evaluated locally on CPU. Again, resource allocation and task execution are automatically handled by Parsl. Hamiltonian arithmetic is incredibly powerful in the context of free energy calculations because algorithms such as hamiltonian replica exchange, thermodynamic integration, and umbrella sampling can be expressed in an elegant manner. Section 5.5 discusses this in more detail.

5.4 Sampling

In the Born-Oppenheimer philosophy, we explore the phase space of a molecule or a material and generate samples using molecular dynamics simulations. Those samples are then used to evaluate time averages of some property of interest in order to predict physical observables. In `psiflow`, such simulations are executed within `i-PI`¹⁹, a versatile and efficient code which supports an impressive number of features. We mention the most important ones below

- **molecular dynamics in various ensembles**: most notably NVE, NVT, and fully anisotropic NPT. There exist a variety of thermostat and barostat options, the default being Langevin. Together with the ability to combine arbitrary hamiltonians, this also includes biased molecular dynamics simulations using e.g. harmonic restraints (umbrella sampling).
- **path-integral molecular dynamics (PIMD)**: allows for the simulation of the quantum behavior of light atomic nuclei. This is important for many systems involving hydrogen atoms at relatively low temperatures (\leq room temperature). Importantly, these simulations can also be performed in any of the aforementioned ensembles.
- **geometry optimizations**: `i-PI` can be used to optimize the geometry of a molecule or a material using a variety of optimization algorithms.
- **replica exchange** (parallel tempering): replica exchange can dramatically improve the sampling efficiency and ergodicity whenever nontrivial free energy barriers are present in the phase space of the system. In this

¹⁹<https://ipi-code.org/>

approach, one considers replicas of the system at various temperatures and/or pressures, and optionally even with different hamiltonians.

- **(multiple walker) metadynamics:** simple but powerful method to overcome free energy barriers when a suitable collective variable is known for the system of interest.

In psiflow, molecular dynamics simulations are performed using the `Walker` class. Walkers define the thermodynamic ensemble in which a given system should be embedded, which includes a full definition of the interaction energy, its boundary conditions (isolated, constant temperature, ...), and any potential couplings to other walkers (as is the case for replica exchange or multiple-walker algorithms).

5.5 Free energy calculations

In the context of free energy calculations, it is very powerful to be able to represent hamiltonians symbolically and do simple arithmetic with them. It allows to express sampling algorithms in code that resembles the mathematical formalisms behind them. We demonstrate this based on two examples:

- **thermodynamic integration:** many applications of thermodynamic integration aim to determine the free energy F of a given system with hamiltonian H based on a reference hamiltonian H_0 of which the free energy F_0 can be determined analytically. In many cases, the reference hamiltonian represents a (high-dimensional) quadratic interaction energy, because it often provides a good approximation to H and is analytically solvable. In those cases, it is necessary to consider interpolations of the interaction energy between H and H_0 :

```
harmonic = Harmonic(...) # H_0
mace = MACE(...)        # H

for alpha in np.linspace(0, 1, 10): # H_0 → H in 10 steps
    H = harmonic + alpha * (mace - harmonic)
```

- **umbrella sampling:** let us again consider the (by now familiar) example of a proton exchange reaction in a formic acid dimer. A series of biased simulations is performed in which the QM interaction energy $E(\mathbf{r})$ is biased along a well-chosen reaction coordinate q using harmonic restraints:

$$E_q(\mathbf{r}) = E(\mathbf{r}) + \frac{K}{2}(q(\mathbf{r}) - q_0)^2 \quad (2)$$

Because `Hamiltonian` instances support basic arithmetic, we can construct $E_q(\mathbf{r})$ as a sum of a MACE hamiltonian and a bias hamiltonian precisely as in Equation 2:

```
bias = PlumedHamiltonian(...) # simple PLUMED bias
mace = MACE(...)             # e.g. from pretrained model

potential = mace + bias      # represents biased PES!
```

without having to worry about specific implementation details.

In addition, it is very easy to assemble all walkers from a given free energy calculation technique and perform on-the-fly learning with them as to ensure that the ML potential is properly trained in all regions of phase space that will be sampled.

5.6 Conclusions

The adoption of ML methods in molecular modeling has increased the complexity of many simulation pipelines. For example, on-the-fly learning algorithms contain a large number of tasks which require a mixture of CPU and GPU compute nodes. The use of post-HF-level training data, for example using transfer learning as in **Paper III**, additionally requires somewhat more specialized CPU hardware with large amounts of memory and fast local node storage. **Paper IV** presents a powerful active learning strategy for discovering transition paths but it does introduce significantly more human effort. In addition, the use of ML potentials always entails certain risks related to overfitting and incorrect extrapolation, and robust validation of production runs is usually a good idea. Psiflow was originally built as a computational tool which streamlines and facilitates these types of calculations. However, the asynchronous execution paradigm appears to be quite powerful in computational sciences, and we believe it has great potential even apart from the research in Chapters 3 & 4. In particular for free energy calculations, we anticipate that psiflow has the ability to relieve much of the technical burden and make the techniques accessible to a much larger audience.

6

Conclusions and Perspectives

While the physical laws which govern the behavior of electrons and nuclei are known, it takes more than that to predict the properties of materials and molecules with computer simulations. Chapter 1 presents a concise overview that explains how the theoretical laws of quantum and statistical mechanics give rise to a computational framework for performing molecular simulation (Figure 1.4). Chapter 2 highlights the two main computational challenges that are associated with their application to complex molecular systems, namely the computation of the interaction energy $E(\mathbf{r})$ and the need for ergodic sampling of the Boltzmann distribution $p(\mathbf{r})$. Conventional strategies which aim to overcome these challenges include classical force fields and enhanced sampling algorithms, and they have been adopted in various branches of physics, chemistry, and materials science. They also laid the basis for the first part of the research in this work; **Paper I** presents a computational study on the phase transition dynamics in a metal-organic framework using classical force fields and metadynamics. Its key methodological innovation is the development and implementation of an anisotropic pressure control algorithm that does not require computation of the virial stress. It allowed us to perform the largest dynamic simulations of metal-organic frameworks to date, and investigate the size-dependent dynamics of a structural phase transition.

Deep learning has been a major driver for much of the computational research over the past five years, and this thesis is no exception. The computational challenges outlined in Chapter 2 can be tackled by neural networks, and in particular graph neural networks. Fast and accurate prediction of the

interaction energy $E(\mathbf{r})$ can be achieved by training GNNs to QM reference data, but there are two main challenges associated with this. First, DFT is often not a very accurate level of theory, and in some cases leads to qualitatively wrong predictions. Second, it is nontrivial to generate training samples in complex phase spaces in which sampling is highly nonergodic (e.g. slow modes and large free energy barriers). **Paper II** presents a massively parallel on-the-fly learning strategy to generate training data in an efficient manner, whereas **Paper III** develops an approach in which we increase the accuracy of DFT-level models using transfer learning. In both cases, we demonstrate the efficacy and scalability of the methods to nanoporous systems; metal-organic frameworks and Brønsted-acidic zeolites. We showcased how the methods continue to work in complex phase spaces with multiple minima and/or nonnegligible entropic contributions, with case studies on the structural phase transition in MIL-53(Al) and the protonation/alkoxide formation of isobutene in chabazite. Aside from predicting interaction energies, neural networks have also been employed in improving the ergodicity of molecular dynamics simulations and stimulate transitions between local free energy minima. A particularly attractive line of research which has gathered a lot of attention is to learn reaction coordinates in order to drive transitions between multiple metastable minima. In **Paper IV**, we present smooth basin classification (SBC); a new framework for construction such collective variables based on the internal features GNN-based ML potentials. Importantly, SBC exceeds the state of the art in terms of data efficiency for alanine dipeptide, learns permutationally invariant collective variables by construction, and works for both reactive and nonreactive transitions. When combined, **Papers II-IV** constitute an integrated machine learning approach to tackle the computational challenges outlined in Chapter 2 for complex molecular systems.

In that sense, the work in Chapters 3 & 4 still deserves further investigation regarding the transferability of the proposed approaches. An interesting case study would be to study hydrocarbon conversion in Brønsted-acidic zeolites as a whole, i.e. including a large number of zeolite topologies and a complete enumeration of all hydrocarbons up to a given size. In theory, it would be possible to train high-level energy corrections simultaneously with a large classifier that is able to recognize all molecules and reaction intermediates in any of the pores. In this way, we can maximally benefit from the trans-

ferability of the learned feature representations across systems with similar chemical composition.

The generation of accurate training data requires accurate and computationally feasible QM methods which are applicable to complex molecular systems. In the vast majority of real-world scenarios, we are still unable to estimate the interaction energy $E(r)$ to within 1 kJ/mol of its true value. Some macroscopic properties of interest (reaction rates, adsorption isotherms) are highly sensitive to the interaction energy and therefore require such precision. Quantum chemistry methods struggle to capture static electron correlation effects, they are generally too expensive for systems of realistic size, and they require very expensive hardware (high-memory nodes). Post-HF calculations such as RPA or CCSD(T) are incredibly inconvenient in practice because very few HPCs can afford the required hardware. To move forward, it is absolutely essential to develop alternatives to these methods. For example, quantum Monte Carlo methods that use neural networks as wave functions have emerged as a powerful and more scalable alternative [12,14].

On a more abstract level, it is instructive to analyze the primary drivers of the research in this thesis – i.e. the key enabling technologies and formalisms that were required to develop and/or execute the work. Upon inspection, we find that all four papers find their origins in recent advances in theoretical and applied computer science. **Paper I** succeeds in visualizing the large-scale dynamics of anisotropic phase transitions thanks to massive hardware-level parallelization (GPU computing started around early 2010s). **Paper II** succeeds in simulating the mechanical behavior of metal-organic frameworks at DFT-level accuracy thanks to efficient graph neural networks (which were conceived around 2017 [125]). **Paper III** succeeds in computing reaction enthalpy and entropies at post-HF accuracy thanks to transfer learning (which gained much attention with large pretrained image models based on e.g. ILSVRC²⁰, around 2012). Finally, **Paper IV** successfully drives and characterizes transitions between free energy minima by machine learning a classification model to distinguish basins of attraction (deep-learning-based classification with the cross-entropy loss appeared around the late 1990s).

²⁰ImageNet Large Scale Visual Recognition Challenge

Long story short; computer science drives computational science²¹.

In trying to predict *future* developments in computational physics and chemistry, we may therefore have our intuition guided by the *current* developments in computer science:

- **scale:** the near-exponential growth in compute capacity allows for larger and larger ML models which are trained on larger and larger datasets. In the context of this thesis, this trend is visible in the emergence of universal interaction potentials which have been trained to large databases which cover almost the entire periodic table [126] (e.g. MACE-MP0). Such universal models effectively provide a pretrained baseline which is qualitatively correct but quantitatively inaccurate. The methods proposed in Chapters 3 & 4 also benefit from such models since their internal feature representation is very diverse in terms of chemical environments, and we show in Chapter 4 that they can be readily used as base model for collective variable learning. Every year, such universal potentials grow increasingly accurate while becoming more and more computationally efficient.
- **transformers & generative AI:** the landmark paper *Attention Is All You Need* [127] appeared in 2017 and has already been cited 138,034 times as of October 2024. It introduced the *transformer* neural network architecture, which can be considered as foundational for the recent breakthroughs in generative AI for text, image, video, and music. The impact of generative AI to molecular simulation has so far remained rather limited, but we expect this to increase dramatically in the coming years²². For example, denoising diffusion has been used as a pretraining step for universal ML potentials and it seems to improve the accuracy significantly [126]. Their application for thermodynamic integration calculations has also been demonstrated recently [128].

Lastly, we mention that even with a combination of (1) powerful hardware, (2) the latest graph neural network architectures, and (3) training data with energies and forces at or near chemical accuracy, we are not guaranteed that molecular simulations become sufficiently powerful for all materials and molecules. A well-known example is the prediction of the three-dimensional

²¹Google DeepMind published six Nature papers over the span of twelve months!

²²Interestingly, a variety of startups have appeared over the past 18 months which aim to commercialize their potential

structure of a protein (as in Figure 2.4); direct atomic simulation generally remains infeasible due to the enormous time scales of protein folding events, and alternative methods are necessary. In the case of protein structure prediction, an efficient solution which has been awarded the 2024 Nobel prize in Chemistry is AlphaFold. It relies on a machine learning model using transformers and incorporates a significant amount of experimental training data based on X-ray crystallography.

Part II

Peer-Reviewed Publications, Preprints,
and Conference Contributions



Selected Publications and Preprints

Paper I

Large-Scale Molecular Dynamics Simulations Reveal New Insights Into the Phase Transition Mechanisms in MIL-53(AI)

Sander Vandenhaute, Sven M. J. Rogge, Veronique Van Speybroeck

Frontiers in Chemistry, **2021**, 9, 718920

Sander Vandenhaute implemented and validated the anisotropic Monte Carlo barostat and performed all simulations.

Reprinted with permission.
Copyright (2021), The Authors



Large-Scale Molecular Dynamics Simulations Reveal New Insights Into the Phase Transition Mechanisms in MIL-53(AI)

Sander Vandenhoute, Sven M. J. Rogge and Veronique Van Speybroeck*

Center for Molecular Modeling (CMM), Ghent University, Ghent, Belgium

Soft porous crystals have the ability to undergo large structural transformations upon exposure to external stimuli while maintaining their long-range structural order, and the size of the crystal plays an important role in this flexible behavior. Computational modeling has the potential to unravel mechanistic details of these phase transitions, provided that the models are representative for experimental crystal sizes and allow for spatially disordered phenomena to occur. Here, we take a major step forward and enable simulations of metal-organic frameworks containing more than a million atoms. This is achieved by exploiting the massive parallelism of state-of-the-art GPUs using the OpenMM software package, for which we developed a new pressure control algorithm that allows for fully anisotropic unit cell fluctuations. As a proof of concept, we study the transition mechanism in MIL-53(AI) under various external pressures. In the lower pressure regime, a layer-by-layer mechanism is observed, while at higher pressures, the transition is initiated at discrete nucleation points and temporarily induces various domains in both the open and closed pore phases. The presented workflow opens the possibility to deduce transition mechanism diagrams for soft porous crystals in terms of the crystal size and the strength of the external stimulus.

OPEN ACCESS

Edited by:

Leonard R. MacGillivray,
The University of Iowa, United States

Reviewed by:

Jingxiang Yang,
New York University, United States
Paulo Filho Marques De Oliveira,
University of São Paulo, Brazil
Jonas Baltrusaitis,
Lehigh University, United States

*Correspondence:

Veronique Van Speybroeck
veronique.vanspeybroeck@
ugent.be

Specialty section:

This article was submitted to
Solid State Chemistry,
a section of the journal
Frontiers in Chemistry

Received: 01 June 2021

Accepted: 13 August 2021

Published: 27 August 2021

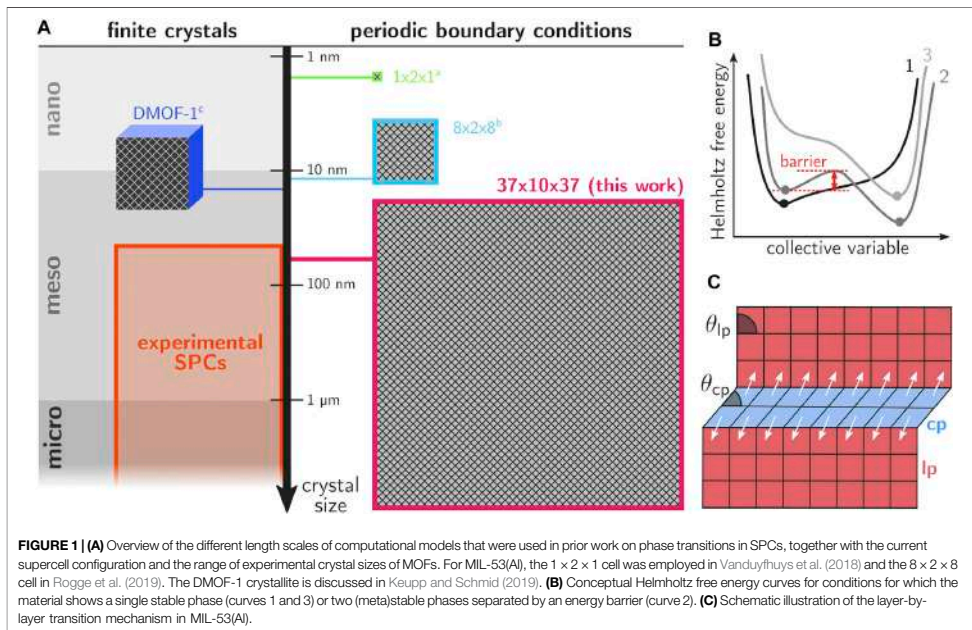
Citation:

Vandenhoute S, Rogge SMJ and
Van Speybroeck V (2021) Large-Scale
Molecular Dynamics Simulations
Reveal New Insights Into the Phase
Transition Mechanisms in MIL-53(AI).
Front. Chem. 9:718920.
doi: 10.3389/fchem.2021.718920

Keywords: soft porous crystals, phase transitions, transition mechanism, phase nucleation, phase propagation, molecular modeling, crystal size

1 INTRODUCTION

Over the last decades, metal-organic frameworks (MOFs) have emerged as an extraordinary class of materials as a result of their unique building block concept from metal nodes and organic linkers, giving rise to an almost limitless number of materials (Chui et al., 1999; Li et al., 1999; Férey, 2001; Yaghi et al., 2003; Kitagawa et al., 2004). MOFs may display anomalous responses to external triggers and exhibit for example negative linear compressibility, negative thermal expansion, negative gas adsorption, or large-amplitude structural transformations while retaining their structural integrity (Coudert, 2015; Balestra et al., 2016; Bennett et al., 2017; Liu et al., 2018; Burtch et al., 2019; Coudert and Evans, 2019; Evans et al., 2019). Such stimuli-responsive behavior has been explored in various research fields including adsorption, separation, catalysis, and drug delivery (Horike et al., 2009; Li et al., 2009; Horcajada et al., 2012; Schneemann et al., 2014). More in particular, the terminology “soft porous crystals” (SPCs) was coined by Kitagawa and colleagues referring to those materials that exhibit bistable or even multistable behavior, possessing the ability to undergo large structural transformations upon exposure to external stimuli while maintaining their long-range structural



order (Horike et al., 2009). Enormous experimental and computational research efforts have been undertaken to understand and ultimately predict or tune this functional behavior. However, most of our insights so far originate from thermodynamic considerations, while the mechanistic details of such large-amplitude phase transformations are yet to be resolved. For example, it is still unclear how the transformation nucleates and how this nucleation is affected by the presence of defects or the size of the crystal. Experimentally, a series of *in situ* experiments have established that the dynamic response of these materials is strongly affected by the presence of defects and the crystal size, whereby crystal downsizing was found to suppress their ability to morph (Sakata et al., 2013; Miura et al., 2017; Krause et al., 2018; Ehrling et al., 2019; Kundu et al., 2019; Wannapaiboon et al., 2019; Bon et al., 2020; Krause et al., 2020; Ehrling et al., 2021).

Another major point of discussion regarding the transition mechanism concerns the degree to which the transition occurs in a collective, cooperative way. In many literature references, the hypothesis of collective behavior was assumed, where the entire framework transforms cooperatively (Schneemann et al., 2014; Wieme et al., 2018; Vanduyfhuys et al., 2018). Insights into the thermodynamics of the phase transformations have been obtained from atomistic simulations by constructing the Helmholtz free energy as function of the state variables governing the observed behavior (Rogge et al., 2015; Vanduyfhuys et al., 2018; Hobday and Kieslich, 2021).

Figure 1B shows a series of conceptual free energy curves in terms of a collective variable that is able to uniquely distinguish between all (meta)stable or activated states. For many frameworks the unit cell volume proved to be an appropriate collective variable (Demuyne et al., 2018), however, this may not always be the case. For flexible linkers, more advanced procedures were necessary to find collective variables that allowed to describe the phase transformation. The study of such frameworks is beyond the scope of the current contribution. Herein, we will focus on MIL-53(AI), the prototypical flexible framework where transitions may be induced by temperature, pressure and gas adsorption (Demuyne et al., 2017; Demuyne et al., 2018). Experimentally determined structural parameters are presented in **Supplementary Section S1**. From the perspective of these Helmholtz free energy curves, collective behavior implies that any small barrier in the thermodynamic potential of a single unit cell would translate into a huge barrier for the entire system and the framework would only transition in case one of the minima in the thermodynamic potential would disappear (curves 1 and 3 in **Figure 1B**).

Nowadays, there are various indications that transitions in SPCs do not occur in a fully cooperative way. Early on, Triguero et al. proposed a theoretical layer-by-layer transition model for winerack frameworks like MIL-53(AI), where collective transformations between the large pore (lp) and closed pore (cp) phases only occur within a single layer at a time (Triguero et al., 2011). Their findings were based on a model

that exploits the characteristic geometry of winerack frameworks, which are composed of rhombus-shaped cells in the directions perpendicular to the aluminum oxide chain. The shape of the rhombus cells can be characterized by the angle θ , which varies between 79° for the lp phase and 40° for the cp phase (see **Figure 1C**). The edge lengths do not change appreciably because the corresponding benzene-1,4-dicarboxylate (BDC) linkers do not allow for large variations in length. These specific geometrical features impose significant constraints on the possible deformation modes of the crystal if its crystallinity is to be preserved; as only those deformations are allowed which preserve the rhombus shape. With these assumptions, cells in a layer must deform coherently in order to preserve the lattice integrity. As such, it was concluded that the phase transformation occurs in an avalanche manner; once the lp-to-cp nucleation occurs at one point, the entire layer quickly transforms to the cp phase. Further lp-to-cp transformations were predicted to occur most likely in layers neighboring already transformed layers. More recent work has constructed energy versus volume curves for MIL-53(AI) using very accurate many-body dispersion calculations within the random-phase approximation (Wieme et al., 2018). Even at 0 K, a non-negligible barrier of about 8 kJ/mol was found between the lp and cp phase, indicating the existence of metastable lp crystals even at very low temperatures. This conclusion is corroborated by Mendt et al. who observed a fraction of lp material for MIL-53(AI) even at 9 K (Mendt et al., 2010).

Computational modeling can help to unravel such mechanistic details, provided that the employed models are realistic representations of experimentally observed MOFs, with dimensions comparable to real crystallites and with the explicit inclusion of the external surface and potential defects. Unfortunately, even with the massive amount of computing resources that is now available, such simulations are not yet feasible (Van Speybroeck et al., 2021). At present, the attainable length scales within the field of nanostructured materials are limited to a few nanometers and common molecular dynamics (MD) runs extend from the picoscale (for ab initio based methods) to the nanoscale range (for force field based methods). In the MOF field, the largest force field based simulations were performed on systems of about 10,000 atoms, with a length scale of tens of nanometers in one direction. Such spatial dimensions are still below experimental crystal sizes which may extend well into the micrometer range (see **Figure 1A**). As a consequence of these model limitations, some effects which are inherently related to (long-range) spatial disorder are not properly accounted for in current computational models (Van Speybroeck et al., 2021).

Recent work has attempted to simulate phase transformations in SPCs in a more realistic manner (see **Figure 1A**). Schmid et al. performed the first finite size simulations on nanocrystallites of DMOF-1(Zn) by setting up a crystallite model containing roughly 250,000 atoms, with explicit inclusion of the external surface (Keupp and Schmid, 2019). Such simulations abandon the periodic boundary conditions (PBCs) that are commonly employed in MOF computational modeling. A series of simulations were performed in a temperature ramp between

300 K and 500 K to observe the thermal opening, and constrained simulations were performed to observe the mechanical closing of the DMOF-1 system. They concluded that the transition nucleates near the external surface, where an interface between the cp and lp phase occurs which then travels dynamically through the lattice. The nucleation was found to be distinctively dependent on the surface-to-volume ratio. Such observations would not have been possible based on simulations with PBCs and small-sized unit cells, as spatial disorder is in that case not allowed and the external surface not taken into account. Simultaneously, Rogge et al. performed large scale force-field-based MD simulations on a series of MOFs, including MIL-53(AI), DMOF-1(Zn) and CoBDP using PBCs, with supercells that contained over 10,000 atoms (Rogge et al., 2019). The latter supercells were referred to as mesocells and their behavior was contrasted with nanocells having substantially smaller dimensions (see **Figure 1A**). Although the mesocells are still about an order of magnitude smaller than experimental crystal sizes, new physicochemical phenomena were observed. For the nanocell, the lp-to-cp transition occurs cooperatively and any form of spatial disorder is prevented by the enforced PBCs. On the other hand, when mesocells are exposed to pressures of 40 MPa at 300 K, part of the system undergoes an lp-to-cp transition whereas the rest temporarily remains in its original lp phase. Temporary interfacial defects were observed near the lp-cp phase boundary, which traverse the lattice until the transformation is complete. The phenomenon was also thermodynamically investigated by constructing the Helmholtz free energy profiles as a function of the volume, which strongly depend on the possibility of phase coexistence occurring in the larger mesocells. The latter phenomenon has not yet been confirmed experimentally, as this requires dedicated in situ cells where the temporal behavior of the material could be followed with very high spatiotemporal resolution.

From this discussion, it is clear that a full mechanistic understanding of the phase transformations in finite MOF crystals requires models that are substantially larger than the ones considered so far, as these would allow us to include spatial disorder at various length scales. Ultimately, simulations of finite-sized crystals containing millions of atoms and having sizes comparable to small experimental crystals would become feasible. In this work, we take an important step in this direction by exploiting the massively parallel architecture of state-of-the-art GPUs in order to simulate these phase transitions at unprecedented length scales. GPU acceleration is known to increase the attainable time and length scale of MD simulations by roughly two orders of magnitude as compared to a CPU, which is achieved by offloading the evaluation of interatomic forces and possibly the time integration onto the GPU. While most molecular mechanics engines nowadays provide some level of GPU acceleration, they are rarely compatible with the specialized force fields and thermodynamic ensembles that are required in common SPC simulation pipelines. In this work, we succeeded in porting the necessary force fields and sampling protocols to the highly extensible and GPU-oriented OpenMM software package (Eastman et al., 2017). In particular, this included the

development and implementation of an extension of the existing pressure control algorithm towards fully anisotropic unit cell fluctuations. The proposed workflow enabled us to go far beyond currently accessible atomistic force field based MD simulations on MOFs. Herein, we perform simulations on models containing 1,040,440 atoms and a unit cell size of $54.9 \times 6.6 \times 45.1 \text{ nm}^3$, thereby entering the range of small experimental crystals (Figure 1A). Visualization and analysis of the transition dynamics become increasingly challenging for such large systems. Herein, we present a dynamic two-dimensional lattice representation that provides a fundamentally new perspective of the framework dynamics. Based on the presented workflow, we are able to give more insight into operative transition mechanisms for different values of the external stimulus.

The remainder of the paper is organized as follows: Section 2 presents the methodological advances that were necessary to achieve this goal, with in particular the derivation and implementation of the proposed pressure control algorithm and the automated dynamic lattice representation. Section 3 then discusses the analysis of the observed transition mechanisms; the concluding discussion and perspectives for future work are provided in Section 4.

2 METHODS

The phase transition mechanism is investigated using MD simulations on a massive scale, with framework models containing up to a million atoms and thousands of pores, and over timescales on the order of 10 ns. Such simulation sizes are unprecedented for SPCs, and require state-of-the-art GPUs in order to become computationally feasible. While GPU-accelerated MD has been around for over a decade, accurate computational models of the large-amplitude transition in winerack-type frameworks additionally require (i) an efficient implementation of a number of specialized force field interactions that are not commonly found elsewhere, and (ii) full flexibility of the simulation cell, as the box vector lengths and angles vary significantly during the transition (Rogge et al., 2018). In this work, we chose to employ the OpenMM software package because of its highly versatile and extensible API, and its flexibility in supporting custom interactions, thereby satisfying the first requirement. The second requirement—fully flexible unit cells—is more complicated because OpenMM does not natively support fully anisotropic unit cell fluctuations. Moreover, the implementation of existing anisotropic barostats in OpenMM would require an extensive rewrite of the entire codebase because they require the virial stress at every step in order to operate properly. To avoid this, we chose to extend an existing isotropic pressure control algorithm with fully anisotropic fluctuations, because it is relatively easy to implement on GPUs and is compatible with custom interaction potentials. The theoretical framework of the proposed method as well as a series of validation experiments are presented in Section 2.1.

The underlying physical mechanisms of the transition are exposed by visualizing trajectories using a dynamic two-dimensional lattice representation. This provides a

fundamentally new perspective on the transition dynamics and the corresponding induced disorder. An overview of the automated and GPU-accelerated visualization procedure is given in Section 2.2.

2.1 Pressure Control

For systems that are subject to periodic boundary conditions, constant pressure MD is achieved by dynamically adjusting the unit cell parameters throughout the simulation in order to maintain equilibrium with an externally applied pressure. The vast majority of current pressure control algorithms [see e.g., Berendsen et al. (1984), MTK (Martyna et al., 1994), Langevin (Feller et al., 1995), and Parrinello-Rahman (Parrinello and Rahman, 1981; Nosé and Klein, 1983)] determine the appropriate change in unit cell parameters based on the instantaneous virial stress of the system, which is both expensive to compute and not always readily available (e.g., in polarizable force fields). To avoid the virial computation, an alternative algorithm has been developed which instead performs Monte Carlo trial moves in the unit cell volume in order to sample the desired isothermal-isobaric phase space distribution (Chow and Ferguson, 1995; Åqvist et al., 2004). Here, we extend this Monte Carlo pressure control algorithm towards fully anisotropic cell fluctuations, as these are critical for the description of phase transformations in (soft) condensed matter and SPCs in particular (Rogge et al., 2018).

2.1.1 Derivation

The Monte Carlo pressure control algorithm works by performing Metropolis sampling in the unit cell degrees of freedom during a MD simulation, with trial moves being performed every 10 to 100 steps. Whereas previous work only considered isotropic moves, i.e., trial moves which attempt to scale the unit cell isotropically while leaving its shape unaltered, we here present a more general approach that considers trial moves in all unit cell degrees of freedom. We begin the derivation by considering a molecular system that is periodic in all three dimensions, with a triclinic unit cell that contains N atoms. A microstate of this system is determined by the cartesian coordinates of its N particles $\mathbf{r} = (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)$ and, in addition, three linearly independent box vectors \mathbf{a} , \mathbf{b} , and \mathbf{c} which determine the periodicity of the system. These box vectors can be arranged along the rows of a 3×3 cell matrix \mathbf{h} , in which case the unit cell volume is given by $V = \det(\mathbf{h})$. The configurational partition function in the isothermal-isobaric ensemble for this system is then defined as (Martyna et al., 1994; Tuckerman, 2010):

$$\Delta(N, P, T) = C \iint e^{-\beta U(\mathbf{r})} e^{-\beta PV} \det(\mathbf{h})^{-2} d\mathbf{h} d\mathbf{r} \quad (1)$$

whereby $\beta = (kT)^{-1}$ represents the inverse temperature, P the externally applied pressure, and C a constant that is otherwise irrelevant. The differentials $d\mathbf{h}$ and $d\mathbf{r}$ that appear in Eq. 1 represent integrations over the nine components of the matrix \mathbf{h} and the $3N$ components of the coordinates \mathbf{r} , respectively. Most MD engines require cell matrices to be in lower triangular form

due to efficiency considerations, and it is therefore necessary to first rewrite Eq. 1 in terms of lower triangular cell matrices \mathbf{h}_{Δ} :

$$\mathbf{h}_{\Delta} = \begin{bmatrix} a_x & 0 & 0 \\ b_x & b_y & 0 \\ c_x & c_y & c_z \end{bmatrix} \quad (2)$$

This may be achieved by eliminating global rotations of the coordinate system. As explained in **Supplementary Section S2**, this enables us to transform the original nine-dimensional integration over \mathbf{h} into a six-dimensional integration over \mathbf{h}_{Δ} :

$$\Delta(N, P, T) = C \iint e^{-\beta U(\mathbf{r})} e^{-\beta PV} (b_y c_z^2)^{-1} d\mathbf{h}_{\Delta} d\mathbf{r} \quad (3)$$

Eq. 3 represents the desired phase space distribution that we wish to approximate. To derive a specific expression for the acceptance probability of a given trial move in a Metropolis random walk, we first need to remove the \mathbf{h}_{Δ} -dependence in the integration limits of the particle coordinates; a similar procedure is necessary when deriving acceptance probabilities in regular isothermal-isobaric Monte Carlo (Tuckerman, 2010). This may be achieved by defining normalized particle coordinates s_i for each particle i (in components):

$$r_{i,x} = a_x s_{i,x} \quad (4)$$

$$r_{i,y} = b_y s_{i,y} \quad (5)$$

$$r_{i,z} = c_z s_{i,z} \quad (6)$$

and after performing the substitution in Eq. 3:

$$\Delta(N, P, T) = C \iint e^{-\beta U(s, \mathbf{h}_{\Delta})} e^{-\beta PV} (b_y c_z^2)^{-1} V^N d\mathbf{h}_{\Delta} ds \quad (7)$$

whereby both integrations are now performed over fixed domains. Eq. 7 now allows us to use Metropolis sampling to generate a Markov chain of unit cell matrices that exhibits the correct phase space distribution (Metropolis et al., 1953). A trial move $\mathbf{h}_{\Delta,o} \rightarrow \mathbf{h}_{\Delta,n}$ is generated by sampling random displacements from a uniform distribution in all six components of the unit cell. The move is then accepted with a probability derived from Eq. 7:

$$\Pr[\mathbf{h}_{\Delta,o} \rightarrow \mathbf{h}_{\Delta,n}] = \min \left[1, \exp \left(\sum_{i=0}^3 c_i \right) \right] \quad (8)$$

with

$$c_0 = -\beta \Delta U \quad (9)$$

$$c_1 = -\beta P \Delta V \quad (10)$$

$$c_2 = N \log \left(\frac{V_n}{V_o} \right) \quad (11)$$

$$c_3 = \log \left(\frac{(b_y c_z^2)_n}{(b_y c_z^2)_o} \right) \quad (12)$$

and where ΔU , ΔV denote the differences in potential energy and volume between the new unit cell $\mathbf{h}_{\Delta,n}$ and the old unit cell $\mathbf{h}_{\Delta,o}$. The amplitude of the displacements is dynamically adjusted such that, on average, about 50% of the trial moves are accepted.

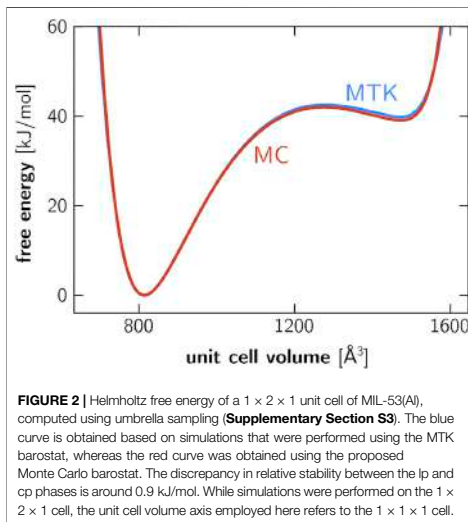


FIGURE 2 | Helmholtz free energy of a $1 \times 2 \times 1$ unit cell of MIL-53(A), computed using umbrella sampling (**Supplementary Section S3**). The blue curve is obtained based on simulations that were performed using the MTK barostat, whereas the red curve was obtained using the proposed Monte Carlo barostat. The discrepancy in relative stability between the lp and cp phases is around 0.9 kJ/mol. While simulations were performed on the $1 \times 2 \times 1$ cell, the unit cell volume axis employed here refers to the $1 \times 1 \times 1$ cell.

2.1.2 Validation

The isothermal-isobaric ensemble is characterized by a phase space distribution in accordance with Eq. 1, and the here proposed pressure control algorithm should succeed in generating phase space trajectories that are distributed accordingly. To critically verify that this is indeed the case, we performed a number of validation experiments on strongly anisotropic systems.

First, we considered a simple harmonic crystal in which only short-range covalent interactions are present. The force constants and geometry of the system were deliberately chosen as to induce a large degree of anisotropy, with fluctuations in unit cell parameters that are sufficiently large. MD trajectories obtained at different pressures were then evaluated based on the convergence of the average stress tensor and by following the rigorous ensemble validation protocol outlined in Shirts (2013) and Merz and Shirts (2018). As reported in **Supplementary Section S3**, all trajectories passed the stress tensor convergence and ensemble validation tests.

A second validation experiment was performed specifically on the MIL-53(A) framework. Previous work has established the importance of anisotropic fluctuations in the relative stability of the lp and cp phases (Rogge et al., 2018), and it is imperative that the Monte Carlo based pressure control algorithm yields the same relative stabilities as virial-based barostats (e.g., MTK or Langevin). To demonstrate their equivalence, we computed the Helmholtz free energy as a function of the unit cell volume using umbrella sampling, employing either the proposed Monte Carlo barostat or the existing MTK barostat (**Supplementary Section S4**). As observed in **Figure 2**, both approaches are equivalent as

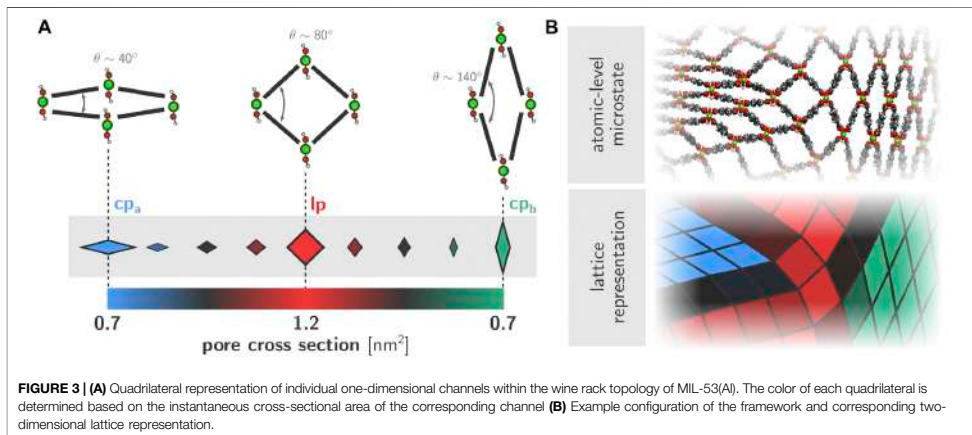


FIGURE 3 | (A) Quadrilateral representation of individual one-dimensional channels within the wine rack topology of MIL-53(A). The color of each quadrilateral is determined based on the instantaneous cross-sectional area of the corresponding channel **(B)** Example configuration of the framework and corresponding two-dimensional lattice representation.

the predicted relative differences in cp-lp stability deviate by less than two percent.

Both validation experiments provide strong evidence for the correctness of the derivation and implementation of the proposed pressure control algorithm. It is worth mentioning that throughout the derivation, no modifications to the equations of motion of individual particles were necessary. This suggests that the effect of Monte Carlo pressure control on the overall dynamics of the system is expected to be negligible, which has indeed been confirmed for the isotropic variant based on the invariance of diffusion coefficients (Chow and Ferguson, 1995). Lastly, an important advantage over virial-based approaches is that the proposed method works equally well for simulations in which forces and energies are computed in single precision, whereas virial-based barostats may in that case experience overall drifts in e.g. the average density (as discussed in Harger and Ren (2019)). This is particularly important for GPU acceleration, as it is well known that the floating point performance of GPUs in single precision is much higher as compared to double precision (Le Grand et al., 2013).

2.2 Visualization

A physically accurate and visually clear representation of the obtained trajectories is vital in order to understand the physical mechanisms that govern the phase transition behavior. Previous work has established how both the transition itself as well as the various forms of induced disorder are essentially two-dimensional phenomena that are translationally invariant along the direction of the aluminum chains (Rogge et al., 2019). As such, we chose to visualize the framework dynamics using a two-dimensional representation of the lattice, where the aluminum chains are represented by vertices and the adjoining organic ligands by edges (Figure 3). The position of the vertices is obtained by projecting the corresponding aluminum chain onto the cross-sectional plane. This two-dimensional representation of

the framework structure is then constructed for each snapshot in a given trajectory. In order to further highlight the state of each of the pores in the system, we color each quadrilateral based on its instantaneous cross-sectional area. More specifically, the lp and both of the cp phases are filled using full colors (respectively red, blue, and green). For pores that are transitioning between phases and therefore contain cross-sectional areas that are between the pure lp and cp values, we used different shades of gray as indicated in the colorbar in Figure 3A. As such, the framework dynamics and local phase behavior are visualized and exposed in a tangible manner.

2.3 Computational Details

2.3.1 Force Fields

All simulations on the MIL-53(AI) framework were performed using an ab initio derived force field. The covalent interactions were obtained based on quantum mechanical input data using the QuickFF protocol (Vanduyfhuys et al., 2015). The electrostatic interactions were included by considering Gaussian charge distributions around each atom. The magnitude of the charge distributions was obtained using the Minimal Basis Iterative Stockholder (MBIS) scheme (Verstraelen et al., 2016), and the radius of the distributions were obtained using the scheme by Chen and Martínez et al (2007). Dispersion interactions were modelled using MM3-type interactions (Allinger et al., 1989). The force field has been validated extensively in previous work (Vanduyfhuys et al., 2018). The OpenMM input files for the force field were prepared using the OpenYAFF conversion tool (Vandenhoute, 2021b).

2.3.2 Molecular Dynamics Simulations

Large-scale MD simulations on the transition mechanism of MIL-53(AI) were performed using OpenMM 7.5.0 (Eastman et al., 2017), supplemented with the implementation of the new barostat (available online (Vandenhoute, 2021a)). The

simulations were performed using a leapfrog Langevin integrator with a friction coefficient of 0.1 ps^{-1} (Zhang et al., 2019) and a timestep of 0.5 fs. Monte Carlo trial moves in the unit cell degrees of freedom were performed every five steps. With this trial move frequency, it takes roughly 10 ns of simulation time for a full lp–cp transition. Particle positions and unit cell vectors were sampled every 5 ps. Dispersion interactions were smoothly truncated at 11 Å, and supplemented with analytical tail corrections. Electrostatic interactions were evaluated using the particle mesh Ewald (PME) method, with a splitting parameter α of 0.32 \AA^{-1} and a reciprocal space cutoff of 0.35 \AA^{-1} . All simulations were performed on a single NVIDIA V100 GPU with 32 GB of memory, achieving a simulation speed of about 0.5 ns/day in mixed precision mode, in which case forces and energies are computed in single precision and the time integration is performed in double precision, thereby achieving an optimal tradeoff between accuracy and computational efficiency (Eastman et al., 2017). We note that no additional free energy calculations were performed beyond the validation in Figure 2, as these are not required to understand the different transition mechanisms. Furthermore, the unit cell volume is no longer a suitable collective variable for systems of this size given that the statistical fluctuations in V are almost entirely absent due to the very large number of particles.

3 RESULTS

We investigated the transition mechanism based on large-scale MD simulations at constant temperature and pressure. In order to eliminate PBC artefacts as much as possible and ensure that phase separation and/or coexistence is maximally allowed by the model, we performed the simulations on a $37 \times 10 \times 37$ unit cell containing 1,040,440 atoms and just over 2,500 individual pores. Such unit cell sizes are unprecedented in computational research on MOFs and SPCs in particular, and represent a significant step forward with respect to the state-of-the-art (as visualized in Figure 1A). To identify the pertinent features in the framework dynamics, we simulated the lp-to-cp transition for this system at different thermodynamic conditions. An overview of all simulations performed with their specific control variables is given in Supplementary Table S2. Specifically, we considered three different temperatures (200, 300 and 500 K) and two different pressures for each temperature (100 and 300 MPa). While both pressures are much higher than the experimentally observed transition pressure for this material—which is estimated at about 13–18 MPa at room temperature (Yot et al., 2014)—these pressures ensured that the phase transition proceeds sufficiently fast as to make the simulations computationally feasible on single-GPU systems. Alternatively, enhanced sampling techniques may be used to speed up the transition, for example by biasing the dynamics of the framework along the θ angle indicated in Figure 3. As this is a first case study of the large-scale dynamics of winerack-type frameworks, we chose not to pursue this direction and consider regular unbiased dynamics at elevated transition pressures instead. Nevertheless, the transition times of the employed $37 \times 10 \times 37$ system are

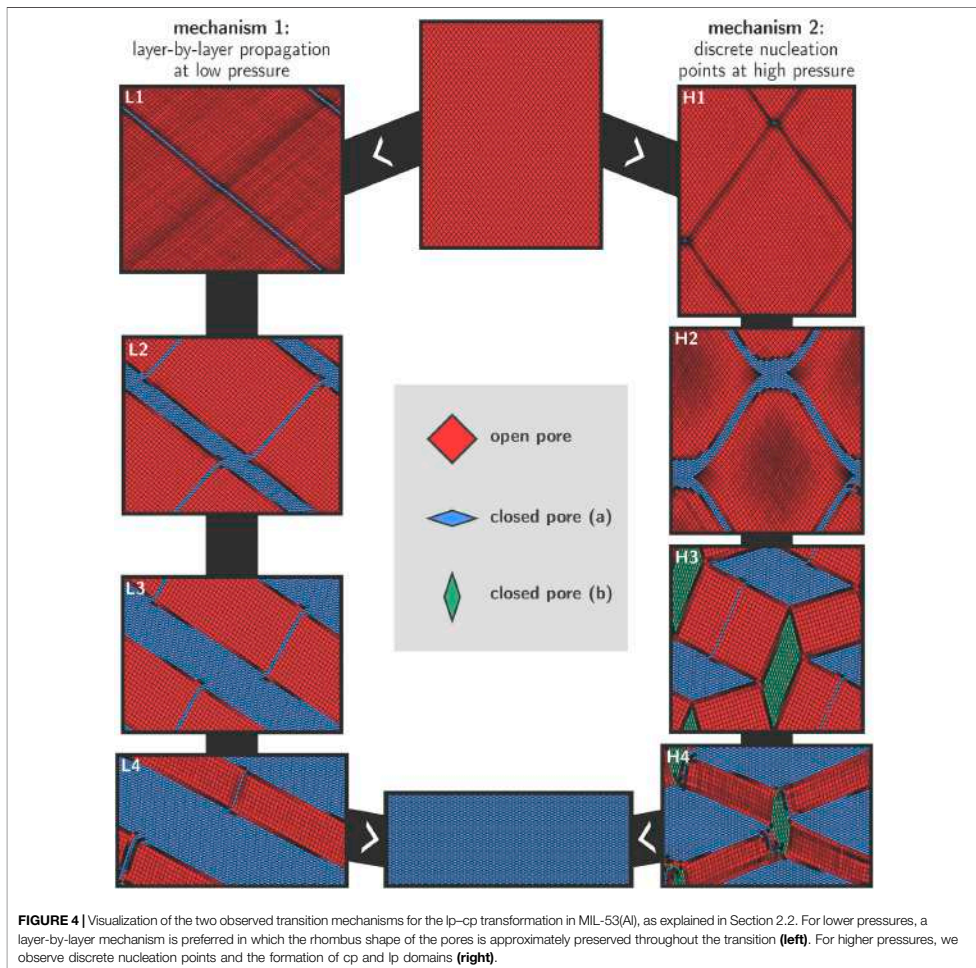
observed to be roughly three orders of magnitude larger as compared to a fully cooperative transition in a $1 \times 2 \times 1$ cell (~ 10 ns versus ~ 10 ps).

Figure 4 shows snapshots of the transition mechanisms at 300 K in the low- and high-pressure regimes, as obtained from simulations at 100 MPa (left) and 300 MPa (right). The final aim is to unravel more details about the transition mechanisms at various conditions. Before discussing the results in detail, we first provide some general remarks on the transition based on topological considerations. Considering the winerack-type structure of the framework and the very limited intrinsic flexibility of the BDC linkers in between aluminum chains, physically feasible transition mechanisms need to preserve the approximately rhombus shape for most of the pores in the system because the energy required to strongly deform the cells from their rhombus shape is expected to be very high. In addition, interatomic forces were modelled using a classical force field and as such we implicitly enforce strict covalent bond integrity throughout the entire transition. Deviations from the rhombus shape might eventually be possible provided that the mechanical energy supplied to the system is sufficiently high. Our simulations show that essentially two different transition mechanisms are active depending on the magnitude of the applied external pressure. At lower pressures, a layer-by-layer transition mechanism is observed whereas at higher pressures discrete nucleation points emerge within the lattice, which ultimately give rise to domain formation during the transformation. These results are schematically shown in Figure 4, where the left column shows the snapshots for the 100 MPa transition, and the right column shows the snapshots for the 300 MPa transition. While both transitions were recorded at 300 K, we observed entirely similar behavior at 200 K and at 500 K.

3.1 Layer-By-Layer Transition at Low Pressures

For all transitions with relatively low values of the external stimulus (100 MPa); we observe a diagonal layer-by-layer mechanism in agreement with the conceptual model proposed in Triguero et al. (2011). This mechanism is initiated in one layer in which all cells switch cooperatively from the lp to the cp phase (Figure 4, panel L1). The formation of the initial cp layer within an lp bulk phase creates two-dimensional phase boundaries that extend diagonally across the material. It may be noted that a second gray-shaded diagonal appears. However, in this pressure regime, this region does not increase substantially during the simulation.

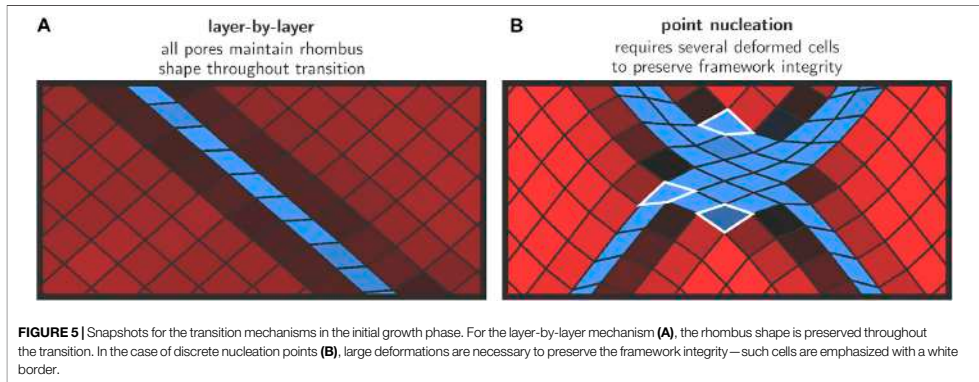
To obtain more insight into the deformation mode of the various cells in and around the layer where the transformation was initiated, an enlarged representation of this area in the lattice is shown in Figure 5A. It can be seen that all unit cells maintain their rhombus-like shape and that only the angle θ changes substantially from ca. 78° to ca. 42° . The linkers approximately maintain their length. This is in line with our earlier hypothesis where we anticipated that only small deformations from a rhombus shape are energetically feasible at low pressures.



Once a cp layer has been formed, the transition proceeds gradually in the direction normal to the nucleation layer, whereby lp layers at the lp/cp phase boundary collapse and are appended by the growing cp region. This effectively leads to a propagation of the lp/cp phase boundary in the direction normal to the nucleation layer (Figure 4, panels L2–L4). This process continues until all layers have transformed to the cp phase, at which point the transition is complete and a minimum in free energy is encountered. Such a layer-by-layer mechanism does not require individual pores to deviate from their rhombus shape, and their equilateral form is preserved at all times during the

transition. The only internal coordinate that varies throughout the entire process is the angle θ .

As already indicated earlier, it is clear from the visualization in panels L1 to L4 of Figure 4 that a second diagonal having the cp phase is formed parallel to the growth direction and roughly orthogonal to the existing growing cp region. This second cp region does not propagate further throughout the material, and does not otherwise mediate the transition in a significant way. We therefore hypothesize that it is a result of the rather large pressures that had to be applied in order to make the transition computationally feasible. The applied pressure of



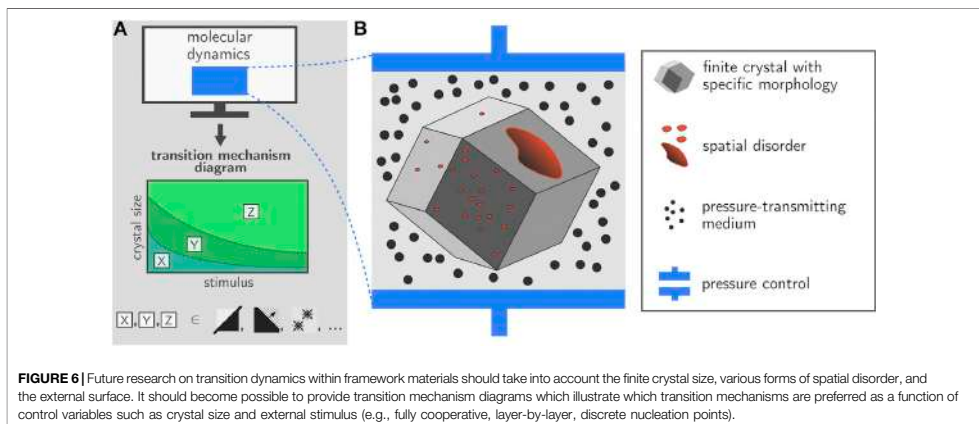
100 MPa in the left panels of **Figure 4** is still significantly higher than the lowest possible transition pressure which was previously established at around 30 MPa (Vanduyfhuys et al., 2018), and the appearance of a second cp diagonal therefore appears to be a way for the framework to release the excessive strain within the lattice.

3.2 Transition via Discrete Nucleation Points at High Pressures

At significantly higher pressures, an alternative transition pathway was observed where the transition is initiated at various nucleation points and where also various domains are formed in the crystal. The snapshots resulting from a MD simulation at 300 K and 300 MPa are shown in panels H1 to H4 in **Figure 4**. Similar observations were made for simulations performed at pressures above 300 MPa.

Early on in the simulation (panel H1), discrete nucleation points are formed throughout the framework. Immediately after their formation, these nucleation points become connected along various diagonals (panel H2). To obtain more insight into the deformation modes of the individual cells in this high-pressure regime, **Figure 5B** shows an enlarged representation of the deformed region around a nucleation point, in which we observe how some cells are strongly deformed from their original rhombus-like shape. This is possible because the energy required to induce such large deformations is now mechanically supplied to the system due to the high external pressure. This is in stark contrast with the layer-by-layer mixed phase configurations which were topologically allowed and did not require this type of deformations.

Also visible in panel H2 are the gray-shaded areas in between the interconnected nuclei, indicating that pores in those regions



are significantly smaller than their full l_p counterparts. We may regard these regions as being squeezed by nuclei on either sides due to the large structural difference between the cp and l_p phase. At this stage, continued growth of the cp regions further increases the strain, until the point where domain formation occurs in which alternate cp (both with acute and obtuse θ angles) and l_p domains are present (panel H3). As such, all three possible phases of the framework are found to coexist temporarily during this phase of the transition. The blue cp regions continue to grow at 300 K and 300 MPa, at the expense of the l_p and alternative cp domains, which ultimately disappear at which point the transition is complete.

In spite of the enormous size of the mesocell, artefacts of the employed PBCs are still visible in the snapshots in **Figure 4**. For example, in panel H3, the structure and geometry of the different domains is necessarily organized such that the required periodicity is satisfied. Clearly, a full understanding of the transition dynamics requires computational models in which the implicit assumption of exact long-range order—as is the case when using PBCs—is completely eliminated, for example by using finite crystallites. In addition, we note that the dynamics of the system is not purely hamiltonian. Specifically, the temperature of the system was controlled using a Langevin thermostat, which implies that the equations of motion are modified with a stochastic contribution, the magnitude of which is determined by the friction coefficient (see Section 2.3). In addition to the thermostat, the dynamics of the system are also perturbed by the barostat, which performs small but frequent changes to the unit cell parameters. While we do not expect either the temperature or pressure control to induce drastic changes in the transition dynamics, they do prohibit a thorough analysis of the kinetics of the transition, as e.g., the transition time is dependent on the frequency of barostat trial moves.

4 CONCLUSION AND PERSPECTIVES

The simulations performed here constitute an important step forward towards a full mechanistic understanding of various phase transformations in realistic MOF particles. The implementation of the proposed pressure control algorithm in the GPU-accelerated OpenMM library as presented here allows to include fully anisotropic unit cell fluctuations in large-scale simulations containing millions of atoms. Its area of application is not limited to the breathing transition in winerack frameworks but is very general, and also includes e.g., crystalline-to-amorphous transitions. As such, we have taken a major leap forward in terms of the size of the systems that can be simulated in the field of MOFs. Based on the current simulations, we demonstrate that the transition mechanism may be critically dependent on various elements such as the size of the crystal but also the strength of the external stimulus, which is in this case the external pressure. For future simulations, it would be interesting to extend the protocol towards other stimuli such as temperature but eventually also guest molecules, and unravel whether similar transition mechanisms are observed. Such simulations might enable to construct a transition mechanism diagram that indicates the expected transition mechanism as a function of various control variables. Such a

hypothetical diagram is illustrated in **Figure 6A**. The icons X,Y,Z indicate various mechanisms such as collective behavior, layer-by-layer behavior, or transitions starting from discrete nucleation points. Based on the simulations performed in this paper and literature data, we are already able to fix a number of points on the transition mechanism diagram in terms of the external pressure. However, more simulations are necessary to provide the complete transition mechanism diagram.

Despite the new insights obtained in this work, further methodological steps are necessary to enable the simulation of systems that are truly representative of realistic MOF particles, i.e., with crystal sizes similar to experimentally observed crystals, and with explicit inclusion of defects and the crystal surface. In what follows, we give some reflections on future perspectives in this direction.

First of all, the current approach will have to be extended towards finite nanocrystals with morphologies that are representative of experimentally observed crystallites and with length scales going beyond 50 nm. Subsequent embedding of this crystallite in a medium will then allow to apply an external pressure to the crystal in a very natural manner (**Figure 6**). Experimentally, both mercury as well as silicone oil have been used as effective pressure transducers in order to detect pressure-induced structural transitions in several winerack frameworks (Beurroies et al., 2010; Yot et al., 2012; Yot et al., 2014, 2016; Ramaswamy et al., 2017; Henke et al., 2018; Krause et al., 2018; Wahiduzzaman et al., 2018; Wieme et al., 2019; Yot et al., 2019) and other flexible frameworks such as ZIF-4 or DUT-8 (Kavoosi et al., 2017; Krylov et al., 2020); the simulation setup as described in **Figure 6B** would mirror such experiments.

Second, the current simulations were performed using an all-atom classical force field which does not allow any bond breakage to occur. It is not excluded that during a phase transition, bonds may temporarily break at specific locations within the framework as it is well known that linkers in MOFs may have a labile nature especially when being exposed to guest particles that allow to temporarily stabilize detached linkers (Hajek et al., 2018; Caratelli et al., 2019). To account for such effects, it is necessary to employ reactive force fields or more complex machine learning potentials (MLPs) that are trained based on quantum mechanical calculations in order to capture such effects. However, the systematic construction of MLPs for the complex systems under study here is highly nontrivial. To the best of our knowledge, only one MLP has so far been constructed for MOFs, by the group of Behler on MOF-5 (Eckhoff and Behler, 2019).

Finally, while current simulations allowed to deduce qualitatively new mechanistic details on the phase transition in MIL-53(Al), a next step would additionally aim to determine thermodynamic and kinetic properties associated with the transition, including its nucleation and growth. From a thermodynamic point of view, we could resort to the construction of Helmholtz free energy diagrams in terms of an appropriate collective variable. As noted in Section 2.3, the unit cell volume can no longer be regarded as an appropriate collective variable due to the absence of significant statistical fluctuations, and other variables will have to be considered such as the opening angle θ . The determination of kinetic properties, such as propagation rate constants for the phase boundaries, is highly

challenging for these systems and may require specialized sampling protocols.

Previous considerations clearly illustrate the complexity associated with a full understanding of phase transformations in SPCs. The problem at hand is a prototypical example of a spatiotemporal process, where the dynamics of the MOF lattice is affected by spatial heterogeneities at various length and time scales (Van Speybroeck et al., 2021). A full understanding of the spatiotemporal response of MOFs will require a close partnership between the modeling and experimental community, whereby dedicated experimental *in situ* methods are necessary to track intermediate metastable states during their dynamic response towards external stimuli, and where theoreticians will have to explore new modeling avenues to tackle processes at various length and time scales.

DATA AVAILABILITY STATEMENT

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

AUTHOR CONTRIBUTIONS

SV, SR, and VVS initiated the discussion, designed and wrote the paper, and were involved in the discussions of the results. SV derived the extension of the anisotropic barostat in OpenMM and performed all simulations included in this work.

REFERENCES

- Allinger, N. L., Yuh, Y. H., and Lii, J. H. (1989). Molecular Mechanics. The MM3 Force Field for Hydrocarbons. 1. *J. Am. Chem. Soc.* 111, 8551–8566. doi:10.1021/ja00205a001
- Åqvist, J., Wennerström, P., Nervall, M., Bjelic, S., and Brandsdal, B. O. (2004). Molecular Dynamics Simulations of Water and Biomolecules with a Monte Carlo Constant Pressure Algorithm. *Chem. Phys. Lett.* 384, 288–294. doi:10.1016/j.cplett.2003.12.039
- Balestra, S. R. G., Bueno-Perez, R., Hamad, S., Dubbeldam, D., Ruiz-Salvador, A. R., and Calero, S. (2016). Controlling thermal Expansion: A Metal–Organic Frameworks Route. *Chem. Mater.* 28, 8296–8304. doi:10.1021/acs.chemmater.6b03457
- Bennett, T., Cheetham, A., Fuchs, A., and Coudert, F.-X. (2017). Interplay between Defects, Disorder and Flexibility in Metal–Organic Frameworks. *Nat. Chem.* 9, 11–16. doi:10.1038/nchem.2691
- Berendsen, H. J. C., Postma, J. P. M., van Gunsteren, W. F., DiNola, A., and Haak, J. R. (1984). Molecular Dynamics with Coupling to an External bath. *J. Chem. Phys.* 81, 3684–3690. doi:10.1063/1.448118
- Beuroires, I., Boulhout, M., Llewellyn, P. L., Kuchta, B., Férey, G., Serre, C., et al. (2010). Using Pressure to Provoked the Structural Transition of Metal–Organic Frameworks. *Angew. Chem. Int. Ed.* 49, 7526–7529. doi:10.1002/anie.201003048
- Bon, V., Brunner, E., Pöppel, A., and Kaskel, S. (2020). Unraveling Structure and Dynamics in Porous Frameworks via Advanced *In Situ* Characterization Techniques. *Adv. Funct. Mater.* 30, 1907847. doi:10.1002/adfm.201907847
- Burth, N., Baxter, S., Heinen, J., Bird, A., Schneemann, A., Dubbeldam, D., et al. (2019). Negative thermal Expansion Design Strategies in a Diverse Series of

FUNDING

VVS acknowledges funding from the European Union's Horizon 2020 research and innovation program (Consolidator ERC Grant Agreement 647755-DYNPOR), as well as the Research Board of Ghent University (BOF). SV and SR wish to thank the Fund for Scientific Research–Flanders (FWO) for an aspirant doctoral and junior postdoctoral fellowship, respectively (grant nos. 11H6821N (SV) and 12T3519N (SR)). The computational resources and services used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation – Flanders (FWO) and the Flemish Government – department EWI.

ACKNOWLEDGMENTS

Peter Kenneth Eastman of Stanford University, Department of Chemistry is acknowledged for his help with the implementation of the anisotropic Monte Carlo barostat in OpenMM.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fchem.2021.718920/full#supplementary-material>

Metal–Organic Frameworks. *Adv. Funct. Mater.* 29, 1904669. doi:10.1002/adfm.201904669

- Caratelli, C., Hajek, J., Meijer, E. J., Waroquier, M., and Van Speybroeck, V. (2019). Dynamic Interplay between Defective UiO-66 and Protic Solvents in Activated Processes. *Chem. Eur. J.* 25, 15315–15325. doi:10.1002/chem.201903178
- Chen, J., and Martínez, T. J. (2007). QTPIE: Charge Transfer with Polarization Current Equalization. A Fluctuating Charge Model with Correct Asymptotics. *Chem. Phys. Lett.* 438, 315–320. doi:10.1016/j.cplett.2007.02.065
- Chow, K.-H., and Ferguson, D. M. (1995). Isothermal-isobaric Molecular Dynamics Simulations with Monte Carlo Volume Sampling. *Comput. Phys. Commun.* 91, 283–289. doi:10.1016/0010-4655(95)00059-O
- Chui, S. S.-Y., Lo, S. M.-F., Charmant, J. P., Orpen, A. G., and Williams, I. D. (1999). A Chemically Functionalizable Nanoporous Material. *Science* 283, 1148–1150. doi:10.1126/science.283.5405.1148
- Coudert, F.-X., and Evans, J. (2019). Nanoscale Metamaterials: Meta-MOFs and Framework Materials with Anomalous Behavior. *Coord. Chem. Rev.* 388, 48–62. doi:10.1016/j.ccr.2019.02.023
- Coudert, F.-X. (2015). Responsive Metal–Organic Frameworks and Framework Materials: Under Pressure, Taking the Heat, in the Spotlight, with Friends. *Chem. Mater.* 27, 1905–1916. doi:10.1021/acs.chemmater.5b00046
- Demuyne, R., Rogge, S. M. J., Vanduyfhuys, L., Wieme, J., Waroquier, M., and Van Speybroeck, V. (2017). Efficient Construction of Free Energy Profiles of Breathing Metal–Organic Frameworks Using Advanced Molecular Dynamics Simulations. *J. Chem. Theor. Comput.* 13, 5861–5873. doi:10.1021/acs.jctc.7b01014
- Demuyne, R., Wieme, J., Rogge, S. M. J., Dedecker, K. D., Vanduyfhuys, L., Waroquier, M., et al. (2018). Protocol for Identifying Accurate Collective Variables in Enhanced Molecular Dynamics Simulations for the Description of Structural Transformations in Flexible Metal–Organic Frameworks. *J. Chem. Theor. Comput.* 14, 5511–5526. doi:10.1021/acs.jctc.8b00725

- Eastman, P., Swails, J., Chodera, J. D., McGibbon, R. T., Zhao, Y., Beauchamp, K. A., et al. (2017). OpenMM 7: Rapid Development of High Performance Algorithms for Molecular Dynamics. *PLoS Comput. Biol.* 13, 1–17. doi:10.1371/journal.pcbi.1005659
- Eckhoff, M., and Behler, J. (2019). From Molecular Fragments to the Bulk: Development of a Neural Network Potential for MOF-5. *J. Chem. Theor. Comput.* 15, 3793–3809. doi:10.1021/acs.jctc.8b01288
- Ehrling, S., Miura, H., Senkovska, I., and Kaskel, S. (2021). From Macro- to Nanoscale: Finite Size Effects on Metal–Organic Framework Switchability. *Trends Chem.* 3, 291–304. doi:10.1016/j.trechm.2020.12.012
- Ehrling, S., Senkovska, I., Bon, V., Evans, J. D., Petkov, P., Krupskaya, Y., et al. (2019). Crystal Size versus Paddle Wheel Deformability: Selective Gas Adsorption Transitions of the Switchable Metal–Organic Frameworks DUT-8(Co) and DUT-8(Ni). *J. Mater. Chem. A* 7, 21459–21475. doi:10.1039/C9TA06781G
- Evans, J., Dürholt, J., Kaskel, S., and Schmid, R. (2019). Assessing Negative Thermal Expansion in Mesoporous Metal–Organic Frameworks by Molecular Simulation. *J. Mater. Chem. A* 7, 24019–24026. doi:10.1039/C9TA06644F
- Feller, S. E., Zhang, Y., Pastor, R. W., and Brooks, B. R. (1995). Constant Pressure Molecular Dynamics Simulation: The Langevin Piston Method. *J. Chem. Phys.* 103, 4613–4621. doi:10.1063/1.470648
- Férey, G. (2001). Microporous Solids: From Organically Templated Inorganic Skeletons to Hybrid Frameworks: A Review in Chemistry. *Chem. Mater.* 13, 3084–3098. doi:10.1021/cm011070n
- Hajek, J., Caratelli, C., Demuynek, R., De Wispelaere, K., Vanduyfhuys, L., Waroquier, M., et al. (2018). On the Intrinsic Dynamic Nature of the Rigid UiO-66 Metal–Organic Framework. *Chem. Sci.* 9, 2723–2732. doi:10.1039/C7SC04947A
- Harger, M., and Ren, P. (2019). Virial-based Berendsen Barostat on GPUs Using AMOEBA in Tinker-OpenMM. *Res. Chem.* 1, 100004. doi:10.1016/j.jrechem.2019.100004
- Henke, S., Wharmby, M. T., Kieslich, G., Hante, I., Schneemann, A., Wu, Y., et al. (2018). Pore Closure in Zeolitic Imidazolate Frameworks under Mechanical Pressure. *Chem. Sci.* 9, 1654–1660. doi:10.1039/C7SC04952H
- Hobday, C. L., and Kieslich, G. (2021). Structural Flexibility in Crystalline Coordination Polymers: A Journey along the Underlying Free Energy Landscape. *Dalton Trans.* 50, 3759–3768. doi:10.1039/D0DT04329J
- Horcajada, P., Gref, R., Baati, T., Allan, P. K., Maurin, G., Couvreur, P., et al. (2012). Metal-organic Frameworks in Biomedicine. *Chem. Rev.* 112, 1232–1268. doi:10.1021/cr200256v
- Horike, S., Shimomura, S., and Kitagawa, S. (2009). Soft Porous Crystals. *Nat. Chem.* 1, 695–704. doi:10.1038/nchem.444
- Kavosi, N., Bon, V., Senkovska, I., Krause, S., Atzori, C., Bonino, F., et al. (2017). Tailoring Adsorption Induced Phase Transitions in the Pillared-Layer Type Metal–Organic Framework DUT-8(Ni). *Dalton Trans.* 46, 4685–4695. doi:10.1039/C7DT00015D
- Keupp, J., and Schmid, R. (2019). Molecular Dynamics Simulations of the ‘breathing’ Phase Transformation of MOF Nanocrystallites. *Adv. Theor. Simul.* 2, 1900117. doi:10.1002/adts.201900117
- Kitagawa, S., Kitaura, R., and Noro, S.-i. (2004). Functional Porous Coordination Polymers. *Angew. Chem. Int. Ed.* 43, 2334–2375. doi:10.1002/anie.200300610
- Krause, S., Bon, V., Senkovska, I., Többehns, D. M., Wallacher, D., Pillai, R. S., et al. (2018). The Effect of Crystallite Size on Pressure Amplification in Switchable Porous Solids. *Nat. Commun.* 9, 1573. doi:10.1038/s41467-018-03979-2
- Krause, S., Reuter, F. S., Ehrling, S., Bon, V., Senkovska, I., Kaskel, S., et al. (2020). Impact of Defects and Crystal Size on Negative Gas Adsorption in DUT-49 Analyzed by *In Situ* 129Xe NMR Spectroscopy. *Chem. Mater.* 32, 4641–4650. doi:10.1021/acs.chemmater.0c01059
- Krylov, A., Senkovska, I., Ehrling, S., Maliuta, M., Krylova, S., Slyusareva, E., et al. (2020). Single Particle Raman Spectroscopy Analysis of the Metal–Organic Framework DUT-8(Ni) Switching Transition under Hydrostatic Pressure. *Chem. Commun.* 56, 8269–8272. doi:10.1039/D0CC02491K
- Kundu, T., Wahiduzzaman, M., Shah, B. B., Maurin, G., and Zhao, D. (2019). Solvent Induced Control over Breathing Behavior in Flexible Metal–Organic Frameworks for Natural Gas Delivery. *Angew. Chem. Int. Ed.* 58, 8073–8077. doi:10.1002/anie.201902738
- Le Grand, S., Gotz, A. W., and Walker, R. C. (2013). SPFP: Speed without Compromise – a Mixed Precision Model for GPU Accelerated Molecular Dynamics Simulations. *Comput. Phys. Commun.* 184, 374–380. doi:10.1016/j.cpc.2012.09.022
- Li, H., Eddaoudi, M., O’Keeffe, M., and Yaghi, O. M. (1999). Design and Synthesis of an Exceptionally Stable and Highly Porous Metal–Organic Framework. *Nature* 402, 276–279. doi:10.1038/46248
- Li, J.-R., Kuppler, R., and Zhou, H.-C. (2009). Selective Gas Adsorption and Separation in Metal–Organic Frameworks. *Chem. Soc. Rev.* 38, 1477–1504. doi:10.1039/B802426J
- Liu, Z., Gao, Q., Chen, J., Deng, J., Lin, K., and Xing, X. (2018). Negative Thermal Expansion in Molecular Materials. *Chem. Commun.* 54, 5164–5176. doi:10.1039/C8CC01153B
- Martyna, G. J., Tobias, D. J., and Klein, M. L. (1994). Constant Pressure Molecular Dynamics Algorithms. *J. Chem. Phys.* 101, 4177–4189. doi:10.1063/1.467468
- Mendt, M., Jee, B., Stock, N., Ahnfeldt, T., Hartmann, M., Himsel, D., et al. (2010). Structural Phase Transitions and thermal Hysteresis in the Metal–Organic Framework Compound MIL-53 as Studied by Electron Spin Resonance Spectroscopy. *J. Phys. Chem. C* 114, 19443–19451. doi:10.1021/jp107487g
- Merz, P. T., and Shirts, M. R. (2018). Testing for Physical Validity in Molecular Simulations. *PLoS ONE* 13, 1–22. doi:10.1371/journal.pone.0202764
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.* 21, 1087–1092. doi:10.1063/1.1699114
- Miura, H., Bon, V., Senkovska, I., Ehrling, S., Watanabe, S., Ohba, M., et al. (2017). Tuning the Gate-Opening Pressure and Particle Size Distribution of the Switchable Metal–Organic Framework DUT-8(Ni) by Controlled Nucleation in a Micromixer. *Dalton Trans.* 46, 14002–14011. doi:10.1039/C7DT02809A
- Nosé, S., and Klein, M. (1983). Constant Pressure Molecular Dynamics for Molecular Systems. *Mol. Phys.* 50, 1055–1076. doi:10.1080/00268978300102851
- Parrinello, M., and Rahman, A. (1981). Polymorphic Transitions in Single Crystals: A New Molecular Dynamics Method. *J. Appl. Phys.* 52, 7182–7190. doi:10.1063/1.328693
- Ramaswamy, P., Wieme, J., Alvarez, E., Vanduyfhuys, L., Itié, J.-P., Fabry, P., et al. (2017). Mechanical Properties of a Gallium Fumarate Metal–Organic Framework: A Joint Experimental-Modelling Exploration. *J. Mater. Chem. A* 5, 11047–11054. doi:10.1039/C7TA01559C
- Rogge, S. M. J., Caroes, S., Demuynek, R., Waroquier, M., Van Speybroeck, V., and Ghysels, A. (2018). The Importance of Cell Shape Sampling to Accurately Predict Flexibility in Metal–Organic Frameworks. *J. Chem. Theor. Comput.* 14, 1186–1197. doi:10.1021/acs.jctc.7b01134
- Rogge, S. M. J., Vanduyfhuys, L., Ghysels, A., Waroquier, M., Verstraeten, T., Maurin, G., et al. (2015). A Comparison of Barostats for the Mechanical Characterization of Metal–Organic Frameworks. *J. Chem. Theor. Comput.* 11, 5583–5597. doi:10.1021/acs.jctc.5b00748
- Rogge, S. M. J., Waroquier, M., and Van Speybroeck, V. (2019). Unraveling the Thermodynamic Criteria for Size-dependent Spontaneous Phase Separation in Soft Porous Crystals. *Nat. Commun.* 10, 4842. doi:10.1038/s41467-019-12754-w
- Sakata, Y., Furukawa, S., Kondo, M., Hirai, K., Horike, N., Takashima, Y., et al. (2013). Shape-memory Nanopores Induced in Coordination Frameworks by crystal Downsizing. *Science* 339, 193–196. doi:10.1126/science.1231451
- Schneemann, A., Bon, V., Schwedler, I., Senkovska, I., Kaskel, S., and Fischer, R. A. (2014). Flexible Metal–Organic Frameworks. *Chem. Soc. Rev.* 43, 6062–6096. doi:10.1039/C4CS00010J
- Shirts, M. R. (2013). Simple Quantitative Tests to Validate Sampling from Thermodynamic Ensembles. *J. Chem. Theor. Comput.* 9, 909–926. doi:10.1021/ct300688p
- Triguero, C., Coudert, F.-X., Boutin, A., Fuchs, A. H., and Neimark, A. V. (2011). Mechanism of Breathing Transitions in Metal–Organic Frameworks. *J. Phys. Chem. Lett.* 2, 2033–2037. doi:10.1021/jz2008769
- Tuckerman, M. E. (2010). *Statistical Mechanics Theory and Molecular Simulation*. Oxford graduate texts. New York: Oxford University Press.
- Van Speybroeck, V., Vandenhaute, S., Hoffman, A. E. J., and Rogge, S. M. J. (2021). Towards Modeling Spatiotemporal Processes in Metal–Organic Frameworks. *Trends Chem.* 3 (8), 605–619. doi:10.1016/j.trechm.2021.04.003
- Vandenhaute, S. (2021a). Implementation of the Proposed Anisotropic Barostat in OpenMM. Available at: <https://github.com/svandenhaute/openmm/ca2b47d8c1855f8f3ec6119d7077bdb8cc50a3b> (Accessed June 1, 2021).

- Vandenhaute, S. (2021b). OpenYAFF. Available at: <https://github.com/svandenhaute/openyaff/commit/d2b4265ae652bb89ad525bda93ea8a12b6e80a4> (Accessed June 1, 2021).
- Vanduyfhuys, L., Rogge, S. M. J., Wieme, J., Vandenbrande, S., Maurin, G., Waroquier, M., et al. (2018). Thermodynamic Insight into Stimuli-Responsive Behaviour of Soft Porous Crystals. *Nat. Commun.* 9, 204. doi:10.1038/s41467-017-02666-y
- Vanduyfhuys, L., Vandenbrande, S., Verstraelen, T., Schmid, R., Waroquier, M., and Van Speybroeck, V. (2015). QuickFF: A Program for a Quick and Easy Derivation of Force fields for Metal-Organic Frameworks from Ab Initio Input. *J. Comput. Chem.* 36, 1015–1027. doi:10.1002/jcc.23877
- Verstraelen, T., Vandenbrande, S., Heidar-Zadeh, F., Vanduyfhuys, L., Van Speybroeck, V., Waroquier, M., et al. (2016). Minimal Basis Iterative Stockholder: Atoms in Molecules for Force-Field Development. *J. Chem. Theor. Comput.* 12, 3894–3912. doi:10.1021/acs.jctc.6b00456
- Wahiduzzaman, M., Reimer, N., Itié, J.-P., Stock, N., Maurin, G., and Yot, P. G. (2018). Mechanical-pressure Induced Response of the MOF Al-MIL-53-TDC. *Polyhedron* 155, 144–148. doi:10.1016/j.poly.2018.08.045
- Wannapaiboon, S., Schneemann, A., Hante, I., Tu, M., Epp, K., Semrau, A. L., et al. (2019). Control of Structural Flexibility of Layered-Pillared Metal-Organic Frameworks Anchored at Surfaces. *Nat. Commun.* 10, 346. doi:10.1038/s41467-018-08285-5
- Wieme, J., Lejaeghere, K., Kresse, G., and Van Speybroeck, V. (2018). Tuning the Balance between Dispersion and Entropy to Design Temperature-Responsive Flexible Metal-Organic Frameworks. *Nat. Commun.* 9, 4899. doi:10.1038/s41467-018-07298-4
- Wieme, J., Rogge, S. M. J., Yot, P. G., Vanduyfhuys, L., Lee, S.-K., Chang, J.-S., et al. (2019). Pillared-layered Metal Organic Frameworks for Mechanical Energy Storage Applications. *J. Mater. Chem. A* 7, 22663–22674. doi:10.1039/C9TA01586H
- Yaghi, O., O'Keeffe, M., Ockwig, N., Chae, H., Eddaoudi, M., and Kim, J. (2003). Reticular Synthesis and the Design of New Materials. *Nature* 423, 705–714. doi:10.1038/nature01650
- Yot, P. G., Boudene, Z., Macia, J., Granier, D., Vanduyfhuys, L., Verstraelen, T., et al. (2014). Metal-organic Frameworks as Potential Shock Absorbers: the Case of the Highly Flexible MIL-53(AI). *Chem. Commun.* 50, 9462–9464. doi:10.1039/C4CC03853C
- Yot, P. G., Ma, Q., Haines, J., Yang, Q., Ghoufi, A., Devic, T., et al. (2012). Large Breathing of the MOF MIL 47(V^{IV}) under Mechanical Pressure: a Joint Experimental Modelling Exploration. *Chem. Sci.* 3, 1100–1104. doi:10.1039/C2SC00745B
- Yot, P. G., Vanduyfhuys, L., Alvarez, E., Rodriguez, J., Itié, J.-P., Fabry, P., et al. (2016). Mechanical Energy Storage Performance of an Aluminum Fumarate Metal-Organic Framework. *Chem. Sci.* 7, 446–450. doi:10.1039/C5SC02794B
- Yot, P. G., Wahiduzzaman, M., Elkaim, E., Fertey, P., Fabry, P., Serre, C., et al. (2019). Modulation of the Mechanical Energy Storage Performance of the MIL-47(VIV) Metal Organic Framework by Ligand Functionalization. *Dalton Trans.* 48, 1656–1661. doi:10.1039/C8DT04214D
- Zhang, Z., Liu, X., Yan, K., Tuckerman, M. E., and Liu, J. (2019). Unified Efficient Thermostat Scheme for the Canonical Ensemble with Holonomic or Isokinetic Constraints via Molecular Dynamics. *J. Phys. Chem. A* 123, 6056–6079. doi:10.1021/acs.jpca.9b02771

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2021 Vandenhaute, Rogge and Van Speybroeck. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Supplementary Material

1 EXPERIMENTAL PROPERTIES OF MIL-53(AL)

Table S1 gives an overview of the experimentally determined structural parameters of MIL-53(Al), as reported by Yot et al. (2014).

	large pore	closed pore
space group	<i>Imcm</i>	<i>C2/c</i>
<i>a</i> [Å]	16.7322(1)	19.633(1)
<i>b</i> [Å]	16.7322(7)	7.160(1)
<i>c</i> [Å]	6.6295(2)	6.559(1)
β [deg]	–	104.70(1)
volume [Å ³]	1423.8(1)	897.2(6)

Table S1. Experimental properties on the structure of MIL-53(Al) (Yot et al., 2014).

2 INTEGRATION OVER LOWER TRIANGULAR CELL MATRICES

This section demonstrates how we may write the isothermal-isobaric partition function into an integration over the six degrees of freedom of a lower triangular cell instead of the nine degrees of freedom for an arbitrarily oriented cell. To achieve this, we first need to define the transformation between the initial cell matrix \mathbf{h} and the lower triangular cell matrix \mathbf{h}_Δ . This transformation is conveniently obtained by computing the QR decomposition of the transpose of \mathbf{h} :

$$QR(\mathbf{h}^T) = \mathbf{q}\mathbf{h}_\Delta^T \quad (\text{S1})$$

and with \mathbf{q} the orthonormal rotation matrix. We then obtain:

$$\mathbf{h}_\Delta = \mathbf{h}\mathbf{q} \quad (\text{S2})$$

$$= \begin{bmatrix} a_x & 0 & 0 \\ b_x & b_y & 0 \\ c_x & c_y & c_z \end{bmatrix} \quad (\text{S3})$$

and for the volume of the unit cell:

$$V = \det(\mathbf{h}) \quad (\text{S4})$$

$$= \det(\mathbf{h}_\Delta) \quad (\text{S5})$$

$$= a_x b_y c_z \quad (\text{S6})$$

Equation S3 defines the change of variables that must be performed in the integral of the partition function (repeated here for convenience):

$$\Delta(N, P, T) = C \int \int e^{-\beta U(\mathbf{r})} e^{-\beta PV} \det(\mathbf{h})^{-2} d\mathbf{h} d\mathbf{r} \quad (\text{S7})$$

with C an irrelevant constant, N the number of particles in the unit cell, P the hydrostatic pressure, T the temperature, β the inverse temperature, and finally $U(\mathbf{r})$ the potential energy of the system as a function of the cartesian coordinates of all atoms (which contains an implicit dependence on the cell matrix). It is important to keep in mind that while equation S3 represents a simple rotation of the coordinate axes, the rotation matrix \mathbf{q} is itself not a constant and instead depends on \mathbf{h} (the integration variable), and this complicates the transformation of the differentials. The Jacobian of the change of variables in equation S3 has been computed elsewhere (see e.g. Edelman and Rao (2005), equation 3.6) and was found to be:

$$(d\mathbf{h}) = a_x^2 b_y (d\mathbf{h}_\Delta) (\mathbf{q}^T d\mathbf{q}) \quad (\text{S8})$$

Based on the obtained Jacobian, we may now perform the substitution suggested in equation S3:

$$\Delta(N, P, T) = C \int \int e^{-\beta U(\mathbf{r})} e^{-\beta PV} \det(\mathbf{h})^{-2} d\mathbf{h} d\mathbf{r} \quad (\text{S9})$$

$$= C \int \int e^{-\beta U(\mathbf{r})} e^{-\beta PV} \det(\mathbf{h}_\Delta)^{-2} a_x^2 b_y (\mathbf{d}\mathbf{h}_\Delta) (\mathbf{q}^T d\mathbf{q}) d\mathbf{r} \quad (\text{S10})$$

$$= C' \int \int e^{-\beta U(\mathbf{r})} e^{-\beta PV} \det(\mathbf{h}_\Delta)^{-2} a_x^2 b_y (\mathbf{d}\mathbf{h}_\Delta) d\mathbf{r} \quad (\text{S11})$$

$$= C' \int \int e^{-\beta U(\mathbf{r})} e^{-\beta PV} (a_x b_y c_z)^{-2} a_x^2 b_y (\mathbf{d}\mathbf{h}_\Delta) d\mathbf{r} \quad (\text{S12})$$

$$= C' \int \int e^{-\beta U(\mathbf{r})} e^{-\beta PV} (b_y c_z^2)^{-1} (\mathbf{d}\mathbf{h}_\Delta) d\mathbf{r} \quad (\text{S13})$$

whereby C' is introduced to absorb the constant integration over the rotation matrix \mathbf{q} . Equation S13 is identical to the result stated in equation 3 in the main text.

3 VALIDATION USING THE ANISOTROPIC CRYSTAL

The proposed pressure control algorithm is critically validated using a fictitious crystal that is designed to exhibit strong anisotropy. The triclinic unit cell of the crystal is described by the following cell matrix (in angstrom):

$$\mathbf{h}_\Delta = \begin{bmatrix} a_x & 0 & 0 \\ b_x & b_y & 0 \\ c_x & c_y & c_z \end{bmatrix} = \begin{bmatrix} 60 & 0 & 0 \\ 20 & 70 & 0 \\ 20 & 20 & 80 \end{bmatrix} \quad (\text{S14})$$

The resulting box vector lengths and angles are given in Figure S1A. We then consider 27 identical atoms and place them on a regular $3 \times 3 \times 3$ grid, with grid lines parallel to either \mathbf{a} , \mathbf{b} , or \mathbf{c} . This is schematically depicted in Figure S1B. In this geometry, the average distance between atoms is over 20 Å, which is of course much larger than the average atomic separation found in real materials. However, this is intentionally chosen to improve the convergence of the internal stress tensor σ_{int} , which has a volume dependence in its denominator (Rogge et al., 2015):

$$\sigma_{\text{int}} = \frac{1}{V} \left[\sum_{i=1}^N \frac{\mathbf{p}_i \otimes \mathbf{p}_i}{m_i} - \mathbf{h}_\Delta^T \frac{\partial U}{\partial \mathbf{h}_\Delta} \right] \quad (\text{S15})$$

where \otimes denotes the exterior product between two vectors: $(\mathbf{u} \otimes \mathbf{v})_{ij} = u_i v_j$. Nearest-neighbor interactions are introduced in the system in order to stabilize the crystalline phase shown in Figure S1B. Only two types of interactions are necessary to achieve this: harmonic bonds and harmonic angles. To enforce anisotropic behavior (beyond the strongly triclinic shape of the unit cell), different force constants and rest values are used depending on the direction of the bond or the plane of the angle; these are summarized in Figure S1C. No other interactions were present in the system, i.e. the constituting atoms only interact through the aforementioned short-range covalent interactions. We then proceeded by performing a number of isothermal-isobaric MD simulations, at different pressures. The resulting trajectories were analyzed in terms of the average value of σ_{int} (section 3.1) as well as the volume distributions (section 3.2).




Figure S1. (A) Description of the (initial) unit cell parameters of the anisotropic crystal. (B) Simplified representation of the three-dimensional structure of the crystal. The unit cell contains 27 atoms that are positioned on a regular $3 \times 3 \times 3$ grid, with grid lines parallel to the box vectors. (C) Parameters for the covalent interactions. Rest values of the interactions are obtained as the actual values in the initial geometry described in (B). For the angles, this implies that both ϕ_0 and $\pi - \phi_0$ are used as rest values, depending on the value of the angle.

3.1 Convergence of the internal stress tensor

Samples that are distributed according to the isothermal-isobaric ensemble satisfy the so-called pressure virial theorem, which states that the average value of σ_{int} equals the externally applied stress σ_{ext} (Tuckerman, 2010):

$$\langle \sigma_{\text{int}} \rangle = \sigma_{\text{ext}} \quad (\text{S16})$$

and for a hydrostatic pressure P , this implies

$$\langle \sigma_{\text{int}} \rangle = P \mathbf{1} \quad (\text{S17})$$

$$= \begin{bmatrix} P & 0 & 0 \\ 0 & P & 0 \\ 0 & 0 & P \end{bmatrix} \quad (\text{S18})$$

We may verify equation S18 by computing the internal stress σ_{int} at regular intervals during a molecular dynamics simulation and monitoring its running average. In doing so, it is useful to decompose the total internal stress σ_{int} into an isotropic and anisotropic contribution:

$$\sigma_{\text{int}} = \frac{\text{tr}(\sigma_{\text{int}})}{3} \mathbf{1} + \left(\sigma_{\text{int}} - \frac{\text{tr}(\sigma_{\text{int}})}{3} \mathbf{1} \right) \quad (\text{S19})$$

$$= \sigma_{\text{iso}} + \sigma_{\text{a}} \quad (\text{S20})$$

whereby naturally

$$\text{tr}(\sigma_{\text{a}}) = 0 \quad (\text{S21})$$

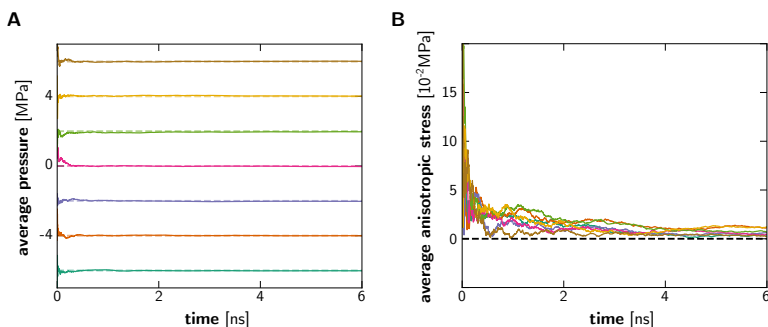


Figure S2. (A) Running average of the internal pressure $P_{\text{int}} = \text{tr}(\sigma_{\text{iso}})/3$ as a function of simulation time, for seven different pressures (-6 MPa to 6 MPa in steps of 2 MPa). Dashed lines indicate the applied pressure in each of the simulations. (B) Root-mean-square of the components of the running average of σ_{a} as a function of simulation time, with colors corresponding to those in (A). All running averages are converging to zero.

Here, it is important to note that the derivative of U with respect to the cell matrix parameters (the second term in equation S15) is not always computed analytically by molecular mechanics engines – such as OpenMM – because it is computationally expensive to evaluate and rather tedious to implement. In those cases, it is necessary to compute the derivative in the virial stress expression (equation S15) numerically, by applying small perturbations in each of the six components and computing the resulting change in potential energy. For numerical reasons, perturbations were performed in the six components of the *symmetric* cell matrix instead of the lower triangular components. Both the running averages of the pressure and the mean squared error of the components of σ_a are displayed in Figure S2 for a variety of pressures. We indeed observe that σ_{int} converges to a diagonal matrix, as required by equation S18.

3.2 Ensemble validation

By performing simulations at different temperatures and pressures, we may obtain histograms of observables (e.g. the potential energy or the unit cell volume) for each temperature and pressure. These histograms may then be compared with each other in order to validate whether they are distributed according to the correct Boltzmann distribution (in our case, the (N, P, T) distribution function of Equation S7). This type of ensemble validation was first proposed by Michael R. Shirts (Shirts, 2013), and we refer the reader to his work for a detailed analytical derivation. There, it is shown that the log likelihood ratio of two volume histograms obtained at different pressures satisfies a linear relation that is a function of the temperature and the difference in applied pressure. Here, we report the results of these volume distribution tests in Figure S3. This was obtained using the `physical-validation` Python package¹. Our results confirm that the proposed pressure control algorithm does indeed generate the correct volume distributions.

P_0 [MPa]	P_1 [MPa]	analytical slope [nm^{-3}]	estimated slope [nm^{-3}]	estimated ΔP [MPa]
-6	-4	-0.4829	-0.4796 +/- 0.0052	-1.99 +/- 0.02
-4	-2	-0.4829	-0.4795 +/- 0.0052	-1.99 +/- 0.02
-2	0	-0.4829	-0.4784 +/- 0.0051	-1.98 +/- 0.02
0	2	-0.4829	-0.4808 +/- 0.0052	-1.99 +/- 0.02
2	4	-0.4829	-0.4805 +/- 0.0050	-1.99 +/- 0.02
4	6	-0.4829	-0.4820 +/- 0.0050	-2.00 +/- 0.02

Figure S3. Results of the ensemble validation procedure (Shirts, 2013). A total of seven simulations were performed at pressures between -6 MPa and 6 MPa, at 300K . Volume distributions were used to determine a maximum likelihood estimate of the pressure difference between two simulations. For each pair of simulations, the estimated pressure difference is in near-perfect agreement with the actual difference in applied pressure.

¹ https://github.com/shirtsgroup/physical_validation/commit/7d3ae2ea71ebaf2cac7b4ab7d922be2c0dfb6639

4 UMBRELLA SAMPLING PROTOCOL

The MIL-53(AI) framework exhibits a large-amplitude structural transition which may be characterized by computing the Helmholtz free energy F as a function of the unit cell volume V using umbrella sampling (see Demuynck et al. (2017)). The interaction energy between particles was modelled using the same force field that was used in the main text. We computed the free energy variation along V using a reference implementation of the MTK barostat using YAFF², and a reference implementation of the proposed Monte Carlo pressure control algorithm in YAFF. The calculations were performed on a small $1 \times 2 \times 1$ supercell of the MIL-53(AI) framework. This is done in order to limit the number of particles and avoid the thermodynamic limit (in which case differences between ensembles disappear). Umbrellas were introduced by applying a harmonic restraint to the unit cell volume V :

$$U_H(V) = K(V - V_0)^2 \quad (\text{S22})$$

with $V = \det \mathbf{h}$. The center of the harmonic restraint, V_0 , was varied between 1480 \AA^3 (volume of the cp phase) and 3120 \AA^3 (volume of the lp phase), in steps of 20 \AA^3 . The force constant K was set at $0.0114 \text{ kJ}/[\text{mol} \cdot \text{\AA}^6]$, which resulted in a sufficient overlap of the volume distributions between neighboring umbrellas. Resulting trajectories were analyzed using `pymbar` (Shirts and Chodera, 2008) in order to obtain the free energy profile as a function of V . The volume axis shown in Figure 2 in the main text refers to the volume of the $1 \times 1 \times 1$ cell, i.e. half of the $1 \times 2 \times 1$ cell used in the calculations.

² available at <https://github.com/molmod/yaff/commit/d33861ab81bbe621364a8f58f7b77f65d03206f7>

5 OVERVIEW OF PERFORMED SIMULATIONS

Several simulations were performed in order to identify the main transition mechanisms and understand how they depend on the periodicity. Table S2 gives an overview of all simulations. The $37 \times 10 \times 37$ system is topologically square, in the sense that the horizontal and vertical direction contain the same number of layers. The $29 \times 10 \times 37$ system is slightly elongated in one direction, with 29 layers horizontally and 37 layers vertically. Due to the difference in the number of layers in the horizontal and vertical direction, this model does not allow a single diagonal to transition from the lp to the cp phase because the 29×37 periodicity is not compatible with a single cp layer structure. Because of this, layer-by-layer transitions are prohibited in an artificial manner.

supercell	temperature [K]	pressure [MPa]	number of atoms
$37 \times 10 \times 37$	200	100	1,040,440
$37 \times 10 \times 37$	200	300	1,040,440
$37 \times 10 \times 37$	200	500	1,040,440
$37 \times 10 \times 37$	300	100	1,040,440
$37 \times 10 \times 37$	300	300	1,040,440
$37 \times 10 \times 37$	300	500	1,040,440
$37 \times 10 \times 37$	500	100	1,040,440
$37 \times 10 \times 37$	500	300	1,040,440
$37 \times 10 \times 37$	500	500	1,040,440
$29 \times 10 \times 37$	200	100	815,480
$29 \times 10 \times 37$	200	300	815,480
$29 \times 10 \times 37$	200	500	815,480
$29 \times 10 \times 37$	300	100	815,480
$29 \times 10 \times 37$	300	300	815,480
$29 \times 10 \times 37$	300	500	815,480
$29 \times 10 \times 37$	500	100	815,480
$29 \times 10 \times 37$	500	300	815,480
$29 \times 10 \times 37$	500	500	815,480

Table S2. Overview of the performed simulations.

REFERENCES

- Demuyne, R., Rogge, S. M. J., Vanduyfhuys, L., Wieme, J., Waroquier, M., and Van Speybroeck, V. (2017). Efficient construction of free energy profiles of breathing metal–organic frameworks using advanced molecular dynamics simulations. *J. Chem. Theory Comp.* 13, 5861–5873. doi:10.1021/acs.jctc.7b01014
- Edelman, A. and Rao, N. R. (2005). Random matrix theory. *Acta Numer.* 14, 233–297. doi:10.1017/S0962492904000236
- Rogge, S., Vanduyfhuys, L., Ghysels, A., Waroquier, M., Verstraelen, T., Maurin, G., et al. (2015). A comparison of barostats for the mechanical characterization of metal–organic frameworks. *J. Chem. Theory Comput.* 11, 5583–5597. doi:10.1021/acs.jctc.5b00748
- Shirts, M. R. (2013). Simple quantitative tests to validate sampling from thermodynamic ensembles. *J. Chem. Theory Comput.* 9, 909–926. doi:10.1021/ct300688p
- Shirts, M. R. and Chodera, J. D. (2008). Statistically optimal analysis of samples from multiple equilibrium states. *J. Chem. Phys.* 129, 124105. doi:10.1063/1.2978177
- Tuckerman, M. E. (2010). *Statistical mechanics theory and molecular simulation*. Oxford graduate texts (New York: Oxford University Press)
- Yot, P. G., Boudene, Z., Macia, J., Granier, D., Vanduyfhuys, L., Verstraelen, T., et al. (2014). Metal–organic frameworks as potential shock absorbers: the case of the highly flexible MIL-53(AI). *Chem. Commun.* 50, 9462–9464. doi:10.1039/C4CC03853C

Paper II

Machine learning potentials for metal-organic frameworks using an incremental learning approach

Sander Vandenhoute, Maarten Cools-Ceuppens, Simon DeKeyser, Toon Verstraelen, Veronique Van Speybroeck

npj Computational Materials, 2023, 9, 19

Sander Vandenhoute developed and implemented the on-the-fly learning algorithm and performed all the simulations.

Reprinted with permission.
Copyright (2023), The Authors

ARTICLE OPEN



Machine learning potentials for metal-organic frameworks using an incremental learning approach

Sander Vandenhoute¹, Maarten Cools-Ceuppens¹, Simon DeKeyser¹, Toon Verstraelen¹ and Veronique Van Speybroeck¹✉

Computational modeling of physical processes in metal-organic frameworks (MOFs) is highly challenging due to the presence of spatial heterogeneities and complex operating conditions which affect their behavior. Density functional theory (DFT) may describe interatomic interactions at the quantum mechanical level, but is computationally too expensive for systems beyond the nanometer and picosecond range. Herein, we propose an incremental learning scheme to construct accurate and data-efficient machine learning potentials for MOFs. The scheme builds on the power of equivariant neural network potentials in combination with parallelized enhanced sampling and on-the-fly training to simultaneously explore and learn the phase space in an iterative manner. With only a few hundred single-point DFT evaluations per material, accurate and transferable potentials are obtained, even for flexible frameworks with multiple structurally different phases. The incremental learning scheme is universally applicable and may pave the way to model framework materials in larger spatiotemporal windows with higher accuracy.

npj Computational Materials (2023)9:19; <https://doi.org/10.1038/s41524-023-00969-x>

INTRODUCTION

Metal-organic frameworks (MOFs) have become one of the most intriguing classes of materials in the last decades as they may exhibit anomalous responses upon exposure to external stimuli such as temperature, pressure, or guest sorption; notable examples are negative thermal expansion, negative gas adsorption, and large-amplitude structural transitions. Apart from being academic peculiarities, MOFs have now also found their way into technologically important applications such as gas storage, sensing, separation, or catalysis.^{1–9} Recently, it became clear that the dynamic response of MOFs is largely affected by spatial heterogeneities at various length scales, varying from the subnanometer scale for isolated defects to the mesoscale when correlated nanoregions are present and to the microscale when accounting for the finite crystal size.^{10–14} It led to the terminology of spatiotemporal processes within MOFs, referring to the entanglement between the dynamic response of the material and the existing spatial disorder.^{15,16} Modeling spatiotemporal processes in realistic MOFs at length and time scales comparable to experimental conditions is a very ambitious goal, and requires a range of innovative methods to translate atomic-scale information into macroscopic insight. One key ingredient is the method used to describe the interatomic forces. Current modeling efforts on MOFs either rely on density functional theory (DFT), which may obtain an accurate description of framework-guest interactions when complemented with proper dispersion models to describe longer-ranged interactions, or on classical force fields, which are simple analytical functions that ignore the quantum mechanical description of the electrons. Classical force fields have become tractable within the MOF field as they are now starting to bridge the length scale gap towards experimentally observed MOF crystallites^{17,18}. However, they are less accurate than DFT-based methods and exhibit only limited transferability; force fields derived under certain thermodynamic conditions are not necessarily applicable to other operating conditions. In addition, they are generally unable to describe bond formation or breakage. For these reasons, DFT-based methods are in principle preferred but

their applicability is limited to nanometer-sized structural models and time scales up to about 100 ps, even on state-of-the-art computing infrastructure. Ideally, it would be possible to combine the best of both worlds, whereby interatomic forces are evaluated at an accuracy comparable to DFT but with a computational efficiency similar to classical force fields. Machine learning potentials (MLPs) may offer such a hybrid alternative; they may learn the potential energy surface (PES) at a given level of quantum chemical accuracy, and can then be used to accelerate subsequent energy evaluations by multiple orders of magnitude^{19,20}. MLPs have been successfully developed for a variety of different materials and molecules. Broadly speaking, these models may be categorized into either kernel regression methods, which determine the interaction energy by comparing a given configuration to a set of reference configurations^{21,22}, or neural network potentials, which directly determine a high-dimensional representation of the PES based on thousands or even millions of parameters²³.

Since 2019, a number of studies have emerged in which neural network potentials have been employed to study e.g., mechanical or diffusion properties for framework materials^{24–26}. Although they demonstrate the potential of MLPs in the field of modeling MOFs, they required a large set of expensive DFT molecular dynamics (MD) trajectories in order to generate the necessary amount of training data (over 10,000 snapshots in all three cases). Such approaches become computationally intractable for systems where large portions of phase space need to be sampled. This is the case for more complex frameworks with large unit cells such as HKUST-1(Cu) or MOF-808(Zr), but also flexible frameworks with multiple stable phases pose a real challenge in terms of generating training data in a computationally efficient manner (e.g., MIL-53(Al) in Fig. 1). For these frameworks, all relevant regions in phase space need to be properly sampled, including not only the stable phases of the material within the thermodynamic conditions of interest but also activated regions and all points along important transition paths. This is very difficult to achieve using equilibrium first principles MD because simulation

¹Center for Molecular Modeling, Ghent University, Technologiepark 46, 9052 Zwijnaarde, Belgium. ✉email: veronique.vanspeybroeck@ugent.be




Fig. 1 Overview of the atomic structures of the frameworks. UiO-66(Zr) is a mechanically rigid framework with a single stable phase; MIL-53(Al) is a flexible material.

times are limited when using DFT to evaluate the interatomic forces at each timestep. In addition, configurations obtained from equilibrium MD typically sample regions close to free energy minima and fail to capture important transition states in between phases, even though these are essential to include in the training data.

In this work, we present an incremental learning scheme for the construction of accurate MLPs with thermodynamic transferability without requiring large amounts of DFT input data. We employ NequIP²⁷, a newly proposed message passing neural network (MPNN) that exhibits extraordinary data efficiency due to its use of equivariant feature representations for the atomic environments (see below)²⁸. The key observation in our learning scheme is that a proper sampling of the phase space can be achieved without performing extensive quantum mechanical-based molecular dynamics simulations. Instead, inspired by other active learning approaches^{29–31}, we propose to sample the phase space using the MLP itself, in combination with on-the-fly training whenever it encounters unknown regions. In this sense, systematically improved MLPs are generated. Key to our approach is the use of metadynamics (MTD) in order to enforce the sampling towards unexplored regions. This allows us to ensure that both in- and out-of-equilibrium configurations are automatically included during training, even when large free energy barriers are present. Importantly, we show that our incremental learning algorithm is highly efficient and does not require a kind of uncertainty metric that determines which samples to include in the training data, as is often the case for other active learning approaches proposed elsewhere^{30,32}. We demonstrate the accuracy and transferability of models constructed with the incremental learning scheme based on two representative frameworks (Fig. 1): the mechanically rigid UiO-66(Zr) framework, which consists of twelve-fold coordinated inorganic bricks and organic benzene-1,4-dicarboxylate (BDC) ligands³³, and the flexible MIL-53(Al) framework, which consists of one-dimensional (periodic) aluminum hydroxide chains connected with the same organic ligands³⁴. Incremental learning addresses the main disadvantages of traditional data generation with an iterative learning approach, in which efficient and parallel enhanced sampling molecular dynamics simulations are combined with on-the-fly learning. Based on a single atomic structure and a definition of one or more collective variables, the algorithm proceeds with simultaneously exploring and learning the phase space of the material; the approach is schematically shown in Fig. 2. The algorithm starts by constructing a first generation MLP, which is trained based on a small set of initial configurations that are obtained by applying random perturbations to the particle positions and strain components of the initial configuration. This first generation MLP is then used in a short (~1 ps) multiple walker metadynamics (MTD) simulation in order to explore the phase

space around the initial configuration. The final configuration of each of the walkers is extracted and subjected to a new DFT calculation to obtain the energy and forces. The latter are added to the training and validation sets after which a next generation MLP is obtained by training the model for a short amount of time. Even though the sampling was performed with an initially inaccurate MLP, it still suffices to explore a meaningful region of phase space and generate almost decorrelated samples. The sampling time per iteration may be kept relatively short as it serves the purpose to gradually explore larger portions of phase space. After training, the model is considered to have learned an incremental region in phase space, and may then be passed on to the next iteration in which it continues the MTD sampling. The state of the bias potential is retained from the previous iteration as to ensure that the walkers explore a slightly different region of phase space in each iteration. By continuously alternating the enhanced sampling and training steps, the entire phase space of the collective variable is explored and learned on-the-fly without the need to perform expensive DFT-based molecular dynamics simulations. Importantly, our approach ensures that atomic configurations for which a QM evaluation is performed are always separated by a ~1 ps MTD trajectory, which implies that they are at most only weakly correlated. As such, we are guaranteed that there is little redundancy between the QM-evaluated configurations in the training data even though we did not rely on any specialized uncertainty measures as found in other active learning schemes^{30,32}.

The incremental learning approach as proposed here is most efficient when combined with MLP models which are easy to train and which are able to learn a large number of atomic environments and chemical species. While existing MLPs for MOFs used strictly invariant architectures^{24–26}, we employ the neural equivariant interatomic potential (NequIP)²⁷ because it achieves a data efficiency similar to kernel regression methods while maintaining a cost of evaluation that is independent of the size of the training set. Like many other message passing neural networks (MPNNs)^{35,36}, NequIP computes the total potential energy of a given configuration as a sum of individual atomic contributions, whereby each contribution is the result of a series of convolutions taken over neighboring atomic environments as to represent the physical interactions between nearby atoms. Equivariant MPNNs differ from more established MPNNs such as SchNet³⁷ or DeePMD³⁸ because they complement rotationally invariant, scalar-like features with equivariant features in the representation of atomic environments³⁹ (Fig. 3a). Inspired by tensor field networks²⁹, NequIP assigns each of the features to a specific irreducible representation of the rotational group SO(3). Each irreducible representation is characterized by a rotation order $\ell \in \mathbb{N}$, which defines how features transform when the

cartesian coordinates at the input are rotated. The first two rotation orders ($\ell = 0$ and $\ell = 1$) correspond to scalar and vector features respectively, and are illustrated in Fig. 3a. Effectively, features for which $\ell > 0$ are better suited to represent directional information of the atomic environment throughout the layers of the network, and are the primary reason for the exceptional data

efficiency of equivariant MPNNs as compared to invariant MPNNs and even kernel regression methods²⁷. To demonstrate the importance of equivariant features, we performed separate incremental learning runs for the invariant ($\ell_{\max} = 0$) and equivariant ($\ell_{\max} = 1$) networks, and monitored the mean absolute error (MAE) on the force predictions for an independent

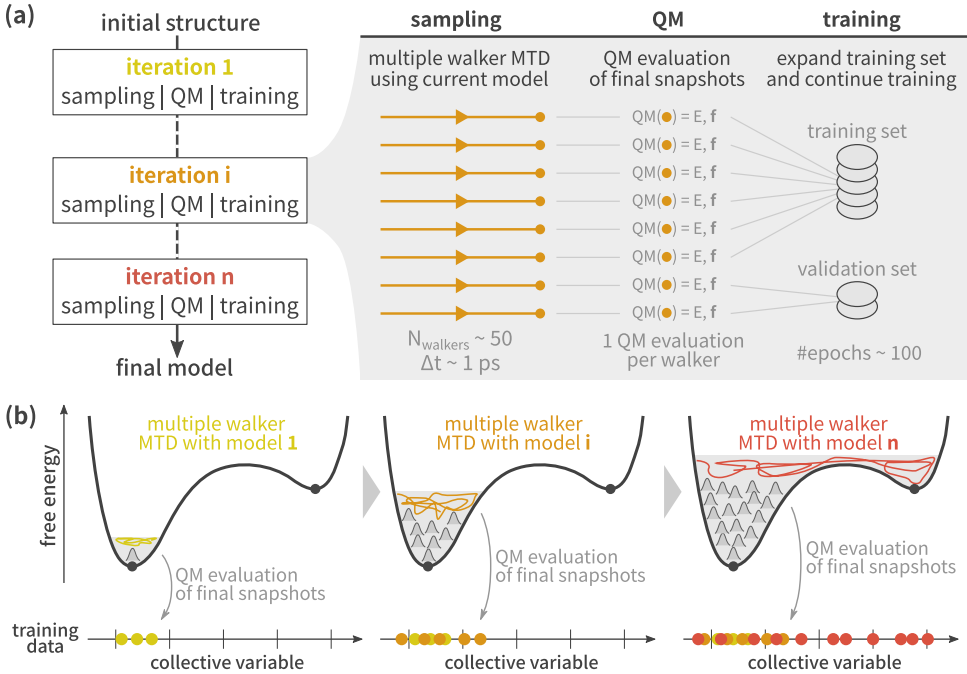


Fig. 2 Overview of the incremental learning approach as a combination of enhanced sampling and on-the-fly training. **a** Each iteration employs the current MLP in multiple walker metadynamics to explore increasingly larger regions in phase space, after which the obtained samples are evaluated at the QM level and the MLP is retrained. **b** Metastable regions and transition paths are automatically explored and included in the training data.

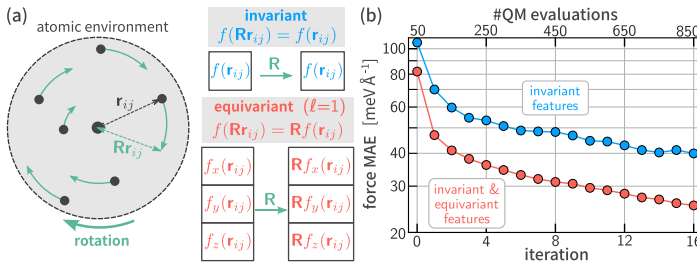


Fig. 3 Comparison between invariant and equivariant feature representations. **a** Features in equivariant neural networks are associated with a specific irreducible representation of $SO(3)$, as characterized by the rotation order ℓ . For $\ell = 0$, features do not transform upon rotation of the system, and may be considered as scalars. For $\ell = 1$, features transform according to vectors in \mathbb{R}^3 . **b** Evolution of the test error as a function of the number of iterations in the incremental learning algorithm as well as the total number of QM evaluations that were performed for both the invariant ($\ell_{\max} = 0$ model (blue)) and equivariant ($\ell_{\max} = 1$ model (red)). The test system was UiO-66(Zr). Computational details regarding network hyperparameters, MTD parameters, and test set generation can be found in the Methods section as well as the Supplementary Information.

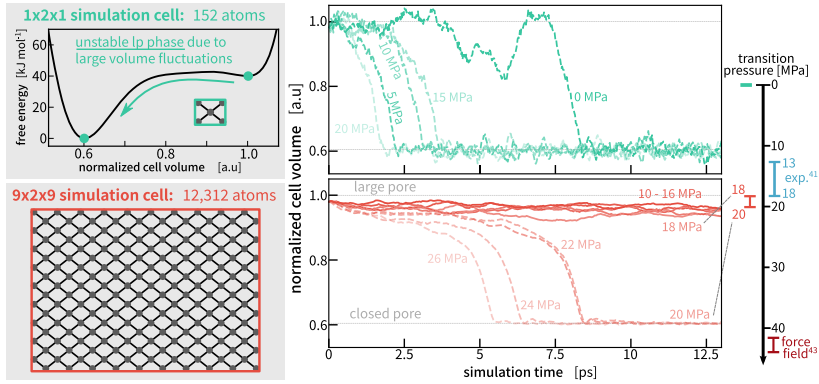


Fig. 5 Estimating the transition pressure via large-scale dynamics. We compare estimates using a regular $1 \times 2 \times 1$ unit cell versus a large-scale $9 \times 2 \times 9$ unit cell, as obtained based on (N, P, T) dynamics. All unit cell volumes were normalized with respect to the lp minimum of the Helmholtz free energy surface at 300 K, which occurs at 2882 Å. Nevertheless, the lp phase is not stable for the $1 \times 2 \times 1$ cell due to the larger volume fluctuations. In contrast, the $9 \times 2 \times 9$ cell allows to determine a transition pressure in the range of 18–20 MPa, which is in good agreement with experiment⁴¹.

trained MLPs at only a fraction of the computational cost when compared to previous learning approaches even though the accuracy is maintained across thermodynamic conditions (see Supplementary Note 4). In fact, the total computational cost of training data generation is now equivalent to or even less than a routine geometry optimization or a very short first principles MD trajectory (excluding the GPU resources required during training).

Mechanical stability and phase transitions

Interestingly, while the network parameters were optimized using a weighted loss function that takes both energies and forces into account, the NequIP models are also capable of providing accurate predictions for the virial stress in the system (see Supplementary Note 4). This is a necessary requirement in order to generate representative (N, P, T) dynamics, and it already suggests that the overall mechanical behavior of the frameworks will be well reproduced by the MLPs. To further investigate this, we compared the final models with the DFT reference in terms of structural and mechanical properties at 0 K. Optimized geometries for UiO-66(Zr) and both phases of MIL-53(Al) are in good agreement with the DFT reference, with RMSD values below 0.05 Å on both atomic positions and box vector components. For each of the optimized geometries, we evaluate the mechanical behavior in terms of the stiffness tensor $C \in \mathbb{R}^{6 \times 6}$ (in Voigt notation), which determines how the stress $\sigma \in \mathbb{R}^6$ within the material changes due to an applied strain $\epsilon \in \mathbb{R}^6$:

$$\sigma = C\epsilon \quad (1)$$

The components of C are computed using either a finite-difference approach (used for the DFT reference), or with an exact second-order hessian matrix (used for the MLPs); both methods are outlined in Supplementary Note 3.

The bottom part of Fig. 4 shows the stiffness tensors for all three optimized structures; each of the squares is color-coded based on the relative deviation between the MLP prediction and the DFT reference, whereas the value of the latter is indicated in black text inside each square. None of the predicted stiffness constants differed by more than 7 GPa from the DFT reference values, demonstrating that our models are capable of reproducing the mechanical properties of the frameworks, as is to be expected based on the force and stress

error assessments. To further demonstrate their thermodynamic transferability, we explicitly performed (N, P, T) MD simulations at a large range of pressures and validated the obtained trajectories with the underlying DFT reference; the results are presented in Supplementary Note 4. The constructed MLPs exhibit essentially the same accuracy as the underlying level of theory while being at least three orders of magnitude faster to evaluate. To further demonstrate this enormous gain in computational efficiency, we will investigate the threshold pressure for the large pore (lp) to closed pore (cp) transition in MIL-53(Al) at 300 K, which was previously estimated at 13–18 MPa based on mercury porosimetry experiments⁴¹. Computational prediction of the transition pressure is difficult because *ab initio* (N, P, T) MD simulations of MIL-53(Al) typically exhibit large volume fluctuations due to the small unit cell size (with $\Delta V/V$ on the order of 10%), and these may act as premature triggers for the phase transition⁴². Fig. 5 visualizes the transition dynamics for a typical $1 \times 2 \times 1$ cell used in *ab initio* simulations. The framework exhibits lp-to-cp transitions for any nonnegative pressure, suggesting that the lp phase is unstable at 300 K in spite of a clear lp minimum on the Helmholtz free energy surface of the material. In addition, the absence of any correlation between the timing of the transitions and the magnitude of the applied pressure further suggests that transitions at this scale are driven entirely by unit cell volume fluctuations⁴². Fortunately, the extraordinary computational efficiency of MLPs allows us to consider much larger supercells of the same framework, such as the $9 \times 2 \times 9$ cell shown in Fig. 5. Because the ensemble standard deviation of physical observables decreases according to the square root of the number of particles, the fluctuations in unit cell volume at this scale are an order of magnitude smaller and therefore no longer able to trigger premature transitions of the framework. The obtained transition pressure of 18–20 MPa as shown in Fig. 5 is further confirmed by a full evaluation of the pressure-versus-volume equation of state at 300 K as presented in Supplementary Note 5, and is in good agreement with the experimental result. Previous computational estimates existed but were obtained using a classical force field that was parameterized based on DFT input data from the optimized lp geometry only —i.e., without taking into account either the cp phase or the

Table 1. Test error performance of an MLP trained on a set of 10 aluminum- or zirconium frameworks (not including UiO-66(Zr) and MIL-53(Al)) at 600 K.

Test errors for Zr- and Al-based MOFs			
Metal	Framework	Forces MAE [meV Å ⁻¹]	Stress MAE [MPa]
Al	CAU-13	25.7	66.2
Al	A-520	40.3	100.3
Al	DUT-4	20.4	36.7
Al	MIL-68	55.0	67.7
Al	NOTT-300	22.8	22.9
Zr	UiO-66(CDC)	24.8	74.7
Zr	UiO-67	17.2	3.7
Zr	MIL-140	22.6	58.5
Zr	MOF-808	40.0	18.0
Zr	BD899C1	25.9	65.2
Al	MIL-53	24.7	107.2
Zr	UiO-66(BDC)	21.3	31.7

Computational details such as NequIP hyperparameters, sampling parameters, and test set generation are given in Supplementary Note 6.

transition region—which resulted in a large disagreement with experiment⁴³.

Towards universal interaction potentials

Overall, our results demonstrate that the physical interactions in a given framework may be learned by equivariant MPNNs based on only a minimal amount of QM evaluations. Naturally, we can exploit this efficiency even further and examine whether we can construct a single model for the prediction of physical interactions in several different frameworks. This kind of transferability has already been shown for other systems such as small organic molecules, and can also be anticipated for MOFs because of their modular building block structure. As a proof of principle, we considered a set of 10 well-known aluminum- and zirconium-based frameworks that are similar to but different from MIL-53(Al) and UiO-66(Zr) in either the topology of the framework or the organic ligand. A more detailed description of the frameworks under consideration in terms of their building blocks is given in Supplementary Note 5. We used incremental learning to explore and learn the phase space of this set of 10 materials, which resulted in a training set of about 3100 configurations. We evaluated the test error performance for the frameworks included during training as well as the UiO-66(Zr) and MIL-53(Al) frameworks that were left out; the results are shown in Table 1. Even though it was trained on a dataset with a significantly larger variety in atomic environments, the model still achieves relatively low force and stress MAEs, even for the frameworks not included during training. In Supplementary Note 6, the performance of the model is further investigated and compared with UFF4MOF⁴⁴; an established universal interaction potential for MOFs. While the model is slightly less accurate when compared to the results in Fig. 4, it still outperforms UFF4MOF by a large margin. In addition, it should be emphasized that the entire training procedure only required about 3100 QM evaluations in total or about 310 QM evaluations per material, which further demonstrates the efficacy of incremental learning in combination with equivariant MPNNs.

DISCUSSION

In this work, we propose an efficient approach for the construction of accurate and transferable MLPs for framework materials. Even

for systems with multiple phases, we show that about 1000 QM evaluations are sufficient to construct accurate equivariant MPNNs. This increased computational efficiency is important for future research, as it is now possible to employ more advanced QM methods (e.g., hybrid functionals or beyond) during MLP training and in this way allow for a more accurate description of dynamic phenomena in these materials. In addition, the ability to construct a single potential for the description of multiple frameworks is highly promising, especially because we observed that the number of QM evaluations per material actually decreases with increasing variety in the training set (from about 1000 to only about 300). Nevertheless, further research in this area is still necessary. For example, it is still unclear how a maximally diverse training set of different frameworks should be assembled, i.e., which distribution of building blocks, topology, and/or pore sizes are necessary to guarantee transferability towards as many materials as possible. In addition, it remains to be seen whether equivariant MPNNs like NequIP will maintain their accuracy across the MOF design space, including e.g., large mesoporous systems. In those cases, message passing architectures may have to be complemented with more recent models that are targeted towards a more accurate description of long-range interactions^{45–47}. Finally, future work should extend the applicability of MLPs towards guest-loaded and even disordered frameworks in order to fully unleash their potential in the design of next-generation MOF technologies.

METHODS

Density functional theory calculations

QM energy evaluations were performed using the CP2K simulation package⁴⁸, version 8.2. We employ the PBE⁴⁹ functional with Grimme D3 dispersion corrections⁵⁰ and a hybrid basis set including both TZVP Gaussian basis functions and plane waves⁵¹; GTH pseudopotentials were used to smoothen the electron density near the nuclei. The plane wave cutoff energy was set to 1000 Ry for all materials as to guarantee that force and stress calculations were fully converged. Additional computational details regarding DFT calculations are available in Supplementary Note 1.

Molecular dynamics

MLP sampling was performed using YAFF⁵² at conditions of constant temperature and constant pressure, with a timestep of 0.5 fs. The temperature was controlled using a Langevin thermostat⁵³ with a time constant of 100 fs, and the pressure was controlled using a Langevin barostat with a time constant of 1 ps⁵⁴. PLUMED version 2.7.2 was used as a plugin to add the metadynamics bias in the simulation⁵⁵.

Machine learning potentials

All the MLP models in this work were constructed using NequIP version 0.5.4. All models employed a cutoff radius of 5 Å for the atomic environments and used four interaction layers, as this was found to strike the optimal balance between accuracy and computational efficiency. Because neighboring atoms exchange information about their atomic environments in each layer, the effective interaction radius of a single atom is at most $4 \times 5 \text{ Å} = 20 \text{ Å}$ because that is the largest distance over which information on a given atomic environment can travel. While this was found sufficient for all frameworks considered in this work, an accurate description of frameworks having larger pores may require dedicated long-range interactions. Feature representations were restricted to rotation orders of either $l = 1$ or $l = 0$, and the sizes of the network layers were chosen such that the total number of network parameters was around 200,000. The loss

function consists of a weighted average of potential energy and force errors, and was optimized using Adam⁵⁶ with a learning rate of 0.005. Additional parameters are presented in Supplementary Note 2.

Workflow management and sampling parameters

All of the computational steps in the learning algorithm are managed using the Snakemake workflow management system, version 7.8⁵⁷. QM evaluations and metadynamics simulations are performed in a massively parallel manner on CPU nodes, whereas MLP training is performed on a single GPU (Nvidia A100 40GB). In this way, each iteration in the algorithm takes about one to two hours in real time. For models that were trained on a single material, we used 50 parallel MTD walkers to explore the phase space, 45 of which were used to construct the training set and the remaining 5 for the validation set. In contrast to traditional multiple walker MTD, each walker maintains its own bias potential as to encourage them to sample different regions in phase space. Hills were added at a pace of 50 fs, with a width $\sigma = 100 \text{ \AA}^3$ and a height of 5 kJ/mol. Models were initialized by training on a small set of 50 structures obtained by applying uniform perturbations in the atomic coordinates and strain components (with respective amplitudes of 0.08 \AA and 0.01) starting from the initial structure.

Additional validation

To verify the validity of the transition dynamics as generated by the MLP, we validated the 0 MPa phase transition simulation on the $1 \times 2 \times 1$ unit cell of MIL-53(Al) from Fig. 5 by post hoc calculation of the DFT energy and forces at regular intervals over the entire transition. In this way, we obtained an average force MAE of only 6.7 meV \AA^{-1} and an average energy MAE of 0.1 meV per atom, which is exceptionally low. In addition to the $9 \times 2 \times 9$ configuration presented in Fig. 4, transition pressures were also determined based on alternative supercell configurations ($3 \times 2 \times 3$ and $6 \times 2 \times 6$), all of which yielded the same 18–20 MPa estimate.

DATA AVAILABILITY

All datasets that were used and/or generated in this work are publicly available via Zenodo^{58,59}.

CODE AVAILABILITY

An automated implementation of the entire algorithm using Snakemake is available online⁵⁹, together with the necessary input files for CP2K and NequIP as well as a variety of scripts to perform molecular dynamics simulations, geometry optimizations, and extended Hessian calculations. In addition, we provide a highly modular and scalable implementation of the incremental learning approach in `psiflow`, a Python library available at github.com/svandenoute/psiflow.

Received: 17 March 2022; Accepted: 25 January 2023;

Published online: 06 February 2023

REFERENCES

- Burch, N. C., Jusuja, H. & Walton, K. S. Water stability and adsorption in metal-organic frameworks. *Chem. Rev.* **114**, 10575 (2014).
- Redfern, L. R. & Farha, O. K. Mechanical properties of metal-organic frameworks. *Chem. Sci.* **10**, 10666 (2019).
- Horcajada, P. et al. Metal-organic frameworks in biomedicine. *Chem. Rev.* **112**, 1232 (2012).
- Rogge, S. M. J. et al. Metal-organic and covalent organic frameworks as single-site catalysts. *Chem. Soc. Rev.* **46**, 3134 (2017).
- Lee, J. et al. Metal-organic framework materials as catalysts. *Chem. Soc. Rev.* **38**, 1450 (2009).
- Freund, R. et al. The current status of mof and cof applications. *Angew. Chem. Int. Ed.* **60**, 23975 (2021).

- Ma, N. & Horike, S. Metal-organic network-forming glasses. *Chem. Rev.* **122**, 4163 (2022).
- Lin, J.-B. et al. A scalable metal-organic framework as a durable physisorbent for carbon dioxide capture. *Science* **374**, 1464 (2021).
- Hanikel, N. et al. Evolution of water structures in metal-organic frameworks for improved atmospheric water harvesting. *Science* **374**, 454 (2021).
- Furukawa, H., Müller, U. & Yaghi, O. M. "heterogeneity within order" in metal-organic frameworks. *Angew. Chem. Int. Ed.* **54**, 3417 (2015).
- Cheetham, A. K., Bennett, T. D., Coudert, F.-X. & Goodwin, A. L. Defects and disorder in metal organic frameworks. *Dalton Trans.* **45**, 4113 (2016).
- Krause, S. et al. The effect of crystallite size on pressure amplification in switchable porous solids. *Nat. Commun.* **9**, 1573 (2018).
- Ehrling, S., Miura, H., Senkova, I. & Kaskel, S. From macro- to nanoscale: finite size effects on metal-organic framework switchability. *Trends Chem.* **3**, 291 (2021).
- Ehrling, S. et al. Crystal size versus paddle wheel deformability: selective gated adsorption transitions of the switchable metal-organic frameworks dut-8(co) and dut-8(ni). *J. Mater. Chem. A* **7**, 21459 (2019).
- Van Speybroeck, V., Vandenhoute, S., Hoffman, A. E. & Rogge, S. M. Towards modeling spatiotemporal processes in metal-organic frameworks. *Trends Chem.* **3**, 605–619 (2021).
- Evans, J. D., Bon, V., Senkova, I., Lee, H.-C. & Kaskel, S. Four-dimensional metal-organic frameworks. *Nat. Commun.* **11**, 2690 (2020).
- Vandenhoute, S., Rogge, S. M. J. & Van Speybroeck, V. Large-scale molecular dynamics simulations reveal new insights into the phase transition mechanisms in mil-53(al). *Front. Chem.* **9**, 2296 (2021).
- Keupp, J. & Schmid, R. Molecular dynamics simulations of the "breathing" phase transformation of mof nanocrystallites. *Adv. Theory Simul.* **2**, 1900117 (2019).
- Friederich, P., Häse, F., Proppe, J. & Aspuru-Guzik, A. Machine-learned potentials for next-generation matter simulations. *Nat. Mater.* **20**, 750 (2021).
- Unke, O. T. et al. Machine learning force fields. *Chem. Rev.* **121**, 10142 (2021).
- Christensen, A. S., Bratholm, L. A., Faber, F. A. & Anatole von Lilienfeld, O. Fchl revisited: faster and more accurate quantum machine learning. *J. Chem. Phys.* **152**, 044107 (2020).
- Deringer, V. L. et al. Gaussian process regression for materials and molecules. *Chem. Rev.* **121**, 10073 (2021).
- Behler, J. Four generations of high-dimensional neural network potentials. *Chem. Rev.* **121**, 10037 (2021).
- Eckhoff, M. & Behler, J. From molecular fragments to the bulk: Development of a neural network potential for mof-5. *J. Chem. Theory Comput.* **15**, 3793 (2019).
- Achar, S. K., Wardzala, J. J., Bernasconi, L., Zhang, L. & Johnson, J. K. Combined deep learning and classical potential approach for modeling diffusion in uio-66. *J. Chem. Theory Comput.* **18**, 3593 (2022).
- Yu, Y., Zhang, W. & Mei, D. Artificial neural network potential for encapsulated platinum clusters in mof-808. *J. Phys. Chem. C* **126**, 1204 (2022).
- Batzner, S. et al. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **13**, 2453 (2022).
- Thomaset, N. et al. Tensor field networks: rotation- and translation-equivariant neural networks for 3d point clouds. Preprint at <https://arxiv.org/abs/1802.08219> (2018).
- Sivaraman, G. et al. Machine-learned interatomic potentials by active learning: amorphous and liquid hafnium dioxide. *NPJ Comput. Mater.* **6**, 104 (2020).
- Schran, C. et al. Machine learning potentials for complex aqueous systems made simple. *Proc. Natl Acad. Sci. USA* **118**, e2110077118 (2021).
- Wang, W., Yang, T., Harris, W. H. & Gómez-Bombarelli, R. Active learning and neural network potentials accelerate molecular screening of ether-based solvate ionic liquids. *Chem. Commun.* **56**, 8920 (2020).
- Vandermause, J. et al. On-the-fly active learning of interpretable bayesian force fields for atomistic rare events. *NPJ Comput. Mater.* **6**, 20 (2020).
- Cavka, J. H. et al. A new zirconium inorganic building brick forming metal organic frameworks with exceptional stability. *J. Am. Chem. Soc.* **130**, 13850 (2008).
- Loiseau, T. et al. A rationale for the large breathing of the porous aluminum terephthalate (mil-53) upon hydration. *Chem. Eur. J.* **10**, 1373 (2004).
- Schütt, K. T., Unke, O. T. & Gastegger, M. Equivariant message passing for the prediction of tensorial properties and molecular spectra. Preprint at <https://arxiv.org/abs/2102.03150> (2021).
- Batatia, I., Kovács, D. P., Simm, G. N. C., Ortner, C. & Csányi, G. Mace: higher order equivariant message passing neural networks for fast and accurate force fields. Preprint at <https://arxiv.org/abs/2206.07697> (2022).
- 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 (2018).
- Wang, H., Zhang, L., Han, J. & Weinan, E. Deepmd-kit: a deep learning package for many-body potential energy representation and molecular dynamics. *Comput. Phys. Commun.* **228**, 178 (2018).
- Musil, F. et al. Physics-inspired structural representations for molecules and materials. *Chem. Rev.* **121**, 9759 (2021).

40. Nazarian, D., Ganesh, P. & Sholl, D. S. Benchmarking density functional theory predictions of framework structures and properties in a chemically diverse test set of metal-organic frameworks. *J. Mater. Chem. A* **3**, 22432 (2015).
41. Yot, P. G. et al. Metal-organic frameworks as potential shock absorbers: the case of the highly flexible mil-53(al). *Chem. Commun.* **50**, 9462 (2014).
42. Rogge, S. M. J. et al. A comparison of barostats for the mechanical characterization of metal-organic frameworks. *J. Chem. Theory Comput.* **11**, 5583 (2015).
43. Rogge, S. M. J., Waroquier, M. & Van Speybroeck, V. Unraveling the thermodynamic criteria for size-dependent spontaneous phase separation in soft porous crystals. *Nat. Commun.* **10**, 4842 (2019).
44. Coupry, D. E., Addicoat, M. A. & Heine, T. Extension of the universal force field for metal-organic frameworks. *J. Chem. Theory Comput.* **12**, 5215 (2016).
45. Staacke, C. G. et al. On the role of long-range electrostatics in machine-learned interatomic potentials for complex battery materials. *ACS Appl. Energy Mater.* **4**, 12562 (2021).
46. Grisafi, A. & Ceriotti, M. Incorporating long-range physics in atomic-scale machine learning. *J. Chem. Phys.* **151**, 204105 (2019).
47. Lewis, A. M., Grisafi, A., Ceriotti, M. & Rossi, M. Learning electron densities in the condensed phase. *J. Chem. Theory Comput.* **17**, 7203 (2021).
48. VandeVondele, J. et al. Quickstep: fast and accurate density functional calculations using a mixed Gaussian and plane waves approach. *Comput. Phys. Commun.* **167**, 103 (2005).
49. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**, 3865 (1996).
50. Grimme, S., Antony, J., Ehrlich, S. & Krieg, H. A consistent and accurate ab initio parametrization of density functional dispersion correction (dft-d) for the 94 elements h-pu. *J. Chem. Phys.* **132**, 154104 (2010).
51. Lippert, B. G., Hutter, J. & Parrinello, M. A hybrid gaussian and plane wave density functional scheme. *Mol. Phys.* **92**, 477 (1997).
52. Verstraelen, T., Vanduyfhuys, L., Vandenbrande, S. & Rogge, S. Yaff, yet another force field. <http://molmod.ugent.be/software/>.
53. Bussi, G. & Parrinello, M. Accurate sampling using langevin dynamics. *Phys. Rev. E* **75**, 056707 (2007).
54. Feller, S. E., Zhang, Y., Pastor, R. W. & Brooks, B. R. Constant pressure molecular dynamics simulation: the Langevin piston method. *J. Chem. Phys.* **103**, 4613 (1995).
55. Tribello, G. A., Bonomi, M., Branduardi, D., Camilloni, C. & Bussi, G. Plumed 2: new feathers for an old bird. *Comp. Phys. Commun.* **185**, 604 (2014).
56. Kingma, D. P. & Ba, J. Adam: a method for stochastic optimization. Preprint at <https://arxiv.org/abs/1412.6980> (2017).
57. Mölderer, F. et al. Sustainable data analysis with snakemake. *F1000Research* **10**, 33 (2021).
58. Vandenhaute, S., Cools-Ceuppens, M., Verstraelen, T. & Van Speybroeck, V. Machine learning potentials for metal-organic frameworks with thermodynamic transferability: training data. <https://doi.org/10.5281/zenodo.6359970> (2022).
59. Vandenhaute, S., Cools-Ceuppens, M., DeKeyser, S., Verstraelen, T. & Van Speybroeck, V. Machine learning potentials for metal-organic frameworks using an incremental learning approach: workflow and data. <https://doi.org/10.5281/zenodo.7539133> (2023).

ACKNOWLEDGEMENTS

V.V.S. and T.V. acknowledge funding from the Research Board of Ghent University (BOF). S.V. and M.C.C. wish to thank the Research Foundation – Flanders (FWO) for doctoral fellowships (grant nos. 11H6821N and 11D0420N respectively). The resources and services used in this work were provided by VSC (Flemish Supercomputer Center), funded by the Research Foundation – Flanders (FWO) and the Flemish Government.

AUTHOR CONTRIBUTIONS

S.V., M.C.C., S.D., T.V., and V.V.S. initiated the discussion. S.V. implemented the algorithm, trained the models, and performed all simulations included in this work. S.V., M.C.C., T.V., and V.V.S. were involved in the discussions of the results, designed, and wrote the paper.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41524-023-00969-x>.

Correspondence and requests for materials should be addressed to Veronique Van Speybroeck.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons

Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023

Paper III

The Operando Nature of Isobutene Adsorbed in Zeolite H-SSZ-13 Unraveled by Machine Learning Potentials Beyond DFT Accuracy

Sander Vandenhaute[†], Massimo Bocus[†], Veronique Van Speybroeck

[†]: these authors contributed equally

Angewandte Chemie Int. Ed. 64, e202413637, 2025

Sander Vandenhaute developed and implemented the on-the-fly and transfer learning algorithms.

Reprinted with permission.
Copyright (2025), The Authors



The Operando Nature of Isobutene Adsorbed in Zeolite H-SSZ-13 Unraveled by Machine Learning Potentials Beyond DFT Accuracy

Massimo Bocus[†], Sander Vandenhoute[†], and Veronique Van Speybroeck*

Abstract: Unraveling the nature of adsorbed olefins in zeolites is crucial to understand numerous zeolite-catalyzed processes. A well-grounded theoretical description critically depends on both an accurate determination of the potential energy surface (PES) and a reliable account of entropic effects at operating conditions. Herein, we propose a transfer learning approach to perform random phase approximation (RPA) quality enhanced sampling molecular dynamics simulations, thereby approaching chemical accuracy on both the determination and exploration of the PES. The proposed methodology is used to investigate isobutene adsorption in H-SSZ-13 as prototypical system to estimate the relative stability of physisorbed olefins, carbenium ions and surface alkoxide species (SAS) in Brønsted-acidic zeolites. We show that the *tert*-butyl carbenium ion formation is highly endothermic and no entropic stabilization is observed compared to the physisorbed complex within H-SSZ-13. Hence, its predicted concentration and lifetime are negligible, making a direct experimental observation unlikely. Yet, it remains a shallow minimum on the free energy surface over the whole considered temperature range (273–873 K), being therefore a short-lived reaction intermediate rather than a transition state species.

Introduction

Zeolites are microporous, crystalline aluminosilicate materials. The charge imbalance created by the defective substitution of Si⁴⁺ with Al³⁺ can be compensated by the presence of an acidic proton (the Brønsted acid site, or BAS for short) on one of the oxygen atoms in the Al first coordination sphere. The BAS acidity, the strong confine-

ment effects of the micro-sized pores and their remarkable stability have made Brønsted-acidic zeolites pivotal catalysts in the traditional petrochemical industry as well as in novel processes for the conversion of biomass-derived sustainable feedstocks.^[1,2] Most acid-catalyzed reactions in zeolites are triggered by an initial transfer of the BAS from the framework to the reacting substrate, potentially creating a positively charged intermediate. Therefore, understanding the chemical nature and reactivity of the latter is of crucial interest to obtain a comprehensive mechanistic view of the chemical process under investigation.^[3]

In this context, a particularly challenging case study is the adsorption of olefins in zeolites (Figure 1a). When unsaturated hydrocarbons approach the BAS, the double C=C bond interacts with the electrophilic proton forming a so-called π -complex. A complete proton transfer from the zeolite to the molecule creates a positively-charged carbenium ion, which can then chemisorb on one of the oxygen atoms in the first coordination sphere of the Al site to form a surface alkoxide species (SAS). The persistent nature of each of these adsorbed species at operating conditions has been the topic of extensive theoretical and experimental studies.^[3] It is generally accepted that their relative stability is affected by a series of factors, including temperature, zeolite topology and substitution degree of the double bond which, in general, create a subtle balance between enthalpic and entropic contributions. Enthalpically, SAS would be preferred but their stability may decrease due to steric hindrance in small concave environments.^[4,5] Carbenium ions are thought to be entropically stabilized, however such stabilization may strongly depend on the zeolite topology. Thus far, disentangling entropic and enthalpic effects has remained challenging for both theory and experiment. In this contribution, we study isobutene in H-SSZ-13 (CHA topology) to reveal the persistent nature of its various adsorbed states. This case is selected as isobutene is the smallest alkene with a tertiary carbon atom and the limited unit cell size of H-SSZ-13 allows for the highest level of theory calculations.

A large number of theoretical works have investigated isobutene chemisorption in zeolites, initially using gas-phase clusters^[8–15] and then moving to more realistic, fully periodic models. These can be inquired with static approaches, only relying on a limited number of points on the potential energy surface (PES),^[4,16–29] or with dynamic approaches, that better mimic operating conditions.^[5,30,31] A more detailed overview of the theoretical investigations in periodic models performed thus far is taken up in Table S1 of the Supplementary Information. On the one hand, seminal work

[*] Dr. M. Bocus,[†] S. Vandenhoute,[†] Prof. Dr. V. Van Speybroeck
 Center for Molecular Modeling, Ghent University, Technologiepark
 46, 9052 Zwijnaarde (Belgium)
 E-mail: veronique.vanspeybroeck@ugent.be

[†] These authors contributed equally.

© 2024 The Author(s). Angewandte Chemie International Edition published by Wiley-VCH GmbH. This is an open access article under the terms of the Creative Commons Attribution Non-Commercial NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

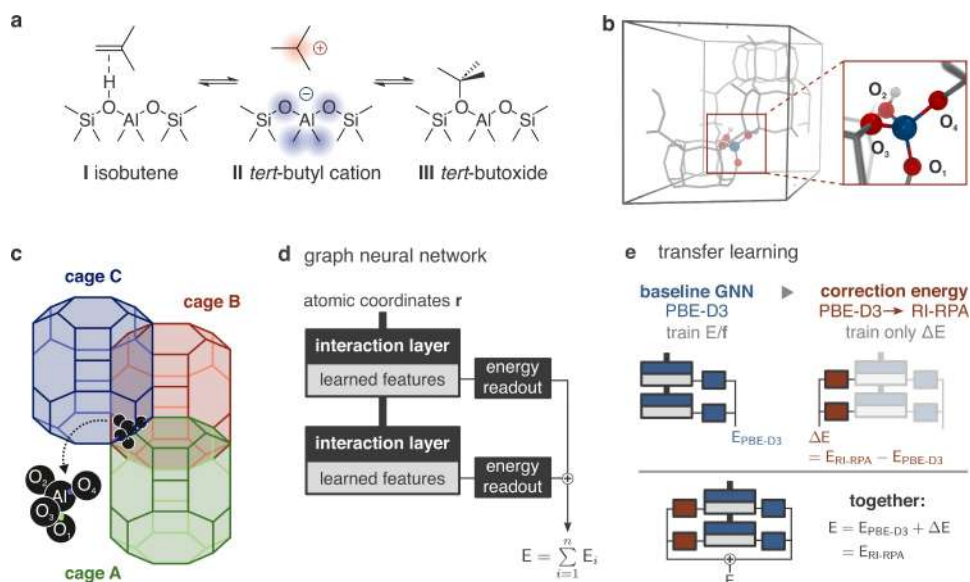


Figure 1. a. The three investigated intermediates formed upon isobutene adsorption in Brønsted acidic zeolites. b. The conventional H–SSZ–13 unit cell used in this work. The inset shows the nomenclature of the four oxygen atoms in the first coordination sphere of the Al defect adopted herein. Al is depicted in blue, O in red and H in white, while the rest of the framework is shown in grey sticks for the sake of clarity. c. Schematic depiction of the three nonequivalent H–SSZ–13 cages. The location of the Al tetrahedron is highlighted with grey spheres. d. Simplified representation of the MACE architecture;⁶⁰ the atomic coordinates are parsed in atomic environments which are passed through a series of interaction layers; this yields a set of features based on which the atomic energies are computed. e. Summary of the two-stage procedure proposed in this contribution: the interaction layers and the first set of energy readouts (blue) are trained based on PBE-D3-level reference data—energy and forces—obtained using active learning.¹⁷ A second set of readouts (red) are trained on a small number of RI-RPA energies. See the Supplementary Sections 2.5 and 2.6 for more information.

of Sauer and co-workers clearly shows that standard generalized gradient approximation density functional theory (GGA DFT), with or without empirical dispersion corrections, is unreliable when modeling species with a net charge separation such as carbenium ions in zeolites.^[20,22,29] High-level (at least MP2) single point corrections based on cluster models increase the carbenium ion energy by tens of $\text{kJ}\cdot\text{mol}^{-1}$ with respect to the neutral alkene. On the other hand, typical industrial processes involving olefins—such as fluid catalytic cracking^[32] (FCC)—are conducted at very high temperatures ($\sim 800\text{ K}$). Static simulations with harmonic thermal corrections are not well-suited to model mobile and weakly bounded species in the zeolite voids.^[33,34] Therefore, studies conducted with *operando* methodologies based on constant temperature molecular dynamics (MD) suggest that thermal effects can remarkably stabilize the mobile carbenium ion, making it a prominent intermediate at reaction conditions.^[5,30] This effect may however be highly dependent on the zeolite topology.^[5] While MD-based approaches seem attractive to unravel the nature of adsorbed species at operating conditions, they require hundreds of thousands energy and forces evaluations,

making it impossible to go beyond GGA DFT when evaluating the PES.

The extent to which carbenium ions are relevant intermediates in high temperature processes such as FCC has therefore remained unclear. Various experimental attempts have been performed to identify the predominant intermediate(s); because some studies report data in favor of transient tertiary carbenium ions^[25,35–38] while others did not detect them^[39–43] or did not provide definitive evidence,^[44] the experimental results remain inconclusive.^[3] From a theoretical viewpoint, various studies have reconciled enhanced sampling molecular dynamics with high-level calculations in zeolites, using for instance corrections based on single-point cluster evaluations,^[31] machine learning assisted free energy perturbation theory,^[45–47] or similar methods.^[48] These approaches are however only successful when the high- and low-level method sufficiently overlap in phase space.^[49] Furthermore, the studies were still based on expensive *ab initio* molecular dynamics simulations.

To fully solve these issues, we harness the power of state-of-the-art machine learning potentials (MLPs) and propose an efficient approach based on transfer

learning.^[50–53] Starting from a machine learning model trained on a production level of theory, we augment its accuracy towards a higher-level electronic structure method with outstanding data efficiency. We showcase the power of this methodology by investigating the conversion of isobutene in the H–SSZ–13 zeolite with RI-RPA^[54,55] (random phase approximation using a resolution-of-identity approach) quality enhanced sampling molecular dynamics (ESMD) simulations. Our results, arising from a total of 60 ns RI-RPA quality ESMD, show that within H–SSZ–13 the *tert*-butylation is not entropically stabilized compared to the physisorbed neutral isobutene within the temperature range 323–823 K. By approaching chemical accuracy in both the enthalpy and entropy estimates of the reaction, physisorbed isobutene is predicted to be the only species present in significant amounts within the zeolite pores while the concentrations of *tert*-butylation and *tert*-butoxide are negligible. The carbenium ion remains an elusive intermediate, with a computed lifetime that is however too short to be observed with currently available experimental techniques.

Results and Discussion

The proposed transfer learning approach is used to investigate the adsorption of isobutene in the H–SSZ–13 zeolite. The framework presents a relatively small unit cell where all tetrahedral atoms are equivalent by symmetry. We use the same equilibrated conventional unit cell of our previous investigation on the proton hopping reaction.^[56] A single Si atom is substituted with Al (Si/Al=35), breaking the unit cell symmetry and creating four nonequivalent oxygen atoms in the first coordination of the Al defect and three nonequivalent cages in which the isobutene can be located (Figure 1b,c). The isobutene conversion in the *tert*-butylation and the surface *tert*-butoxide is captured by means of two collective variables (CVs, Figure 2a). CVs are functions of the atomic coordinates which should be able to describe the progress of the reaction under investigation. One CV describes the proton transfer from the zeolite to isobutene with a difference in coordination numbers (CNs). The first CN is between the primary carbon atoms (C_p) and all hydrogen atoms in the system, while the second between the four zeolite oxygen atoms in the first coordination sphere of the Al defect (O_z) and, again, all hydrogen atoms. The other CV is a CN between the tertiary carbon atom (C_t) and O_z (additional details are reported in Supplementary Section 2.3). Two assumptions will be made to keep the complexity of the problem manageable: (1) the formation of the primary surface isobutoxide species can be ignored. The reaction is known to proceed through a high energy barrier^[22,24] and does not lead to the formation of a carbenium ion, the main focus of this study. (2) It will be assumed that the reactivity of isobutene is entirely independent of inter-cage diffusion and that it does not prefer to reside in any specific cage. An integrated reaction-diffusion model is beyond the scope of this study and does not impact its conclusions.

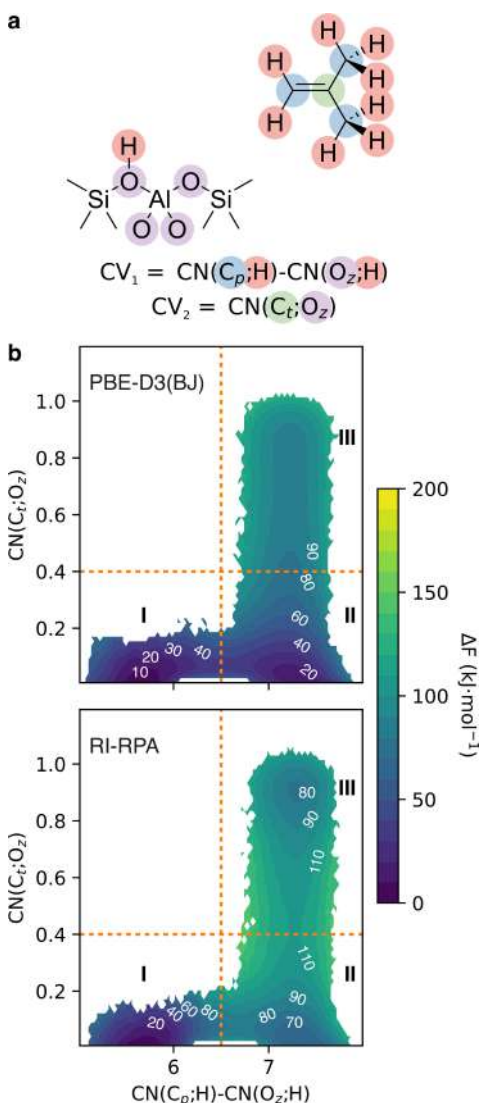


Figure 2. a. Graphical depiction of the atoms included in the CNs that constitute the two final CVs used in the US simulations. b. The computed FES changes remarkably when moving from the PBE-D3(BJ) to the RI-RPA PES. Showcased as examples are the 623 K FESs for isobutene conversion in cage B. The orange lines highlight the FES partition in the regions corresponding to isobutene (I), to the *tert*-butylation (II) and to the *tert*-butoxide (III). A complete overview of the computed FESs is provided in Supplementary Section 4.2.

An accurate computational estimate of both enthalpic and entropic contributions to the reaction free energy requires millions of energy and force evaluations. We employ state of the art graph neural networks (GNNs) to learn a computationally efficient representation of the *ab initio* potential energy surface in order to be able to simulate the dynamics of the system on such a long time scale. The incorporation of equivariant feature representations into GNNs has reduced the required training set sizes to about a few thousand structures for a chemical reaction in zeolites,^[56] each of which is to be labeled with the correct *ab initio* energy and atomic forces.^[7] The key challenge is that we cannot generate this amount of training data for the required (post-HF) level of theory. First, (RI-)MP2 or RPA energy evaluations require at least two orders of magnitude more memory and compute time than regular GGA DFT. Second, the calculation of the nuclear gradients introduces additional memory requirements and is moreover very hard to converge with respect to basis set and grid size. To overcome these issues, we developed a transfer learning approach which (1) reduces the required number of post-HF single-point evaluations by an order of magnitude, and (2) only requires the post-HF energy during training (and not the nuclear gradients, i.e. forces). The approach is schematically illustrated in Figure 1d,e, and is founded on the empirical observation that the interaction layers within the GNN implicitly learn an informative feature representation for the chemical environment of each atom in the system.^[52,53,57] In particular, we observe that if we train a baseline GNN model to PBE-D3(BJ) energy and forces, the resulting interaction layers can be used to train a small additional energy correction ΔE which captures the difference between the learned PBE-D3(BJ) baseline output and the target post-HF energy. Importantly, the prediction of this ΔE requires only a small number of additional weights ($\sim 1\%$ extra), and can be learned based on just a few hundred *energy-only* post-HF single points—see Supplementary Section 2.6 for more information.

Based on our benchmarks, we determined RI-RPA/cc-QT(Al,Si)Z as our target level of theory (see Supplementary Section 2.2), whereas the baseline level of theory was PBE-D3(BJ). We employ a combination of active learning and enhanced sampling to train a baseline GNN model to PBE-D3(BJ) reference data using psiflow.^[7,58,59] Quantum mechanical evaluations are performed with CP2K,^[60] using either a MOLOPT/TZVP or correlation-consistent cc-QT-(Al,Si)Z basis set for respectively the PBE-D3(BJ) and RI-RPA levels of theory.^[61–63] Umbrella sampling simulations are used to sample the phase space in the active learning loop as well as in the production runs to derive all free energy surfaces (FESs), which are computed with the weighted histogram analysis (WHAM) method as implemented in the ThermoLIB library.^[64] All computational details are reported in Supplementary Section 2, whereas Supplementary Section 3 presents an extensive validation of the entire workflow. In particular, we show that (1) the baseline GNN perfectly reproduces the PBE-D3(BJ) reaction free energy profile as obtained from brute force AIMD umbrella sampling, and (2) that our transfer learning

approach allows us to achieve an RMSE of only 0.3 meV/atom on the target RI-RPA potential energy surface. Nevertheless, the entire active learning workflow requires only 3682 PBE-D3(BJ) single-point calculations, with an additional ~ 200 *energy-only* RI-RPA calculations to parameterize the ΔE prediction. Hereafter, all results of PBE-D3(BJ) quality are obtained using the vanilla baseline GNN. The results of RI-RPA quality are obtained with the same GNN augmented with an additional set of readout layers that predicts ΔE (Figure 1e).

As previously mentioned, the introduction of the Al substitution breaks the CHA unit cell symmetry and creates 4 nonequivalent O atoms in the first coordination sphere and 3 nonequivalent cages (Figure 1b,c). In principle, isobutene can be located in any of the 3 cages and the BAS/*tert*-butoxide can be bound to any of the 4 oxygen atoms. While some combinations can be excluded by simple steric considerations (for instance, when the BAS is on O₃ it is exclusively accessible from cage C), we tested all remaining paths at 623 K to select the most favorable one(s) (Supplementary Section 4.1). At the end, the two most promising ones were selected: in the first, isobutene is located in cage B and the BAS/*tert*-butoxide are formed on O₄. This path exhibits the lowest free energy barriers for both reactive steps. However, based on our previous investigation of the proton hopping reaction in H-SSZ-13,^[56] O₄ is also the least populated site—which is in line with a higher activity. Hence, we also selected an additional path in which isobutene resides in cage C, the BAS is located on O₃ and the *tert*-butoxide is formed on O₂. This path presents moderate free energy barriers and, moreover, O₃ was also computed to be the most populated site.

The two paths are investigated with both PBE-D3(BJ) and RI-RPA quality MLP US simulations over the temperature range 323–823 K in steps of 100 K. While a complete overview of all the computed FESs and free energy barriers is reported in Supplementary Section 4.2, we showcase the 623 K FESs of cage B in Figure 2b to illustrate the main differences between PBE-D3(BJ) and RI-RPA, as these can be generalized to both cages and all temperatures. Starting from the protonation reaction (**I**→**II**), a more accurate level of theory compensates—as expected—the deficiencies of GGA DFT. Indeed, the raw free energy barrier increases from ~ 30 kJ·mol⁻¹ to ~ 80 kJ·mol⁻¹ and the *tert*-butylation is lifted in energy with respect to isobutene, going from $\sim +20$ kJ·mol⁻¹ to $\sim +70$ kJ·mol⁻¹. This shift reflects the one observed for the 0 K electronic energy and agrees with the reference results of Plessow et al.^[28] (Supplementary Section 3.1). Yet, the *tert*-butylation remains a minimum on the FES at all temperatures, indicating that it is a transient reaction intermediate rather than a transition state. Quite some difference is also seen in the FES region concerning the *tert*-butoxide formation (**II**→**III**). While PBE-D3(BJ) does not predict the *tert*-butoxide to be a minimum, the RI-RPA FES has a clear transition state between the two intermediates. The overall stability of the alkoxide is similar for both levels of theory ($\sim +80$ kJ·mol⁻¹), in line with the 0 K results (Supplementary Section 3.1).

From the computed 2-dimensional FESs a CV-independent free energy for an intermediate *A* can be retrieved by integration over the respective CV region:

$$F_A = -k_B T \log \left(\int_A \exp \left(-\frac{F(q_1, q_2)}{k_B T} \right) dq_1 dq_2 \right) \quad (1)$$

q_1 and q_2 are the two CVs, T is the reaction temperature and k_B Boltzmann's constant. The boundaries used to define the three regions associated with isobutene, *tert*-butylation and *tert*-butoxide are highlighted with orange dotted lines in Figure 2b. This allowed us to compute the free energy differences between the three intermediates for both cages and both levels of theory over the whole temperature range. Remark that for PBE-D3(BJ) the surface *tert*-butoxide (**III**) is not a minimum on the FES, *i.e.* there is no free energy barrier separating it from **II**. Hence, its computed free energy is arbitrary as the state boundaries are ill-defined. However, it can still provide a useful mean of comparison with the RI-RPA results where the state is metastable. The results are shown in Figure 3a–d, where it can be seen how the free energy differences have in all cases an almost perfect linear behavior over the considered temperature range. This allowed us to rely on the simple relation

$$\Delta F_{A-B} = \Delta U_{A-B} - T\Delta S_{A-B} \quad (2)$$

and perform linear regression to estimate the internal energy (ΔU_{A-B}) and entropy (ΔS_{A-B}) variations of the reactions. Given the linear behavior of the data points, these variations can be approximated as temperature-independent over the considered temperature range.

Starting from the formation of the *tert*-butylation (**I**→**II**), PBE-D3(BJ) predicts a small positive internal energy variation in both the considered cages (+16 kJ·mol⁻¹). The main cause for the lower free energy difference observed in cage B is the entropy variation, which goes from being slightly negative in cage C to moderately positive in cage B. Moving to RI-RPA has a dramatic effect on the results: first, the internal energy difference between **I** and **II** drastically increases to more than +55 kJ·mol⁻¹ in both cages. Additionally, the entropy variation decreases by ~10 J·K⁻¹·mol⁻¹ in both cages, becoming moderately negative for cage C and close to zero for cage B. In any case, our results indicate that there exists no entropic stabilization of weakly bounded carbenium ions in H-SSZ-13 compared to the physisorbed complex, and their previous observation in MD simulations^[5] is therefore exclusively related to the deficiencies of GGA DFT. This is made clear by observing the equilibrium population of the three intermediates (Figure 3e,f). Especially for cage B, PBE-D3(BJ) predicts a significant concentration of *tert*-butylation above 400 K, which also becomes easily observable in unbiased MD simulations given the underestimated free energy barrier. On the other hand, the RI-RPA results clearly foresee neutral isobutene as the only species present in the zeolite for any temperature, with the concentrations of *tert*-butylation and *tert*-butoxide being nearly zero.

To support these conclusions, we used the MLPs to run 1 ns long unbiased MD simulations at 323, 623 and 823 K both in cage B and C (a complete overview of the results is reported in Supplementary Section 4.4). When using either isobutene or the *tert*-butylation as initial structure, PBE-D3(BJ) allows frequent jumps between the two states, which become more frequent when increasing the temperature (Supplementary Figure S22). With RI-RPA, on the other hand, the *tert*-butylation always deprotonates within 5 ps and neutral isobutene is the only observed species for all temperature. Only in one simulation at 823 K we observed a transient formation of *tert*-butylation lasting for few ps, in line with the results of the enhanced sampling simulations. In order to better understand the differences in behavior between isobutene and the *tert*-butylation, we performed additional simulations in cage B and C at 623 K with half-quadratic potential walls to restrain the system in the isobutene or *tert*-butylation region (**I** and **II** in Figure 2b, respectively). Interestingly, the *tert*-butylation has a large mobility in cage B, because the Al defect is located at half the cage height. Hence, it explores a similar space volume as isobutene (Figure 4). Conversely, cage C has the Al defect located at its base. While isobutene shifts between a BAS-bound π -complex and a free Van der Waals complex, the *tert*-butylation remains in tight proximity to the negative Al. This reflects in a larger entropic penalty, which is in line with the lower ΔS_{I-II} for cage C derived from enhanced sampling simulations. When comparing PBE-D3(BJ) and RI-RPA (Figure 4) no major differences can be seen, indicating that the entropy decrease of RI-RPA is not related to a reduced mobility of the intermediates. Additionally, we also investigated the global rotations of the molecules within the two cages (Supplementary Figure S24) and the internal rotations of the methyl groups (Supplementary Figure S25). None showed major differences between PBE-D3(BJ) and RI-RPA, suggesting that the entropic differences between the two levels of theory arises from changes in the vibrational modes, which are however more challenging to quantify.

The formation of *tert*-butoxide (**I**→**III**) is predicted to be moderately endothermic by both levels of theory, with an internal energy variation with respect to neutral isobutene of +25 kJ·mol⁻¹ in both cages according to PBE-D3(BJ) and +18 (cage C) or +7 kJ·mol⁻¹ (cage B) according to RI-RPA. The difference could in this case be attributed to limitations of RI-RPA, which was shown to modestly over-stabilize the surface-bonded species (Supplementary Section 3.1). As expected, the entropy variation of the reaction is in all cases largely negative, which makes the formation of *tert*-butoxide highly unfavorable at all temperatures (Figure 3e,f).

The final question we tried to answer is whether the *tert*-butylation, despite being clearly present in negligible concentrations in H-SSZ-13, might survive long enough to be detected experimentally. Tuma and Sauer^[22] reported a half life $\tau_{1/2} = \ln 2/k_{II-I}$ of 59 μ s in ferrierite, which is sufficiently long for detection by fast spectroscopies such as UV-vis. To estimate a lifetime for the *tert*-butylation in H-SSZ-13, we computed the backward kinetic constant k_{II-I}

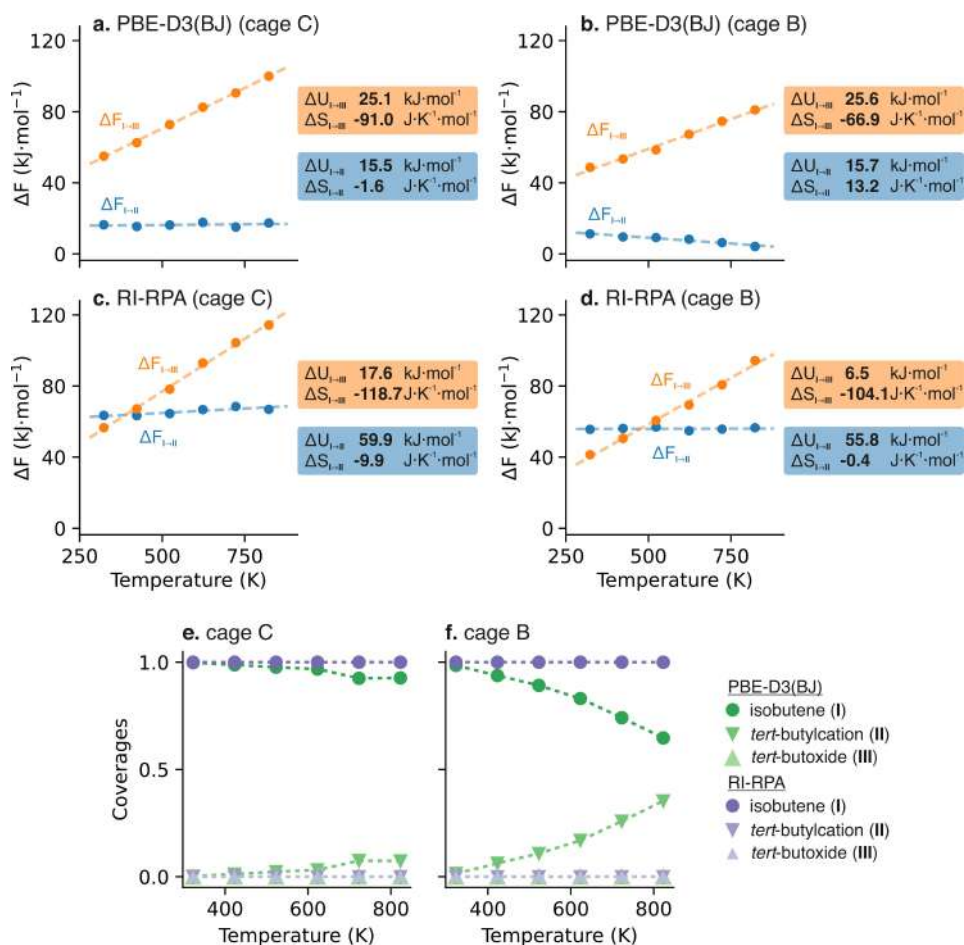


Figure 3. The *tert*-butylation does not benefit from any entropic stabilization with respect to neutral isobutene in the H-SSZ-13 pores. a-d. Free energy difference between isobutene and *tert*-butylation (I→II) and isobutene and *tert*-butoxide (I→III) computed with PBE-D3(BJ) and RI-RPA in cage B and cage C. The internal energy and entropy differences obtained from linear regression are also reported. e, f. Fractional coverages of the 3 intermediates derived from the PBE-D3(BJ) and RI-RPA free energy differences in cage C (e) and cage B (f).

from our simulations within the transition state theory (TST) approximation (Supplementary Section 4.3). Our results indicate that the half life of the intermediate in H-SSZ-13 is extremely short, where for both cage B and C is in the order of hundreds of fs or few ps at most, in agreement with the results of the unbiased MD simulations (Supplementary Section 4.4). This, joined with its extremely low concentration, would make its experimental observation highly unlikely.

Conclusions

In this contribution, we propose a transfer learning approach based on state-of-the-art machine learning interatomic potentials to perform RI-RPA quality ESMC simulations. We select a relevant and challenging case study to showcase the capability of the methodology, namely the chemisorption of isobutene in the H-SSZ-13 zeolite. The reaction proceeds through a weakly bounded *tert*-butylation intermediate, whose mobility is hardly captured by static

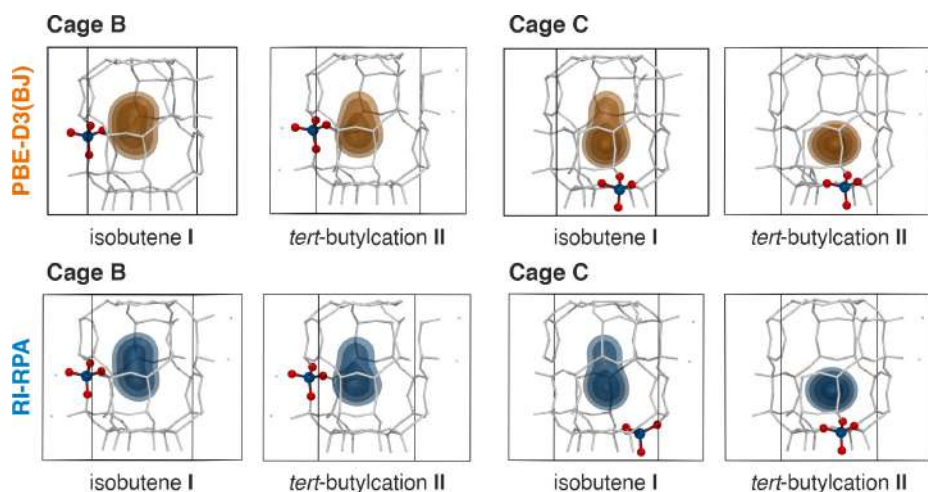


Figure 4. Density maps showing the location of the tertiary isobutene carbon atom during 1 ns long PBE-D3(BJ) and RI-RPA MD simulations of isobutene and *tert*-butylcation at 623 K. The surfaces are obtained with the volmap extension implemented in VMD^[65] (resolution of 0.2 Å). The isovalues are arbitrarily chosen to be (from the exterior) 0.005, 0.01, 0.02 and 0.04. Zeolite atoms are represented with white sticks, only the Al tetrahedron is highlighted with colored spheres (Al in blue, O in red) for the sake of clarity. See also Supplementary Section 4.4.

simulations with harmonic corrections. Additionally, its charged nature makes standard GGA DFT unreliable in capturing its relative stability, with an overstabilization in the order of tens of $\text{kJ}\cdot\text{mol}^{-1}$. Our approach provides an affordable solution to both issues, allowing to explicitly include temperature effects with a potential energy surface of arbitrary accuracy.

By performing RI-RPA quality US simulations over a broad range of temperatures, we show that the entropy difference between the neutral isobutene and the *tert*-butylcation is around zero or even slightly negative independently on the level of theory. Yet, moving from PBE-D3(BJ) to RI-RPA drastically increases the endothermic nature of the reaction, leading to equilibrium populations dominated by the neutral isobutene, while the concentrations of the *tert*-butylcation and *tert*-butoxide intermediates are negligible. Yet, the *tert*-butylcation remains a minimum on the FES for all considered temperatures, indicating that it is a transient reaction intermediate rather than a transition state species. This is in line with the available experimental evidences for H-ZSM-5, where the intermediate was captured with nucleophilic moieties.^[25,38] Our estimate for the *tert*-butylcation half life is drastically lower than previous reports from static simulations in ferrierite, in the order of few picoseconds at most, making its direct experimental observation highly unlikely.

Of course, our conclusions strictly hold for H-SSZ-13 only and frameworks with different pore shapes and sizes could lead to different observations.^[4,5] One of the main strengths of our methodology, being based on graph neural networks, is that it provides a per-atom environment-

dependent correction. Hence, as the parent model, it can be highly transferable provided that a sufficiently diverse training set is provided. Moreover, it also has the potential of being used with cluster-based corrections, where only the atoms directly involved in the correction (adsorbed molecule and active site) are modelled with the target level of theory. This could allow to model significantly larger systems (that are impossible to evaluate in full with RI-RPA) with chemical accuracy^[66] and further reduce the gap between theory and experiment.

Acknowledgements

V.V.S. acknowledges funding from the Research Board of Ghent University. S.V. wishes to thank the Research Foundation—Flanders (FWO) for a doctoral fellowship (grant n. 11H6821N). M.B. wishes to thank the Research Foundation—Flanders (FWO) for a junior postdoctoral fellowship (grant n. 1269725N) and for the Excellence of Science (EOS) Project BioFact (EOS ID 30902231). The resources and services used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation - Flanders (FWO) and the Flemish Government. Additionally, we acknowledge the EuroHPC Joint Undertaking for awarding this project access to the EuroHPC supercomputer LUMI, hosted by CSC (Finland) and the LUMI consortium, and the Luxembourg national supercomputer MeluXina through a EuroHPC Regular Access call. M.B. thanks Prof. Louis Vanduyfhuys for the assistance with the use of ThermoLIB and Dr. Frederick

Stein and Prof. Jürg Hutter for their help with the RI-RPA calculations. S.V. wishes to thank Ben Clifford for his guidance during the development of pswflow.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability

An implementation of the modified MACE architecture is available on GitHub (<https://github.com/molmod/transfer-mace>); examples of input files, relevant datasets, analysis scripts, and trained models used to obtain the results presented in this manuscript are available on Zenodo (<https://doi.org/10.5281/zenodo.13992565>). Any additional data is available from the authors upon request.

Keywords: Carbocations · Machine Learning · Molecular Dynamics · Random Phase Approximation · Zeolites

- [1] J. Cejka, A. Corma, S. Zones (Editors), *Zeolites and Catalysis: Synthesis, Reactions and Applications*, Wiley-VCH Verlag GmbH & Co. KGaA **2010**.
- [2] T. Ennaert, J. Van Aelst, J. Dijkmans, R. De Clercq, W. Schutyser, M. Dusselier, D. Verboeckend, B. F. Sels, *Chem. Soc. Rev.* **2016**, *45*, 584.
- [3] R. Zhao, G. L. Haller, J. A. Lercher, *Microporous Mesoporous Mater.* **2023**, *358*, 112390.
- [4] M. L. Sarazen, E. Iglesia, *Proc. Natl. Acad. Sci. USA* **2017**, *114*, E3900, <https://doi.org/10.1073/pnas.1619557114>.
- [5] P. Cnudde, M. Waroquier, V. Van Speybroeck, *Catalysis Science, Technology* **2023**, *13*, 4857.
- [6] I. Batatia, D. P. Kovacs, G. Simm, C. Ortner, G. Csanyi, MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields, in S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Editors), *Advances in Neural Information Processing Systems*, volume 35, Curran Associates, Inc. **2022** pages 11423–11436.
- [7] S. Vandenhoute, M. Cools-Ceuppens, S. DeKeyser, T. Verstraelen, V. Van Speybroeck, *npj Computational Materials* **2023**, *9*, 1, <https://doi.org/10.1038/s41524-023-00969-x>.
- [8] P. Viruela-Martin, C. Zicovich-Wilson, A. Corma, *J. Phys. Chem.* **1993**, *97*, 13713.
- [9] A. Rigby, M. Frash, *J. Mol. Catal. A* **1997**, *126*, 61.
- [10] V. Kazansky, *Catal. Today* **1999**, *51*, 419.
- [11] M. Boronat, P. Viruela, A. Corma, *Phys. Chem. Chem. Phys.* **2001**, *3*, 3235.
- [12] M. Boronat, A. Corma, *Appl. Catal. A* **2008**, *336*, 2.
- [13] H. Fang, A. Zheng, S. Li, J. Xu, L. Chen, F. Deng, *J. Phys. Chem. C* **2010**, *114*, 10254.
- [14] H. Fang, A. Zheng, J. Xu, S. Li, Y. Chu, L. Chen, F. Deng, *J. Phys. Chem. C* **2011**, *115*, 7429.
- [15] Y. Chu, B. Han, A. Zheng, X. Yi, F. Deng, *J. Phys. Chem. C* **2013**, *117*, 2194.
- [16] P. Sinclair, A. Vries, P. Sherwood, et al., *J. Chem. Soc. Faraday Trans.* **1998**, *94*, 3401.
- [17] M. Boronat, C. M. Zicovich-Wilson, P. Viruela, A. Corma, *J. Phys. Chem. B* **2001**, *105*, 11169.
- [18] C. Tuma, J. Sauer, *Angew. Chem. Int. Ed.* **2005**, *44*, 4769.
- [19] V. Nieminen, M. Sierka, D. Y. Murzin, J. Sauer, *J. Catal.* **2005**, *231*, 393.
- [20] C. Tuma, J. Sauer, *Phys. Chem. Chem. Phys.* **2006**, *8*, 3955.
- [21] B. A. De Moor, M.-F. Reyniers, M. Sierka, J. Sauer, G. B. Marin, *J. Phys. Chem. C* **2008**, *112*, 11796.
- [22] C. Tuma, T. Kerber, J. Sauer, *Angew. Chem. Int. Ed.* **2010**, *49*, 4678.
- [23] N. Rosenbach Jr, A. P. dos Santos, M. Franco, C. J. Mota, *Chem. Phys. Lett.* **2010**, *485*, 124.
- [24] C. M. Nguyen, B. A. De Moor, M.-F. Reyniers, G. B. Marin, *J. Phys. Chem. C* **2012**, *116*, 18236.
- [25] W. Dai, C. Wang, X. Yi, A. Zheng, L. Li, G. Wu, N. Guan, Z. Xie, M. Dyballa, M. Hunger, *Angew. Chem. Int. Ed.* **2015**, *54*, 8783.
- [26] G. A. Ferguson, L. Cheng, L. Bu, S. Kim, D. J. Robichaud, M. R. Nimlos, L. A. Curtiss, G. T. Beckham, *J. Phys. Chem. A* **2015**, *119*, 11397.
- [27] T. J. Goncalves, P. N. Plessow, F. Studt, *ChemCatChem* **2019**, *11*, 4368.
- [28] P. N. Plessow, F. Studt, *J. Phys. Chem. Lett.* **2020**, *11*, 4305.
- [29] O. Ren, M. Rybicki, J. Sauer, *J. Phys. Chem. C* **2020**, *124*, 10067.
- [30] P. Cnudde, K. De Wispelaere, J. Van der Mynsbrugge, M. Waroquier, V. Van Speybroeck, *J. Catal.* **2017**, *345*, 53.
- [31] K. De Wispelaere, P. N. Plessow, F. Studt, *ACS Physical Chemistry Au* **2022**, *2*, 399.
- [32] E. T. Vogt, B. M. Weckhuysen, *Chem. Soc. Rev.* **2015**, *44*, 7342.
- [33] G. Collinge, S. F. Yuk, M.-T. Nguyen, M.-S. Lee, V.-A. Glezakou, R. Rousseau, *ACS Catal.* **2020**, *10*, 9236.
- [34] V. Van Speybroeck, M. Bocus, P. Cnudde, L. Vanduyfhuys, *ACS Catal.* **2023**, *13*, 11455.
- [35] A. G. Stepanov, K. I. Zamaraev, J. M. Thomas, *Catal. Lett.* **1992**, *13*, 407.
- [36] A. G. Stepanov, K. I. Zamaraev, *Catal. Lett.* **1993**, *19*, 153.
- [37] A. G. Stepanov, V. Van Speybroeck, *J. Catal.* **2017**, *345*, 53.
- [38] A. G. Stepanov, M. V. Luzgin, V. N. Romannikov, V. N. Side-inkov, K. I. Zamaraev, *J. Catal.* **1996**, *164*, 411.
- [39] J. F. Haw, B. R. Richardson, I. S. Oshiro, N. D. Lazo, J. A. Speed, *J. Am. Chem. Soc.* **1989**, *111*, 2052.
- [40] M. Aronson, R. Gorte, W. Farneth, D. White, *J. Am. Chem. Soc.* **1989**, *111*, 840.
- [41] N. Lazo, B. Richardson, P. Schettler, J. White, E. Munson, J. Haw, *J. Phys. Chem.* **1991**, *95*, 9420.
- [42] J. F. Haw, J. B. Nicholas, T. Xu, L. W. Beck, D. B. Ferguson, *Acc. Chem. Res.* **1996**, *29*, 259.
- [43] H. Ishikawa, E. Yoda, J. N. Kondo, F. Wakabayashi, K. Domen, *J. Phys. Chem. B* **1999**, *103*, 5681.
- [44] K. Onda, K. Tanabe, H. Noguchi, K. Domen, A. Wada, *J. Phys. Chem. B* **2003**, *107*, 11391.
- [45] T. Bučko, M. Gešvandnerová, D. Rocca, *J. Chem. Theory Comput.* **2020**, *16*, 6049, <https://doi.org/10.1021/acs.jctc.0c00486>.
- [46] J. Rey, C. Chizallet, D. Rocca, T. Bučko, M. Badawi, *Angew. Chem. Int. Ed.* **2024**, *63*, e202312392, <https://doi.org/10.1002/anie.202312392>.
- [47] J. Rey, M. Badawi, D. Rocca, C. Chizallet, T. Bučko, *Catalysis Science, Technology* **2024**, *14*, 5314.
- [48] F. Berger, M. Rybicki, J. Sauer, *ACS Catal.* **2023**, *13*, 2011.
- [49] B. Herzog, M. Chagas da Silva, B. Casier, M. Badawi, F. Pascale, T. Bučko, S. Lebeuge, D. Rocca, *J. Chem. Theory Comput.* **2022**, *18*, 1382.
- [50] R. Ramakrishnan, P. O. Dral, M. Rupp, O. A. Von Lilienfeld, *J. Chem. Theory Comput.* **2015**, *11*, 2087.
- [51] V. Zaverkin, D. Holzmüller, L. Bonferraro, J. Kästner, *Phys. Chem. Chem. Phys.* **2023**, *25*, 5383.
- [52] J. Falk, L. Bonati, P. Novelli, M. Parrinello, M. Pontil, Transfer learning for atomistic simulations using GNNs and kernel

- mean embeddings, <https://doi.org/10.48550/arXiv.2306.01589> **2024**.
- [53] A. Erlebach, M. Šipka, I. Saha, P. Nachtigall, C. J. Heard, L. Grajciar, *Nat. Commun.* **2024**, *15*, 4215, <https://doi.org/10.1038/s41467-024-48609-2>.
- [54] M. Del Ben, J. Hutter, J. VandeVondele, *J. Chem. Theory Comput.* **2013**, *9*, 2654.
- [55] M. Del Ben, O. Schütt, T. Wentz, P. Messmer, J. Hutter, J. VandeVondele, *Comput. Phys. Commun.* **2015**, *187*, 120.
- [56] M. Bocus, R. Goeminne, A. Lamaire, M. Cools-Ceuppens, T. Verstraelen, V. Van Speybroeck, *Nat. Commun.* **2023**, *14*, 1008, <https://doi.org/10.1038/s41467-023-36666-y>.
- [57] S. Vandenhaute, T. Braeckvelt, P. Dobbelaere, M. Bocus, V. V. Speybroeck, Rare Event Sampling using Smooth Basin Classification, <https://doi.org/10.48550/arXiv.2404.03777> **2024**.
- [58] S. Vandenhaute, Psiflow - interatomic potentials using online learning, <https://github.com/molmod/psiflow> **2023**.
- [59] Y. Babuji, A. Woodard, Z. Li, D. S. Katz, B. Clifford, R. Kumar, L. Lacinski, R. Chard, J. M. Wozniak, I. Foster, M. Wilde, K. Chard, Parsl: Pervasive Parallel Programming in Python, in *Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing*, HPDC'19, Association for Computing Machinery, New York, NY, USA **2019** page 25–36.
- [60] T. D. Kühne, M. Iannuzzi, M. Del Ben, V. V. Rybkin, P. Seewald, F. Stein, T. Laino, R. Z. Khaliullin, O. Schütt, F. Schiffmann, et al., *J. Chem. Phys.* **2020**, *152*, <https://doi.org/10.1063/5.0007045>.
- [61] J. P. Perdew, K. Burke, M. Ernzerhof, *Phys. Rev. Lett.* **1996**, *77*, 3865.
- [62] S. Grimme, J. Antony, S. Ehrlich, H. Krieg, *J. Chem. Phys.* **2010**, *132*, <https://doi.org/10.1063/1.3382344>.
- [63] S. Grimme, S. Ehrlich, L. Goerigk, *J. Comput. Chem.* **2011**, *32*, 1456.
- [64] L. Vanduyfhuys, ThermoLIB, <https://molmod.ugent.be/software/thermolib> **2023**.
- [65] W. Humphrey, A. Dalke, K. Schulten, *J. Mol. Graphics* **1996**, *14*, 33.
- [66] J. A. Pople, *Rev. Mod. Phys.* **1999**, *71*, 1267.
- Manuscript received: July 19, 2024
Accepted manuscript online: October 31, 2024
Version of record online: November 25, 2024

Supporting Information:

**Unraveling the nature of adsorbed isobutene in
H-SSZ-13 with operando simulations at the top
of Jacob's ladder**

Massimo Bocus,^{†,‡} Sander Vandenhaute,^{†,‡} and Veronique Van Speybroeck^{*,†}

[†]*Center for Molecular Modeling, Ghent University, Technologiepark 46, 9052 Zwijnaarde, Belgium*

[‡]*These authors contributed equally*

E-mail: veronique.vanspeybroeck@ugent.be

Contents

1	Literature overview	S-3
2	Extended computational details	S-4
2.1	Zeolite model	S-4
2.2	First-principles evaluations	S-4
2.3	Collective variable	S-6
2.4	Enhanced sampling molecular dynamics	S-7
2.5	Active learning	S-8
2.6	Transfer learning	S-11
3	Methodology validation	S-14
3.1	RI-RPA accuracy	S-14
3.2	Active learning	S-16
3.3	Transfer learning	S-16
4	Extended results	S-24
4.1	Preferred reaction path exploration	S-24
4.2	Computed free energy surfaces	S-28
4.3	<i>tert</i> -butylation half life	S-35
4.4	Standard molecular dynamics	S-35
	References	S-39

1 Literature overview

Table S1: Summary of the theoretical works that have investigated isobutene physi- and chemisorption in periodic zeolite models throughout the years.

Year	Reference	LOT	PES exploration	Framework
1998	Sinclair et al. <i>J. Chem. Soc. Faraday Trans.</i> ¹	(MP2/6-31G*//HF/3-21G*):MM	static	CHA
2001	Boronat et al. <i>J. Phys. Chem. B</i> ²	HF/6-31G**	static	TON
2002	Clark et al. <i>Stud. Surf. Sci. Catal.</i> ³	QM-Pot (B3LYP/T(O)DZP)	static	FAU
2003	Rozanska et al. <i>J. Phys. Chem. B</i> ⁴	PW91/PW	static	CHA, TON, MOR
2004	Boronat et al. <i>J. Am. Chem. Soc.</i> ⁵	ONIOM (B3PW91/6-31G(d,p):MNDO)	static	MOR
2005	Tuma & Sauer <i>Angew. Chem. Int. Ed.</i> ⁶	PBE/PW	static	FER
2005	Nieminen et al. <i>J. Catal.</i> ⁷	QM-Pot (B3LYP/TZVP)	static	FER
2006	Tuma & Sauer <i>PCCP</i> ⁸	MP2/TZVP:PBE/PW + CBS extrapolation	static	FER
2008	De Moor et al. <i>J. Phys. Chem. C</i> ⁹	QM-Pot (MP2/TZVP//B3LYP/T(O)D)	static	FAU, FER
2010	Tuma et al. <i>Angew. Chem. Int. Ed.</i> ¹⁰	MP2/TZVP:PBE/PW + CBS extrapolation + CCSD(T) corrections	static	FER
2010	Rosenbach et al. <i>Chem. Phys. Lett.</i> ¹¹	ONIOM MP2/6-31G(d,p):MNDO	static	FAU
2012	Nguyen et al. <i>J. Phys. Chem. C</i> ¹²	PW91-D/PW//PW91/PW	static	FAU, MOR, MFI, TON
2015	Dai et al. <i>Angew. Chem. Int. Ed.</i> ¹³	BEEF-vdw/PW	static	CHA, MFI, FER, MOR
2015	Ferguson et al. <i>J. Phys. Chem. A</i> ¹⁴	PBE-vdw/PW	static	MFI
2017	Sarazen & Iglesia <i>PNAS</i> ¹⁵	revPBE-D3/PW	static	TON
2017	Cnudde et al. <i>J. Catal.</i> ¹⁶	PBE-D3/DZVP	metadynamics	MFI
2020	Plessow & Studt <i>J. Phys. Chem. Lett.</i> ¹⁷	DLPNO(CCSD(T)/cc-pVDZ:PBE-D2/PW	static	CHA
2020	Ren et al. <i>J. Phys. Chem. C</i> ¹⁸	MP2:PBE-D2 + Δ CCSD(T)	static	FER
2022	De Wispelaere et al. <i>J. Phys. Chem. Au</i> ¹⁹	PBE-D3/PW + DLPNO-CCSD(T)/def2-TZVPP corrections	US	CHA
2024	This work	RI-RPA/cc-QT(Si,Al)Z	US	CHA

2 Extended computational details

2.1 Zeolite model

The reaction was modelled in a conventional CHA unit cell consisting of 36 T Si atoms. One of these is replaced with an Al atom and a proton (the Brønsted acid site, BAS) on one of its bridging oxygen atoms, leading to a Si/Al of 35. The Al introduction breaks the unit cell symmetry, such that the three cages around the substitution are no longer equivalent. Moreover, the four oxygen atoms in the first coordination sphere of the Al defect are also nonequivalent. The BAS and the *tert*-butoxide species can then be bonded to any of these oxygen atoms and the isobutene molecule can be placed in any of the three cells (see Figure 1 b,c in the main manuscript). The equilibrium unit cell parameters were obtained from a previous investigation from some of the authors²⁰ (13.7733 Å, 13.7824 Å, 14.9983 Å, 90.06°, 90.05°, 119.91°) and kept fixed for all simulations, given the material's rigidity and the expected limited effect in minor volume changes on the free energy differences between the intermediates.

2.2 First-principles evaluations

First-principles calculations have been performed using the CP2K software^{21,22} (version 2023.1), relying on a combination of atom-centered and plane-waves basis sets. The latter was selected with an energy cutoff of 900 Ry to ensure well-converged results. GTH pseudopotentials²³ were used to smooth the electron density in proximity of the atomic core. 3 main levels of theory were employed to generate and/or validate our results:

- **A production level of theory**, which is used to generate an initial model from the active learning procedure (Section 2.5). We here selected the well-known dispersion-corrected PBE-D3(BJ)²⁴⁻²⁶ generalized gradient approximation (GGA) functional with a MOLOPT-TZVP quality basis set, due to its ubiquitous role as 'standard level of theory' for zeolite-catalyzed reactions.

- **A target level of theory** for the transfer learning scheme. To reach the highest possible accuracy for fully periodic simulations, we resorted to the random phase approximation^{27,28} (RPA) in combination with a cc-QZ basis set for H, C, O and a cc-TZ basis set for Si and Al. We will refer to this level of theory as RI-RPA/cc-QT(Al,Si)Z, as the resolution of identity (RI) approximation²⁷ with an RI-TZ auxiliary basis set was used to increase the simulation efficiency. The RI-RPA energies remarkably improve the PBE-D3(BJ) results, especially by counteracting the massive stabilization of the tert-butylcation intermediate (see Section 3.1). We originally selected 12 quadrature points for the minimax grid but we noticed that this was not providing perfectly converged energies for structures with a small bandgap (mostly associated with the *tert*-butylcation state). Hence, we selected all structures with a minimax gap below 2×10^{-2} Ha and re-evaluated them with 20 quadrature points, which provided better converged energies.
- **A test level of theory.** Since performing an extensive amount of *first-principles* simulations with RI-RPA is extremely demanding from a computational standpoint, we selected a second, cheaper level of theory to validate the transfer learning procedure. In an attempt to deviate as much as possible from dispersion-corrected GGA DFT, we opted for the local density approximation²³ (LDA) with a TZVP basis set.

To provide an idea of the differences in computational cost for the various methodologies, a single energy + forces evaluation on one of our structures takes about 1 minute on 64 cores (AMD Epyc Rome 7H12) while an energy-only RI-RPA evaluation requires about 1 hour on 256 cores (two nodes, each with two AMD Epyc Rome 7H12 64-core processors) and about 1 TB RAM.

2.3 Collective variable

The protonation and *tert*-butoxide formation reactions were studied by two separate collective variables (CVs). Both are based on coordination numbers (CNs), which are defined as:

$$\text{CN}(A; B) = \sum_{a \in A} \sum_{b \in B} \frac{1 - \left(\frac{r_{ab}}{r_0}\right)^6}{1 - \left(\frac{r_{ab}}{r_0}\right)^{12}} \quad (1)$$

Where A and B are two groups of atoms. The protonation is described by $\text{CN}(\text{C}_p; \text{H}) - \text{CN}(\text{O}_z; \text{H})$, where $\text{CN}(\text{C}_p; \text{H})$ is the CN between the primary isobutene carbon atoms and all hydrogen atoms in the system (including the BAS, $r_0 = 1.4 \text{ \AA}$) and $\text{CN}(\text{O}_z; \text{H})$ the CN between the zeolite oxygen atoms in the first coordination sphere of the Al defect and all hydrogen atoms ($r_0 = 1.4 \text{ \AA}$). This CV goes from ~ 7 for isobutene (8 C_p -H bonds and 1 O_z -H bond) to ~ 9 for the *tert*-butylation and *tert*-butoxide (9 C_p -H bonds and 0 O_z -H bonds). The *tert*-butoxide formation is described with $\text{CN}(\text{C}_t; \text{O}_z)$ between the tertiary carbon of isobutene and the four zeolite oxygens in the first coordination sphere of the Al defect ($r_0 = 2.0 \text{ \AA}$). This goes from ~ 0 for isobutene and *tert*-butylation to ~ 1 for *tert*-butoxide.

A single CV was also devised to model the two reactions at once for the active learning procedure. We took inspiration from a previous work from some of the authors²⁹ to propose the following linear combination:

$$\text{CV} = \frac{2}{\sqrt{5}} \text{CN}(\text{C}_t; \text{O}_z) + \frac{1}{\sqrt{5}} (\text{CN}(\text{C}_p; \text{H}) - \text{CN}(\text{O}_z; \text{H})) \quad (2)$$

The coefficient of the linear combination are chosen such that they represent a rotation in the space defined by $\text{CN}(\text{C}_t; \text{O}_z)$ and $\text{CN}(\text{C}_p; \text{H}) - \text{CN}(\text{O}_z; \text{H})$ with a subsequent projection on the x axis (Figure S1).

It is important to notice that this is *not* a ‘good’ CV as it does not properly separate the intermediates—i.e. it is not orthogonal to the transition state lines (Figure S1b). Nonetheless, a single CV remarkably simplifies the phase space exploration during active learning. For the

subsequent production US simulations with the final MLP models, we relied on the two partial CVs $CN(C_t;O_z)$ and $CN(C_p;H)-CN(O_z;H)$ as they are better suited to describe the protonation and alkoxide formation reactions, respectively.

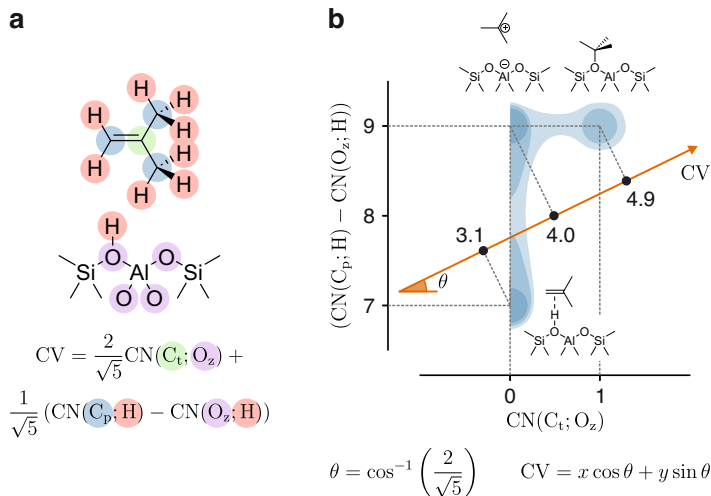


Figure S1: *a*. Summary of the atoms involved in the CNs used as CVs in this work. *b*. A 1-dimensional CV can be obtained by a rigid rotation of the 2-dimensional space defined by $CN(C_t;O_z)$ and $CN(C_p;H)-CN(O_z;H)$ with a subsequent projection on the x axis. The figure depicts a fictitious 2-dimensional free energy surface for the isobutene conversion.

2.4 Enhanced sampling molecular dynamics

The aforementioned CVs are used to perform enhanced sampling molecular dynamics (ESMD) simulations, where a bias is applied along the CV to force the system into exploring high-energy regions of the PES. The ESMD technique of choice for both the active learning procedure as well as for the final production runs is umbrella sampling^{30,31} (US). In US, the CV range of interest is explored with quadratic bias potentials

$$V(CV) = \frac{1}{2}\kappa(CV - CV_0)^2 \quad (3)$$

centered at CV_0 and with a force constant of κ . The final free energy profiles, free energy differences and related uncertainties are obtained with our in-house developed ThermoLIB library.^{32,33}

2.5 Active learning

The training set for the final MLP model is obtained using an active learning loop as recently proposed by some of us³⁴ and implemented in `psiflow`.³⁵ As explained in Section 2.1, the HSSZ-13 unit cell has three nonequivalent cages and four nonequivalent oxygen atoms in the first coordination sphere of the Al defect. To maximise the phase space explored by our simulations, we prepared an initial dataset of 24 manually-built structures, divisible in four groups of six (Figure S2). Each group contains a sequence of reactive intermediates, namely (1) isobutene in the middle of a cage, (2) isobutene π -complex, (3) protonation transition state, (4) *tert*-butyl cation in the middle of a cage, (5) *tert*-butyl cation in proximity of the Al defect and, finally, (6) surface *tert*-butoxide species. The four groups encode a different reactive path, where the isobutene is in one of the cages with the BAS on an accessible oxygen atom. The *tert*-butoxide is formed on the same oxygen on which the BAS was. The four cases we considered are (1) cage A, O₁, (2) cage B, O₂, (3) cage C, O₃ and (4) cage B, O₄. Remark that the structures were build manually by simply adjusting the bond lengths to reasonable values, analogously as what is commonly done to begin conventional DFT-based geometry optimizations or transition state searches, as it is unnecessary to start the active learning loop from ideal geometries. To generate an initial dataset, we apply small random perturbations to the atomic coordinates (normally distributed, with a standard deviation of 0.05 Å). Each structure was subjected to four random perturbations, resulting in a total of 96 structures in the initial dataset. The initial MLP that is used to start the active learning loop is obtained by training a MACE model to this initial dataset (with a 90:10 train/validation split). The most important network hyperparameters are provided in Table S2.

During active learning, we use the MLP from the previous iteration to perform a set of biased molecular dynamics simulations in order to generate samples, evaluate them with DFT to label

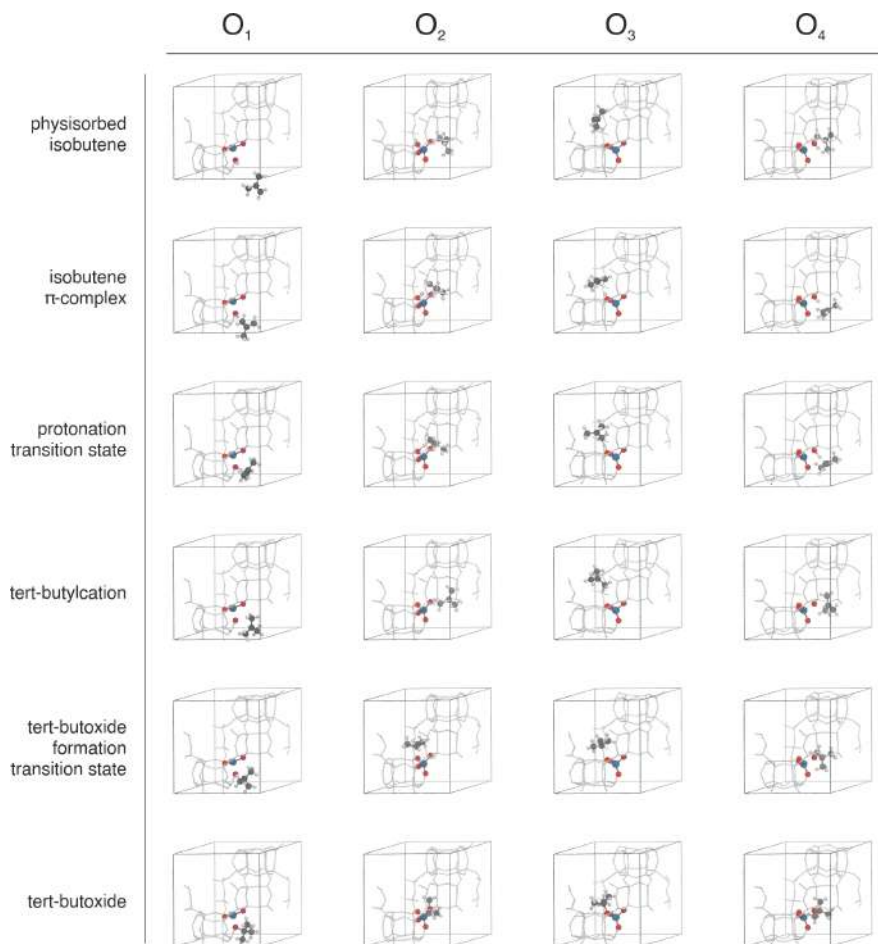


Figure S2: Schematic depiction of the 24 manually-built structures used as initial dataset for the active learning loop. Al is in blue, O in red, C in grey, and H in white, while the rest of the zeolite framework is depicted as a wireframe for the sake of clarity.

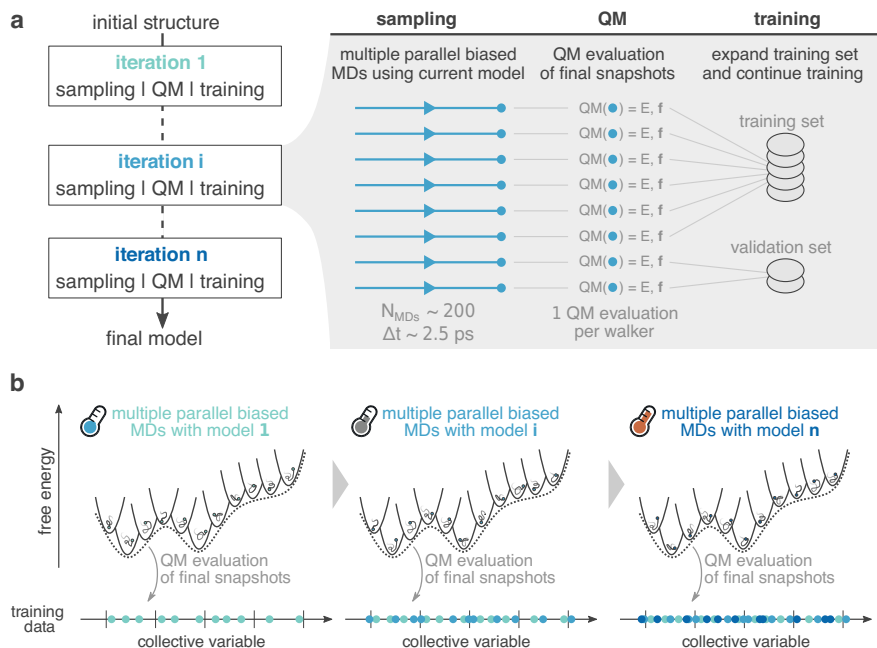


Figure S3: The active learning scheme used to generate the PBE-D3(BJ) training set relies on performing multiple parallel biased MD simulations using the current ML model. The biases are quadratic potentials distributed along the CV range of interest, analogously to a regular US simulation. The last structure of each MD is evaluated with CP2K and added to the training or validation sets. A new model is then trained on the extended dataset and used for the subsequent round of MD simulations (a). This strategy ensures a fast and uniform coverage of the whole CV region of interest (b). The figure is adapted with permission from Ref.³⁴

them with the PBE-D3(BJ) energies and forces, and retrain the MLP (Figure S3). A single such simulation is referred to as a *walker* because it traces a path through phase space. A total of 200 walkers are defined by creating 50 umbrellas equally spaced in the CV range [2.3, 4.4] for each of the four aforementioned groups of structures (note that the CV is invariant with respect to the specific BAS / isobutene / *tert*-butoxide locations, see Section 2.3). The umbrellas force constant was set equal to $3000 \text{ kJ}\cdot\text{mol}^{-1}$. Each walker is propagated for 2.5 ps with a timestep of 0.5 fs in the NVT ensemble using a Langevin thermostat³⁶ for temperature control. At the end

of the simulation, the last state for each of the walkers is evaluated with CP2K and added to the original dataset. A new model is then trained and the loop restarts (Figure S3). 20 loops were executed in total, providing a final dataset of ~ 3500 structures. A temperature ramp was used to bring the temperature from 173 to 1073 K during the first 5 iterations.

A number of safety checks were in place to guarantee that the sampled structures were chemically relevant. First, a walker was stopped and reset if the simulation temperature exceeded the target temperature by more than 1000 K or if any of the interatomic distances was smaller than 0.8 Å. Second, a walker was reset if the error between the current MACE model and CP2K was higher than 10 meV per atom or $200 \text{ meV}\cdot\text{Å}^{-1}$. Third, sampled states for which the SCF converged but for which the error between the current model and the DFT reference was larger than a certain threshold were discarded. This error threshold was set to 20 meV per atom for the energy and $500 \text{ meV}\cdot\text{Å}^{-1}$ for the forces. In this way, we avoid that unphysical geometries appear in the training set, which can corrupt the training/validation metrics and decrease the overall accuracy of the model.

After 20 iterations, the complete dataset is used to train a final model with an increased scheduler patience of 20 epochs and an early stopping patience of 40 epochs. In this way, the learning rate is gradually reduced if no improvement in the loss is observed for 20 epochs, allowing to better approach the local minimum in the hyperparameter space. Our final model, trained on 3682 PBE-D3(BJ) evaluated structures (3313 in the training and 369 in the validation set) achieves a root mean squared error (RMSE) of 0.3 meV per atom and $32 \text{ meV}\cdot\text{Å}^{-1}$ on the validation set.

2.6 Transfer learning

We now present a more detailed overview of the transfer learning approach as presented in Figure 1e in the main manuscript. When applying graph neural networks (GNNs) in the context of atomic systems, the atoms are considered as nodes in a graph. An edge is created between any two nodes whenever their distance is smaller than a predefined cutoff distance (6.5 Å in

Table S2: Summary of the network configuration adopted in the active learning loop. All other parameters are kept at their default values.

cutoff	6.5 Å
maximum rotation order (max_L)	2
correlation	3
number of channels	16
highest spherical harmonics order (max_e11)	3
number of radial basis functions	10
energy:forces weight in loss	200:1
batch size	16
learning rate	0.005

our case, see Table S2). Message passing GNNs represent the chemical environment of node i with a feature vector \mathbf{f}_i . The node features \mathbf{f}_i are updated through a series of message passing or *interaction* layers, which allows neighboring atoms to exchange information regarding their chemical environments. The total potential energy of the system is decomposed into per-atom contributions E_i , and each per-atom contribution is obtained using a simple readout function R on the node features \mathbf{f}_i :

$$E = \sum_{i=1}^n E_i \quad (4)$$

$$= \sum_{i=1}^n R(\mathbf{f}_i) \quad (5)$$

For MACE in particular, the readout function is layer-dependent; the first $N - 1$ readouts are strictly linear functions of the (scalar part of the) node features, whereas the last readout contains a single nonlinearity.³⁷ By construction, Equation 5 implies that the node features \mathbf{f}_i are informative for the interactions that involve atom i . For example, a GNN which is trained on PBE-D3(BJ) energy and force targets ($E_{\text{PBE-D3(BJ)}}$, $-\nabla E_{\text{PBE-D3(BJ)}}$) will yield node features which are informative for the PBE interaction energy in the system.

Our goal is to predict a systematic energy correction to $E_{\text{PBE-D3(BJ)}}$:

$$\Delta E = E_{\text{RI-RPA}} - E_{\text{PBE-D3(BJ)}} \quad (6)$$

using a minimal number of additional RI-RPA energy evaluations since these are expensive to evaluate. Even more important is the fact that nuclear gradients at the RI-RPA level are numerically noisy and often computationally too demanding to compute. As such, we would ideally learn to predict ΔE by training on a limited number of *energy-only* targets. As an ansatz, we consider the fact that the energy correction between PBE-D3(BJ) and RI-RPA is systematic, and therefore learnable based on the (existing) PBE-trained GNN, which was trained to a much larger volume of $(E_{\text{PBE-D3(BJ)}}, \nabla E_{\text{PBE-D3(BJ)}})$ targets. The most straightforward way to proceed is then to reuse the PBE-trained GNN but with a second readout layer which predicts a per-atom energy correction ΔE_i based on the original node features \mathbf{f}_i (see also Figure 1d in the main manuscript):

$$\Delta E = \sum_{i=1}^n \Delta E_i \quad (7)$$

$$= \sum_{i=1}^n \Delta R(\mathbf{f}_i) \quad (8)$$

In analogy with the conventional MACE readout layers R , we consider ΔR as a simple layer-dependent function which is applied to the (scalar part of the) node features \mathbf{f}_i after each layer. Contrary to R which is nonlinear only in the last interaction layer, we consider ΔR nonlinear for all interaction layers. For each interaction layer, it consists of two dense layers and two swish nonlinearities.

As such, the new readout function ΔR contains only about 1,000 trainable weights, as opposed to the 100,000 trainable weights in the full GNN. This makes its training highly data efficient, to the extent that it can be performed based on energy-only targets (i.e. $E_{\text{RI-RPA}} - E_{\text{PBE-D3(BJ)}}$, without gradients). The general procedure is then to first train the GNN to PBE-D3(BJ) energies and forces using the active learning approach as outlined in Section 2.5, after which the ΔR readouts are trained using a much smaller dataset which is labeled with the correct energy difference, i.e. $\Delta E = E_{\text{RI-RPA}} - E_{\text{PBE-D3(BJ)}}$. Importantly, during training of ΔR , the GNN itself and the original PBE-D3(BJ) readout R are kept fixed such that the network does not forget its PBE-D3(BJ) training; only the weights in ΔR are optimized.

Starting from the complete PBE-D3 dataset from the active learning loop, 3 grids were defined with 200, 400, and 800 equally-spaced points along the reaction CV (in the range 2.3–4.3) and we extracted the closest structure to each point from the dataset. Sometimes this led to duplicate structures, which were discarded providing 3 datasets of 200, 400 and 788 structures. We also used the whole dataset of 3656 structures for reference (this number is smaller than the number of structures in the complete PBE-D3 dataset as, for few structures, the RPA SCF did not converge).

After few tests, we noticed that the final results are mostly dependent on the weight decay and the size of the additional readout layers. For this reason, we tested both a weight decay of 0.0 and 1E-6 together with ΔR layer sizes of 8 and 16. The models were retrained on the four datasets achieving in every case sub-meV·atom⁻¹ errors (Table S3). Of course, this does not necessarily mean that the models can now be used to perform RPA-quality enhanced sampling simulations. For this reason, we extensively tested the reliability of the transfer learning procedure against an extremely challenging case study, i.e. going from PBE-D3 to LDA (Section 3.3).

Table S3: RMSE on the validation set energies (meV·atom⁻¹) obtained for various RI-RPA dataset sizes (90:10 train-alidation split). The two tested hyperparameters are the weight decay and the layer sizes within ΔR .

weight decay	ΔR layers	dataset size			
		200	400	788	3656
1E-6	[8,8]	0.6	0.5	0.5	0.7
1E-6	[16,16]	0.6	0.6	0.5	0.8
0.0	[8,8]	0.6	0.5	0.5	0.7
0.0	[16,16]	0.6	0.5	0.5	0.7

3 Methodology validation

3.1 RI-RPA accuracy

To confirm that the selected target level of theory (RI-RPA/cc-QT(Al,Si)Z) produces results which are significantly better than the production level of theory (PBE-D3(BJ)/MOLOPT-TZVP), we

performed a single point evaluation on the equilibrium structures for the π -complex, the *tert*-butylation and the *tert*-butoxide species in H-SSZ-13 reported by Plessow *et al.*¹⁷ The study adopts a chemically-accurate approach originally proposed by Tuma and Sauer,⁸ where a periodic structure is optimized with GGA DFT and high-level corrections are applied based on clusters extracted from the periodic structure. In particular, Plessow *et al.* use DLPNO-CCSD(T)/cc-PVDZ calculations with complete basis set (CBS) corrections on a 46T cluster obtained from a periodic PBE-D3 optimization.

The results are shown in Figure S4, where it can be seen how the energy difference between the π -complex and *tert*-butylation is perfectly reproduced while the *tert*-butoxide is modestly overstabilized by RI-RPA ($\sim 10 \text{ kJ}\cdot\text{mol}^{-1}$). For completeness, we also tested the RI-MP2^{27,28,38} method with the same basis set as RPA and the auxiliary density matrix method³⁹ (ADMM) with a tz2p auxiliary basis set to speed up the calculation of the Hartree-Fock exchange. The results are in line with the RPA ones, though some additional overstabilization of the *tert*-butoxide species is observed (Figure S4). This is in line with the results reported by Ren *et al.*¹⁸ in H-FER, where the cluster correction is split in two parts: first, a MP2 correction was done on a large cluster and second an additional CCSD(T) correction is performed on a smaller 5T cluster. The final CCSD(T) nudge to the energy is reported to be $+0.1 \text{ kJ}\cdot\text{mol}^{-1}$ for the *tert*-butylation and $+8.0 \text{ kJ}\cdot\text{mol}^{-1}$ for the *tert*-butoxide. Therefore, not surprisingly, both RPA and MP2 lack a certain amount of electron correlation required to reach CCSD(T) quality results for the *tert*-butoxide intermediate. Finally, we tested the hybrid functional PBE0-D3(BJ)/MOLOPT-TZVP+ADMM which, despite a small improvement, retains a major overstabilization of the *tert*-butylation (Figure S4).

These tests indicate that RI-RPA is the best option for fully periodic high quality results, as it completely recovers the deficiencies of GGA and hybrid DFT in overstabilizing charged intermediates at the price of a modest overstabilization of the surface-bounded ones.




Figure S4: 0 K electronic energy differences between the tert-butylcation and tert-butoxide referred to the isobutene π -complex, as obtained by Plessow et al.¹⁷ (DLPNO-CCSD(T)/cc-PVDZ:PBE-D3/PW) and by us (RI-RPA/cc-QT(Al,Si)Z, RI-MP2/cc-QT(Al,Si)Z+ADMM, PBE0-D3(BJ)/MOLOPT-TZVP+ADMM, PBE-D3(BJ)/MOLOPT-TZVP).

3.2 Active learning

It is essential for the validity of our results that the final MLP trained on the active learning dataset can reproduce with very high fidelity the original first-principles PES. To assess this, we performed a brute-force independent PBE-D3(BJ) US simulation at 623 K using CP2K (85 umbrellas, 50 ps simulation time) having the isobutene in cage B, the BAS on O₂ and the tert-butoxide on O₄. The same simulation was then repeated with the trained MLP from active learning. The results are shown in Figure S5. It can clearly be seen that the final MACE model perfectly matches the 623 K brute-force CP2K FEP with deviations in the order of ± 2 kJ·mol⁻¹, which are well contained within the statistical uncertainties of the profiles.

3.3 Transfer learning

Validating the transfer learning procedure poses more challenges, as it is impossible to perform any type of active learning with RI-RPA given the enormous computational cost that is associated with RPA energy and forces evaluations. To overcome this limitation, we devised a challenging




Figure S5: Comparison between the 623 K FEPs obtained from umbrella sampling simulations using CP2K and our final MACE model. The FEP minimum is set to 0 for both profiles.

case study that provides solid evidence for the validity of our approach. Instead of setting our initial target to RI-RPA, we went in the opposite direction in terms of accuracy and selected LDA (see Section 2.2). While PBE and RPA have been shown to explore relatively similar regions of phase space for hydrocarbon adsorption in zeolites,⁴⁰ this is much less likely for LDA where large changes in the predicted electron density are expected. As such, transfer learning from PBE to LDA is expected to be significantly more challenging as compared to transfer learning from PBE to RPA.

First, we trained a MACE model based on the active learning cycle outlined in Section 2.5, with LDA CP2K evaluations. All settings are kept identical with respect to the PBE-D3(BJ) active learning loop. After 20 loops, we obtained a dataset of 3389 structures which allowed us to train a final model achieving an RMSE of 0.3 meV per atom and 28 meV·Å⁻¹ on energies and forces of the validation set, respectively. Since these values are basically identical to the PBE-D3(BJ) results (Section 2.5), we can assume that the model is a faithful reproduction of the LDA PES region of interest and can be used as reference. Remark that, coming from the active learning loop, *this model should cover the most important regions of the LDA phase space, which do not necessarily*

overlap with the PBE-D3(BJ) ones.

We then moved to the actual transfer learning. To this end, we re-evaluated the whole PBE-D3(BJ) dataset from the active learning procedure with LDA. Few SCFs did not converge, leading to a final dataset of 3679 structures. 3 grids were defined with 200, 400 and 800 equally-spaced points along the reaction CV (in the range 2.3–4.3) and we extracted the closest structure to each point from the dataset. Sometimes this led to duplicate structures, which were discarded providing 3 datasets of 200, 400 and 788 structures. We also used the whole dataset of 3679 structures for reference. As a reminder, this dataset consists of structures collected from the PBE-D3(BJ) active learning loop which are currently labeled with the energy difference between LDA and PBE-D3(BJ).

Similarly to RI-RPA (Supplementary Section 2.6), the final results are mostly dependent on the weight decay and the size of the additional readout layers. For this reason, we tested both a weight decay of 0.0 and 1E-6 together with two additional readout layers of size [8,8] and [16,16]. The models were retrained on the four datasets achieving in every case sub-meV·atom⁻¹ errors (Table S4).

Table S4: RMSE on the validation set energies (meV·atom⁻¹) obtained for various LDA-level dataset sizes (90:10 train-validation split). The two tested hyperparameters are the weight decay and the layer sizes within ΔR .

weight decay	ΔR layers	Dataset size			
		200	400	788	3679
1E-6	[8,8]	0.7	0.8	0.5	0.5
1E-6	[16,16]	0.6	0.8	0.5	0.4
0.0	[8,8]	0.7	0.8	0.5	0.4
0.0	[16,16]	0.6	0.7	0.5	0.3

Because of the limited difference in achieved error between the [8,8] and [16,16] models, we decided to test both (weight decay 0.0) against the LDA model from active learning. The FEPs are reported in Figure S6 and Figure S7 for the [8,8] and [16,16] models, respectively. It can be seen how increasing the amount of training data has a beneficial effect in both cases, where sub 5 kJ·mol⁻¹ differences can be obtained starting from 788 training structures. The size of the

model seems also to have an impact, where the [16,16] model performs generally slightly better than the [8,8] model. This suggests that the proposed approach is suited to adequately capture the PES differences even between very different methods, a task which remained challenging for other methodologies such as machine learning free energy perturbation.⁴⁰

With these results supporting our transfer learning protocol, we repeated the same simulations with RI-RPA. Also in that case, we reevaluated the full PBE-D3 dataset for benchmark purposes. First, we compared the [8,8] models trained on various amounts of training data (Figure S9). Additionally, we assessed consistency between the results by comparing the [8,8] and [16,16] models trained on the whole dataset (Figure S10). As it can be seen, the model appears to have much less difficulty in reproducing the RI-RPA FEPs than in the case of LDA, as excellent convergence is already observed with 200 training structures. This is to be anticipated given that RPA is essentially a second-order correction energy based on PBE orbitals, whereas LDA is an entirely different functional with a different ground-state density. Furthermore, previous work on hydrocarbon adsorption in zeolites by Herzog and co-workers⁴⁰ has shown that for these systems, there exists significant overlap between the phase spaces explored by PBE-D3 and RPA. In view of these results, we opted to use the [8,8] model trained on 200 structures for the production runs as reported Section 4 and in the main manuscript, which achieved a validation RMSE on the energy of 0.6 meV per atom. It should be noted that the validation set contained only 20 structures – based on a 90:10 split of the 200 structures – so the validation error is not necessarily representative for the actual performance of the model in production simulations. Therefore, we explicitly validated the test error of the model during the production enhanced sampling simulations at 623 K by extracting snapshots from all umbrellas. Two snapshots were extracted per umbrella; one at 25 ps and one at 50 ps such that they are fully decorrelated. There are 50 umbrellas for the protonation and 50 for the alkoxide formation, leading to a total of 400 structures. These were evaluated at the RI-RPA level of theory, and compared with the model's prediction. As such, we obtain a test RMSE on the energy of 0.36 meV per atom, which roughly corresponds to a maximum deviation of $4.2 \text{ kJ}\cdot\text{mol}^{-1}$ on the energy of a full structure.

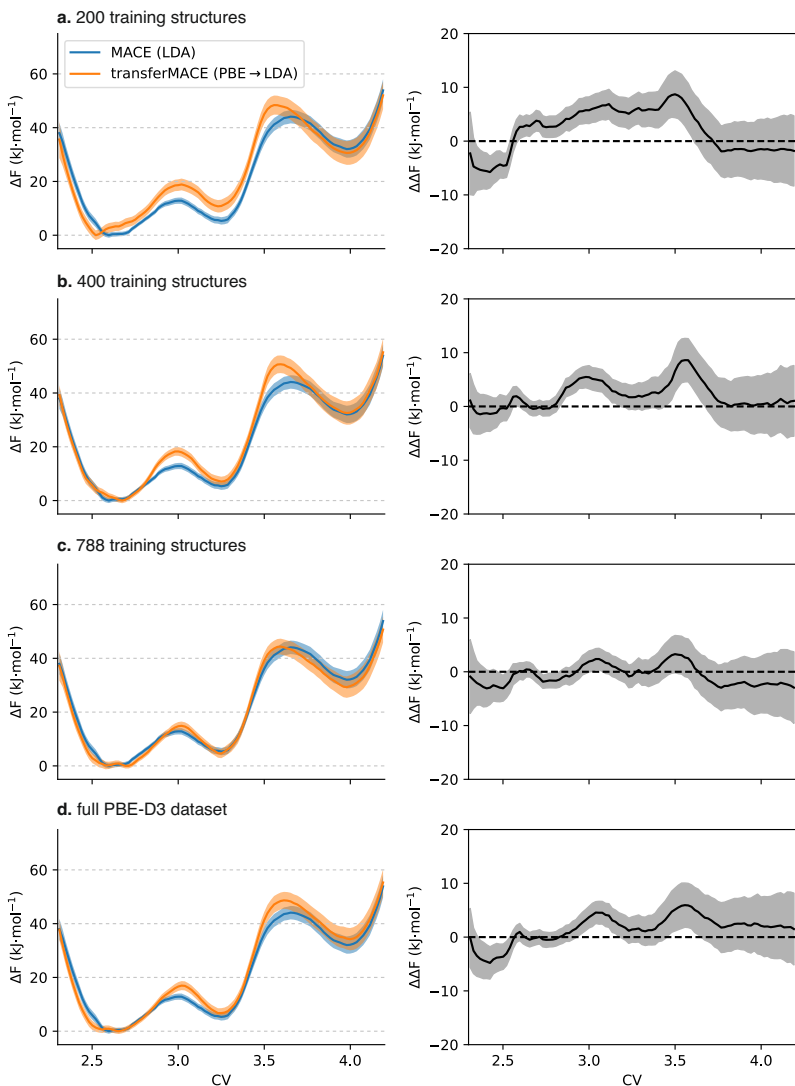


Figure S6: Comparison between the LDA FEPs (623 K) obtained from a MACE model trained on active learning results and [8,8] transferMACE models based on the PBE-D3 model from active learning trained on 200 (a), 400 (b), 788 (c) structures and the full PBE-D3 dataset (d).

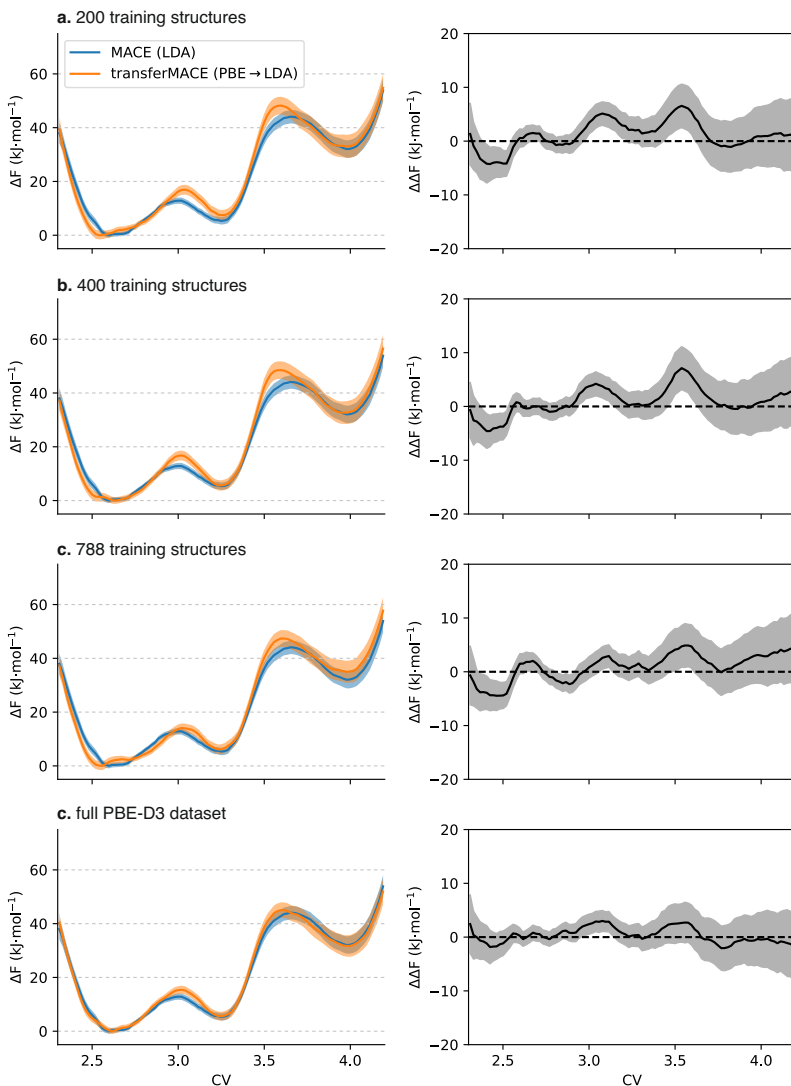


Figure S7: Comparison between the LDA FEPs (623 K) obtained from a MACE model trained on active learning results and [16,16] transferMACE models based on the PBE-D3 model from active learning trained on 200 (a), 400 (b), 788 (c) structures and the full PBE-D3 dataset (d).




Figure S8: Atomic structures of the *tert*-butoxide and *tert*-butylcation configurations, with atoms color-coded according to ΔE_i .

By construction, interatomic potentials based on GNNs are trained to decompose the total potential energy of the system into per-atom contributions (see Equation 5). The same is true for the delta energies ΔE_i ; they represent a partitioning of the energy difference between RI-RPA and PBE-D3(BJ) into atomic contributions. As such, it should be large in absolute value whenever atoms find themselves in chemical environments for which PBE-D3(BJ) diverges from RI-RPA, and it should be approximately zero whenever both levels of theory would largely yield the same interaction energy for that environment. For the chemisorption of isobutene in H-CHA, this implies that the correction energies ΔE_i are expected to be zero for most of the framework atoms (because RI-RPA and PBE-D3(BJ) are expected to agree rather well for those chemical environments), and mostly nonzero for the atoms in the active site as well as the isobutene molecule. We explicitly verified this by visualizing the distribution of ΔE_i throughout the structure in Figure S8, for both the *tert*-butoxide and *tert*-butylcation configurations. Interestingly, this supports the conclusion that we succeeded in having the network understand what the pertinent differences between GGA and second-order post-HF are – all based on just a few hundred energy-only singlepoint evaluations.

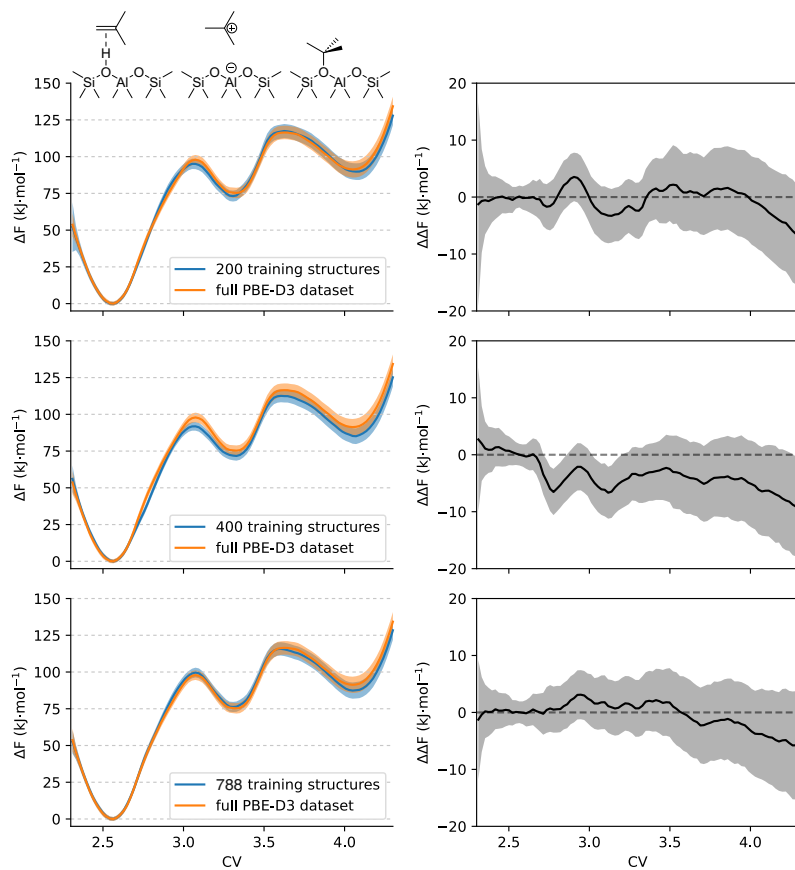


Figure S9: Comparison between the RI-RPA FEPs (623 K) obtained from [8,8] transferMACE models based on the PBE-D3 model from active learning trained on 200, 400 and 788 structures compared with a model trained on the full PBE-D3 dataset.




Figure S10: Comparison between the RI-RPA FEPs (623 K) obtained from the [8,8] and [16,16] transferMACE models based on the PBE-D3 model from active learning trained on the full PBE-D3 dataset.

4 Extended results

4.1 Preferred reaction path exploration

Having 4 nonequivalent oxygen atoms and 3 nonequivalent cages would, in principle, give rise to 12 possible reactive paths. However, some of these are clearly unfeasible. For instance, when the BAS is on O_3 it is pointing towards the inside of cage C, making it inaccessible from cage A or B. This reduces the number of cases to 7, which are summarized in Table S5.

Table S5: Summary of the reaction paths that are in principles possible within the H-SSZ-13 unit cell. For the protonation reaction, the oxygen atom is where the BAS is located while the cage is where the isobutene is located. For the *tert*-butoxide formation reaction, the oxygen atom is where the alkoxide is formed while the cage is where the *tert*-butylcation is located.

oxygen atom	cage
O_1	A
O_1	B
O_2	B
O_2	C
O_3	C
O_4	A
O_4	B

Note that all isobutene intermediates are too hindered to easily diffuse through the 8-member

rings during the simulations, hence the necessity of considering O₂-cage B and O₂-cage C as separate cases. We tested all these paths for both the protonation reaction (where BAS is located on the indicated oxygen atom while the isobutene in the indicated cage) and the *tert*-butoxide formation (where *tert*-butoxide is located on the indicated oxygen atom while the *tert*-butylation in the indicated cage). Note that the two reactive events are now described by two separate CVs, as shown in Figure S1. We started by performing a 50,000 steps MOVINGRESTRAINT simulation (as defined in PLUMED). In the first 5,000, the system goes from a κ (see Equation 3) of 500 to one of 3000 kJ·mol⁻¹ while fixed at a CV value of 5.3 (for the protonation reaction, corresponding to BAS + isobutene) or 1.0 (for the alkoxide formation, corresponding to the *tert*-butoxide species). Then, in the remaining time, the bias gradually moves to a CV value of 7.8 (protonation, *tert*-butylation) or 0.0 (alkoxide formation, *tert*-butylation). These simulations do not only give a first indication of whether a reaction path is feasible, but they also generate a useful set of initial structures for subsequent umbrella sampling simulations. All MOVINGRESTRAINT simulations turned out to be successful, except for the cases in which the BAS/*tert*-butoxide are on O₁ and O₄ with the mobile species in cage B and A, respectively. This is possibly due to a combination of steric hindrance and unfavorable BAS orientation that makes the reaction unlikely to occur.

With these two paths excluded, we are left with 5 possible ones that were all further investigated with US simulations at 623 K (Figure S11 and Figure S12). The results show that PBE-D3(BJ) and RI-RPA are in general qualitative agreement and the same conclusions can be drawn for both. To limit the amount of simulations, we considered two paths: the first one involves protonation of isobutene in cage B from a BAS on O₄ with subsequent formation of the *tert*-butoxide again on O₄. This path is kinetically favorable, with barriers of both steps that are significantly lower than the other pathways. However, a possible reason for this is the high reactivity of the BAS on O₄, which reflects in a low stability. Indeed, when investigating the proton hopping reaction in H-SSZ-13, we found that its occupancy is very low and only becomes significant at higher temperatures. Therefore, we decided to consider a second path as well, where isobutene is in cage C and gets protonated by the BAS on O₃ while forming a *tert*-butoxide species on O₂ after-

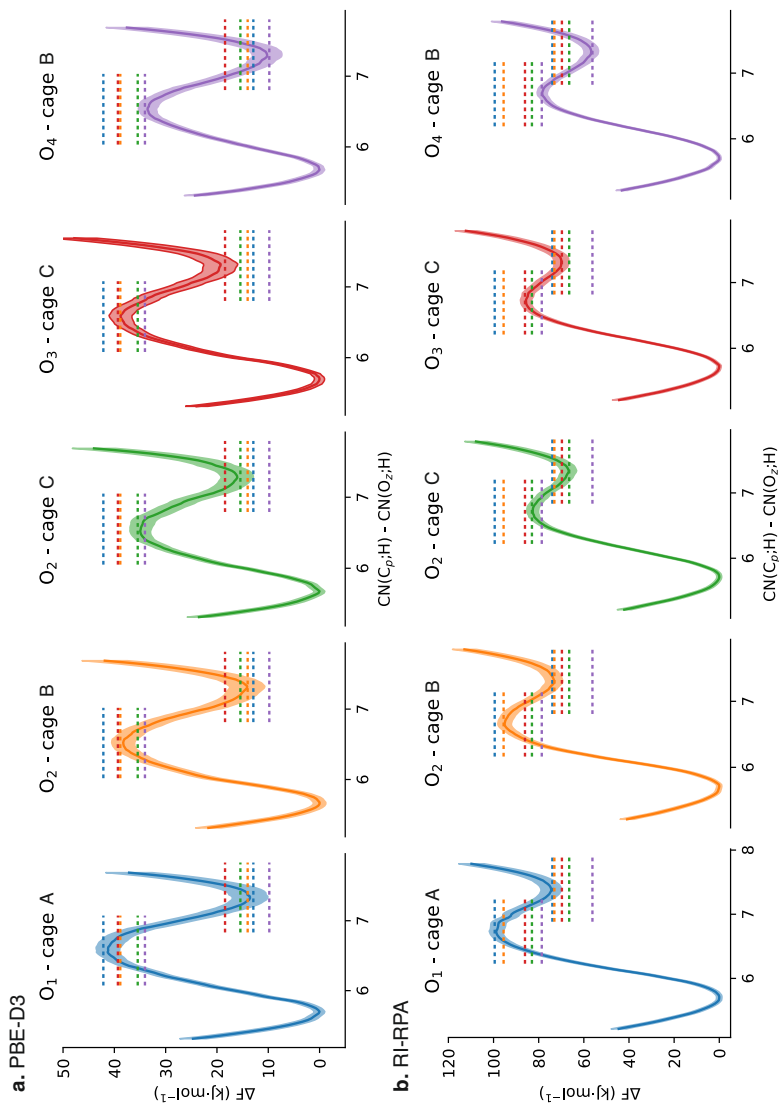


Figure S11: 623 K FEPs obtained for all the possible protonation pathways at a PBE-D3(BJ) (a) and RI-RPA (b) levels of theory. The colored horizontal lines indicate the height of transition state and product in all profiles and are intended to guide the reader's eye.

wards. This path still has reasonably moderate barriers while O₃ was predicted to have a very high occupancy according to our previous work.

4.2 Computed free energy surfaces

As explained in Section 4.1, two paths were investigated: one in which isobutene is in cage B, the BAS is on O₂ and the *tert*-butoxide is also formed on O₂, and a second path in which isobutene is in cage C, the BAS is on O₃ and the *tert*-butoxide is formed on O₂. To this end, we performed US simulations in the temperature range 323–823 K in steps of 100 K. For each temperature, two sets of 1-dimensional US simulations was performed, one along CN(C_p;H)-CN(O₂;H) and a second along CN(C_t;O₂). The two simulations were then combined in a single FES *via* 2-dimensional WHAM analysis, imposing a κ of zero along the unbiased direction (see Equation 3).

The resulting 2-dimensional FESs together with their error estimate are reported in Figure S13–Figure S20. Remark that for the error estimate, we ignored points for which the biased probability was 5×10^{-4} as these are generated by random fluctuation of the system at the edges of the profile and produce such large error bars that they would compromise any subsequent analysis of the profiles.

Additionally, we computed reaction rates within the transition state theory (TST) approximation.³² First, we cropped the portion of the FES not related to the reaction under investigation (III for I↔II and I for II↔III) and we projected the resulting cropped FES along the CV of interest (CN(C_p;H)-CN(O₂;H) for I↔II and CN(C_t;O₂) for II↔III). The projection formula is simply:

$$F(q) = -k_{\text{B}}T \ln \left(\int_{-\infty}^{+\infty} \exp \left(-\frac{F(q, q')}{k_{\text{B}}T} \right) dq' \right). \quad (9)$$

This gives a 1-dimensional free energy profile (FEP) from which the reaction rate can be computed as:

$$k_{\text{TST}} = \sqrt{\frac{k_{\text{B}}T}{2\pi}} \langle |\nabla_x q| \rangle_{q=q^*} \frac{\exp \left(-\frac{F(q^*)}{k_{\text{B}}T} \right)}{\int_{\mathcal{A}} \exp \left(-\frac{F(q)}{k_{\text{B}}T} \right) dq} \quad (10)$$

In which q is the CV, k_B is Boltzmann's constant, $\langle |\nabla_x q| \rangle_{q=q^*}$ is the averaged gradient of the CV with respect to the mass-weighted atomic coordinates computed atop the TS (q^*) and $\int_{\mathcal{A}} \exp\left(-\frac{F(q)}{k_B T}\right) dq$ is the integrated free energy of the reactant basin \mathcal{A} . For convenience, the kinetic constant can be transformed in a *phenomenological free energy barrier* using the relation:

$$\Delta F^\ddagger = -k_B T \ln(k_{\text{TST}} \beta h), \quad (11)$$

A summary of all phenomenological barriers between all intermediates in both cages and all temperatures is reported in Table S6–Table S9. Note that for PBE-D3(BJ) the *tert*-butylation **II** and *tert*-butoxide **III** are not separated by a clear transition state (Figure S12a) and therefore no barrier can be computed between the two species.

Table S6: Phenomenological free energy barrier for isobutene protonation (**I**↔**II**) in cage B at a PBE-D3(BJ) level of theory.

Reaction	T (K)	$\Delta F_{\text{forward}}^\ddagger$ (kJ·mol ⁻¹)	$\Delta F_{\text{backward}}^\ddagger$ (kJ·mol ⁻¹)
I ↔ II	323	16 ± 1	5 ± 1
	423	17 ± 1	8 ± 2
	523	19 ± 1	10 ± 2
	623	22 ± 1	14 ± 2
	723	24 ± 1	18 ± 2
	823	28 ± 2	23 ± 2

Table S7: Phenomenological free energy barrier for isobutene protonation (**I**↔**II**) in cage C at a PBE-D3(BJ) level of theory.

Reaction	T (K)	$\Delta F_{\text{forward}}^\ddagger$ (kJ·mol ⁻¹)	$\Delta F_{\text{backward}}^\ddagger$ (kJ·mol ⁻¹)
I ↔ II	323	18 ± 1	2 ± 1
	423	20 ± 1	4 ± 2
	523	23 ± 1	7 ± 2
	623	27 ± 1	10 ± 2
	723	28 ± 1	14 ± 2
	823	32 ± 2	17 ± 2

Table S8: Phenomenological free energy barrier for isobutene protonation (**I**↔**II**) and *tert*-butoxide formation (**II**↔**III**) in cage B at a RI-RPA level of theory.

Reaction	T (K)	$\Delta F_{\text{forward}}^{\ddagger}$ (kJ·mol ⁻¹)	$\Delta F_{\text{backward}}^{\ddagger}$ (kJ·mol ⁻¹)
I ↔ II	323	59 ± 1	3 ± 1
	423	61 ± 1	5 ± 1
	523	65 ± 1	8 ± 2
	623	65 ± 1	11 ± 2
	723	70 ± 2	15 ± 2
	823	74 ± 2	19 ± 2
II ↔ III	323	24 ± 2	39 ± 2
	423	29 ± 2	35 ± 2
	523	35 ± 2	32 ± 2
	623	43 ± 2	30 ± 2
	723	51 ± 2	28 ± 2
	823	62 ± 2	26 ± 2

Table S9: Phenomenological free energy barrier for isobutene protonation (**I**↔**II**) and *tert*-butoxide formation (**II**↔**III**) in cage C at a RI-RPA level of theory.

Reaction	T (K)	$\Delta F_{\text{forward}}^{\ddagger}$ (kJ·mol ⁻¹)	$\Delta F_{\text{backward}}^{\ddagger}$ (kJ·mol ⁻¹)
I ↔ II	323	64 ± 1	0 ± 1
	423	66 ± 2	2 ± 2
	523	68 ± 1	4 ± 2
	623	72 ± 1	7 ± 2
	723	76 ± 2	9 ± 2
	823	78 ± 2	13 ± 2
II ↔ III	323	28 ± 1	34 ± 1
	423	35 ± 1	31 ± 2
	523	42 ± 2	28 ± 2
	623	51 ± 2	25 ± 2
	723	60 ± 2	24 ± 2
	823	69 ± 2	22 ± 2




Figure S13: 2-dimensional PBE-D3(B) FESs for the conversion of isobutene (**I**) to *tert*-butylation (**II**) and *tert*-butoxide (**III**) in cage B, computed for various reaction temperatures. The associated uncertainties are shown in Figure S14.




Figure S14: Uncertainties on the FESs shown in Figure S13. Remark that the points with a biased probability lower than 5×10^{-4} are excluded being statistically irrelevant.




Figure S15: 2-dimensional PBE-D3(B) FESs for the conversion of isobutene (**I**) to *tert*-butylation (**II**) and *tert*-butoxide (**III**) in cage C, computed for various reaction temperatures. The associated uncertainties are shown in Figure S16.




Figure S16: Uncertainties on the FESs shown in Figure S15. Remark that the points with a biased probability lower than 5×10^{-4} are excluded being statistically irrelevant.

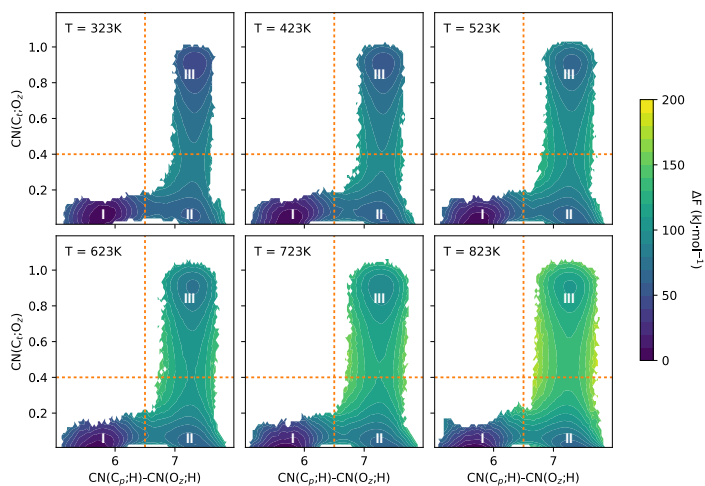


Figure S17: 2-dimensional RI-RPA FESs for the conversion of isobutene (**I**) to *tert*-butylcation (**II**) and *tert*-butoxide (**III**) in cage B, computed for various reaction temperatures. The associated uncertainties are shown in Figure S18.




Figure S18: Uncertainties on the FESs shown in Figure S17. Remark that the points with a biased probability lower than 5×10^{-4} are excluded being statistically irrelevant.




Figure S19: 2-dimensional RI-RPA FESs for the conversion of isobutene (**I**) to *tert*-butylcation (**II**) and *tert*-butoxide (**III**) in cage C, computed for various reaction temperatures. The associated uncertainties are shown in Figure S20.




Figure S20: Uncertainties on the FESs shown in Figure S19. Remark that the points with a biased probability lower than 5×10^{-4} are excluded being statistically irrelevant.

4.3 *tert*-butylation half life

The half life of the *tert*-butylation is roughly estimated with the relation $10^{\tau_{1/2}} = \ln 2/k_{\text{II} \rightarrow \text{I}}$. The backward kinetic constant $k_{\text{II} \rightarrow \text{I}}$ is estimated from the US simulations within the transition state theory approximation as explained in the previous section. The results are shown in Figure S21.




Figure S21: *tert*-butylation half life estimates as obtained from ESMD simulations.

4.4 Standard molecular dynamics

To study the behavior of isobutene and the *tert*-butylation without the US biases, we performed 1 ns long unbiased MD simulations of the two species in both cages B and C at 323, 623 and 823 K. Figure S22 shows the time evolution of $\text{CN}(\text{C}_p;\text{H})-\text{CN}(\text{O}_2;\text{H})$ in the various simulations, where it can be seen how the *tert*-butylation is regularly sampled by PBE-D3(BJ) even at 323 K while for RI-RPA it immediately decays to isobutene within the first few ps of simulation (insets in the right graphs of Figure S22) and is only formed once for few ps at 823 K.

We then performed additional biased simulations at 623 K to better sample the *tert*-butylation. To this end, we introduced two half-quadratic potential walls, one acting for $\text{CV}_1 < 6.5$ and one acting for $\text{CV}_2 > 0.4$ (both with a κ of $3000 \text{ kJ}\cdot\text{mol}^{-1}$). We used these simulations to evaluate the difference in mobility between the two species. As it can be seen, both levels of theory present very similar mobilities, while the change of cage has a much more pronounced effect. Indeed, in

cage B isobutene and the *tert*-butylcation sample similar volumes because the Al defect is located around the central region of the cage. In cage C, on the other hand, Al is on the base. This causes the *tert*-butylcation to be constrained to the bottom region of the cage while isobutene samples occasionally also the top region. This is consistent with the differences in $\Delta S_{\text{I} \rightarrow \text{II}}$ computed from the FESs (Figure 3 in the main manuscript), which is about $10 \text{ J} \cdot \text{K}^{-1} \cdot \text{mol}^{-1}$ lower in cage C than in cage B. To better quantify this difference we used the simulations to estimate the translational entropy of the adsorbed species (after fixing the center of mass of the zeolite atoms) using the formula:

$$\begin{aligned}
 S_{\text{trans}} &= S_x + S_y + S_z \\
 &= -k_{\text{B}} \int_{-\infty}^{+\infty} p_x(x) \ln(p_x(x)) dx \\
 &\quad - k_{\text{B}} \int_{-\infty}^{+\infty} p_y(y) \ln(p_y(y)) dy \\
 &\quad - k_{\text{B}} \int_{-\infty}^{+\infty} p_z(z) \ln(p_z(z)) dz
 \end{aligned} \tag{12}$$

Where $p_i(i)$ is the probability of observing the molecule's center of mass at a certain i coordinate during the simulation. The results are reported in Table S10 and are in line with the decrease in entropy between cage B and C obtained from the free energy profiles, where the $\Delta S_{\text{trans, I} \rightarrow \text{II}}$ is fundamentally 0 for cage B and significantly negative for cage C, independently on the level of theory.

Table S10: Translational entropy ($\text{J} \cdot \text{K}^{-1} \cdot \text{mol}^{-1}$) obtained from regular MD simulations of isobutene (I, unbiased) and *tert*-butylcation (II, with walls) at 623 K. For isobutene, the states for which $\text{CV}_1 > 6.5$ were ignored.

	Cage B		Cage C	
	PBE-D3(BJ)	RI-RPA	PBE-D3(BJ)	RI-RPA
$S_{\text{trans, I}}$	141.7	140.6	138.1	137.7
$S_{\text{trans, II}}$	141.3	141.7	131.6	130.5
$\Delta S_{\text{trans, I} \rightarrow \text{II}}$	-0.4	1.0	-6.5	-7.2

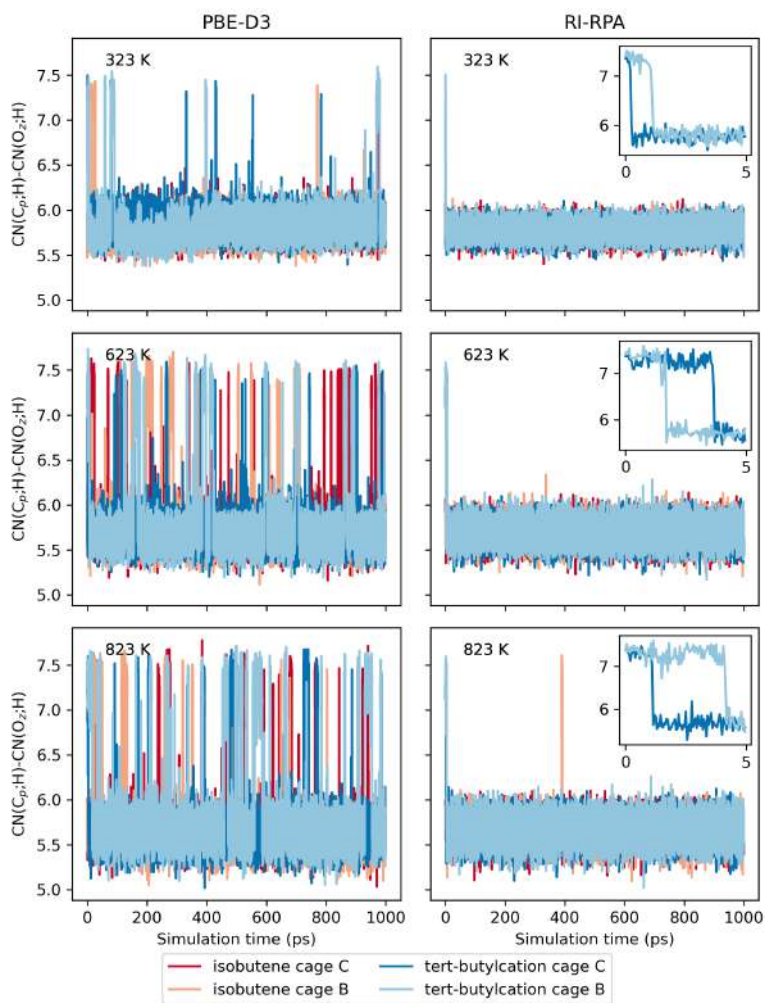


Figure S22: Time evolution of $CN(C_p;H) - CN(O_z;H)$, which describes isobutene protonation, in 1 ns long unbiased MD simulations starting from either isobutene or the *tert*-butylcation in either cage B or C. The inset in the right graphs highlights the deprotonation of the *tert*-butylcation in the first few ps of simulation.

S-37

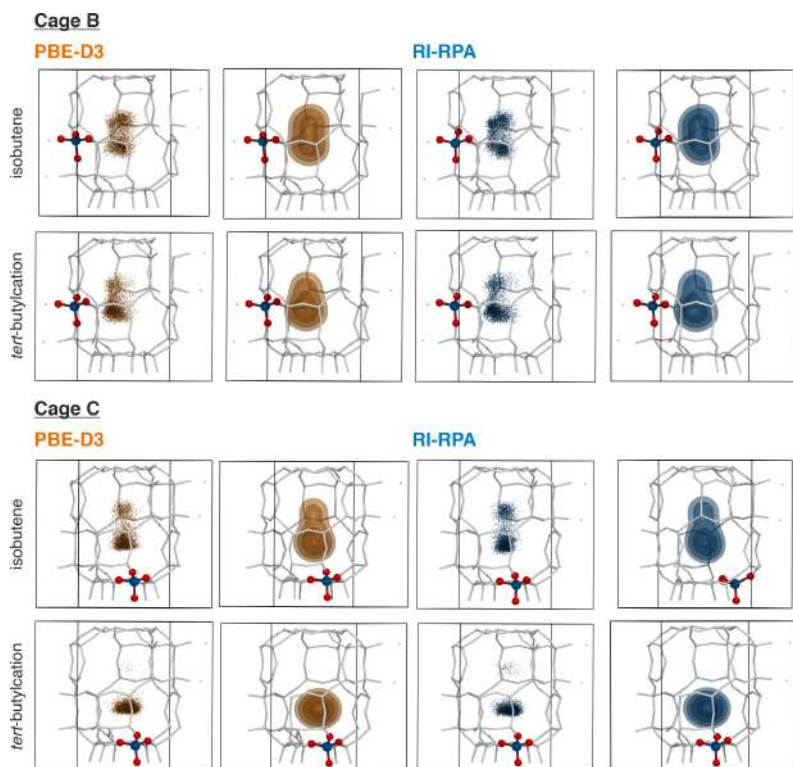


Figure S23: Comparison between the movement of the tertiary isobutene carbon atom during the MD simulations of isobutene and the *tert*-butylation at 623 K. The carbon atom position is either shown as semitransparent dots subsampled every 250 fs or as density maps obtained with the `volmap` extension implemented in VMD⁴¹ (resolution of 0.2 Å). The isolines are arbitrarily placed at (from the exterior) 0.005, 0.01, 0.02 and 0.04. Zeolite atoms are represented with white sticks, only the Al tetrahedron is highlighted with colored spheres (Al in blue, O in red).

References

- (1) Sinclair, P.; Vries, A.; van Santen, R.; others Quantum-chemical studies of alkene chemisorption in chabazite: A comparison of cluster and embedded-cluster models. *Journal of the Chemical Society, Faraday Transactions* **1998**, *94*, 3401–3408.
- (2) Boronat, M.; Zicovich-Wilson, C. M.; Viruela, P.; Corma, A. Influence of the local geometry of zeolite active sites and olefin size on the stability of alkoxide intermediates. *The Journal of Physical Chemistry B* **2001**, *105*, 11169–11177.
- (3) Clark, L. A.; Sierka, M.; Sauer, J. Relative stability of alkoxides and carbocations in zeolites. QM/MM embedding and QM calculations applying periodic boundary conditions. *Studies in Surface Science and Catalysis* **2002**, *142*, 643–649.
- (4) Rozanska, X.; van Santen, R. A.; Demuth, T.; Hutschka, F.; Hafner, J. A periodic DFT study of isobutene chemisorption in proton-exchanged zeolites: dependence of reactivity on the zeolite framework structure. *The Journal of Physical Chemistry B* **2003**, *107*, 1309–1315.
- (5) Boronat, M.; Viruela, P.; Corma, A. Reaction intermediates in acid catalysis by zeolites: prediction of the relative tendency to form alkoxides or carbocations as a function of hydrocarbon nature and active site structure. *Journal of the American Chemical Society* **2004**, *126*, 3300–3309.
- (6) Tuma, C.; Sauer, J. Protonated isobutene in zeolites: tert-butyl cation or alkoxide? *Angewandte Chemie International Edition* **2005**, *44*, 4769–4771.
- (7) Nieminen, V.; Sierka, M.; Murzin, D. Y.; Sauer, J. Stabilities of C₃–C₅ alkoxide species inside H-FER zeolite: a hybrid QM/MM study. *Journal of Catalysis* **2005**, *231*, 393–404.
- (8) Tuma, C.; Sauer, J. Treating dispersion effects in extended systems by hybrid MP2: DFT calculations—protonation of isobutene in zeolite ferrierite. *Physical Chemistry Chemical Physics* **2006**, *8*, 3955–3965.
- (9) De Moor, B. A.; Reyniers, M.-F.; Sierka, M.; Sauer, J.; Marin, G. B. Physisorption and chemisorption of hydrocarbons in H-FAU using QM-Pot (MP2//B3LYP) calculations. *The Journal of Physical Chemistry C* **2008**, *112*, 11796–11812.

- (10) Tuma, C.; Kerber, T.; Sauer, J. The tert-Butyl Cation in H-Zeolites: Deprotonation to Isobutene and Conversion into Surface Alkoxides. *Angewandte Chemie International Edition* **2010**, *49*, 4678–4680.
- (11) Rosenbach Jr, N.; dos Santos, A. P.; Franco, M.; Mota, C. J. The tert-butyl cation on zeolite Y: A theoretical and experimental study. *Chemical Physics Letters* **2010**, *485*, 124–128.
- (12) Nguyen, C. M.; De Moor, B. A.; Reyniers, M.-F.; Marin, G. B. Isobutene Protonation in H-FAU, H-MOR, H-ZSM-5, and H-ZSM-22. *The Journal of Physical Chemistry C* **2012**, *116*, 18236–18249.
- (13) Dai, W.; Wang, C.; Yi, X.; Zheng, A.; Li, L.; Wu, G.; Guan, N.; Xie, Z.; Dyballa, M.; Hunger, M. Identification of tert-Butyl Cations in Zeolite H-ZSM-5: Evidence from NMR Spectroscopy and DFT Calculations. *Angewandte Chemie International Edition* **2015**, *54*, 8783–8786.
- (14) Ferguson, G. A.; Cheng, L.; Bu, L.; Kim, S.; Robichaud, D. J.; Nimlos, M. R.; Curtiss, L. A.; Beckham, G. T. Carbocation stability in H-ZSM5 at high temperature. *The Journal of Physical Chemistry A* **2015**, *119*, 11397–11405.
- (15) Sarazen, M. L.; Iglesia, E. Stability of bound species during alkene reactions on solid acids. *Proceedings of the National Academy of Sciences* **2017**, *114*, E3900–E3908.
- (16) Cnudde, P.; De Wispelaere, K.; Van der Mynsbrugge, J.; Waroquier, M.; Van Speybroeck, V. Effect of temperature and branching on the nature and stability of alkene cracking intermediates in H-ZSM-5. *Journal of Catalysis* **2017**, *345*, 53–69.
- (17) Plessow, P. N.; Studt, F. How accurately do approximate density functionals predict trends in acidic zeolite catalysis? *The Journal of Physical Chemistry Letters* **2020**, *11*, 4305–4310.
- (18) Ren, Q.; Rybicki, M.; Sauer, J. Interaction of C3–C5 alkenes with zeolitic Brønsted sites: π -complexes, alkoxides, and carbenium ions in H-FER. *The Journal of Physical Chemistry C* **2020**, *124*, 10067–10078.
- (19) De Wispelaere, K.; Plessow, P. N.; Studt, F. Toward computing accurate free energies in heterogeneous catalysis: a case study for adsorbed isobutene in H-ZSM-5. *ACS Physical Chemistry Au* **2022**, *2*, 399–406.

- (20) Bocus, M.; Goeminne, R.; Lataire, A.; Cools-Ceuppens, M.; Verstraelen, T.; Van Speybroeck, V. Nuclear quantum effects on zeolite proton hopping kinetics explored with machine learning potentials and path integral molecular dynamics. *Nature Communications* **2023**, *14*, 1008.
- (21) VandeVondele, J.; Krack, M.; Mohamed, F.; Parrinello, M.; Chassaing, T.; Hutter, J. Quickstep: Fast and accurate density functional calculations using a mixed Gaussian and plane waves approach. *Computer Physics Communications* **2005**, *167*, 103–128.
- (22) Kühne, T. D.; Iannuzzi, M.; Del Ben, M.; Rybkin, V. V.; Seewald, P.; Stein, F.; Laino, T.; Khalullin, R. Z.; Schütt, O.; Schiffmann, F.; others CP2K: An electronic structure and molecular dynamics software package-Quickstep: Efficient and accurate electronic structure calculations. *The Journal of Chemical Physics* **2020**, *152*.
- (23) Goedecker, S.; Teter, M.; Hutter, J. Separable dual-space Gaussian pseudopotentials. *Physical Review B* **1996**, *54*, 1703.
- (24) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized gradient approximation made simple. *Physical Review Letters* **1996**, *77*, 3865.
- (25) Grimme, S.; Antony, J.; Ehrlich, S.; Krieg, H. A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu. *The Journal of Chemical Physics* **2010**, *132*.
- (26) Grimme, S.; Ehrlich, S.; Goerigk, L. Effect of the damping function in dispersion corrected density functional theory. *Journal of Computational Chemistry* **2011**, *32*, 1456–1465.
- (27) Del Ben, M.; Hutter, J.; VandeVondele, J. Electron correlation in the condensed phase from a resolution of identity approach based on the Gaussian and plane waves scheme. *Journal of chemical theory and computation* **2013**, *9*, 2654–2671.
- (28) Del Ben, M.; Schütt, O.; Wentz, T.; Messmer, P.; Hutter, J.; VandeVondele, J. Enabling simulation at the fifth rung of DFT: Large scale RPA calculations with excellent time to solution. *Computer Physics Communications* **2015**, *187*, 120–129.
- (29) Bocus, M.; Vanduyfhuys, L.; De Proft, F.; Weckhuysen, B. M.; Van Speybroeck, V. Mech-

- anistic characterization of zeolite-catalyzed aromatic electrophilic substitution at realistic operating conditions. *JACS Au* **2022**, *2*, 502–514.
- (30) Torrie, G. M.; Valleau, J. P. Monte Carlo free energy estimates using non-Boltzmann sampling: Application to the sub-critical Lennard-Jones fluid. *Chemical Physics Letters* **1974**, *28*, 578–581.
- (31) Torrie, G.; Valleau, J. Monte Carlo study of a phase-separating liquid mixture by umbrella sampling. *The Journal of chemical physics* **1977**, *66*, 1402–1408.
- (32) Bailleul, S.; Dedecker, K.; Cnudde, P.; Vanduyfhuys, L.; Waroquier, M.; Van Speybroeck, V. Ab initio enhanced sampling kinetic study on MTO ethene methylation reaction. *Journal of catalysis* **2020**, *388*, 38–51.
- (33) Vanduyfhuys, L. ThermoLIB. <https://molmod.ugent.be/software/thermolib>, 2023.
- (34) Vandenhaute, S.; Cools-Ceuppens, M.; DeKeyser, S.; Verstraelen, T.; Van Speybroeck, V. Machine learning potentials for metal-organic frameworks using an incremental learning approach. *npj Computational Materials* **2023**, *9*, 1–8.
- (35) Vandenhaute, S. Psiflow - interatomic potentials using online learning. <https://github.com/molmod/psiflow>, 2023.
- (36) Allen, M. P.; Tildesley, D. J. *Computer simulation of liquids*; Oxford university press, 2017.
- (37) Kovács, D. P.; Batatia, I.; Arany, E. S.; Csányi, G. Evaluation of the MACE force field architecture: From medicinal chemistry to materials science. *The Journal of Chemical Physics* **2023**, *159*, 044118, DOI: 10.1063/5.0155322.
- (38) Del Ben, M.; Hutter, J.; VandeVondele, J. Second-order Møller–Plesset perturbation theory in the condensed phase: An efficient and massively parallel Gaussian and plane waves approach. *Journal of Chemical Theory and Computation* **2012**, *8*, 4177–4188.
- (39) Guidon, M.; Hutter, J.; VandeVondele, J. Auxiliary density matrix methods for Hartree-Fock exchange calculations. *Journal of Chemical Theory and Computation* **2010**, *6*, 2348–2364.
- (40) Herzog, B.; Chagas da Silva, M.; Casier, B.; Badawi, M.; Pascale, F.; Bucko, T.; Lebegue, S.; Rocca, D. Assessing the accuracy of machine learning thermodynamic perturbation theory:

density functional theory and beyond. *Journal of Chemical Theory and Computation* **2022**, *18*, 1382–1394.

- (41) Humphrey, W.; Dalke, A.; Schulten, K. VMD – Visual Molecular Dynamics. *Journal of Molecular Graphics* **1996**, *14*, 33–38.

Paper IV

Rare Event Sampling using Smooth Basin Classification

Sander Vandenhaute, Tom Braeckvelt, Pieter Dobbelaere, Massimo Bocus,
Veronique Van Speybroeck

arXiv:2404.03777, 2024

Sander Vandenhaute developed and implemented the algorithm and performed all simulations.

Reprinted with permission.
Copyright (2024), The Authors

Rare Event Sampling using Smooth Basin Classification

Sander Vandenhoute, Tom Braeckevelt, Pieter Dobbelaere, Massimo Bocus, and Veronique Van Speybroeck*

Center for Molecular Modeling, Ghent University (BE)

(Dated: May 27, 2024)

The efficiency of atomic simulations of materials and molecules can rapidly deteriorate when large free energy barriers exist between local minima. We propose smooth basin classification, a universal method to define reaction coordinates based on the internal feature representation of a graph neural network. We achieve high data efficiency by exploiting their built-in symmetry and adopting a transfer learning strategy. We benchmark our approach on challenging chemical and physical transformations, and show that it matches and even outperforms reaction coordinates defined based on human intuition.

Dynamic simulations of materials and molecules can provide both qualitative and quantitative information on chemical and physical transformations. While explicit integration of Newton’s equation of motion requires timesteps in the order of a femtosecond, the time required for many transformations to complete is often beyond the micro- or millisecond range because of large energy barriers between initial and final state. Such transformations are therefore almost always rare events, and their simulation can easily require billions of time steps [1, 2].

Over the past decades, a variety of advanced sampling techniques have been proposed in order to improve the sampling efficiency of dynamic simulations. A particularly intuitive approach is to bias the dynamics along certain directions in phase space, effectively lifting the system over obstructive free energy barriers. The bias energy is defined along a certain predefined reaction coordinate or collective variable (CV). The CV is a function of the atomic coordinates and should naturally discriminate between initial and final state without violating the physical symmetries in the system [3]. Specifically, the CV should remain invariant with respect to global translations and rotations, as well as permutations of atoms which are chemically identical. In the vast majority of practical applications, such CVs are manually defined based on chemical and physical intuition, and they usually contain simple invariants such as specific unit cell parameters or coordination numbers of one chemical species with respect to another. However, manual definition can be highly challenging and time-consuming, because the metastable intermediate(s) and/or transition state(s) are not always known in advance, or because the required invariances are difficult to satisfy [4–6].

Considering the numerous achievements of machine learning in atomistic modeling, it is sensible to anticipate that suitable reaction coordinates can also be learned from (unbiased) simulation data, and ideally without requiring human intuition about the transition mechanism(s). A wide variety of such methods have been proposed over the past decade, and we briefly mention a few notable examples while referring the reader to specialized reviews for a more comprehensive discussion [7–9]. Supervised methods learn a continuous discriminant func-

tion which separates free energy minima based on user-defined structural descriptors, for example using (deep) linear discriminant analysis [10–13] or (variational) autoencoders [14–16]. Unsupervised methods do not directly learn a discriminant function but instead identify so-called slow modes, either within a given set of descriptors or entirely end-to-end; see e.g. time-lagged independent component analysis or time-lagged autoencoders [17–19]. While these methods demonstrate the potential of using machine learning for defining reaction coordinates, their widespread adoption has remained rather limited. We hypothesize that this is due to inherent limitations regarding the required volume of simulation data, the difficulty of constructing an appropriate set of input features, or the difficulty in interpreting and validating the obtained CVs.

In this work, we propose smooth basin classification (SBC); a universal method to construct CVs and compute free energy differences for chemical and physical transformations. SBC builds upon the successful development of graph neural networks (GNNs) as effective interatomic potentials by using their learned feature space as ansatz for constructing physically meaningful CVs [20–23]. Our key finding is that GNN-based interaction potentials produce a feature representation for atomic environments which can be used to drive transitions between free energy minima and characterize their relative stability.

Consider a system with two basins of attraction, named A and B. Each point \mathbf{r} in the system’s coordinate space can be uniquely labeled with A or B depending on the trajectory of a geometry optimization; \mathbf{r} belongs to basin A whenever it optimizes to any of the local minima within A, and vice-versa for B (Figure 1, a). In SBC, we aim to construct a mapping between each configuration \mathbf{r} and its associated basin of attraction. Instead of defining this mapping directly in terms of the atomic coordinates – which are not invariant with respect to the symmetry operations in the system – we embed it into the learned feature space of a GNN which has been pre-trained to reproduce the potential energy surface of the system (Figure 1, b). In general, the interaction layers of the GNN parse the chemical environment of atom i into a

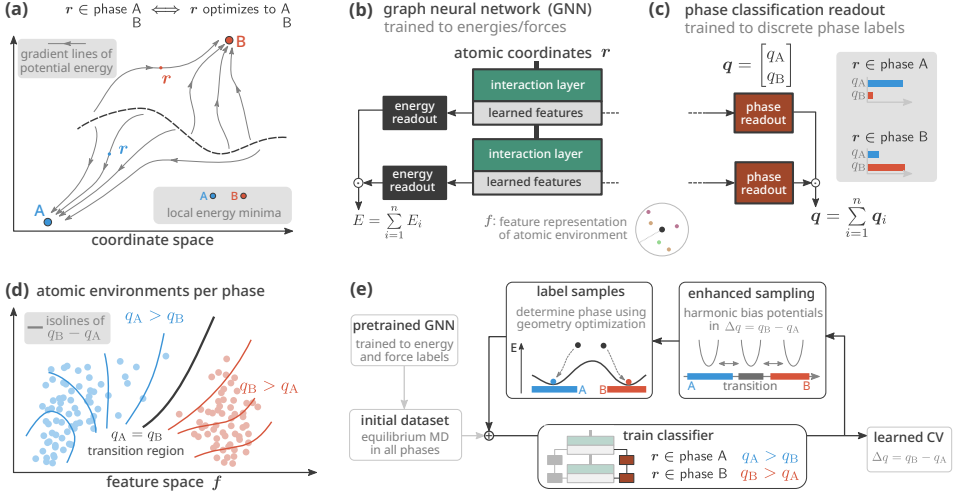


FIG. 1. **(a)** Illustration of the partitioning of the phase space into basins of attraction. Each point x can be uniquely associated to its basin, i.e. its local potential energy minimum. **(b)** Base GNN which is pretrained to reference energies and forces. The interaction layers effectively generate a learned feature representation f for the atomic environments. **(c)** a phase classifier which is trained separately on atomic geometries which are labeled with the associated basin as in **(a)**. **(d)** Illustration of the feature space embedding of atomic environments based on their (local) phase, together with the isolines of Δq . **(e)** Active learning loop for SBC; the log probabilities q are used to bias the dynamics from one basin to another, in order to sample geometries along the transition path(s). The gathered data is labeled with the correct basin of attraction **(a)** after which the classifier is retrained. See the Appendix for more information.

feature vector f_i in this space. Importantly, this feature vector is rigorously invariant with respect to all symmetry operations in the system, and is specifically trained to describe the atomic interactions. Both the per-atom energy E_i and the total energy E are computed from the feature vectors f_i using a simple readout function:

$$r \xrightarrow{\text{GNN}} \{f_i\}_{i=1}^n \xrightarrow{\text{energy readout}} E = \sum_{i=1}^n E_i \quad (1)$$

in which n denotes the total number of atoms in configuration r . Because the features f_i are maximally informative for determining the interaction energy of a given atomic environment, we may assume that they are also highly informative in predicting the basin of attraction. In full analogy with equation 1, we can augment the GNN with an additional classification readout function (Figure 1, c):

$$r \xrightarrow{\text{GNN}} \{f_i\}_{i=1}^n \xrightarrow{\text{phase readout}} q = \sum_{i=1}^n q_i \quad (2)$$

in which we introduce the per-atom log probabilities q_i as an invariant tuple which measures how strongly the model associates a given atomic environment f_i to basin

A or B – as commonly done in a classification setting [24]:

$$q_i = (q_i^A, q_i^B) \quad (3)$$

For example, if the model believes that a given atomic environment i is more likely to appear in phase A as opposed to phase B, it will predict log probabilities q_i such that $q_i^A > q_i^B$. After summing the per-atom log probabilities q_i to obtain the total log probabilities $q = (q^A, q^B)$, the predicted basin of attraction of the total configuration is then determined by the largest component in q . A detailed overview of the phase classification readout can be found in the Appendix.

While the prediction of the basin of attraction is essentially a discrete classification task, the actual log probabilities q_i are smooth functions of the atomic feature vectors f_i and hence the atomic coordinates – which means that they can be used to introduce biasing forces (and biasing stress). Ideally, we would like to design this bias such that it can drive transitions from and to any of the basins. To this end, consider the difference between two log probabilities:

$$\Delta q = q^B - q^A \quad (4)$$

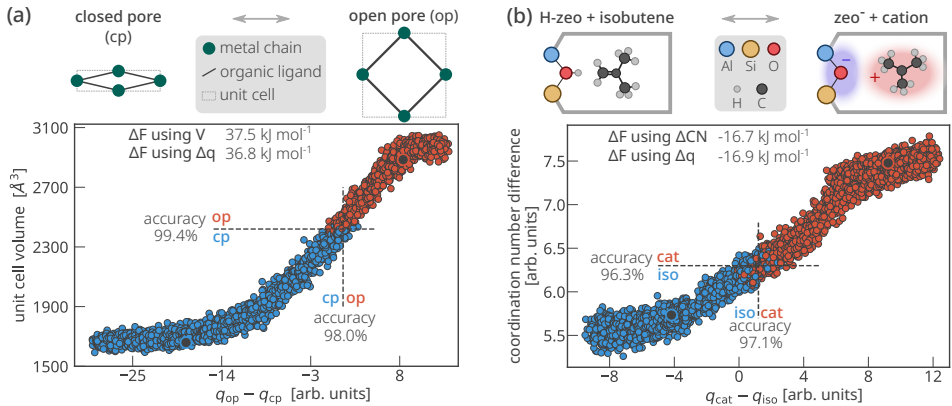


FIG. 2. (a) The nonreactive phase transition in MIL-53(Al). The scatter plot shows the distribution of both the unit cell volume V as well as the log probability difference $\Delta q = q_{op} - q_{cp}$. Each dot represents a single atomic geometry, and its color represents whether it minimizes to the *op* or *cp* phase. Black dots indicate potential energy minima. (b) A proton transfer reaction in a zeolite. The scatter plot shows the distribution of both the coordination number difference as well as the log probability difference $\Delta q = q_{cat} - q_{iso}$, whereby the color again indicates the corresponding basin of attraction. The reported relative free energy differences are computed based on the free energy profiles in the Appendix.

In essence, Δq measures how much more likely the model thinks a given configuration belongs to basin B rather than basin A. It is therefore strictly positive in basin B and strictly negative in basin A provided that the phase readout is properly trained (Figure 1, d). Clearly, the quantity Δq correctly discriminates between two basins of attraction while also being continuously differentiable with respect to the atomic coordinates (using standard backpropagation). It is therefore, by definition, a valid collective variable for the transition between basins A and B.

In what follows, we discuss the application of SBC for prototypical physical and chemical transformations. We use MACE as base GNN [22, 25], and employ active learning as in Figure 1 (e) to train the classifier readout – see the appendix for additional computational details.

Figure 2 (a) shows a nonreactive phase transition in a solid-state system named MIL-53(Al), a flexible metal-organic framework based on a winerack topology [26]. The system exhibits two basins of attraction which are relevant for tuning the material towards gas storage and/or sensing applications; a closed pore (*cp*) and open pore (*op*) phase [27]. In these types of transitions, i.e. those for which the unit cell parameters vary significantly between the different basins, researchers often employ some combination of unit cell parameters as collective variable – in this case, the unit cell volume [28]. Figure 2 (a) demonstrates the performance of the unit cell volume as formal discriminator between the two basins, and compares it with a difference of learned log proba-

bilities as proposed in equation 4. The misclassified geometries are all localized near the transition region (such they will not cause substantial differences in the relative stability). Note that the network learns to distinguish low- and high-volume configurations based on atomic environments only, i.e. without explicitly relying on unit cell parameters as input.

For global phase transformations in solids such as the one in Figure 2 (a), a size-extensive definition of \mathbf{q} as a sum over per-atom log probabilities q_i is intuitively sensible because it gives the model the ability to assign nonzero log probability contributions to all atoms simultaneously. However, in chemical transformations, the rare event is most often a local process whose free energy barrier does not necessarily depend on the size of the surrounding environment (e.g. the catalyst). To demonstrate the applicability of our approach to those cases, we investigate a proton transfer reaction from a Brønsted acid site in a zeolite to an alkene guest (Figure 2, b). The reference collective variable in this case is a coordination number difference which is defined explicitly in terms of the indices of the oxygen atoms of the active site as well as the indices of the carbon atoms of the molecule and all of the hydrogens – see the Supplemental Material (Section S1.2) for more details. Classification accuracies are excellent for both the proposed CV based on chemical intuition as well as the learned log probability difference. Misclassified geometries are localized at or near the transition region, which means that they will not induce discrepancies in the relative stability. Indeed,

	training cost		input descriptors	symmetries			initialized without prior transition path(s)
	snapshots	time [ns]		trans	rot	perm	
time-lagged AE [18]		750 ns	learned	●	●	●	
FEBLAE [29]		1500 ns	learned	●	●	●	
VAMPnets [19]	250,000	250 ns	learned	●	●	●	✓
LINES [30]		100 ns	manual	●	●	●	
LED [31]		100 ns	manual	●	●	●	
Ref. [32]	>300,000		learned	●	●	●	✓
Deep-LDA [11]	10,000		manual	●	●	●	
Deep-TDA [33]		40 ns	learned	●	●	●	✓
Ref. [34]	15,000	15 ns	manual	●	●	●	✓
Ref. [35]	8,000	512 ns	manual	●	●	●	
SBC	3,694	11 ns	learned	●	●	●	✓

TABLE I. Performance of various CV learning methods on alanine dipeptide. Key characteristics are the required amount of input data (lower is better), the built-in symmetry as inductive bias, and whether or not an initial transition path is required for the model to learn the transition. Missing values indicate that the corresponding quantity was not directly reported in either the main manuscript or any of its appendices. Yellow dots indicate that the required invariance was achieved by means of a manual structural alignment procedure – see the corresponding references for more information.

for both systems, free energy calculations based on either the reference or the learned CV yield the same relative stability (to within 1 kJ/mol).

Note that the quantity Δq satisfies the same symmetries as the potential energy E of the system (Figure 1, b/c). As such, the learned collective variable is invariant with respect to global translations and rotations as well as permutations of chemically identical atoms. Achieving permutational symmetry in particular is quite remarkable, and should be contrasted with both common practice CVs and most other ML-based approaches [11, 18, 19, 33] which often use functions and/or input features that depend explicitly on the atomic indices (thereby violating permutation invariance) or specific components of the unit cell vectors (thereby violating rotational as well as periodic invariance). We hypothesize that the Euclidean and permutation symmetry of Δq is the primary reason for the exceptional data efficiency of the method, alongside the effective active learning scheme (Figure 1, e) and the implicit use of equivariant features within the GNN. We demonstrate this explicitly based on a case study on alanine dipeptide, which represents a topical nonreactive conformational change that has served as enhanced sampling benchmark system in many studies [11, 18, 19, 29–32, 34–36]. Table I enumerates and compares the most prominent CV learning methods with SBC for alanine dipeptide. We evaluate a number of key characteristics, including the required amount of input data to train a given CV, its transferability towards symmetrically equivalent geometries, whether or not it requires manual feature selection, and whether or not an initial transition path is required. In essence, Table I shows that SBC outperforms every single method in all of these aspects. It requires the fewest structures and smallest total simulation time; it is

the only method that is rigorously invariant with respect to all symmetry operations; and it does not require any kind of manual input feature selection or initial transition path. Figure 3 demonstrates how SBC is capable of discovering several transition paths during its active learning loop. In Section S1.3 of the Supplemental Material, we further show that the relative free energy difference as computed using (ϕ, ψ) and Δq agrees within 1 kJ/mol.

We conclude with a discussion on the current limitations of SBC as well as directions for future research. Clearly, the success of our method is critically dependent on the availability of ML potentials which can predict the atomic interactions in the system with high fidelity. In this work, we considered both self-developed and foundation models as base GNN in which the SBC model is inserted – the alanine dipeptide and MIL-53(Al) case studies were performed with MACE models trained using psiflow [22, 37], while a MACE foundation model was used for the proton transfer reaction[38]. As such, we can conclude that SBC is fundamentally compatible with both. However, more challenging scenarios might require multi-task training of the entire GNN (including its interaction layers), in which each sample in the training set is labeled with QM reference data and additionally a phase or basin label – this will incentivize the model to design the per-atom feature embedding to be informative for both potential energy prediction as well as basin classification. In addition to relative free energies, the learned CVs can be employed to predict the kinetics of rare events based on e.g. reactive flux and/or path sampling methods [39]. Finally, we mention that the proposed method can be extended towards the more general case of n free energy minima simply by increasing the number of classes in the classification layer.




FIG. 3. Free energy landscape of alanine dipeptide as obtained from umbrella sampling in the conventional backbone dihedral angles ϕ and ψ). The scattered dots indicate the training data distribution as obtained during active learning of the SBC model; the color-coding indicates the Δq values after final training. See the Supplemental Material (Section S1) for additional computational details.

DATA AVAILABILITY STATEMENT

An implementation of SBC in MACE will be made available in a public GitHub repository upon final publication; for peer-review purposes, we also provide the code through the following link: <https://tinyurl.com/mufvj4ur>. All additional input files, all generated datasets, and all of the models which were trained in this manuscript will be made available in a public Zenodo archive upon final publication; for peer-review purposes, these are available upon request as well.

ACKNOWLEDGEMENTS

The authors acknowledge funding from the Research Board of Ghent University (V.V.S) and iBOF-21-085 PERSIST (V.V.S. and T.B.). S.V. and P.D. wish to thank the Research Foundation – Flanders (FWO) for doctoral fellowships (grant nos. 11H6821N and 11O2123N respectively). M.B. acknowledges financial support from the Fund for Scientific Research Flanders and the Excellence of Science (EOS) Project BioFact (EOS ID 30902231). We acknowledge the EuroHPC Joint Undertaking for awarding this project access to the EuroHPC supercomputer LUMI, hosted by CSC (Finland) and the LUMI consortium through a EuroHPC Regular Access call. Part of the simulations were performed on the Luxembourg national supercomputer

MeluXina. The authors gratefully acknowledge the Lux-Provide teams for their expert support.

Appendix on the Model Architecture

In analogy with the prediction of the total potential energy E as a sum of predicted atomic energies E_i , we predict the total log probability \mathbf{q} as a sum of per-atom log probabilities \mathbf{q}_i (see Equation 2 in the main text). In a system with N basins, the log probabilities \mathbf{q}_i are a vector of length N , i.e. there is one per-atom logit for each basin. We use $\mathbf{f}_i^{(j)}$ to denote the GNN features from layer j for atom i . First, we forward the node features into a MLP per layer:

$$\mathbf{h}_i^{(j)} = \text{MLP}^{(j)}(\mathbf{f}_i^{(j)}) \quad (5)$$

Next, we concatenate $\mathbf{h}_i^{(j)}$ for all GNN layers j , and use a final linear layer with weights W to compute the per-atom log probabilities \mathbf{q}_i :

$$\mathbf{q}_i = W \left(\bigoplus_j \mathbf{h}_i^{(j)} \right) \quad (6)$$

Finally, the total log probabilities \mathbf{q} are obtained by summing the individual \mathbf{q}_i :

$$\mathbf{q} = \sum_{i=1}^n \mathbf{q}_i \quad (7)$$

The phase readout layers are analogous to the existing energy readouts found in conventional GNNs [20–22] (Figure 1, c), and essentially consist of a multilayer perceptron (MLP) containing at most two nonlinearities and about ~ 1000 weights in total (compared to $\sim 100,000$ weights in the GNN).

Appendix on Model Training

The predicted log probabilities \mathbf{q} are normalized and trained to atomic geometries which are labeled with the correct basin. The label is obtained by minimizing the energy of the structure, as explained in Figure 1 in the main manuscript. In a classification setting, one often uses the cross-entropy loss function [24]. We can define the phase label \mathbf{p} of an atomic geometry as a one-hot vector which has the same length as \mathbf{q} (i.e. the number of basins), and whose components indicate the correct basin of attraction:

$$p^c = \begin{cases} 1 & \text{if phase} = C \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

Based on Equation 8, we can write the cross-entropy loss as a simple scalar product:

$$\text{CE}(\mathbf{q}) = -\mathbf{p} \cdot \mathbf{q} \quad (9)$$

The classifier is trained while keeping the GNN and energy readouts fixed, similar to conventional transfer learning approaches. Its training data consists of atomic geometries which are labeled with their corresponding basin of attraction (e.g. A or B).

Note that the cross-entropy loss is a function defined on the total log probabilities \mathbf{q} , not on the specific per-atom log probabilities \mathbf{q}_i . This implies that the cross-entropy loss will incentivize the model to predict correct total log probabilities, but otherwise introduces no restrictions regarding the specific per-atom log probabilities \mathbf{q}_i . This makes the model prone to overfitting, especially in the low-data regime. To improve this, we considered a number of regularization strategies including label smoothing, an L2 penalty on the weights, and an L1 penalty on logit gradients with respect to the atomic coordinates. Their efficacy was highly dependent on both the system under study and their specific weights in the loss function, and we chose not to apply them. Instead, we designed a new regularization term which prevents the per-atom log probabilities \mathbf{q}_i from becoming strongly positive for any given atomic environment i .

First, note that the global probability for each phase can be computed based on the predicted log probabilities \mathbf{q} using a softmax function:

$$P[A] = \frac{e^{q^A}}{\sum_c e^{q^c}} \quad (10)$$

for any class A. We can apply a similar reasoning to the per-atom log probabilities \mathbf{q}_i in order to obtain per-atom phase probabilities $P_i[A]$. Next, consider

$$\mathbb{E}[\mathbf{q}_i] = \sum_c P_i[C] q_i^C \quad (11)$$

The quantity $\mathbb{E}[\mathbf{q}_i]$ denotes an expected per-atom logit value. In practice, given the exponential dependence between the probabilities and the log probabilities, there will only be one nonzero contribution in the sum in Equation 11, and the expectation value will essentially reduce to the largest (per-atom) logit value. Note that the predicted per-atom phase (based on the per-atom log probabilities) can differ from the predicted global phase (based on the total log probabilities). Finally, in the loss function, we add an L2 regularization on $\mathbb{E}[\mathbf{q}_i]$, which essentially prevents the model from being overconfident in any of the per-atom log probabilities:

$$\mathcal{L} = CE(\mathbf{q}) + \lambda \sum_{i=1}^n \|\mathbb{E}[\mathbf{q}_i]\|_2 \quad (12)$$

For all experiments in this work, we use $\lambda = 1$.

Furthermore, note that the actual training of the phase readouts can be made very efficient because the GNN node features \mathbf{f}_i can be cached and reused at every step

of the optimization since the weights of the pretrained GNN are fixed. As a result, training is very efficient, and all SBC models in this work can be trained within one hour on a consumer-level GPU.

Appendix on the Active Learning

Figure 1 (e) in the main manuscript outlines the active learning workflow in more detail. We start by training the phase readouts on equilibrium samples from each of the free energy minima, after which we begin with the active learning loop and perform enhanced sampling. We introduce harmonic restraints centered on different values of Δq in order to force the system to transition from one basin to another. In each iteration, we position 32 umbrellas along the entire CV range, at three different temperatures as to ensure that both low- and (relevant) high-energy regions are properly sampled. In each umbrella, the system is randomly initialized in either of the two minima. The samples from the resulting trajectory are labeled according to their basin of attraction, as determined by a simple energy minimization for each sample. After this, the model is retrained on all data, and the procedure is repeated. See the Supplemental Material (Section S1) for a complete overview of all hyperparameters.

Appendix on the Graph Neural Network

All calculations in this work employed MACE as base potential [22], although we should emphasize that SBC is independent of the specific GNN architecture [40]. The interatomic potentials for MIL-53(Al) and alanine dipeptide were trained from scratch using psiflow, at the PBE-D3(BJ) level of theory and with a triple zeta basis set [37]. DFT evaluations were performed with CP2K 2023.1 for MIL-53(Al) in which the basis set was augmented with plane waves up until a cutoff of 1000 Ry, and ORCA 5.0.4 for alanine dipeptide [41, 42]. For the proton transfer reaction in the zeolite catalyst, we employed a foundational MACE model [38]. For the MACE models, we employ node features with $\ell_{\max} = 2$ and 16 channels each, with a cutoff radius of 6.5 Å. All models will be made available upon final publication.

Appendix on the Free Energy Calculations

Using the log probability difference $\Delta q = q_{\text{op}} - q_{\text{cp}}$ as CV (see equation 4), we can employ conventional enhanced sampling techniques and evaluate the free energy profile $F(\Delta q)$. Here, we choose to employ umbrella sam-




FIG. 4. Helmholtz free energy profiles for (a) the phase transition in MIL-53(Al) and (b) the proton transfer reaction, as obtained from umbrella sampling in either the learned collective variable Δq or the reference collective variable.

pling with harmonic restraints on the CV:

$$U_{\text{bias}}(\mathbf{r}, \mathbf{h}) = \frac{K}{2} [\Delta q(\mathbf{r}, \mathbf{h}) - q_0]^2 \quad (13)$$

in which we emphasize that the bias energy contains an implicit dependence on both the atomic coordinates \mathbf{r} and the unit cell parameters \mathbf{h} through the learned atomic environment features \mathbf{f}_i [23]. We use the multistate Bennett acceptance ratio (MBAR) equations to combine samples from multiple umbrellas in a statistically optimal manner and obtain free energy profiles as shown in Figure 4 [43–46]. Note that while relative stabilities are intrinsic to the system, the actual free energy profiles depend on the particular choice of CV, so the learned and reference free energy profiles do not need to have the same shape (although there exist ways to transform one profile into another [47]).

* veronique.vanspeybroeck@ugent.be

- [1] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Elsevier Science, 2023).
- [2] M. Tuckerman, *Statistical Mechanics: Theory and Molecular Simulation*, Oxford Graduate Texts (OUP Oxford, 2010).
- [3] B. Peters, Chapter 11 - landau free energies and restricted averages, in *Reaction Rate Theory and Rare Events Simulations*, edited by B. Peters (Elsevier, Amsterdam, 2017) pp. 273–310.
- [4] E. Grifoni, G.-M. Piccini, and M. Parrinello, Microscopic description of acid-base equilibrium, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 4054 (2019).
- [5] N. Ahalawat and J. Mondal, Assessment and optimization of collective variables for protein conformational landscape: GB1 β -hairpin as a case study, *J. Chem. Phys.* **149**, 094101 (2018).
- [6] H. Niu, P. M. Piaggi, M. Invernizzi, and M. Parrinello, Molecular dynamics simulations of liquid silica crystallization, *Proc. Natl. Acad. Sci. U.S.A.* **115**, 5348 (2018).
- [7] M. Chen, Collective variable-based enhanced sampling and machine learning, *Eur. Phys. J. B* **94**, 211 (2021).
- [8] S. Bhakat, Collective variable discovery in the age of machine learning: reality, hype and everything in between, *RSC Adv.* **12**, 25010 (2022).
- [9] S. Mehdi, Z. Smith, L. Herron, Z. Zou, and P. Tiwary, Enhanced sampling with machine learning, *Annual Review of Physical Chemistry* <https://doi.org/10.1146/annurev-physchem-083122-125941> (2024).
- [10] O. Elishav, R. Podgaetsky, O. Meikler, and B. Hirshberg, Collective variables for conformational polymorphism in molecular crystals, *J. Phys. Chem. Lett.* **14**, 971 (2023).
- [11] L. Bonati, V. Rizzi, and M. Parrinello, Data-driven collective variables for enhanced sampling, *J. Phys. Chem. Lett.* **11**, 2998 (2020).
- [12] D. Mendels, G. Piccini, and M. Parrinello, Collective variables from local fluctuations, *J. Phys. Chem. Lett.* **9**, 2776 (2018).
- [13] D. Ray, E. Trizio, and M. Parrinello, Deep learning collective variables from transition path ensemble, *J. Chem. Phys.* **158**, 204102 (2023).
- [14] M. Štíplka, A. Erlebach, and L. Grajciar, Constructing collective variables using invariant learned representations, *J. Chem. Theory Comput.* **19**, 887 (2023).
- [15] Z. Zou and P. Tiwary, Enhanced sampling of crystal nucleation with graph representation learnt variables, *The Journal of Physical Chemistry B* **128**, 3037 (2024).
- [16] D. Wang and P. Tiwary, State predictive information bottleneck, *J. Chem. Phys.* **154**, 134111 (2021).
- [17] G. Pérez-Hernández, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, Identification of slow molecular order pa-

- rameters for markov model construction, *J. Chem. Phys.* **139**, 015102 (2013).
- [18] C. Wehmeyer and F. Noé, Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics, *J. Chem. Phys.* **148**, 241703 (2018).
- [19] A. Marđt, L. Pasquali, H. Wu, and F. Noé, Vampnets for deep learning of molecular kinetics, *Nat. Commun.* **9**, 5 (2018).
- [20] N. Thomas, T. Smidt, S. Kearnes, L. Yang, L. Li, K. Kohlhoff, and P. Riley, Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds (2018), arXiv:1802.08219 [cs.LG].
- [21] S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt, and B. Kozinsky, E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials, *Nat. Commun.* **13**, 2453 (2022).
- [22] D. P. Kovács, I. Batatia, E. S. Arany, and G. Csányi, Evaluation of the MACE force field architecture: From medicinal chemistry to materials science, *J. Chem. Phys.* **159**, 044118 (2023).
- [23] A. R. Tan, J. C. B. Dietschreit, and R. Gomez-Bombarelli, Enhanced sampling of robust molecular datasets with uncertainty-based collective variables (2024), arXiv:2402.03753 [cs.LG].
- [24] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)* (Springer-Verlag, Berlin, Heidelberg, 2006).
- [25] I. Batatia, D. P. Kovacs, G. Simm, C. Ortner, and G. Csanyi, Mace: Higher order equivariant message passing neural networks for fast and accurate force fields, in *Advances in Neural Information Processing Systems*, Vol. 35, edited by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Curran Associates, Inc., 2022) pp. 11423–11436.
- [26] C. Serre, F. Millange, C. Thouvenot, M. Noguès, G. Marsolier, D. Louër, and G. Férey, Very large breathing effect in the first nanoporous chromium(III)-based solids: MIL-53 or $\text{Cr}^{\text{III}}(\text{OH})\cdot\{\text{O}_2\text{C}-\text{C}_6\text{H}_4-\text{CO}_2\}\cdot\{\text{HO}_2\text{C}-\text{C}_6\text{H}_4-\text{CO}_2\text{H}\}_x\cdot\text{H}_2\text{O}_x$, *Journal of the American Chemical Society* **124**, 13519 (2002).
- [27] L. Vanduyfhuys, S. M. J. Rogge, J. Wieme, S. Vandebrende, G. Maurin, M. Waroquier, and V. Van Speybroeck, Thermodynamic insight into stimuli-responsive behaviour of soft porous crystals, *Nat. Commun.* **9**, 204 (2018).
- [28] R. Demuyne, J. Wieme, S. M. J. Rogge, K. D. Dedecker, L. Vanduyfhuys, M. Waroquier, and V. Van Speybroeck, Protocol for identifying accurate collective variables in enhanced molecular dynamics simulations for the description of structural transformations in flexible metal-organic frameworks, *Journal of Chemical Theory and Computation* **14**, 5511 (2018).
- [29] Z. Belkacemi, P. Gkeka, T. Lelièvre, and G. Stoltz, Chasing collective variables using autoencoders and biased trajectories, *J. Chem. Theory Comput.* **18**, 59 (2022).
- [30] R. E. Odstrcil, P. Dutta, and J. Liu, Lines: Log-probability estimation via invertible neural networks for enhanced sampling, *J. Chem. Theory Comput.* **18**, 6297 (2022).
- [31] P. R. Vlachas, J. Zavadlav, M. Praprotnik, and P. Koumoutsakos, Accelerated simulations of molecular systems through learning of effective dynamics, *J. Chem. Theory Comput.* **18**, 538 (2022).
- [32] L. Sun, J. Vandermause, S. Batzner, Y. Xie, D. Clark, W. Chen, and B. Kozinsky, Multitask machine learning of collective variables for enhanced sampling of rare events, *J. Chem. Theory Comput.* **18**, 2341 (2022).
- [33] E. Trizio and M. Parrinello, From enhanced sampling to reaction profiles, *The Journal of Physical Chemistry Letters* **12**, 8621 (2021).
- [34] L. Bonati, G. Piccini, and M. Parrinello, Deep learning the slow modes for rare events sampling, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2113533118 (2021).
- [35] L. Müllender, A. Rizzi, M. Parrinello, P. Carloni, and D. Mandelli, Effective data-driven collective variables for free energy calculations from metadynamics of paths (2023), arXiv:2311.05571 [physics.comp-ph].
- [36] P. Kang, E. Trizio, and M. Parrinello, Computing the committor with the committor: an anatomy of the transition state ensemble (2024), arXiv:2401.05279 [physics.comp-ph].
- [37] S. Vandenhoute, M. Cools-Ceuppens, S. DeKeyser, T. Verstraelen, and V. Van Speybroeck, Machine learning potentials for metal-organic frameworks using an incremental learning approach, *npj Comput. Mater.* **9**, 19 (2023).
- [38] I. Batatia, P. Benner, Y. Chiang, A. M. Elena, D. P. Kovács, J. Riebesell, X. R. Advincula, M. Asta, W. J. Baldwin, N. Bernstein, A. Bhowmik, S. M. Blau, V. Cărare, J. P. Darby, S. De, F. D. Pia, V. L. Deringer, R. Elijošius, Z. El-Machachi, E. Fako, A. C. Ferrari, A. Genreith-Schriever, J. George, R. E. A. Goodall, C. P. Grey, S. Han, W. Handley, H. H. Heenen, K. Hermansson, C. Holm, J. Jaafar, S. Hofmann, K. S. Jakob, H. Jung, V. Kapil, A. D. Kaplan, N. Karimitari, N. Kroupa, J. Kullgren, M. C. Kuner, D. Kuryla, G. Liepuoniute, J. T. Margraf, I.-B. Magdău, A. Michaelides, J. H. Moore, A. A. Naik, S. P. Niblett, S. W. Norwood, N. O'Neill, C. Ortner, K. A. Persson, K. Reuter, A. S. Rosen, L. L. Schaaf, C. Schran, E. Sivonxay, T. K. Stenczel, V. Svahn, C. Sutton, C. van der Oord, E. Varga-Umbrich, T. Vegge, M. Vondrák, Y. Wang, W. C. Witt, F. Zills, and G. Csányi, A foundation model for atomistic materials chemistry (2023), arXiv:2401.00096 [physics.chem-ph].
- [39] B. Peters, *Reaction Rate Theory and Rare Events Simulations*, edited by B. Peters (Elsevier, Amsterdam, 2017).
- [40] A. Duval, S. V. Mathis, C. K. Joshi, V. Schmidt, S. Miret, F. D. Malliaros, T. Cohen, P. Lio, Y. Bengio, and M. Bronstein, A hitchhiker's guide to geometric gnns for 3d atomic systems (2023), arXiv:2312.07511 [cs.LG].
- [41] T. D. Kühne, M. Iannuzzi, M. Del Ben, V. V. Rybkin, P. Seewald, F. Stein, T. Laino, R. Z. Khalullin, O. Schütt, F. Schiffmann, D. Golze, J. Wilhelm, S. Chulkov, M. H. Bani-Hashemian, V. Weber, U. Borstnik, M. TAILLEFUMIER, A. S. Jakobovits, A. Lazzaro, H. Pabst, T. Müller, R. Schade, M. Guidon, S. Andermatt, N. Holmberg, G. K. Schenter, A. Hehn, A. Bussy, F. Belleflamme, G. Tabacchi, A. Glöb, M. Lass, I. Bethune, C. J. Mundy, C. Plessl, M. Watkins, J. Vandevondele, M. Krack, and J. Hutter, Cp2k: An electronic structure and molecular dynamics software package - quickstep: Efficient and accurate electronic structure calculations, *J. Chem. Phys.* **152**, 194103 (2020).
- [42] F. Neese, F. Wennmoths, U. Becker, and C. Riplinger, The orca quantum chemistry program package, *J. Chem. Phys.* **152**, 224108 (2020).

- [43] M. R. Shirts and J. D. Chodera, Statistically optimal analysis of samples from multiple equilibrium states, *J. Chem. Phys.* **129**, 124105 (2008).
- [44] M. R. Shirts, Reweighting from the mixture distribution as a better way to describe the multistate bennett acceptance ratio (2017), arXiv:1704.00891 [cond-mat.stat-mech].
- [45] M. R. Shirts and A. L. Ferguson, Statistically optimal continuous free energy surfaces from biased simulations and multistate reweighting, *J. Chem. Theory Comput.* **16**, 4107 (2020).
- [46] M. Shirts, K. Beauchamp, L. Naden, J. Chodera, jaimergp, S. Martiniani, C. Stern, M. Henry, J. Fass, R. Gowers, R. T. McGibbon, B. Dice, C. Jones, D. L. Dotson, F. York, I. Pulido, M. Thompson, and T. Burgin, *choderalab/pymbar: 4.0.2 optional jax build* (2023).
- [47] S. Bailleul, K. Dedecker, P. Cnudde, L. Vanduyfluys, M. Waroquier, and V. Van Speybroeck, Ab initio enhanced sampling kinetic study on mto ethene methylation reaction, *Journal of Catalysis* **388**, 38 (2020).

Free Energy Calculations using Smooth Basin Classification: Supplemental Material

Sander Vandenhoute¹, Tom Braeckevelt¹, Pieter Dobbelaere¹, Massimo Bocus¹, and Veronique Van Speybroeck¹

¹Center for Molecular Modeling, Ghent University (BE)

May 24, 2024

S1 Additional Computational Details

All molecular dynamics simulations were performed using Langevin temperature and/or pressure control. Bias potentials on traditional CVs were added via PLUMED [1]; bias potentials on Δq were added manually using a modified MACE-ASE interface. Table S1 provides an overview with the most important hyperparameters for each case study. Note that there is little system-specific tuning necessary – all systems use three different temperatures, similar sampling times, similar bias strengths, and identical SBC parameters.

S1.1 MIL-53(Al)

MIL-53(Al) is a metal-organic framework which exhibits a closed pore (*cp*) and large pore (*lp*) phase (Figure S1). In previous work, we showed that the *op*-to-*cp* transition can proceed spontaneously at room temperature conditions and pressures above approximately 30 MPa [3]. Since both phases are characterized by a different unit cell shape, the transition between them necessarily requires anisotropic unit cell dynamics during active learning. As a result, the geometry optimizations which are used to label each of the samples with either a *cp* or an *op* label relax both atomic coordinates as well as unit cell components. In spite of the (small) unit cell shape change during the transition, the unit cell volume is a good collective variable for this transition, as can be seen in Figure 2 in the main manuscript; the volume appears to be a near-perfect discriminator for predicting the basin of attraction. The reference free energy profile was computed with umbrella sampling in the unit cell volume; we applied harmonic bias potentials centered at 64 equidistant volume points between $V = 1450 \text{ \AA}^3$ and $V = 3300 \text{ \AA}^3$, and with a force constant of $0.0045 \text{ kJ}/(\text{mol} \cdot \text{ \AA}^6)$.




Figure S1 : Atomic geometry of the closed pore (*cp*) and large pore (*lp*) of MIL-53(Al), a flexible metal-organic framework with the winerack topology [2]. It consists of an aluminum hydroxide chain which is interconnected with terephthalate organic linkers.




Figure S2 : Unit cell of chabazite with an aluminum active site. A proton transfer reaction can occur from the framework to an isobutene guest molecule inside the largest cage.




Figure S3 : Atomic geometry of alanine dipeptide.

S1.2 isobutene @ CHA

We consider an isobutene guest in the largest cage of chabazite, a relatively small zeolite framework (Figure S2). The active site of the reaction (technically a Brønsted acid site) is created by replacing a silicon atom with an aluminum atom, and adding a proton to one of its neighboring oxygens. The reference collective variable as mentioned in Figure 2 (b) in the main manuscript is a simple coordination number difference. With $r_0 = 1.4 \text{ \AA}$ and A, B specific groups of atoms, we have:

$$CN(A, B) = \sum_{\substack{i \in A \\ j \in B}} \frac{1 - \left(\frac{r_{ij}}{r_0}\right)^6}{1 - \left(\frac{r_{ij}}{r_0}\right)^{12}} \quad (1)$$

The coordination number difference is then defined between the coordination of the three outer isobutene carbons and all hydrogens, and the coordination of the four active site oxygens and all hydrogens. The reference free energy profile was computed with umbrella sampling in the coordination number difference; we applied harmonic bias potentials centered at 64 equidistant points between $CV = 5.2$ and $CV = 7.9$, and with a force constant of 1500 kJ/mol.

	GNN	model		ensemble	temperatures	bias potentials		sampling times [ps]		
		SBC layer sizes				#	kappa [eV ⁻¹]	1	2	3
MIL-53(Al) 152 atoms, triclinic	system-specific	(16, 8)		NPT (P = 0 MPa)	150 K 300 K 500 K	32	0.1	96	384	384
isobutene @ CHA 121 atoms, triclinic	foundation	(16, 8)		NVT	150 K 300 K 600 K	32	0.1	192	480	960
alanine dipeptide 22 atoms	system-specific	(16, 8)		NVT	300 K 600 K 900 K	32	0.1/0.2	192	480	960

Table S1 : Overview of the SBC and active learning hyperparameters for each of the case studies performed in this work. SBC requires no tuning (the layer sizes and nonlinearities are identical for each of the three cases, even if the underlying GNNs were different), whereas the sampling requires some fine-tuning related to the temperature(s) and the sampling times for each of the three iterations.

S1.3 alanine dipeptide

The two basins of attraction (referred to in the main text as $c7eq$ and $c7ax$) are shown in Figure S3. The reference free energy surface as a function of ϕ and ψ was computed using 2D umbrella sampling. We used an extensive 24×24 grid of umbrella centers, with a force constant of 400 kJ/(mol·rad²). The relative free energies were obtained by integrating over the two minima on the free energy surface, which yielded a value of -4.8 kJ/mol (indicating that $c7eq$ is more stable than $c7ax$). Based on the learned log probabilities, we performed 1D umbrella sampling, using 64 umbrellas and with a force constant of about 20 kJ/mol. After integrating over the minima, we obtained a relative free energy difference of -4.5 kJ/mol, which is in excellent agreement with the reference value.

References

- [1] Massimiliano Bonomi et al. "PLUMED: A portable plugin for free-energy calculations with molecular dynamics". In: *Comp. Phys. Comm.* 180.10 (2009), pp. 1961–1972. issn: 0010-4655. doi: <https://doi.org/10.1016/j.cpc.2009.05.011>.
- [2] Christian Serre et al. "Very Large Breathing Effect in the First Nanoporous Chromium(III)-Based Solids: MIL-53 or Cr^{III}(OH)₂C₆H₄-CO₂·{HO₂C-C₆H₄-CO₂H}_x·H₂O_y". In: *Journal of the American Chemical Society* 124.45 (Nov. 2002), pp. 13519–13526. issn: 0002-7863. doi: 10.1021/ja0276974.
- [3] Sander Vandenhaute et al. "Machine learning potentials for metal-organic frameworks using an incremental learning approach". In: *npj Comput. Mater.* 9.1 (Feb. 2023), p. 19. issn: 2057-3960. doi: 10.1038/s41524-023-00969-x.

B

Appendices

B.1 Empirical Scaling of Quantum Chemistry Methods

This Section presents more details regarding the empirical scaling test as shown in Figure 2.2. We use ORCA, a highly optimized free but closed-source software package which implements a variety of very powerful approximations [23]. Examples are the domain local pair natural orbital localization method (DLPNO) for drastic improvements to the scaling of correlated methods [129], the chain of spheres (COSX) approximation for computing Hartree-Fock exchange [130], and the resolution of identity (RI) approximation for faster evaluation of overlap integrals [131]. While these are great approximations, it is critical to emphasize that they complicate the evaluation of the gradients (i.e. the atomic forces) considerably. For example, ORCA 6 – which has been released in 2024 – still does not support the calculation of forces when using DLPNO even though that approximation has been around for about a decade. Given that the gradients of $E(\mathbf{r})$ are required to integrate Newton’s equation of motion through time, the CCSD(T) and MP2 timings as reported in Figure 2.2 can not be considered as realistic in the context of dynamic simulations where forces are required. Finally, we should emphasize that many systems are not embedded in vacuum but apply periodic boundary conditions (PBCs) in order to mimic the effect of bulk material. While PBCs are not fundamentally incompatible with techniques such as RI or DLPNO, they drive up the computational cost and overall complexity of the algorithms significantly [132–134]. As such, while correlated methods such as MP2 or CCSD(T) are feasible on isolated clusters of hundreds of atoms, the same cannot be said for periodic systems. The periodicity implicitly restricts

the computational scaling of two-electron integrals to at least cubic or even quartic with the number of atoms.

For Figure 2.2, we consider a cluster of water molecules of varying size, ranging from one molecule to 31 molecules. All calculations were performed on a single node, with 2x AMD EPYC Rome 7H12 and 512 GB RAM. Each calculation was executed on four cores of an otherwise empty node to obtain objective scaling information. The following production-level input files were used:

- **PBE:**

```
! DFT PBE D3 cc-pVTZ def2-ECP TRAH VeryTightSCF DEFGRID3
```

- **PBE0:**

```
! DFT PBE0 D3 cc-pVTZ RIJCOSX def2-ECP TRAH VeryTightSCF  
DEFGRID3
```

- **DLPNO-MP2-F12:**

```
! RHF DLPNO-MP2-F12 cc-pVTZ-F12 cc-pVTZ-F12-CABS cc-pVTZ/C cc-  
pVTZ/JK RIJCOSX def2-ECP TRAH VeryTightSCF DEFGRID3
```

```
%method
```

```
RI on
```

```
end
```

```
%basis
```

```
AuxJ "cc-pVTZ/JK"
```

```
end
```

- **DLPNO-CCSD(T)-F12:**

```
! RHF DLPNO-CCSD(T)-F12 cc-pVTZ-F12 cc-pVTZ-F12-CABS RIJCOSX  
cc-pVTZ/JK cc-pVTZ/C def2-ECP TRAH VeryTightSCF DEFGRID3
```

```
%method
```

```
RI on
```

```
end
```

```
%basis
```

```
AuxJ "cc-pVTZ/JK"
```

```
end
```

See the ORCA manual for further information regarding the employed basis sets and approximations.

B.2 Extended Hessian calculations

(adopted and modified with permission from paper 2)

In voigt notation, the elasticity tensor $\mathbf{C} \in \mathbb{R}^{6 \times 6}$ determines how the stress $\boldsymbol{\sigma} \in \mathbb{R}^6$ within the material changes due to an applied strain $\boldsymbol{\epsilon} \in \mathbb{R}^6$:

$$\boldsymbol{\sigma} = \mathbf{C}\boldsymbol{\epsilon} \quad (1)$$

If we write the periodic box vectors of the system along the rows of a unit cell matrix $\mathbf{g} \in \mathbb{R}^{3 \times 3}$, then the strain of a given atomic configuration may be computed as:

$$\boldsymbol{\epsilon} = \frac{1}{2}(\mathbf{g}_0^{-1}\mathbf{g}\mathbf{g}^T\mathbf{g}_0^{-T} - \mathbf{1}) \quad (2)$$

where \mathbf{g}_0 is the optimized unit cell. Conversely, a given strain $\boldsymbol{\epsilon}$ may be converted into a target unit cell \mathbf{g} using

$$\mathbf{g} = \mathbf{g}_0\sqrt{2\boldsymbol{\epsilon} + \mathbf{1}} \quad (3)$$

where the square root may be computed based on the eigenvalue decomposition of $2\boldsymbol{\epsilon} + \mathbf{1}$. Each of the components of \mathbf{C} is determined by applying a series of small strains, optimizing the positions while keeping the (strained) unit cell fixed, and computing the observed stress.




Figure B.1: Linear stress-strain behavior of UiO-66(Zr), computed at the PBE-D3 level of theory. The slope of the curve determines the coefficient C_{11} .

An example of such a fit for UiO-66(Zr) is shown in Figure B.1. For both the optimization as well as the stress computation, it helps to tighten the SCF convergence bound to 10^{-8} Ha. In theory, it is possible to express the stiffness tensor at 0 K in terms of second-order derivatives of the potential energy, thereby avoiding the need for multiple optimizations or finite-difference approximations. Because this expression is very ill-conditioned, it is not possible to use such an expression for stiffness constants derived using either DFT or classical force fields simply because second-order derivatives computed using finite-difference schemes do not possess the required precision. Neural network potentials on the other hand are often implemented using computational frameworks which support automatic differentiation, and it is therefore possible to compute higher-order derivatives at arbitrary precision. This paragraph derives the expression for \mathbf{C} in terms of such derivatives of the potential energy surface $E(\mathbf{r}, \mathbf{g})$. For convenience, we use regular indexing for the stress and strain tensors as it is more convenient in this context. In that case, Equation 1 is equivalent to the following expression

$$C_{ijkl} = \frac{\partial \sigma_{ij}}{\partial \epsilon_{kl}} \quad (4)$$

Alternatively, it can be written as the second-order derivative of the energy [135],

$$C_{ijkl} = \frac{1}{V} \frac{\partial^2 E}{\partial \epsilon_{ij} \partial \epsilon_{kl}} \quad (5)$$

using the fact that

$$\sigma_{ij} = \frac{1}{V} \frac{\partial E}{\partial \epsilon_{ij}} \quad (6)$$

where V is the volume and the system is at its equilibrium at 0 K, i.e. $\sigma_{ij} = 0$. All the necessary information to compute those second-order derivatives is contained in the *extended* Hessian of the system:

$$\mathbf{H}_{\text{ext}} = \begin{bmatrix} \mathbf{H}_{\text{ff}} & \mathbf{H}_{\text{f}\epsilon} \\ \mathbf{H}_{\text{e}\text{f}} & \mathbf{H}_{\text{e}\epsilon} \end{bmatrix} \quad (7)$$

where the submatrices, denoted by f and ϵ , correspond respectively to the $3N$ fractional coordinates of the system and the 9 unit cell components. Because

first-order derivatives vanish at the equilibrium position, we have for the second-order taylor expansion of the potential energy:

$$E = \begin{bmatrix} \mathbf{x}_f^T & \mathbf{x}_\epsilon^T \end{bmatrix} \begin{bmatrix} \mathbf{H}_{ff} & \mathbf{H}_{f\epsilon} \\ \mathbf{H}_{\epsilon f} & \mathbf{H}_{\epsilon\epsilon} \end{bmatrix} \begin{bmatrix} \mathbf{x}_f \\ \mathbf{x}_\epsilon \end{bmatrix} \quad (8)$$

At constant strain \mathbf{x}_ϵ , the fractional coordinates \mathbf{x}_f which minimize the energy E can be found by setting the derivative of Equation 8 to zero and solving for \mathbf{x}_f :

$$\mathbf{x}_f = -\mathbf{H}_{ff}^{-1} \mathbf{H}_{f\epsilon} \mathbf{x}_\epsilon \quad (9)$$

If we substitute this back in Equation 8, we find a relation between the total energy of the system for a given strain and at optimized coordinates:

$$E = \mathbf{x}_\epsilon^T (\mathbf{H}_{\epsilon\epsilon} - \mathbf{H}_{\epsilon f} \mathbf{H}_{ff}^{-1} \mathbf{H}_{f\epsilon}) \mathbf{x}_\epsilon \quad (10)$$


giving rise to the effective hessian:

$$\mathbf{H}_{\text{eff}} = \mathbf{H}_{\epsilon\epsilon} - \mathbf{H}_{\epsilon f} \mathbf{H}_{ff}^{-1} \mathbf{H}_{f\epsilon} \quad (11)$$

such that the final stiffness constants are given by

$$C_{ijkl} = \frac{1}{V} (\mathbf{H}_{\text{eff}})_{ij,kl} \quad (12)$$

Numerically, this arguably more elegant approach will yield the same values for \mathbf{C} as obtained through the finite-difference scheme outlined above, provided that the extended hessian is computed with sufficient accuracy. This is the case for ML potentials due to the ability to perform automatic differentiation for the total energy E with respect to input coordinates \mathbf{r} and box vectors \mathbf{g} . This is precisely how the ML predictions in Figure 3.7 were obtained. DFT reference values were still obtained using the finite-difference approach because energy derivatives obtained from DFT codes such as CP2K are too noisy.



List of Publications and Conference Contributions

Publications and Preprints

- 1 Water motifs in zirconium metal-organic frameworks induced by nanoconfinement and hydrophilic adsorption sites**
Aran Lamaire, Jelle Wieme, Sander Vandenhoute, Ruben Goeminne, Sven M. J. Rogge, Veronique Van Speybroeck
Nature Communications (accepted), **2024**
- 2 The Operando Nature of Isobutene Adsorbed in Zeolite H-SSZ-13 Unraveled by Machine Learning Potentials Beyond DFT Accuracy**
Sander Vandenhoute[†], Massimo Bocus[†], Veronique Van Speybroeck
[†]: these authors contributed equally
Angewandte Chemie Int. Ed. (accepted) e202413637, **2024**
- 3 Rare Event Sampling using Smooth Basin Classification**
Sander Vandenhoute, Tom Braeckvelt, Pieter Dobbelaere, Massimo Bocus, Veronique Van Speybroeck
arXiv:2404.03777, **2024**
- 4 Machine learning potentials for metal-organic frameworks using an incremental learning approach**
Sander Vandenhoute, Maarten Cools-Ceuppens, Simon DeKeyser, Toon

Verstraelen, Veronique Van Speybroeck
npj Computational Materials, **2023**, *9*, 19

5 **Accurately Determining the Phase Transition Temperature of CsPbI₃ via Random-Phase Approximation Calculations and Phase-Transferable Machine Learning Potentials**

Tom Braeckevelt, Ruben Goeminne, [Sander Vandenhautte](#), Sander Borgmans, Toon Verstraelen, Julian A. Steele, Maarten B. J. Roeffaers, Johan Hofkens, Sven M. J. Rogge, Veronique Van Speybroeck
Chemistry of Materials, **2022**, *34*, 8561–8576

6 **Large-Scale Molecular Dynamics Simulations Reveal New Insights Into the Phase Transition Mechanisms in MIL-53(Al)**

[Sander Vandenhautte](#), Sven M. J. Rogge, Veronique Van Speybroeck
Frontiers in Chemistry, **2021**, *9*, 718920

7 **Towards modeling spatiotemporal processes in metal–organic frameworks**

Veronique Van Speybroeck, [Sander Vandenhautte](#), Alexander E. J. Hoffman, Sven M. J. Rogge
Trends in Chemistry, **2021**, *3*, 605-619

Presentations at Conferences & Workshops

The following list enumerates only those presentations for which I was the main presenter.

1 **(invited) Pushing the Limits of Computational Chemistry**

[Sander Vandenhautte](#), Tom Braeckevelt, Veronique Van Speybroeck
LUMI-BE User Day

November 2024

Charleroi, BE

2 **Learned Rare Event Sampling at the Top of Jacob's Ladder**

[Sander Vandenhautte](#), Massimo Bocus, Veronique Van Speybroeck

Advances in catalytic reactivity simulations under operando conditions

October 2024

Varigotti, IT

- 3 **Scalable Molecular Simulation with psiflow**
Sander Vandenhoute, Veronique Van Speybroeck
ParslFest
September 2024
Chicago, US
- 4 **(invited) Tutorial: Scalable Molecular Simulation with psiflow**
Sander Vandenhoute, Veronique Van Speybroeck
Machine Learning Interatomic Potentials: Theory and Practice
November 2023
Helsinki, FI
- 5 **(invited) Free Energy Calculations at the Top of Jacob's Ladder**
Sander Vandenhoute, Massimo Bocus, Veronique Van Speybroeck
Machine Learning Interatomic Potentials: Theory and Practice
November 2023
Helsinki, FI
- 6 **Scalable Molecular Simulation with psiflow**
Sander Vandenhoute, Veronique Van Speybroeck
ParslFest
October 2023
Chicago, US
- 7 **Free Energy Calculations at the Top of Jacob's Ladder**
Sander Vandenhoute, Massimo Bocus, Veronique Van Speybroeck
EuroHPC User Day
October 2023
Brussels, BE
- 8 **On-the-fly Learning for Activated Processes**
Sander Vandenhoute, Maarten Cools-Ceuppens, Simon DeKeyser, Toon Verstraelen, Veronique Van Speybroeck
Swiss Equivariant Learning Workshop
July 2022
Lausanne, CH
- 9 **Large-Scale Molecular Simulations for Metal-Organic Frameworks**

Sander Vandenhaute, Sven M. J. Rogge, Toon Verstraelen, Veronique Van Speybroeck
The Influence of Crystal Size and Morphology on Framework Materials
Februari 2022
online

- 10 **Systematic Coarse-Graining in Metal-Organic Frameworks**
Sander Vandenhaute, Sven M. J. Rogge, Veronique Van Speybroeck
EuroMAT
September 2021
online

Bibliography

- [1] E. M. Landau Lev D.; Lifshitz, *Quantum Mechanics: Non-Relativistic Theory* (Pergamon Press, 1965)
- [2] E. M. Landau Lev D.; Lifshitz, *Statistical Physics: Part I* (Pergamon Press, 1969)
- [3] M. Tuckerman, *Statistical Mechanics: Theory and Molecular Simulation* (OUP Oxford, 2023)
- [4] M. Born and R. Oppenheimer, Zur Quantentheorie der Molekeln, *Annalen Der Physik* **389**, 457 (1927)
- [5] P. Atkins and R. Friedman, *Molecular Quantum Mechanics* (OUP Oxford, 2011)
- [6] A. Szabo and N. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover Publications, 1996)
- [7] F. Jensen, *Introduction to Computational Chemistry* (Wiley, 2017)
- [8] B. S. Daan Frenkel, *Understanding Molecular Simulation: From Algorithms to Applications* (Elsevier, 2023)
- [9] J. Gibbs, *Elementary Principles in Statistical Mechanics: Developed with Especial Reference to the Rational Foundations of Thermodynamics* (C. Scribner's Sons, 1902)
- [10] J. Izaac and J. Wang, *Computational Quantum Mechanics* (Springer International Publishing, 2019)
- [11] F. Verstraete, T. Nishino, U. Schollwöck, M. C. Bañuls, G. K. Chan, and M. E. Stoudenmire, Density matrix renormalization group, 30 years on, *Nature Reviews Physics* **5**, 273 (2023)
- [12] D. Pfau, S. Axelrod, H. Sutterud, I. von Glehn, and J. S. Spencer, Accurate computation of quantum excited states with neural networks, *Science* **385**, eadn137 (2024)
- [13] R. Li et al., A computational framework for neural network-based variational Monte Carlo with Forward Laplacian, *Nature Machine Intelligence* **6**, 209 (2024)

- [14] N. Gao and S. Günemann, Neural Pfaffians: Solving Many Many-Electron Schrödinger Equations, (2024)
- [15] S. Guo, M. A. Watson, W. Hu, Q. Sun, and G. K.-L. Chan, N-Electron Valence State Perturbation Theory Based on a Density Matrix Renormalization Group Reference Function, with Applications to the Chromium Dimer and a Trimer Model of Poly(p-Phenylenevinylene), *Journal of Chemical Theory and Computation* **12**, 1583 (2016)
- [16] V. Fock, Näherungsmethode zur Lösung des quantenmechanischen Mehrkörperproblems, *Zeitschrift Für Physik* **61**, 126 (1930)
- [17] D. R. Hartree, The Wave Mechanics of an Atom with a Non-Coulomb Central Field. Part I. Theory and Methods, *Mathematical Proceedings of the Cambridge Philosophical Society* **24**, 89 (1928)
- [18] C. Møller and M. S. Plesset, Note on an Approximation Treatment for Many-Electron Systems, *Phys. Rev.* **46**, 618 (1934)
- [19] D. Bohm and D. Pines, A Collective Description of Electron Interactions. I. Magnetic Interactions, *Phys. Rev.* **82**, 625 (1951)
- [20] J. Čížek, On the Correlation Problem in Atomic and Molecular Systems. Calculation of Wavefunction Components in Ursell-Type Expansion Using Quantum-Field Theoretical Methods, *The Journal of Chemical Physics* **45**, 4256 (1966)
- [21] P. Hohenberg and W. Kohn, Inhomogeneous Electron Gas, *Phys. Rev.* **136**, B864 (1964)
- [22] W. Kohn and L. J. Sham, Self-Consistent Equations Including Exchange and Correlation Effects, *Phys. Rev.* **140**, A1133 (1965)
- [23] F. Neese, F. Wennmohs, U. Becker, and C. Riplinger, The ORCA quantum chemistry program package, *The Journal of Chemical Physics* **152**, 224108 (2020)
- [24] J. P. Perdew, K. Burke, and M. Ernzerhof, Generalized Gradient Approximation Made Simple, *Phys. Rev. Lett.* **77**, 3865 (1996)
- [25] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu, *J. Chem. Phys.* **132**, 154104 (2010)

- [26] A. Tkatchenko and M. Scheffler, Accurate Molecular Van Der Waals Interactions from Ground-State Electron Density and Free-Atom Reference Data, *Phys. Rev. Lett.* **102**, 73005 (2009)
- [27] C. Adamo, M. Cossi, and V. Barone, An accurate density functional method for the study of magnetic properties: the PBE0 model, *Journal of Molecular Structure: THEOCHEM* **493**, 145 (1999)
- [28] E. T. Whittaker, XVIII.—On the Functions which are represented by the Expansions of the Interpolation-Theory, *Proceedings of the Royal Society of Edinburgh* **35**, 181 (1915)
- [29] R. O. Dror, R. M. Dirks, J. Grossman, H. Xu, and D. E. Shaw, Biomolecular Simulation: A Computational Microscope for Molecular Biology, *Annual Review of Biophysics* **41**, 429 (2012)
- [30] C. Chipot and A. Pohorille, editors, *Free Energy Calculations: Theory and Applications in Chemistry and Biology*, 1st ed. (Springer Berlin, Heidelberg, Berlin, Heidelberg, 2007)
- [31] J. Hénin, T. Lelièvre, M. R. Shirts, O. Valsson, and L. Delemotte, Enhanced Sampling Methods for Molecular Dynamics Simulations [Article v1.0], *Living Journal of Computational Molecular Science* **4**, 1583 (2022)
- [32] M. R. Shirts and J. D. Chodera, Statistically optimal analysis of samples from multiple equilibrium states, *The Journal of Chemical Physics* **129**, 124105 (2008)
- [33] A. S. J. S. Mey and M. K. Gilson, Predicting Binding Free Energies: Frontiers and Benchmarks, *Chemical Reviews* **120**, 3281 (2020)
- [34] A. P. Bhati, S. Wan, D. W. Wright, and P. V. Coveney, Rapid, Accurate, Precise, and Reliable Relative Free Energy Prediction Using Ensemble Based Thermodynamic Integration, *Journal of Chemical Theory and Computation* **15**, 1265 (2019)
- [35] M. R. Shirts and D. L. Mobley, An Introduction to Best Practices in Free Energy Calculations, *Methods in Molecular Biology* **2114**, 121 (2020)
- [36] D. L. Mobley and P. V. Klimovich, Perspective: Alchemical free energy calculations for drug discovery, *The Journal of Chemical Physics* **137**, 230901 (2012)

- [37] J. D. Chodera and D. L. Mobley, Entropy-Enthalpy Compensation: Role and Ramifications in Biomolecular Ligand Recognition and Design, *Annual Review of Biophysics* **42**, 121 (2013)
- [38] G. M. Torrie and J. P. Valleau, Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling, *Journal of Computational Physics* **23**, 187 (1977)
- [39] A. Laio and M. Parrinello, Escaping free-energy minima, *Proceedings of the National Academy of Sciences* **99**, 12562 (2002)
- [40] M. Invernizzi and M. Parrinello, Rethinking metadynamics: From bias potentials to probability distributions, *The Journal of Physical Chemistry Letters* **11**, 2731 (2020)
- [41] B. B. Ye, S. Chen, and Z.-G. Wang, GCMe: Efficient Implementation of the Gaussian Core Model with Smeared Electrostatic Interactions for Molecular Dynamics Simulations of Soft Matter Systems, *Journal of Chemical Theory and Computation* **20**, 6870 (2024)
- [42] T. Verstraelen, S. Vandenbrande, F. Heidar-Zadeh, L. Vanduyfhuys, V. Van Speybroeck, M. Waroquier, and P. W. Ayers, Minimal Basis Iterative Stockholder: Atoms in Molecules for Force-Field Development, *Journal of Chemical Theory and Computation* **12**, 3894 (2016)
- [43] J. A. McCammon, B. R. Gelin, and M. Karplus, Dynamics of folded proteins, *Nature* **267**, 585 (1977)
- [44] A. Rahman, Correlations in the Motion of Atoms in Liquid Argon, *Phys. Rev.* **136**, A405 (1964)
- [45] T. Darden, D. York, and L. Pedersen, Particle mesh Ewald: An Nlog(N) method for Ewald sums in large systems, *The Journal of Chemical Physics* **98**, 10089 (1993)
- [46] J. Shimada, H. Kaneko, and T. Takada, Performance of fast multipole methods for calculating electrostatic interactions in biomacromolecular simulations, *Journal of Computational Chemistry* **15**, 28 (1994)
- [47] K. Vanommeslaeghe and A. D. MacKerell Jr, CHARMM additive and polarizable force fields for biophysics and computer-aided drug design, *Biochim. Biophys. Acta* **1850**, 861 (2015)

- [48] J. Wang, R. M. Wolf, J. W. Caldwell, P. A. Kollman, and D. A. Case, Development and testing of a general amber force field, *Journal of Computational Chemistry* **25**, 1157 (2004)
- [49] T. R. Underwood and H. C. Greenwell, The Water-Alkane Interface at Various NaCl Salt Concentrations: A Molecular Dynamics Study of the Readily Available Force Fields, *Scientific Reports* **8**, 352 (2018)
- [50] L. Vanduyfhuys, S. Vandenbrande, T. Verstraelen, R. Schmid, M. Waroquier, and V. Van Speybroeck, QuickFF: A program for a quick and easy derivation of force fields for metal-organic frameworks from ab initio input, *Journal of Computational Chemistry* **36**, 1015 (2015)
- [51] J. H. Cavka, S. Jakobsen, U. Olsbye, N. Guillou, C. Lamberti, S. Bordiga, and K. P. Lillerud, A New Zirconium Inorganic Building Brick Forming Metal Organic Frameworks with Exceptional Stability, *J. Am. Chem. Soc.* **130**, 13850 (2008)
- [52] S. Vandenhaute, S. M. J. Rogge, and V. Van Speybroeck, Large-Scale Molecular Dynamics Simulations Reveal New Insights Into the Phase Transition Mechanisms in MIL-53(Al), *Frontiers in Chemistry* **9**, (2021)
- [53] S. Páll, A. Zhmurov, P. Bauer, M. Abraham, M. Lundborg, A. Gray, B. Hess, and E. Lindahl, Heterogeneous parallelization and acceleration of molecular dynamics simulations in GROMACS, *The Journal of Chemical Physics* **153**, 134110 (2020)
- [54] P. Eastman et al., OpenMM 8: Molecular Dynamics Simulation with Machine Learning Potentials, *The Journal of Physical Chemistry B* **128**, 109 (2024)
- [55] G. J. Martyna, D. J. Tobias, and M. L. Klein, Constant pressure molecular dynamics algorithms, *The Journal of Chemical Physics* **101**, 4177 (1994)
- [56] S. E. Feller, Y. Zhang, R. W. Pastor, and B. R. Brooks, Constant pressure molecular dynamics simulation: The Langevin piston method, *The Journal of Chemical Physics* **103**, 4613 (1995)
- [57] M. Parrinello and A. Rahman, Polymorphic transitions in single crystals: A new molecular dynamics method, *Journal of Applied Physics* **52**, 7182 (1981)

- [58] H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, and J. R. Haak, Molecular dynamics with coupling to an external bath, *The Journal of Chemical Physics* **81**, 3684 (1984)
- [59] M. E. Tuckerman, *Statistical Mechanics Theory and Molecular Simulation* (Oxford University Press, New York, 2010)
- [60] C. Bishop, *Pattern Recognition and Machine Learning* (Springer (India) Private Limited, 2013)
- [61] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016)
- [62] A. D. White, Deep Learning for Molecules and Materials, *Living Journal of Computational Molecular Science* **3**, 1499 (2021)
- [63] J. Schwichtenberg, *Physics from Symmetry* (Springer International Publishing, 2015)
- [64] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature* **521**, 436 (2015)
- [65] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, *Attention Is All You Need*, in *Advances in Neural Information Processing Systems*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Vol. 30 (Curran Associates, Inc., 2017), p.
- [66] A. Krizhevsky, I. Sutskever, and G. E. Hinton, *Imagenet Classification with Deep Convolutional Neural Networks*, in *Advances in Neural Information Processing Systems*, edited by F. Pereira, C. Burges, L. Bottou, and K. Weinberger, Vol. 25 (Curran Associates, Inc., 2012), p.
- [67] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, *Neural Message Passing for Quantum Chemistry*, in *Proceedings of the 34th International Conference on Machine Learning - Volume 70* (JMLR.org, Sydney, NSW, Australia, 2017), pp. 1263–1272
- [68] K. T. Schütt, H. E. Sauceda, P.-J. Kindermans, A. Tkatchenko, and K.-R. Müller, SchNet – A deep learning architecture for molecules and materials, *J. Chem. Phys.* **148**, 241722 (2018)

- [69] K. T. Schütt, O. T. Unke, and M. Gastegger, Equivariant message passing for the prediction of tensorial properties and molecular spectra, *Corr* (2021)
- [70] J. Gastegger, C. Yeshwanth, and S. Günnemann, *Directional Message Passing on Molecular Graphs via Synthetic Coordinates*, in *Advances in Neural Information Processing Systems*, edited by M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Vol. 34 (Curran Associates, Inc., 2021), pp. 15421–15433
- [71] J. Gastegger, F. Becker, and S. Günnemann, *Gemnet: Universal Directional Graph Neural Networks for Molecules*, in *Conference on Neural Information Processing Systems (Neurips)* (2021)
- [72] T. Cohen and M. Welling, *Group Equivariant Convolutional Networks*, in *Proceedings of the 33rd International Conference on Machine Learning*, edited by M. F. Balcan and K. Q. Weinberger, Vol. 48 (PMLR, New York, New York, USA, 2016), pp. 2990–2999
- [73] B. Anderson, T.-S. Hy, and R. Kondor, *Cormorant: Covariant Molecular Neural Networks*, in *Proceedings of the 33rd International Conference on Neural Information Processing Systems* (Curran Associates Inc., Red Hook, NY, USA, 2019)
- [74] N. Thomas, T. Smidt, S. Kearnes, L. Yang, L. Li, K. Kohlhoff, and P. Riley, Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds, (2018)
- [75] R. Kondor, Z. Lin, and S. Trivedi, Clebsch–Gordan Nets: a Fully Fourier Space Spherical Convolutional Neural Network, (2018)
- [76] M. Weiler, M. Geiger, M. Welling, W. Boomsma, and T. Cohen, 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data, (2018)
- [77] J. Gastegger, J. Groß, and S. Günnemann, *Directional Message Passing for Molecular Graphs*, in *International Conference on Learning Representations* (2020)
- [78] S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt, and B. Kozinsky, E(3)-equivariant graph

neural networks for data-efficient and accurate interatomic potentials, *Nat. Commun.* **13**, 2453 (2022)

- [79] G. Simeon and G. De Fabritiis, *TensorNet: Cartesian Tensor Representations for Efficient Learning of Molecular Potentials*, in *Proceedings of the 37th International Conference on Neural Information Processing Systems* (Curran Associates Inc., New Orleans, LA, USA, 2024)
- [80] W. Greiner and B. Müller, *Quantum Mechanics: Symmetries* (Springer-Verlag, 1989)
- [81] M. Geiger et al., Euclidean neural networks: e3nn, (2022)
- [82] I. Batatia, D. P. Kovacs, G. Simm, C. Ortner, and G. Csanyi, *MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields*, in *Advances in Neural Information Processing Systems*, edited by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Vol. 35 (Curran Associates, Inc., 2022), pp. 11423–11436
- [83] S. Chmiela, A. Tkatchenko, H. E. Sauceda, I. Poltavsky, K. T. Schütt, and K.-R. Müller, Machine learning of accurate energy-conserving molecular force fields, *Science Advances* **3**, e1603015 (2017)
- [84] P. Eastman et al., SPICE, A Dataset of Drug-like Molecules and Peptides for Training Machine Learning Potentials, *Scientific Data* **10**, 11 (2023)
- [85] R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. von Lilienfeld, Quantum chemistry structures and properties of 134 kilo molecules, *Scientific Data* **1**, 140022 (2014)
- [86] X. Fu, Z. Wu, W. Wang, T. Xie, S. Keten, R. Gomez-Bombarelli, and T. Jaakkola, Forces are not Enough: Benchmark and Critical Evaluation for Machine Learning Force Fields with Molecular Simulations, (2023)
- [87] A foundation model for atomistic materials chemistry, (2024)
- [88] S. Vandenhoute, M. Cools-Ceuppens, S. DeKeyser, T. Verstraelen, and V. Van Speybroeck, Machine learning potentials for metal-organic frameworks using an incremental learning approach, *Npj Computational Materials* **9**, 19 (2023)

- [89] T. Loiseau, C. Serre, C. Huguenard, G. Fink, F. Taulelle, M. Henry, T. Bataille, and G. Férey, A Rationale for the Large Breathing of the Porous Aluminum Terephthalate (MIL-53) Upon Hydration, *Chem. Eur. J.* **10**, 1373 (2004)
- [90] A. Laio and M. Parrinello, Escaping free-energy minima, *Proceedings of the National Academy of Sciences* **99**, 12562 (2002)
- [91] P. G. Yot et al., Metal-organic frameworks as potential shock absorbers: the case of the highly flexible MIL-53(Al), *Chem. Commun.* **50**, 9462 (2014)
- [92] S. M. J. Rogge, M. Waroquier, and V. Van Speybroeck, Unraveling the thermodynamic criteria for size-dependent spontaneous phase separation in soft porous crystals, *Nature Communications* **10**, 4842 (2019)
- [93] A. U. Ortiz, A. Boutin, A. H. Fuchs, and F.-X. Coudert, Anisotropic Elastic Properties of Flexible Metal-Organic Frameworks: How Soft are Soft Porous Crystals?, *Physical Review Letters* **109**, 195502 (2012)
- [94] J. Kirkpatrick et al., Pushing the frontiers of density functionals by solving the fractional electron problem, *Science* **374**, 1385 (2021)
- [95] M. Bocus, S. Vandenhaute, and V. Van Speybroeck, Unraveling the nature of adsorbed isobutene in H-SSZ-13 with operando simulations at the top of Jacob's ladder, *Chemrxiv* (2024)
- [96] P. N. Plessow and F. Studt, How Accurately Do Approximate Density Functionals Predict Trends in Acidic Zeolite Catalysis?, *The Journal of Physical Chemistry Letters* **11**, 4305 (2020)
- [97] W. E and E. Vanden-Eijnden, Transition-Path Theory and Path-Finding Algorithms for the Study of Rare Events, *Annual Review of Physical Chemistry* **61**, 391 (2010)
- [98] S. Vandenhaute, T. Braeckvelt, P. Dobbelaere, M. Bocus, and V. V. Speybroeck, Rare Event Sampling using Smooth Basin Classification, (2024)
- [99] O. Elishav, R. Podgaetsky, O. Meikler, and B. Hirshberg, Collective Variables for Conformational Polymorphism in Molecular Crystals, *J. Phys. Chem. Lett.* **14**, 971 (2023)

- [100] L. Bonati, V. Rizzi, and M. Parrinello, Data-Driven Collective Variables for Enhanced Sampling, *J. Phys. Chem. Lett.* **11**, 2998 (2020)
- [101] D. Mendels, G. Piccini, and M. Parrinello, Collective Variables from Local Fluctuations, *J. Phys. Chem. Lett.* **9**, 2776 (2018)
- [102] L. Bonati, E. Trizio, A. Rizzi, and M. Parrinello, A unified framework for machine learning collective variables for enhanced sampling simulations: mlcolvar, *J. Chem. Phys.* **159**, 14801 (2023)
- [103] E. Grifoni, G.-M. Piccini, and M. Parrinello, Microscopic description of acid–base equilibrium, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 4054 (2019)
- [104] N. Ahalawat and J. Mondal, Assessment and optimization of collective variables for protein conformational landscape: GB1 $b\eta$ -hairpin as a case study, *J. Chem. Phys.* **149**, 94101 (2018)
- [105] H. Niu, P. M. Piaggi, M. Invernizzi, and M. Parrinello, Molecular dynamics simulations of liquid silica crystallization, *Proc. Natl. Acad. Sci. U.S.A.* **115**, 5348 (2018)
- [106] R. E. Odstrcil, P. Dutta, and J. Liu, LINES: Log-Probability Estimation via Invertible Neural Networks for Enhanced Sampling, *J. Chem. Theory Comput.* **18**, 6297 (2022)
- [107] G. Pérez-Hernández, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, Identification of slow molecular order parameters for Markov model construction, *J. Chem. Phys.* **139**, 15102 (2013)
- [108] C. Wehmeyer and F. Noé, Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics, *J. Chem. Phys.* **148**, 241703 (2018)
- [109] A. Mardt, L. Pasquali, H. Wu, and F. Noé, VAMPnets for deep learning of molecular kinetics, *Nat. Commun.* **9**, 5 (2018)
- [110] S. Bhakat, Collective variable discovery in the age of machine learning: reality, hype and everything in between, *RSC Adv.* **12**, 25010 (2022)
- [111] M. Chen, Collective variable-based enhanced sampling and machine learning, *Eur. Phys. J. B* **94**, 211 (2021)

- [112] Z. Belkacemi, P. Gkeka, T. Lelièvre, and G. Stoltz, Chasing Collective Variables Using Autoencoders and Biased Trajectories, *J. Chem. Theory Comput.* **18**, 59 (2022)
- [113] P. R. Vlachas, J. Zavadlav, M. Praprotnik, and P. Koumoutsakos, Accelerated Simulations of Molecular Systems through Learning of Effective Dynamics, *J. Chem. Theory Comput.* **18**, 538 (2022)
- [114] L. Sun, J. Vandermause, S. Batzner, Y. Xie, D. Clark, W. Chen, and B. Kozinsky, Multitask Machine Learning of Collective Variables for Enhanced Sampling of Rare Events, *J. Chem. Theory Comput.* **18**, 2341 (2022)
- [115] E. Trizio and M. Parrinello, From Enhanced Sampling to Reaction Profiles, *The Journal of Physical Chemistry Letters* **12**, 8621 (2021)
- [116] L. Bonati, G. Piccini, and M. Parrinello, Deep learning the slow modes for rare events sampling, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2113533118 (2021)
- [117] L. Müllender, A. Rizzi, M. Parrinello, P. Carloni, and D. Mandelli, Effective Data-Driven Collective Variables for Free Energy Calculations from Metadynamics of Paths, (2023)
- [118] Z. Zou and P. Tiwary, Enhanced Sampling of Crystal Nucleation with Graph Representation Learnt Variables, *The Journal of Physical Chemistry B* **128**, 3037 (2024)
- [119] T. D. Kühne et al., CP2K: An electronic structure and molecular dynamics software package - Quickstep: Efficient and accurate electronic structure calculations, *The Journal of Chemical Physics* **152**, 194103 (2020)
- [120] J. J. Mortensen et al., GPAW: An open Python package for electronic structure calculations, *The Journal of Chemical Physics* **160**, 92503 (2024)
- [121] I. Batatia et al., A foundation model for atomistic materials chemistry, (2024)
- [122] Y. Litman et al., i-PI 3.0: A flexible and efficient framework for advanced atomistic simulations, *The Journal of Chemical Physics* **161**, 62504 (2024)

- [123] M. Bonomi et al., Promoting transparency and reproducibility in enhanced molecular simulations, *Nature Methods* **16**, 670 (2019)
- [124] Y. Babuji et al., *Parsl: Pervasive Parallel Programming in Python*, in *Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing* (Association for Computing Machinery, Phoenix, AZ, USA, 2019), pp. 25–36
- [125] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, Neural Message Passing for Quantum Chemistry, (2017)
- [126] J. Riebesell, R. E. A. Goodall, P. Benner, Y. Chiang, B. Deng, A. A. Lee, A. Jain, and K. A. Persson, Matbench Discovery – A framework to evaluate machine learning crystal stability predictions, (2024)
- [127] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, Attention Is All You Need, *Corr* (2017)
- [128] B. Máté, F. Fleuret, and T. Berau, Neural Thermodynamic Integration: Free Energies from Energy-based Diffusion Models, (2024)
- [129] C. Riplinger and F. Neese, An efficient and near linear scaling pair natural orbital based local coupled cluster method, *The Journal of Chemical Physics* **138**, 34106 (2013)
- [130] F. Neese, F. Wennmohs, A. Hansen, and U. Becker, Efficient, approximate and parallel Hartree–Fock and hybrid DFT calculations. A ‘chain-of-spheres’ algorithm for the Hartree–Fock exchange, *Chemical Physics* **356**, 98 (2009)
- [131] M. Feyereisen, G. Fitzgerald, and A. Komornicki, Use of approximate integrals in ab initio theory. An application in MP2 energy calculations, *Chemical Physics Letters* **208**, 359 (1993)
- [132] A. Bussy and J. Hutter, Efficient periodic resolution-of-the-identity Hartree–Fock exchange method with k-point sampling and Gaussian basis sets, *The Journal of Chemical Physics* **160**, 64116 (2024)
- [133] B. Ramberger, T. Schäfer, and G. Kresse, Analytic Interatomic Forces in the Random Phase Approximation, *Phys. Rev. Lett.* **118**, 106403 (2017)
- [134] A. Bussy, O. Schütt, and J. Hutter, Sparse tensor based nuclear gradients for periodic Hartree–Fock and low-scaling correlated wave function methods in the CP2K software package: A massively paral-

lel and GPU accelerated implementation, *The Journal of Chemical Physics* **158**, 164109 (2023)

- [135] R. Golesorkhtabar, P. Pavone, J. Spitaler, P. Puschnig, and C. Draxl, *ElaStic: A tool for calculating second-order elastic constants from first principles*, *Comput. Phys. Commun.* **184**, 1861 (2013)



Part of the simulations were performed on LUMI – we acknowledge the LUMI User Support Team for their expert support.



MELUXINA

HIGH PERFORMANCE
COMPUTING IN LUXEMBOURG

Part of the simulations were performed on the Luxembourg national supercomputer MeluXina. The authors gratefully acknowledge the LuxProvide teams for their expert support.



EuroHPC
Joint Undertaking



We acknowledge EuroHPC Joint Undertaking for providing access to state of the art computing infrastructure in Europe through Regular and Extreme Scale Access projects (EHPC-REG-2023R02-152 and EHPC-EXT-2023E02-065)

We acknowledge the Research Foundation - Flanders (FWO) for a doctoral fellowship (grant no. 11H6821N)

VLAAMS
SUPERCOMPUTER
CENTRUM



Vlaanderen
is supercomputing

We acknowledge VSC for their Tier1 infrastructure

