



HAL
open science

Developping Numerical Tools for the Understanding and Quantification of Chemical Interactions of Adsorbates on Metallic Surfaces

Ruben Staub

► **To cite this version:**

Ruben Staub. Developping Numerical Tools for the Understanding and Quantification of Chemical Interactions of Adsorbates on Metallic Surfaces. Theoretical and/or physical chemistry. Université de Lyon, 2020. English. NNT : 2020LYSEN059 . tel-03510216

HAL Id: tel-03510216

<https://theses.hal.science/tel-03510216>

Submitted on 4 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Numéro National de Thèse : 2020LYSEN059

THÈSE DE DOCTORAT DE L'UNIVERSITÉ DE LYON

opérée par

l'École Normale Supérieure de Lyon

École Doctorale N° 206

École Doctorale de Chimie (Chimie, Procédés, Environnement)

Spécialité : Chimie Théorique

Discipline : Chimie

Soutenue publiquement le 27/10/2020, par :

Ruben STAUB

Développement d'outils de modélisation des surfaces d'alliages et leur application à la cinétique en catalyse hétérogène

Devant le jury composé de :

Daan FRENKEL	Professeur	University of Cambridge	Rapporteur
Michail STAMATAKIS	Professeur	University College London	Rapporteur
Céline CHIZALLET	Chef de projet	IFP Energies nouvelles	Examinatrice
Paul FLEURAT-LESSARD	Professeur	Université de Bourgogne	Examineur
Aurélien GARIVIER	Professeur	ENS de Lyon	Examineur
Franziska HESS	Professeure	TU Berlin	Examinatrice
Stephan STEINMANN	Chargé de recherche CNRS	Laboratoire de Chimie	Directeur de thèse

To the 12-years old me that dreamed of contributing to Science

Science is the art of the appropriate approximation.

— Byron K. Jennings, *On the Nature of Science*



Acknowledgements

In the context of this thesis, I have been fortunate enough to interact with truly wonderful people that helped me all along this thesis.

First of all, I would like to acknowledge the fantastic job that Stephan N. Steinmann accomplished as my supervisor all along this thesis. Always kind, supportive, available and sharp, Stephan constantly encouraged me to pursue my own challenges and approaches (even when it led to explore uncharted territories a bit far from chemical applications) while helping me turn my ideas into publishable material and always providing valuable insights to make it more “chemical”. It is safe to say that Stephan has been the best supervisor I could have dreamed of, and I am truly honored and ever grateful that he accepted me as its very first PhD student, forcing him to pass its accreditation.

I am also extremely grateful to Carine Michel for her long time support and follow-up: ever since I started studying at the ENS de Lyon, I have been fortunate enough to have had Carine review my internship reports, always full of advice and vocational guidance. Carine is the reason I have been able to work on this thesis, as she is the one that introduced me to Stephan in the first place, and I am ever grateful about it, but also for the formidable care she puts into building fantastic working conditions and an inclusive workplace by leading a group dynamic that feels like a community.

Speaking of community, I have been most fortunate to have been surrounded by marvelous ~~eo~~-workers friends in the lab, leading to fruitful discussions and that certainly made this adventure much more enjoyable: Paul, Sarah, Pauline, Kamila, Danish, Adrian, Carles, Alessio, Akif, Jérôme, Masoud, Daniel, Kateryna, Benjamin, Chris, Antton, Romain, Vincenza, Thomas, Baptiste, Caroline, Emanuele, Vincent (both), Christian, David, Tao, Tangui, Elise, Sophie, Pascal, Martin (both), Cyrille, Chantal and 98 more, plus all whom I have forgotten (I have an extremely bad memory for names...)

But also outside of the lab, I have been surrounded by many friends to whom I could complain when my computations did not converge, compare our theses and get challenged in coding contests: Kenny, Loïs, Rédouane, Maxime, Adérito, Raphaël, Mathieu, Gabrielle, Claire,

Acknowledgements

Alice, David, and of course many other that did not directly impact my thesis but were still here.

Going back to the context of this thesis, I wish to mention all those who kindly took a bit of their time to discuss the projects of this thesis. In particular, I am absolutely grateful to Yann Strozecki for inviting me to talk about the Reinforcement Sampling project which led to fruitful discussions with Franck Quessette and the whole ALMOST-DAVID team. I also cannot thank enough Aurélien Garivier for agreeing to discuss the Reinforcement Sampling with me. His extremely valuable guidance toward KL-UCB and his ever relevant remarks certainly helped improve the quality of the UCT-based discussions in this thesis in many ways.

An other source of fruitful interactions came from conferences and symposiums, and so, I am grateful to the organizers of the conferences where I have been invited (CTTC VIII, JTMS 2019, SCF 2018, CP2K Day, and the interdisciplinary ML thematic meeting where I met Aurélien Garivier) and the audience for their questions giving food for thought and forcing to find new formulations. I am also thankful to the organisers of the fantastic TCCM Winter School LTTC 2018, where I met truly wonderful minds (among others: Flavio, Clément and Jean that introduced me to ASE).

I wish to specially acknowledge those who inspired me the initial idea for each project of this thesis:

- Of course, Stephan introduced me to the ALMO formalism and the need for a metal-compatible variant.
- Paul Clabaut first sensitized me to the geometry-based coordination numbers during an exciting discussion where he presented me some cool results he had just obtained, which in the end inspired me to work around this notion.
- Paul Fleurat-Lessard sensitized me to the hidden features of the linear model during a thesis defence where he promoted (as a jury) the use of confidence intervals, piquing my curiosity which, in the end, led to the development of rank-Greville.
- Finally, Olivier Teytaud sensitized me to the MCTS and UCT approach during a presentation at a ski resort organized by the Computer Science department of the ENS de Lyon, which inspired me the Reinforcement Sampling approach.

By all means, I want to warmly thank the referees and jury of this thesis for kindly accepting to dedicate a bit of their time to review this work, and therefore greatly contributing to the final value of this thesis.

Acknowledgements

On a side note, I would also like to thank my teachers all along my education for believing in me, and without whom none of this would have been possible.

Speaking of which, I have been fortunate enough to be trusted with assisting with the teaching of a few courses at the ENS de Lyon, and I am grateful to the teachers in charge (Christophe, Claire, Laureline, Stéphane, Peter, ...) and the students for their enthusiasm and patience, making this experience most enjoyable.

I wish to thank the ENS de Lyon (also for allocating me a grant with which I did this thesis), the administrative staff and the PSMN (our local cluster) staff of course.

How could I not be grateful to my parents who have always believed in me, and supported me in all occasions? They carefully followed my thesis as I explained and vulgarized it to them at each step, leading me to better grasp the notions I was dealing with.

I wish to mention and thank my, sadly late, grandfather for giving me the opportunity to receive the education that led me to this thesis.

Last but not least, I have been most blessed to have been accompanied by my life companion Carine Seraphim during this journey, where she has been ever supportive, caring and involved. Carine have been an ever reassuring presence as we sailed together through the ups and downs encountered all along this thesis. It is safe to say that she have had a tremendous impact, both directly and indirectly, on this thesis in its present form, so I am ever grateful for her supportive assistance (especially during the challenging task of writing this manuscript) and the inspiring discussions that we have had all along this journey that have been most pleasant because of her.



Abstract

Chemical bonding can be described using the density functional theory (DFT) framework as electronic density reorganisation between chemical entities. Relying on Absolutely Localized Molecular Orbitals (ALMO), the density reorganization can be separated into polarization and charge transfer. However, the ALMO formalism is intractable when combined with mixed-states theory, which is required for describing metals. As shown herein, these two theories can be unified into a mean-field approximation called S-ALMO, allowing for a fundamental density-based description of most chemical bonds.

In practice, bonds are most often defined by geometrical considerations. Based on an intuitive and powerful idea (an atom's nearest neighbours must cover its field of view), SANN (Solid Angle based Nearest Neighbours) provides locally adaptive parameter-free coordination numbers. A natural extension of this algorithm, called ASANN, is developed to tackle local anisotropy while remaining parameter-free. We demonstrate that this method effectively provides a fundamental topology-based definition of coordination numbers applicable to close packed bulks, liquids and interfaces.

Even though bonds are topologically defined, they contain chemically relevant information, providing descriptors for describing the associated energy. Lattice-based cluster expansion model Hamiltonians are particularly popular for determining the adsorption energy in heterogeneous catalysis. Fitting the corresponding model Hamiltonian is usually performed on a training set composed of hundreds of structures hand-picked by a chemist. For the model Hamiltonian to be relevant for describing a given reaction, the chemist uses his own expertise and chemical intuition to select a diverse set of chemically relevant structures. Treating the generation of a relevant input as a strategy-based game, a novel active learning scheme for model Hamiltonians is designed based on a UCT (Upper Confidence Tree). Updating the model Hamiltonian on-the-fly is achieved by a novel recursive least-squares algorithm, called rank-Greville, exploiting the rank deficiencies of the occurrence matrix to reduce the scaling of the update. Furthermore, a domain knowledge extension on top of the UCT framework allows to optimize the training set construction, leading to the Reinforcement Sampling scheme. This approach effectively replaces chemical intuition with reproducible reinforcement learning techniques.

Keywords: Energy decomposition analysis, coordination numbers, active learning, rank-deficient recursive least-squares



Résumé

Cette thèse se focalise sur le développement d'outils pour l'étude fondamentale des interactions latérales au travers de multiples projets :

D'un point de vue quantique, les effets énergétiques des interactions chimiques sont décomposables rigoureusement en utilisant des orbitales moléculaires absolument localisées (ALMO). Au cours de cette thèse, une approximation de champ moyen a été développée et implémentée dans CP2K afin d'unifier le formalisme ALMO avec la théorie des états mixtes dans une description DFT, permettant d'étendre l'analyse en composantes énergétiques aux systèmes métalliques.

D'un point de vue purement topologique, un nouvel algorithme de détermination de liaisons chimiques (ASANN) a été développé durant cette thèse. Au travers d'un simple terme correctif, ASANN étend l'algorithme de référence SANN aux systèmes présentant une anisotropie locale, sans introduire aucun paramètre. Les nombres de coordination ainsi produits sont particulièrement adaptés pour la descriptions d'interfaces.

Ces points de vue complémentaires se combinent sous la forme d'un Hamiltonien effectif basé sur les interactions latérales. Un nouvel outil non-stochastique a été développé et implémenté durant cette thèse, utilisant des algorithmes d'apprentissage par renforcement pour l'entraînement automatique de tels modèles destinés à simuler un système à N corps sur surface réactive. Cet outil repose sur l'adaptation d'un UCT avec une pré-exploration guidée par le modèle en cours d'apprentissage. La mise à jour d'un tel modèle linéaire a été optimisée par la formulation d'un nouvel algorithme de résolution des moindres carrés récursifs exploitant les déficiences de rang.

Mots-clés : décomposition de l'énergie, nombres de coordination, résolution des moindres carrés récursifs, systèmes linéaires déficient de rang, apprentissage actif



Contents

Acknowledgements	v
Abstract	ix
Résumé	xi
Table of Contents	xiii
1 Introduction	1
1.1 Motivation	1
1.2 Electronic point of view	2
1.3 Topological point of view	5
1.4 Unification through model Hamiltonians	6
1.5 Linear regression	9
1.6 Training set selection	11
1.7 Thesis diversity	13
Bibliography	14
2 Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions	21
2.1 Introduction	22
2.2 Methodology	23
2.2.1 Ensemble DFT	23
2.2.2 BLW formalism	25
2.2.3 Ensemble BLW	27
2.3 Computational Details	31
2.3.1 Energy Decomposition Analysis	32
2.4 Results and Discussion	35
2.5 Conclusions	39
Bibliography	39
3 Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis	45
3.1 Introduction	46
3.2 Methods	47

Contents

3.3	Computational Details	50
3.4	Results and Discussion	51
3.4.1	Relative Stability of Ice-like Layers	51
3.4.2	Electronic Analysis	53
3.4.3	Energy Decomposition Analysis	55
3.5	Conclusion	67
	Bibliography	67
4	Parameter-Free Coordination Numbers for Solutions and Interfaces	75
4.1	Introduction	76
4.2	Algorithms	78
4.2.1	The SANN algorithm	78
4.2.2	Anisotropically corrected Solid-Angle based Nearest-Neighbors, ASANN algorithm	80
4.3	Computational Details	85
4.4	Results and Discussion	86
4.4.1	Metallic surfaces	86
4.4.2	Coordination numbers at charged interfaces	89
4.4.3	Application to Au NP energy fitting	94
4.5	Conclusion	96
	Bibliography	96
5	Efficient Recursive Least Squares Solver for Rank-Deficient Matrices	103
5.1	Introduction	104
5.2	Derivation	106
5.2.1	General non-orthogonal rank factorization	106
5.2.2	Orthogonal rank factorization	115
5.3	Implementation	116
5.4	Numerical Tests	117
5.4.1	Computational efficiency	118
5.4.2	Numerical stability	122
5.5	Conclusion	126
	Bibliography	126
6	UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling	129
6.1	Introduction	130
6.2	Cluster Expansion Hamiltonian	132
6.3	Presentation of the problem and the strategy	135
6.3.1	Definition of an “adequate” training set	135
6.3.2	Upper Confidence Trees for solving the adequate training set problem	136
6.3.3	Reformulation of the training set construction as a game	140
6.4	Designing a UCT for model Hamiltonians: Reinforcement Sampling	143
6.4.1	Patterns	143
6.4.2	Basic UCT design	148

6.4.3	Pre-exploration	154
6.4.4	Implementation	155
6.5	Results	161
6.5.1	Protocol presentation	161
6.5.2	Conclusion	172
	Bibliography	172
7	Conclusions and Outlook	177
	Bibliography	180
A	Appendix to Chapter 2	183
A.1	Density recovery	183
A.2	Comparison of results with two different Pt pseudo potentials and results from VASP	185
A.3	Further analysis of the deformation energy	185
A.4	Influence of the electronic temperature	185
A.5	Expanded Figures 2, 3 and 4	185
A.6	Effect of K-point sampling	185
B	Appendix to Chapter 3	193
B.1	Additional Analysis	193
B.1.1	Water–Water interaction	193
B.2	Additional Tables and Figures	193
	Bibliography	193
C	Appendix to Chapter 4	205
C.1	Angular correction term derivation	205
C.2	ASANN Theorems	206
C.3	Discussion regarding a continuous version	209
C.4	Additional Figures	213
C.5	Model Hamiltonian parameters	213
C.6	Generalized Coordination Numbers	213
C.7	Ionic crystals	214
C.8	Molecules	216
C.9	ROC data	218
	Bibliography	218
D	Appendix to Chapter 6	221
D.1	Linear regression models	221
D.1.1	Common assumptions	221
D.1.2	Ordinary least squares	223
D.1.3	Generalized least squares	224
D.1.4	Conclusion	224
D.2	Score range for UCB	224
D.2.1	Range without uncertainty	225

Contents

D.2.2	Range with uncertainty	227
D.2.3	Variance without uncertainty	229
D.2.4	Variance with uncertainty	231
D.2.5	Conclusion	231
D.3	Heuristic score justification	231
D.3.1	A first guess: total adsorption energy	231
D.3.2	Relevant heuristic: absolute deviation from non-interacting model . . .	232
	Bibliography	235



1 Introduction

1.1 Motivation

Chemistry is infused with the concept of bonds. From covalent bonds to weak bonds^{1;2}, ionic bonds³ and metallic bonds⁴, chemical bonding plays a major role in almost all sub-fields of Chemistry^{5;6}. But what exactly is a chemical bond? In its most general definition, a chemical bond can be seen as a significant interaction between chemical entities. But in fact, a precise definition would be very much dependent on the context and abstraction level. Indeed, real chemical phenomena are incredibly complex processes that occur across diverse characteristic space and time scales, at the same time⁷. In order to describe them, these phenomena are usually decomposed into macro/meso/micro descriptions and various time-range related effects (long term deviations, short term fluctuations, ...), all interacting together, but studied separately with the appropriate approximations⁸. Total synthesis, atomistics, data science, thermodynamics, kinetics and so forth, all provide different contexts for studying chemical events. Even at the atomic level, Theoretical Chemistry relies on a collection of tools and models for exploring this vast domain of study. In particular, two major descriptions coexist: Quantum Mechanics (QM) and Molecular Mechanics (MM).

In a QM approach, one is interested in the proper description of the electronic structure of a chemical system. In this approach, electrons are treated as quantum objects, through their wavefunction (i.e. the mathematical representation of a quantum state), by evaluating and solving the (approximate) Hamiltonian of a given system. Because of the computational complexity of this resolution, such approach is limited to systems with a relatively few number of atoms. Instead of directly manipulating wavefunctions, one can consider the much simpler, and chemically more intuitive, electronic density. This is particularly popular⁹ in the context of the Density Functional Theory^{10;11} (DFT). Furthermore, DFT allows for the description of relatively larger systems at a quantum level description. One should note that, unlike the

electronic density, chemical bonds are not an observable of the system¹². Hence, just like on experimental grounds, QM enables an unbiased exploration of various bonding situations. However, QM provides a plethora of tools to characterize, quantify and classify chemical bonds¹³⁻¹⁶.

However, in a MM approach, one is focused on the molecular structure only, using classical models to describe the electronic effects: force fields¹⁷. Indeed, a force field is an empirical model aimed at reproducing electronic effects through classical interactions between classical objects (e.g., spheres) representing chemical entities. As a consequence, MM calculations are performed at a rather abstract “high-level” description, dealing with classical nuclei coordinates and, more generally, topology descriptors¹⁸⁻²⁰. Such simplified model is particularly well adapted to describe relatively large systems and/or large time scales that could not be handled by a full quantum description in a reasonable computational time²¹. One should note that, in an MM approach, chemical bonds are treated as additional relations (typically restraints) between coordinates of classically defined chemical entities. As a consequence, an acceptable chemical bond definition, in this context, should inherently be dealing with geometrical considerations.

Interestingly enough, chemical bonds are defined in drastically different ways depending on the approach. Indeed, both approaches rely on different representations and abstractions. So much that covalent bonding is merely a derivable property in the QM approach, while covalent bonds are crucial input data for the MM approach. In the first chapters of this thesis, we explore further the fundamentally dissimilar definitions, depending on the electronic vs. topology-based context, of the same notion of covalent bond that is so characteristic of Chemistry.

1.2 Electronic point of view

From a fundamental point of view, chemical bonding can be described in terms of interaction between chemical entities. Such description is not limited to bonds between atoms, since interactions are defined between molecules, ions, radicals, ... (incorporating many-body effects). These sub-systems will be referred to as fragments. A crucial question now is: how can we identify or even quantify an interaction between fragments?

An emblematic characterization of bonds can be found in Bader’s Atoms-In-Molecule (AIM) analysis¹⁶, where bonding is linked to the topology of the electronic density. Alternatively, we consider instead an energy-based “responsive” scheme, where an interaction is assessed as the response of a fragment to the presence of an other fragment, compared with a hypothetical

situation where the fragments are not interacting (i.e. do not “feel” each other)²². We believe that this approach provides meaningful definitions that can be easily interpreted²³.

In the DFT framework, an interaction between fragments is inevitably linked to an electronic density reorganisation of these fragments. A convenient way to characterize such reorganisation is by evaluating its energy contribution. Indeed, such evaluation is retrieving the strength of the interaction. However, the strength is not the only property that defines a bond. Another important property lies in the nature of the interaction, and can be investigated by discriminating the components of the reorganisation. Indeed, the electronic density reorganisation, caused by the interaction between two fragments, can be naturally decomposed into two main components:

- Electronic density reorganisation within the defined fragments, called polarisation.
- Electronic density reorganisation in-between the fragments, called charge transfer (CT).

This decomposition is not only interesting from a fundamental point of view, or as a tool for rational design. It is also essential for choosing the right approximations and models in order to accurately describe specific processes. For example, if charge transfer effects are found to be significant (and unpredictable from polarization effects), then by definition, a polarizable force-field will likely be not adapted in this context and one would need to rely, at least, on a reactive force field, such as ReaxFF²⁴. Now, a natural question is: how such decomposition can be evaluated/computed in practice?

The most straightforward way to study this decomposition would be to switch on and off one of the components, in order to deduce its contribution. In fact, it is possible to do so with charge transfer between fragments (i.e. disabling it artificially, without changing anything else). The charge transfer contribution (and therefore the reorganization decomposition) can be obtained by forcing the electronic density to remain localized on each fragment (i.e. polarization effects only), and comparing it to the same system where the electronic density is free to be delocalized (i.e. polarization and charge transfer effects). The energy difference between these two systems represents, therefore, the energy contribution of the charge transfer effects, leading to an energy-based decomposition of interactions. This Energy Decomposition Analysis (EDA) requires, therefore, a method for forcing electrons to remain localized on predefined fragments.

The most rigorous method for applying density localization constraints rely on the use of localized orbitals, called Block-Localized Wavefunctions²⁵ (BLW), Absolutely Localized

Chapter 1. Introduction

Molecular Orbitals^{22;26–28} (ALMO) or even Extremely Localized Molecular Orbitals²⁹ (ELMO). The use of localized orbitals also provides significant computational advantages:

- The Basis Set Superposition Error (BSSE) between fragments is naturally reducedⁱ.
- Near linear scaling DFT can be achieved in some settingsⁱⁱ. In this matter, the mixed Gaussian Plane Waves (GPW) approach makes CP2K, the open-source simulation software, particularly relevant for using localized orbitalsⁱⁱⁱ.

Unfortunately, ALMO has originally been formulated for pure states only, i.e., all spin-orbitals are either occupied or empty. In metallic systems the density of states is continuous across the Fermi-level. For Hartree-Fock and Kohn-Sham density functional theory, fractional orbital/band occupations following Fermi-Dirac statistics, are equivalent to the much more involved explicit consideration of pure-state ensembles³¹. The limitation of ALMO to pure states is especially unfortunate for exploring the charge transfer vs. polarization effects involved in heterogeneous catalysis. It is also particularly problematic for electrochemical simulations, with passive species whose charge must remain localized.

The occupied ALMOs are necessarily non-orthogonal between different fragments.^{iv} The self-consistent coupling of occupied (non-orthogonal) ALMOs between fragments implies that the interactions with the surroundings depend on the quantum state of each fragment, leading to non-identical orbitals from one quantum state to an other. It would therefore seem that the ALMO formalism cannot be rigorously combined with mixed state theory in a computationally tractable way. But what about an approximate unification? Is there a meaningful approximation under which the ALMO formalism can be applied to metallic systems?

I suggest in chapter 2 a mean-field approximation where the average interaction is considered for all quantum states independently, leading to a more general orbital occupations rescaling trick, based on a novel and subtly different concept: selfish orbitals. Using selfish orbitals, instead of rescaled orbitals (to obtain rescaled occupations), leads to a mixed-state description that is (approximately) compatible with the ALMO formalism. This approximate unified theory has been integrated into CP2K, under the name S-ALMO. This work extends, therefore, the rigorous CT EDA-based characterization of chemical bonding to metallic systems, under a meaningful approximation.

ⁱThe BSSE comes from the description inconsistencies arising from the superposition of basis set functions, leading to some parts of the electronic density being described by additional basis set functions.

ⁱⁱThis is due to the arising block-diagonal structure of the Kohn-Sham (KS) matrix.

ⁱⁱⁱSince the use of plane waves ensures a fast construction of the locally-projected KS Hamiltonian³⁰

^{iv}This can be explained by additional degree of freedom constrains, in the general case.

This S-ALMO powered EDA was applied, in chapter 3, to water adlayers adsorbed on noble metal surfaces, in order to unravel the properties and origins of the adsorption energies.

1.3 Topological point of view

In practice, however, characterizing chemical bonds through an EDA is only reserved to specific interactions of interest, on systems with a relatively few number of atoms. Indeed, chemical bonds are commonly defined by geometrical considerations instead of density-based analyses^{32–46}. Yet, these geometrically characterized bonds represent valuable and chemically relevant topology-based descriptors^{47–49}. Such geometry-based definition is also particularly relevant to structure analysis, for example applied to Molecular Dynamics (MD) simulations of rather large systems on large time scales that could certainly not be processed by EDA (for studying all possible interactions) in a reasonable amount of time. It now becomes clear that a geometry-based definition of chemical bonds, allowing for fast detection, is essential in this context.

Such detection is performed by a coordination number algorithm, using its own geometry-based definition of covalent bonds. Many algorithms have been proposed^{32–46}, each with its own set of rules for deciding which bonds should be accepted. The most common coordination algorithms rely on predefined cutoffs to define chemically relevant bonds: a covalent bond is detected if and only if two atoms are at a distance shorter than a predefined cutoff distance depending on the nature of the atoms pair⁵⁰. These algorithms are simple, fast and typically used in visualization software³². However, their generalizability can be identified to be impacted by two major issues:

- Non adaptation to local environment: a fixed cutoff definition does not take the local environment into account, reducing its transferability from one connectivity context to another.
- Parameter dependent: the cutoff distances must be defined for each type of bond being evaluated. As a consequence, such algorithm cannot be generalized to connections that were not already pre-defined.

Parameter-free algorithms were developed, using the radial distribution function of atom pairs of interest to define adequate cutoffs^{34–40}. However, this method requires to be trained on a large and similar input dataset in order to extract typical bond distances that are relevant for cutoff definition. Finally, the cutoff distances obtained are globally adapted to the kind of structures being processed, but local variations are averaged out, so that each final cutoff distance is fixed and applied regardless of the local environment.

Only a few algorithms^{33;41;42;46} try to extract connectivity from the local topological environment and therefore define adaptive parameter-free bonds solely based on geometry considerations. These algorithms define geometry-based bonding in terms of proximity, so that an entity is defined to be connected only with its nearest neighbors. This definition is particularly convenient for ideal close-packed materials, where all nearest neighbors are exactly at the same distance. But how are nearest neighbors defined in practice? A state-of-the-art approach is given by the SANN (Solid Angle based Nearest Neighbours) algorithm⁴⁶. This approach is based on an intuitive yet powerful idea: an atom's nearest neighbors must cover its field of view. To be more precise, the central atom is considered a sphere and its neighbors are considered as balls, all with the same radius, defined as the minimal radius such that the union of the intersections between the central sphere and the neighboring balls covers the whole central sphere. The nearest neighbors are defined as the neighbors whose associated ball is intersecting the central sphere. Using this general definition, SANN provides locally adaptive and parameter-free coordination numbers.

However, structures displaying a strong local anisotropy (e.g. edges, protuberance, interfaces) are not well handled by this algorithm⁴⁵. Indeed, the SANN algorithm relies on an isotropic description, and therefore, does not take the local anisotropy into account. A solution was proposed in the form of the Relative Angular Distance (RAD) algorithm⁴⁵, by introducing parameters. So a natural question is: Can nearest neighbors also be defined on structures displaying strong local anisotropy, without introducing parameters?

In chapter 4, I suggest a simple extension of the SANN algorithm, called ASANN, that is compatible with structures displaying a strong local anisotropy, while remaining parameter-free. The ASANN algorithm incorporates into its description a simple^v information about local anisotropy, by using the barycenter of local neighbors. This information is used to estimate the field of view that is effectively covered by neighbors. Then, the nearest neighbors are computed similarly to the SANN algorithm, using a slightly modified version of the original idea: an atom's nearest neighbors must cover its locally relevant field of view. Therefore, using only a basic local anisotropy descriptor, ASANN effectively provides a fundamental geometry-based definition of bonds that is applicable to close packed bulks, liquids and interfaces.

1.4 Unification through model Hamiltonians

These opposite, yet both fundamental, definitions of bonds rely on complementary visions that can actually be combined. Indeed, EDA describes bonds in terms of energetics, whereas coordination number algorithms process bonds as geometry-based descriptors of the local

^vOnly first-order anisotropy is taken into account, i.e. only one privileged direction is considered.

1.4. Unification through model Hamiltonians

connectivity. These two visions are conveniently exploited together in an effective (or model) Hamiltonian. In a model Hamiltonian, geometry-based features are associated with energetics⁴⁷⁻⁴⁹, just like in force fields. As a consequence, a model Hamiltonian can be seen as an embedding of energetics into the space of geometry-based descriptors.

The simplest bond-based model Hamiltonian is a linear model that correlates the energy of a system $E(g)$ to the number of chemical entities (usually atoms) of each type in system g and each possible geometry-defined bond between them:

$$E(g) = \sum_a \beta_a N_a(g) + \sum_{a,b} \beta_{a-b} N_{a-b}(g), \quad (1.1)$$

where a, b are among all defined elements, $N_a(g)$ is the number of a atoms in g , $N_{a-b}(g)$ is the number of $a-b$ bonds in g (for example defined by a coordination number algorithm), and $(\beta_a, \dots, \beta_{a-b}, \dots)$ are the parameters of the (linear) model Hamiltonian, representing single-body terms and the energy contribution of each type of bond. This drastically simplified model can be applied quickly and is extremely easy to interpret. It is such a model that I have applied in Chapter 4 to achieve a first approximation for the energy of CuAu alloy surfaces and nanoparticles. However, this model suffers from two main issues:

- Bonds are treated in a binary fashion (present or not), which cannot account for continuous variations in bond distances (see Figure 1.1 for a schematic representation). This is no longer an issue for a lattice-based representation, where positions (and therefore, distances) are discrete variables.
- Higher order many-body effects are not taken into account. Only singletons and pairs are considered, whereas effects arising from the formation of triplets, and so forth, are neglected.

As a consequence, such a model Hamiltonian would only be acceptable for lattice-based structures where many-body effects arising from triplets (and more) are negligible. These assumptions are quite restrictive in practice. But, in fact, this most basic model can be extended/generalized into an exact framework for describing lattice-based structures⁵¹.

This extension is accomplished by generalizing the concept of bonds between more than two entities. Indeed, just as pair-based (i.e. 2-body) interactions can be described by geometry-based descriptors (i.e., bonds), many-body interactions can be described by geometry-based patterns, representing relations between the coordinates of possibly many objects. Even though these patterns can be defined in 3D, they are much more conveniently applied^{vi} on

^{vi}Due to the related degrees of freedom reduction.

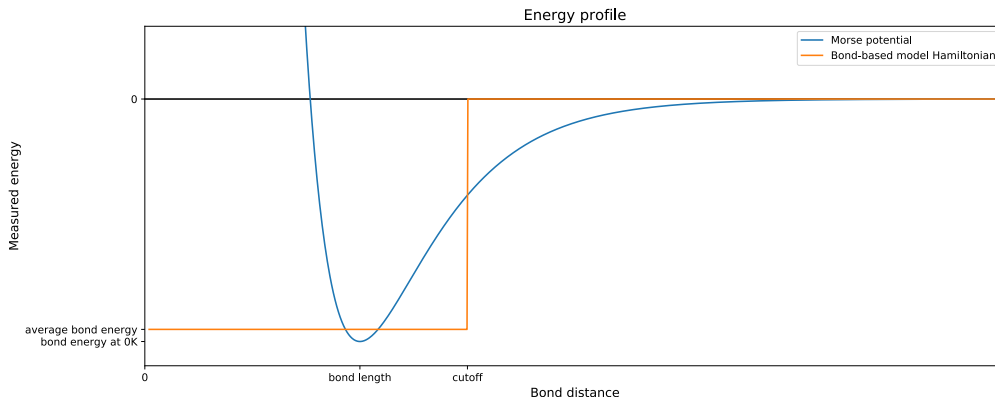


Figure 1.1 – Comparison of the energy profiles of a single bond as the function of the bond length obtained from a Morse potential or a bond-based model Hamiltonian. The bond detection being binary (i.e. found or not) in a typical model Hamiltonian, the corresponding energy profile is a step function.

2D surfaces. In surface science and heterogeneous catalysis it is common to distinguish two types of patterns: the “single-body” patterns, corresponding to chemical entities (adsorbates), which can, in the general case, occupy more than a single site on the surface, and “patterns of patterns”, representing the so called lateral interaction between adsorbates. Therefore, a pattern-based description is particularly relevant for studying reactive surfaces such as catalysts, especially at high coverage, where lateral interactions are important⁵²⁻⁵⁵. Besides, since adsorbates are likely found at specific adsorption sites (typically top, bridge and hollow), a lattice-based representation is particularly well suited. As a consequence, reactive surfaces can quite rigorously be described by an enumerable set of discrete variables representing the occurrences of each possible pattern^{vii}.

Using patterns as a generalization of geometry-based bonds, a unification of energetics with geometry-based descriptors of connectivity can be found in a pattern-based model Hamiltonian. This model is, therefore, now general enough to rigorously describe reactive surfaces. The most common^{56;57} and rigorous pattern-based model Hamiltonian applied to surfaces is the Cluster Expansion^{58;59} based model Hamiltonian:

$$E_{ads}(g) = \sum_a \beta'_a N_a(g) + \sum_{a,b} \beta'_{a-b} N_{a-b}(g) + \sum_{a,b,c} \beta'_{a-b-c} N_{a-b-c}(g) + \dots, \quad (1.2)$$

where a, b, c are now among possible adsorbates at the surface, N_a, N_{a-b}, N_{a-b-c} refer to the number of related patterns present at the surface and the target is now the total adsorption energy $E_{ads}(g)$ of g , and β' are the contributions of patterns to the adsorption energy. The

^{vii}These patterns are indeed enumerable, because of the lattice-like nature of site preference of adsorbates on reactive surfaces

peculiarity of this model Hamiltonian lies in the enumeration of all possible patterns: 1-body terms, 2-body terms, 3-body terms, ... up to terms where all adsorbates are considered, ultimately accounting for all possible many-body effects. Of course, such model would require far too many parameters to fit (possibly infinitely many). Fortunately, only a few number of parameters are non-negligible in this expansion, as the energy contribution of high-order terms and long-range patterns quickly decay to zero. Therefore, this expansion makes it simple to truncate this model Hamiltonian while maintaining an accurate energy description with a reduced number of variables⁶⁰.

To summarize, the CE-based model Hamiltonian is one of the of most general and rigorous, yet simple, way to associate energetics with geometry-based descriptors of connectivity, by defining energy contributions to the associated patterns. The not-so-trivial question is: how can these energy contributions be obtained in practice? Is there a canonical/optimal way to extract this information? To go deeper along this path, it is essential to identify the key components of this learning process. Like any empirical model, a CE-based model Hamiltonian is built from two key components:

- A functional form for the regression. Here, it is a simple linear model.
- A training set containing data to learn from.

Let us then explore the fundamentals of each component and, of course, search for optimal algorithms.

1.5 Linear regression

Performing a linear regression between a target variable Y (also called regressand) and some independent variables (called regressors) X_1, \dots, X_m , using a set of n observations $X \in \mathbb{R}^{n \times m}$ (with $X_{i,j}$ being the value for the j -th regressor on the i -th observation) and associated target values $\mathbf{y} \in \mathbb{R}^n$ is equivalent to finding a set of parameters \mathbf{b} ^{viii} such that:

$$\mathbf{y} = \mathbf{b}^\top X + \mathbf{e}, \tag{1.3}$$

has least squares residuals \mathbf{e} . Among the possibly infinitely many least-squares solutions, the minimum-norm least-squares solution \mathbf{b} is unique⁶¹.

Under the ordinary least-squares assumptions, the minimum-norm least-squares solution

^{viii} \mathbf{b} is an estimator of the true parameters β of the linear model $Y = \beta^\top X + \epsilon$

Chapter 1. Introduction

can be written⁶²:

$$\mathbf{b} = X^+ \mathbf{y}, \quad (1.4)$$

where X^+ is the pseudoinverse (also called Moore-Penrose inverse^{63;64}) of the observation matrix X . Commonly available least-square solvers⁶⁵ usually rely on a decomposition of X (Singular Value Decomposition (SVD), Cholesky decomposition, QR factorization, ...) to efficiently compute the minimum-norm least-squares solution⁶¹ \mathbf{b} .

During the training of a CE-based model Hamiltonian, one might be interested in adding new observations, to further refine the model, without the need to recompute the solution from scratch. In other words, updating the least squares solution b_{n+1} due to the addition of a new observation $\Gamma_{n+1} = (X_{n+1,1}, \dots, X_{n+1,m})$ (with associated target y_{n+1}) can be done by re-using the previous solution b_n , effectively saving computational time. Such updated solution is obtained by a Recursive Least-Squares^{61;66} (RLS) solver and can be written:

$$b_{n+1} = b_n + K \times (y_{n+1} - b_n^\top \Gamma_{n+1}), \quad (1.5)$$

where K is called the Kalman gain vector⁶¹, and $y_{n+1} - b_n^\top \Gamma_{n+1}$ is the predicted residual (or *a priori* error) representing the difference between the new target y_{n+1} and its predicted value using the previous model $b_n^\top \Gamma_{n+1}$.

Commonly available least-squares solvers can be adapted to solve the recursive RLS problem by using specific algorithms^{61;67} to update their internal decomposition of X . Unfortunately, these algorithms were not designed to benefit from an eventual rank deficiency of X (i.e. X has rank $r < \min(n, m)$). Indeed, a rank deficiency occurs when the observation matrix X of a linear system is both:

- Column rank-deficient (i.e. $r < m$). In terms of a linear model, this reads as: the regressors/input space was not fully observed. For example, if the regressor X_2 was null in all observations, no information can be extracted on its corresponding contribution β_2 to the target Y . In other words, some parts of the linear system are underdetermined.
- Row rank-deficient (i.e. $r < n$). This is due to some observations being linearly redundant with others (i.e. at least one observation can be obtained by linear combination of others). For example, if the observation Γ_2 is identical to Γ_1 , a redundant information is assimilated (leading to an eventual inconsistency if the target values are different $y_1 \neq y_2$). In other words, some parts of the linear system are overdetermined.

Real-life applications (e.g. in embedded systems⁶⁶) generate observations that are rarely

rank-deficient, since this would suggest that the regressors were poorly chosen, some of them being redundant. However, rank-deficiency is particularly relevant for a model Hamiltonian being trained, since the relatively small number of observations attainable (due to the computational cost of each new observation) forces to neglect groups of patterns with apparently negligible contributions (i.e. underdetermination of apparently negligible contributions) while non-negligible contributions should be properly sampled to integrate and average out the contribution of eventually non-negligible truncated terms (i.e. overdetermination of non-negligible contributions). So, how can RLS benefit, in terms of computational cost, from rank-deficiency?

In chapter 5, I suggest an RLS algorithm based on a rank decomposition called Rank-Greville, due to its apparent similarity with the Greville^{68;69} algorithm. By maintaining a basic rank factorization, a maximal free family^{ix} \mathcal{F} of observations is maintained (i.e. observations are categorized as redundant or not), leading to the full linear system being decomposed into two separate sub-systems: a purely non-overdetermined system (composed of the observations in \mathcal{F} only) and a purely non-underdetermined system (where all observations are expressed as linear combination of the observations in \mathcal{F}). These sub-systems, with reduced size, can be solved independently and combined, allowing for a full solution update in $O(mr)$ operations instead of $O(mn)$ operations required by commonly available recursive least-squares solvers. This Rank-Greville solver is, therefore, particularly adapted to maintain a constantly up-to-date version of the model Hamiltonian being trained.

1.6 Training set selection

A critical component of any empirical model is the dataset used for its training. The CE-based model Hamiltonian is no exception: it must be trained on a set of geometries that are of relevance to the kind of applications that the model Hamiltonian is destined to. Typically, a CE-based model Hamiltonian is designed for a specific surface-catalyzed reaction and is trained with geometries representing the entities involved in the modelled process, adsorbed on the reactive surface of interest. Unfortunately, this specific dataset is costly to obtain, as the adsorption energy of each geometry needs to be evaluated, with eventual geometry optimization beforehand, usually at a DFT level of theory. Besides, the combinatorial nature of possibly relevant patterns leads to a relatively large number of parameters that, therefore, require to be fitted from a rather large and diverse, yet specialized, training set (a hundred geometries is the typical order of magnitude). As a consequence, the construction of such training set is a tedious and (human and computing) resource consuming task⁷⁰. Therefore, the optimal and automatic choice of geometries to include in such training set is highly desirable⁷⁰⁻⁷².

^{ix}A free family is a set of linearly independent elements.

Chapter 1. Introduction

In practice, the geometries are hand-picked by the chemist, using chemical intuition and/or a pseudo-exhaustive/random sampling. Such sampling is likely sub-optimal and hardly transferable from one model Hamiltonian to the other. Finally, a pseudo-random sampling is neither reproducible (without reporting explicitly all generated geometries) nor truly random. So a natural question is: Can this training set selection task be automated? And even better, can it be done in a way that is more optimal than a pseudo-exhaustive/random sampling?

Design of Experiments (DoE) are commonly used in experimental Chemistry to plan a series of experiments in advance in order to optimize their usefulness⁷³, whereas active learning is a sub-field of machine learning dedicated to the sequential optimization of a training set (i.e. only the next selection is optimized, repetitively, while the model is being trained)⁷⁴. Therefore, both approaches are of relevance for optimizing the training set of a model Hamiltonian. But, before anything, what is an optimal training set?

In chapter 6, an adequate training set is identified by two main properties:

- **Chemical relevance:** an adequate training set should focus on geometries that hold significant chemical information, just like a chemist would use its own chemical intuition to prefer geometries displaying strong lateral interactions. The idea is that stronger interactions should be explored more.
- **Diversity:** an adequate training set should ideally cover the whole chemical space and allow for the description of most situations. Therefore, diversity is essential, since a model has no rigorous predictive power outside the scope of its training set.

These apparently incompatible properties are typical of the classical exploitation vs. exploration dilemma, that can be formulated here as: focusing on chemically relevant geometries while, at the same time, exploring other types of geometries to globally explore the chemical space.

When such dilemma is combined with a huge action space (i.e. here the chemical space), it is commonly solved by relying on Upper Confidence Trees (UCTs). A UCT is a classical tool of reinforcement learning, a sub-field of machine learning dedicated to the design of optimal policies, that has become unavoidable in strategy-game playing algorithms⁷⁵⁻⁷⁸. So, how can a UCT be applied to solve the adequate training set problem in the context of a CE-based model Hamiltonian construction?

In the very same chapter 6, I suggest to reformulate the choice of a geometry (to include in the training set) as a strategy-based game. Using this reformulation, I link the construction of an adequate training set to the learning of an optimal strategy for playing this game. However, this strategy-based game is unlike most board games, so that a conventional UCT is expected to be ill-suited. Therefore, I suggest a custom UCT-based framework, called Reinforcement Sampling, powered by a chemically relevant metric and a powerful pre-exploration step, acting as an additional DoE approach on top of an active learning scheme.

1.7 Thesis diversity

In summary, this thesis focuses on the multiple points of view of the chemical bond (from DFT-based energy effects to a more practical geometry-based description), and their unification in the form of a model Hamiltonian. In particular, the focus is made on methodological developments, guided by a strong algorithmic interest. For each aspect, a state-of-the-art method is explored along with its applicability domain, and a novel extension is proposed by merging various concepts. The developed solutions either exploit the specificity of some applications or extend the applicability domain, geared towards the description of metallic interfaces:

- The ALMO formalism is combined with mixed-states theory, leading to S-ALMO, a mean-field based smearing-compatible generalization, allowing for the energy decomposition analysis of the interaction between adsorbates and metallic surfaces.
- The SANN algorithm is combined with a simple barycenter-based local anisotropy descriptor, leading to ASANN, an extension aware of local anisotropy, allowing for a better description of interfaces.
- A rank factorization is combined with a Greville-like update, leading to Rank-Greville, an RLS algorithm with improved computational complexity on linear systems with rank-deficiency.
- Finally, the model Hamiltonian training is linked to playing a strategy-based game and a UCT approach is combined with diverse add-ons, including a DoE-inspired pre-exploration step, leading to the Reinforcement Sampling framework as an automation tool for building an adequate training set on reactive surfaces.

Consequently, this thesis tackles challenging highly diverse themes. The topics addressed in this thesis have barely any common notation nor shared concepts. Hence, each chapter represents a specific topic and introduces its own context and notation. Therefore, each chapter ought to be self-sufficient in its presentation of the related work, as is also the very

Chapter 1. Introduction

purpose of articles. As a consequence, I believe articles to be relevant as chapters of this thesis, to ensure an unentangled reading. Chapters 2, 3 and 4 are based on published papers, whereas chapter 5 is based on an article being reviewed at the time of writing and chapter 6 is meant to be condensed and submitted by the end of this thesis.

Bibliography

- [1] Dzyaloshinskii, I. E.; Lifshitz, E. M.; Pitaevskii, L. P. The general theory of van der Waals forces. *Advances in Physics* **1961**, *10*, 165–209.
- [2] Parsegian, V. A. *Van der Waals forces: a handbook for biologists, chemists, engineers, and physicists*; Cambridge University Press, 2005.
- [3] Bader, R.; Henneker, W. The ionic bond. *Journal of the American Chemical Society* **1965**, *87*, 3063–3068.
- [4] Anderson, W. P.; Burdett, J. K.; Czech, P. T. What is the metallic bond? *Journal of the American Chemical Society* **1994**, *116*, 8808–8809.
- [5] Pauling, L. *The Nature of the Chemical Bond...*; Cornell university press Ithaca, NY, 1960; Vol. 260.
- [6] Pauling, L. The chemical bond. *Cornell Univ* **1967**, 64.
- [7] Kirchner, B.; Vrabec, J. *Multiscale molecular methods in applied chemistry*; Springer Science & Business Media, 2012; Vol. 307.
- [8] Talanquer, V. Macro, submicro, and symbolic: the many faces of the chemistry “triplet”. *International Journal of Science Education* **2011**, *33*, 179–195.
- [9] Becke, A. D. Perspective: Fifty years of density-functional theory in chemical physics. *The Journal of chemical physics* **2014**, *140*, 18A301.
- [10] Hohenberg, P.; Kohn, W. Inhomogeneous electron gas. *Physical review* **1964**, *136*, B864.
- [11] Kohn, W.; Sham, L. J. Self-consistent equations including exchange and correlation effects. *Physical review* **1965**, *140*, A1133.
- [12] Bader, R. F. Bond paths are not chemical bonds. *The Journal of Physical Chemistry A* **2009**, *113*, 10391–10396.
- [13] Reed, A. E.; Curtiss, L. A.; Weinhold, F. Intermolecular interactions from a natural bond orbital, donor-acceptor viewpoint. *Chem. Rev.* **1988**, *88*, 899–926.

- [14] Tsipis, A. C. DFT challenge of intermetallic interactions: From metallophilicity and metallaromaticity to sextuple bonding. *Coord. Chem. Rev.* **2017**, *345*, 229–262.
- [15] Gonthier, J. F.; Steinmann, S. N.; Wodrich, M. D.; Corminboeuf, C. Quantification of “fuzzy” chemical concepts: a computational perspective. *Chem. Soc. Rev.* **2012**, *41*, 4671–4687.
- [16] Bader, R. F. W. *Atoms in Molecules: A Quantum Theory*. **1990**,
- [17] Tuckerman, M. E.; Martyna, G. J. *Understanding modern molecular dynamics: techniques and applications*. 2000.
- [18] Phillips, J.; Braun, R.; Wang, W.; Gumbart, J.; Tajkhorshid, E.; Villa, E.; Chipot, C.; Skeel, R.; Kale, L.; Schulten, K. Scalable molecular dynamics with NAMD. *JOURNAL OF COMPUTATIONAL CHEMISTRY* **2005**, *26*, 1781–1802.
- [19] Case, D.; Cheatham, T.; Darden, T.; Gohlke, H.; Luo, R.; Merz, K.; Onufriev, A.; Simmerling, C.; Wang, B.; Woods, R. The Amber biomolecular simulation programs. *JOURNAL OF COMPUTATIONAL CHEMISTRY* **2005**, *26*, 1668–1688.
- [20] Brooks, B. R. et al. CHARMM: The Biomolecular Simulation Program. *JOURNAL OF COMPUTATIONAL CHEMISTRY* **2009**, *30*, 1545–1614.
- [21] MacKerell, A. et al. All-atom empirical potential for molecular modeling and dynamics studies of proteins. *JOURNAL OF PHYSICAL CHEMISTRY B* **1998**, *102*, 3586–3616.
- [22] Khaliullin, R. Z.; Head-Gordon, M.; Bell, A. T. An efficient self-consistent field method for large systems of weakly interacting components. *J. Chem. Phys.* **2006**, *124*, 204105.
- [23] Khaliullin, R. Z.; Cobar, E. A.; Lochan, R. C.; Bell, A. T.; Head-Gordon, M. Unravelling the Origin of Intermolecular Interactions Using Absolutely Localized Molecular Orbitals. *J. Phys. Chem. A* **2007**, *111*, 8753–8765.
- [24] van Duin, A. C. T.; Dasgupta, S.; Lorant, F.; Goddard, W. A. ReaxFF: A Reactive Force Field for Hydrocarbons. *The Journal of Physical Chemistry A* **2001**, *105*, 9396–9409.
- [25] Mo, Y.; Gao, J.; Peyerimhoff, S. D. Energy decomposition analysis of intermolecular interactions using a block-localized wave function approach. *J. Chem. Phys.* **2000**, *112*, 5530–5538.
- [26] Stoll, H.; Wagenblast, G.; Preuss, H. On the use of local basis sets for localized molecular orbitals. *Theor. Chim. Acta* **1980**, *57*, 169–178.
- [27] Nagata, T.; Takahashi, O.; Saito, K.; Iwata, S. Basis set superposition error free self-consistent field method for molecular interaction in multi-component systems: Projection operator formalism. *J. Chem. Phys.* **2001**, *115*, 3553–3560.

Chapter 1. Introduction

- [28] Gianinetti, E.; Raimondi, M.; Tornaghi, E. Modification of the Roothaan equations to exclude BSSE from molecular interaction calculations. *Int. J. Quantum Chem.* **1996**, *60*, 157–166.
- [29] Sironi, M.; Genoni, A.; Civera, M.; Pieraccini, S.; Ghitti, M. Extremely localized molecular orbitals: theory and applications. *Theoretical Chemistry Accounts* **2007**, *117*, 685–698.
- [30] Khaliullin, R. Z.; VandeVondele, J.; Hutter, J. Efficient Linear-Scaling Density Functional Theory for Molecular Systems. *J. Chem. Theory Comput.* **2013**, *9*, 4421–4427.
- [31] Yang, W.; Mori-Sanchez, P.; Cohen, A. J. Extension of many-body theory and approximate density functionals to fractional charges and fractional spins. *J. Chem. Phys.* **2013**, *139*, 104114.
- [32] Larsen, A. H. et al. The atomic simulation environment—a Python library for working with atoms. *Journal of Physics: Condensed Matter* **2017**, *29*, 273002.
- [33] Collins, R. On the entropy of the liquid state. *Proceedings of the Physical Society* **1965**, *86*, 199–212.
- [34] Kirkwood, J. G.; Boggs, E. M. The Radial Distribution Function in Liquids. *The Journal of Chemical Physics* **1942**, *10*, 394–402.
- [35] Coulson, C. A.; Rushbrooke, G. S. On the Interpretation of Atomic Distribution Curves for Liquids. *Phys. Rev.* **1939**, *56*, 1216–1223.
- [36] Frank, F. C.; Kasper, J. S. Complex alloy structures regarded as sphere packings. I. Definitions and basic principles. *Acta Cryst* **1958**, *11*, 184–190.
- [37] Brostow, W. Radial distribution function peaks and coordination numbers in liquids and in amorphous solids. *Chemical Physics Letters* **1977**, *49*, 285 – 288.
- [38] O’Keeffe, M. A proposed rigorous definition of coordination number. *Acta Crystallogr A Cryst Phys Diffr Theor Gen Crystallogr* **1979**, *35*, 772–775.
- [39] Varma, S.; Rempe, S. B. Coordination numbers of alkali metal ions in aqueous solutions. *Biophysical Chemistry* **2006**, *124*, 192 – 199.
- [40] Mikolaj, P. G.; Pings, C. J. The Use of the Coordination Number in the Interpretation of Fluid Structure. *Physics and Chemistry of Liquids* **1968**, *1*, 93–108.
- [41] Bernal, J. D. An attempt at a molecular theory of liquid structure. *Trans. Faraday Soc.* **1937**, *33*, 27–40.
- [42] Rahman, A. Liquid Structure and Self-Diffusion. *The Journal of Chemical Physics* **1966**, *45*, 2585–2592.

- [43] Noel, J. K.; Whitford, P. C.; Onuchic, J. N. The Shadow Map: A General Contact Definition for Capturing the Dynamics of Biomolecular Folding and Function. *The Journal of Physical Chemistry B* **2012**, *116*, 8692–8702.
- [44] Taylor, R. Which intermolecular interactions have a significant influence on crystal packing? *CrystEngComm* **2014**, *16*, 6852–6865.
- [45] Higham, J.; Henchman, R. H. Locally adaptive method to define coordination shell. *The Journal of Chemical Physics* **2016**, *145*, 084108.
- [46] van Meel, J. A.; Fillion, L.; Valeriani, C.; Frenkel, D. A parameter-free, solid-angle based, nearest-neighbor algorithm. *The Journal of Chemical Physics* **2012**, *136*, 234107.
- [47] Chen, W.; Schmidt, D.; Schneider, W. F.; Wolverton, C. Ordering and Oxygen Adsorption in Au–Pt/Pt(111) Surface Alloys. *J. Phys. Chem. C* **2011**, *115*, 17915–17924.
- [48] Zhu, B.; Creuze, J.; Mottet, C.; Legrand, B.; Guesmi, H. CO Adsorption-Induced Surface Segregation and Formation of Pd Chains on AuPd(100) Alloy: Density Functional Theory Based Ising Model and Monte Carlo Simulations. *J. Phys. Chem. C* **2016**, *120*, 350–359.
- [49] Vignola, E.; Steinmann, S. N.; Le Mapihan, K.; Vandegehuchte, B. D.; Curulla, D.; Sautet, P. Acetylene Adsorption on Pd–Ag Alloys: Evidence for Limited Island Formation and Strong Reverse Segregation from Monte Carlo Simulations. *J. Phys. Chem. C* **2018**, *122*, 15456–15463.
- [50] Cordero, B.; Gómez, V.; Platero-Prats, A. E.; Revés, M.; Echeverría, J.; Cremades, E.; Barragán, F.; Alvarez, S. Covalent radii revisited. *Dalton Trans.* **2008**, 2832–2838.
- [51] Stamatakis, M. Kinetic modelling of heterogeneous catalytic systems. *Journal of Physics: Condensed Matter* **2014**, *27*, 013001.
- [52] Farkas, A.; Hess, F.; Over, H. Experiment-based kinetic monte carlo simulations: CO oxidation over RuO₂ (110). *The Journal of Physical Chemistry C* **2012**, *116*, 581–591.
- [53] Sutton, J. E.; Lorenzi, J. M.; Krogel, J. T.; Xiong, Q.; Pannala, S.; Matera, S.; Savara, A. Electrons to reactors multiscale modeling: catalytic CO oxidation over RuO₂. *ACS Catalysis* **2018**, *8*, 5002–5016.
- [54] Hess, F.; Over, H. Rate-determining step or rate-determining configuration? the Deacon reaction over RuO₂ (110) studied by DFT-based KMC simulations. *ACS Catalysis* **2017**, *7*, 128–138.
- [55] Piccinin, S.; Stamatakis, M. CO Oxidation on Pd(111): A First-Principles-Based Kinetic Monte Carlo Study. *ACS Catal.* **2014**, *4*, 2143–2152.

Chapter 1. Introduction

- [56] Cao, L.; Li, C.; Mueller, T. The use of cluster expansions to predict the structures and properties of surfaces and nanostructured materials. *Journal of chemical information and modeling* **2018**, *58*, 2401–2413.
- [57] Stampfl, C. Predicting surface phase transitions from ab initio based statistical mechanics and thermodynamics. *Phase Transitions* **2007**, *80*, 311–332.
- [58] Sanchez, J. M.; Ducastelle, F.; Gratias, D. Generalized cluster description of multicomponent systems. *Physica A: Statistical Mechanics and its Applications* **1984**, *128*, 334–350.
- [59] Lerch, D.; Wieckhorst, O.; Hammer, L.; Heinz, K.; Müller, S. Adsorbate cluster expansion for an arbitrary number of inequivalent sites. *Physical Review B* **2008**, *78*, 121405.
- [60] Schmidt, D. J.; Chen, W.; Wolverton, C.; Schneider, W. F. Performance of cluster expansions of coverage-dependent adsorption of atomic oxygen on Pt (111). *Journal of chemical theory and computation* **2012**, *8*, 264–273.
- [61] Bjorck, A. *Numerical Methods for Least Squares Problems*; Other Titles in Applied Mathematics; Society for Industrial and Applied Mathematics, 1996.
- [62] Penrose, R. On best approximate solutions of linear matrix equations. *Mathematical Proceedings of the Cambridge Philosophical Society* **1956**, *52*, 17–19.
- [63] Moore, E. On the reciprocal of the general algebraic matrix. *Bull. Amer. Math. Soc.* **1920**, *26*, 394–395.
- [64] Penrose, R. A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society* **1955**, *51*, 406–413.
- [65] Anderson, E.; Bai, Z.; Bischof, C.; Blackford, L. S.; Demmel, J.; Dongarra, J.; Du Croz, J.; Greenbaum, A.; Hammarling, S.; McKenney, A.; Sorensen, D. *LAPACK Users' Guide*, 3rd ed.; Society for Industrial and Applied Mathematics, 1999.
- [66] Björck, Å. In *Encyclopedia of Optimization*; Floudas, C. A., Pardalos, P. M., Eds.; Springer US: Boston, MA, 2009; pp 1856–1866.
- [67] Daniel, J. W.; Gragg, W. B.; Kaufman, L.; Stewart, G. W. Reorthogonalization and stable algorithms for updating the Gram-Schmidt QR factorization. *Mathematics of Computation* **1976**, *30*, 772–795.
- [68] Greville, T. N. E. Some Applications of the Pseudoinverse of a Matrix. *SIAM Review* **1960**, *2*, 15–22.
- [69] Ben-Israel, A.; Greville, T. N. *Generalized inverses: theory and applications*; Springer Science & Business Media, 2003; Vol. 15.

- [70] van de Walle, A.; Ceder, G. Automating first-principles phase diagram calculations. *Journal of Phase Equilibria* **2002**, *23*, 348.
- [71] Seko, A.; Koyama, Y.; Tanaka, I. Cluster expansion method for multicomponent systems based on optimal selection of structures for density-functional theory calculations. *Physical Review B* **2009**, *80*, 165122.
- [72] Herder, L. M.; Bray, J. M.; Schneider, W. F. Comparison of cluster expansion fitting algorithms for interactions at surfaces. *Surface Science* **2015**, *640*, 104–111.
- [73] Wagner, J. R.; Mount, E. M.; Giles, H. F. In *Extrusion (Second Edition)*, second edition ed.; Wagner, J. R., Mount, E. M., Giles, H. F., Eds.; Plastics Design Library; William Andrew Publishing: Oxford, 2014; pp 291 – 308.
- [74] Settles, B. *Active learning literature survey*; 2009.
- [75] Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M., et al. Mastering the game of Go with deep neural networks and tree search. *nature* **2016**, *529*, 484–489.
- [76] Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A., et al. Mastering the game of go without human knowledge. *nature* **2017**, *550*, 354–359.
- [77] Toyoda, T.; Kotani, Y. Monte Carlo Go Using Previous Simulation Results. Proceedings of the 2010 International Conference on Technologies and Applications of Artificial Intelligence. USA, 2010; p 182–186.
- [78] Drake, P.; Uurtamo, S. Move ordering vs heavy playouts: Where should heuristics be applied in Monte Carlo Go,”. Proceedings of the 3rd North American Game-On Conference. 2007; pp 171–175.



2 Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions

This Chapter is based on the following article: R. Staub, M. Iannuzzi, R. Z. Khaliullin, and S. N. Steinmann *J. Chem. Theory Comput.* **2019**, *15*, 265–275.

The energy decomposition analysis based on block localized wave functions (BLW-EDA) allows to gain physical insight into the nature of chemical bonding, decomposing the interaction energy in (1) a “frozen” term, accounting for the attraction due to electrostatic and dispersion interactions, modulated by Pauli repulsion, (2) the variationally assessed polarization energy and (3) the charge-transfer. This method has so far been applied to gas- and condensed-phase molecular systems. However, its standard version is not compatible with fractionally occupied orbitals (i.e. electronic smearing) and, as a consequence, cannot be applied to metallic surfaces. In this work, we propose a simple and practical extension of BLW-EDA to fractionally occupied orbitals, termed Ensemble BLW-EDA. As illustrative examples, we have applied the developed method to analyze the nature of the interaction of various adsorbates on Pt(111), ranging from physisorbed water to strongly chemisorbed ethylene. Our results show that polarization and charge-transfer both contribute significantly to the adsorption minimum for all studied systems. The energy decomposition analysis provides details with respect to competing adsorption sites (e.g., CO on atop, vs. hollow sites) and elucidates the respective importance of polarization and charge transfer for the increased adsorption energy of H₂S compared to H₂O. Our development will enable a deeper understanding of the impact of charge transfer on catalytic processes in general.

2.1 Introduction

Analyzing the chemical bond in order to understand the driving force and diversity of bonding is almost as old as quantum mechanics.^{1;2} Bader's quantum theory of the atom in a molecule³ and the natural bond orbitals by Weinhold,⁴ are among the most popular approaches, but since the decomposition of the chemical bond energy into different contributions is not unique,⁵ a multitude of other energy decomposition analysis (EDA) schemes have been developed.⁶ All these different schemes come with their advantages and weaknesses, so that application of several tools can either bring contradicting results or provide confidence that at least the trend is captured accurately.⁷ Most of these tools are developed having molecules in mind and their application to the bonding between surfaces and adsorbates is comparatively rare.^{8–11} As one of the few examples, Tonner and co-workers have demonstrated that EDAs on semi-conducting surfaces can provide a deep understanding of coverage effects¹² and elucidate peculiar bonding mechanisms.¹³ Our goal is to extend the EDA based on the block-localized wave function (BLW)^{14;15} to metallic surfaces. The BLW is designed to localize the electrons in Hilbert space in the Mulliken sense, i.e., in terms of atom-centered basis functions. A typical choice is to restrict the expansion to basis functions of atoms of a given molecule, excluding contributions from all other basis functions in the system, thus defining one block per molecule. The BLW – also known in this context as absolutely localized molecular orbitals (ALMO) – is variationally optimized and the different blocks polarize each other. The main advantage of the BLW-EDA compared to other EDAs is that the polarization of fragments in their mutual presence is computed fully variationally. This allows to rigorously separate the polarization energy from charge-transfer. The charge-transfer in BLW-EDA includes bonding orbital interactions between the fragments, which distinguishes it from the charge-transfer obtained by constrained density functional theory (CDFT)^{16;17}, where the charge-transfer is defined in real-space, rather than Hilbert space.¹⁸

The formulation of the theoretical framework exploited for optimizing the BLW goes back to the works of Stoll and co-workers,¹⁹ who aimed at reducing the computational cost of the self-consistent field (SCF) procedure by using ALMOs. The variational optimization of blocks that are localized in Hilbert space has found other applications over the years, exploiting other properties of the BLW unrelated to computational speed-ups. For instance, BLW has been proposed to be used to remove the basis set superposition error (BSSE)^{20;21} and BLW-EDA has been applied to molecular complexes at the DFT^{22–24} or correlated wave function level of theory,^{25;26} quantifying hyperconjugation,^{27;28} strain energies²⁹ and even covalent bonds.^{30;31} Furthermore, ALMOs can provide transferable molecular orbitals (where they are also called ELMOs for extremely localized molecular orbitals)^{32–35} or fragment densities to be used in X-ray structure elucidations.³⁶ Through its variational character, BLW also provides the unique opportunity to directly assess the impact of electron delocalization on the properties of molecules by comparing computed NMR chemical shifts^{37;38} and J-couplings³⁹ for (de-

localized states, which can be seen as the comparison between a single Lewis structure and the true, electron delocalized state. Similarly, the electronic structure of the electron localized state can be analyzed by scalar fields such as the electron localization function (ELF) or the localized orbital locator (LOL) in order to shed more light on the consequences of electron delocalization on the electronic structure.^{40;41} Due to the variational definition of the polarization energy, other applications involve investigating of the impact of polarization on the hydration shell of ionic solutes⁴² and comparison between the BLW polarization energy and polarizable force fields, which has also been exploited to parametrize first principles based force fields.^{43;44}

Heterogeneous catalytic reactions involving metal surfaces or particles are involved in many major industrial processes, such as selective hydrogenations in refineries, ammonia synthesis, steam reforming etc. Furthermore, metal catalysts are key in heterogeneous electrocatalysis, which is promising to improve the efficiency of electrolysis, fuel cells, CO₂ reduction, but also the synthesis of fine chemicals.⁴⁵ Therefore, analysis of the interaction of adsorbates with metallic surfaces can provide valuable insights for the design of novel, more efficient catalysts. Furthermore, when considering metal alloys, the difference between electronic and ensemble effects is a widely discussed topic.⁴⁶ BLW would allow to construct an electron localized (neutral) state of a given secondary element, clearly resolving the electronic (charge-transfer between the two metals) and ensemble effect.

The formulation of BLW is based on the assumption of doubly occupied orbitals, although extensions to spin-unrestricted systems exist,^{47;48} where the system is assumed to have a gap between the highest occupied molecular orbital (HOMO) and the lowest unoccupied molecular orbital (LUMO). Metallic electronic structures are, instead, characterized by a continuum of energy levels over the Fermi energy, leading to partially occupied states.

In this work, we propose an extension of the BLW approach based on a mean field approximation, to which we refer as Ensemble BLW-EDA. In the following, we first present the conventional SCF procedure for metallic systems, followed by the main notions of BLW. Then, we combine the two, extending BLW to metallic systems. Finally, we provide computational details and applications to the prototypical adsorption of molecules (H₂O, H₂S, C₂H₄ and CO) on Pt(111).

2.2 Methodology

2.2.1 Ensemble DFT

The ground-state of a metallic system, i.e., where the density of states around the Fermi level is continuous, cannot be described by a single quantum state. Rather, an ensemble of states is required, both for a physically sound description and for a smooth convergence of the

Chapter 2. Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions

wave function optimization in a self-consistent field (SCF) process. The different quantum states within the ensemble are weighted according to their probability in order to determine observables. Let \hat{A} be an operator, and Ω is the set of all possible quantum states of a given system. Then, the expectation value $\langle A \rangle_\Omega$ of the ensemble is:

$$\langle A \rangle_\Omega = \sum_{\chi \in \Omega} \langle \Psi_\chi | \hat{A} | \Psi_\chi \rangle p_\chi \quad (2.1)$$

where Ψ_χ is the wavefunction associated with the quantum state χ . Ψ_χ spans the ground- and single excited state determinants of the system. p_χ is the probability that the real system is in the quantum state χ , based on the energy of χ .

In the present context, we are interested in the reformulation of the 1-electron density operator $\hat{\rho}$. In the Hartree-Fock approximation, and in the case of orthonormal molecular orbitals ψ_i , $\hat{\rho}$ can be written as:

$$\hat{\rho} = \sum_{\text{occ } \psi_i \in \Psi} |\psi_i\rangle \langle \psi_i| \quad (2.2)$$

Dealing explicitly with ensembles is computationally inefficient and would require dedicated implementations. Since singly excited determinants do not overlap with the ground-state determinant, the same set of molecular orbitals can be used to construct all the relevant quantum states of the system. Therefore, the ensemble 1-electron density ($\langle \hat{\rho} \rangle_\Omega$) can be written in a convenient manner by considering the population of the orbitals (also known as occupation numbers, n_i) instead of the probability of quantum states:

$$\langle \hat{\rho} \rangle_\Omega = \sum_{\chi \in \Omega} p_\chi \hat{\rho}_\chi = \sum_{\chi \in \Omega} p_\chi \sum_{\psi_i \in \Psi_\chi} |\psi_i\rangle \langle \psi_i| = \quad (2.3)$$

$$= \sum_i |\psi_i\rangle \langle \psi_i| n_i = \sum_i |\sqrt{n_i} \psi_i\rangle \langle \sqrt{n_i} \psi_i| = \sum_i |\psi'_i\rangle \langle \psi'_i| \quad (2.4)$$

where i is the sum over all molecular orbitals, i.e., solutions of the Fock equation. In Eq. 2.4 we have introduced the rescaled orbitals:

$$|\psi'_i\rangle = \sqrt{n_i} |\psi_i\rangle \quad (2.5)$$

which are particularly convenient computationally. The occupation number n_i is related to the probability for the orbital ψ_i to be occupied. For finite temperatures, they are determined based on the energy ϵ_i of ψ_i , invoking Fermi-Dirac statistics:

$$n_i = \frac{1}{\exp(\frac{\epsilon_i - \mu}{kT}) + 1} \quad (2.6)$$

where μ is the Fermi level, k is Boltzmann's constant, and T is the (electronic) temperature. With these rescaled orbitals $|\psi'_i\rangle$, the mixed-state 1-electron density R_Ω can be easily computed.

$$R_\Omega = \mathbf{CNC}^\dagger \quad (2.7)$$

where \mathbf{C} is the molecular coefficient matrix and \mathbf{N} is a diagonal matrix of Fermi weights n_i according to Eq. 2.6. Yang and co-workers have generalized this reformulation to any method that can be reformulated based on the non-interacting Greens function, allowing to apply the orbital scaling to virtual orbitals and thus to compute correlation energies with MP2 or RPA for systems with fractional electrons.⁴⁹ With a mild approximation, these scaled orbitals can also be used for coupled cluster computations for systems with fractional electrons.⁵⁰

2.2.2 BLW formalism

The basic idea of the BLW formalism is to express the wave function of the system in terms of blocks of localized orbitals. In the present context, the orbitals are always localized on a subset of atoms, so that the BLW partitions the atom-centered basis functions in mutually exclusive blocks and imposes thereby a localization of the wave function in Hilbert space.

Formally, a block B_m is defined as a set of N_m basis functions $\{\phi_1^m \dots \phi_{N_m}^m\}$, such that each basis function is associated with exactly one block, i.e. $\sum_m N_m = N$, where N is the total number of basis functions of the system. In this work, the set of basis functions $\{\phi_1^m \dots \phi_{N_m}^m\}$ associated with the block B_m is the union of *all* N_j basis set functions $\{\phi_1^j \dots \phi_{N_j}^j\}$ used to describe the j_m atoms of B_m , so that $\sum_m j_m$ equals to the total number of atoms in the system. Similarly, the electrons of the system are assigned to a given block. In our case, all blocks are neutral.

An ALMO is defined as a linear combination of basis functions associated with the same block. Therefore, an orbital ψ_{i^m} pertaining to block B_m is written:

$$|\psi_{i^m}\rangle = \sum_{\nu} C_{\nu,i}^m |\phi_{\nu}^m\rangle \quad (2.8)$$

where \mathbf{C} is the orbital coefficient matrix for the entire system, while \mathbf{C}^m is the matrix of block B_m . We have used greek letters to label atomic orbital basis functions and i^m indexes ALMOs of a given block B_m .

The global orbital coefficient matrix \mathbf{C} , which spans all molecular orbitals and, thus, has

Chapter 2. Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions

dimensions of $N \times N$, has a block-diagonal structure:

$$\mathbf{C} = \begin{pmatrix} \mathbf{C}^1 & 0 & \cdots & 0 \\ 0 & \mathbf{C}^m & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{C}^{\Sigma_m} \end{pmatrix} \quad (2.9)$$

where \mathbf{C}^m is the orbital coefficient matrix restricted to the block B_m . The orbital coefficients $C_{v,i}^m$ are variationally optimized according to the local diagonalisation-based SCF algorithm developed by Stoll¹⁹ and implemented in CP2K as the first stage of the two-step ‘‘ALMO-SCF’’ scheme.⁵¹

In general, locality and orthogonality constraints cannot be both satisfied simultaneously.⁵² Therefore, ALMOs are inherently non-orthogonal between blocks, although they can be kept orthonormal within a block without loss of generality. In this study, we will work under this assumption. Hence, the ALMO overlap matrix (σ) has identity-like diagonal blocks, but non-zero entries for overlaps between blocks.

In practice, the so called reciprocal (or bi-orthogonal) occupied orbitals $|\tilde{\psi}_{i^m}\rangle$ are defined such that $\langle \tilde{\psi}_{i^m} | \psi_{j^l} \rangle = \delta_{i^m j^l}$:

$$|\tilde{\psi}_{i^m}\rangle = \sum_{l,j} |\psi_{j^l}\rangle \sigma_{i^m j^l}^{-1}, \quad \tilde{\mathbf{T}} = \mathbf{T} \sigma^{-1} \quad (2.10)$$

where l goes over all blocks and j^l indexes the occupied orbitals of block B_l , $\tilde{\mathbf{T}}$ is the coefficient matrix of reciprocal occupied orbitals and \mathbf{T} is the occupied part of Eq. 2.9. To be explicit, $|\tilde{\psi}_{i^m}\rangle$ is expanded in terms of *all* basis functions of the system and not only of block B_m .

Reciprocal orbitals enable one to rewrite the Fock equations. As Stoll has demonstrated, the self-consistent solution of projected eigenvalue equations for each block is equivalent to finding variationally optimal ALMOs.

$$\hat{F}^m |\psi_i^m\rangle = \epsilon_i^m |\psi_i^m\rangle, \quad \hat{F}^m = \left(\hat{\mathbb{1}} - \hat{\rho} + \hat{\rho}^{m\dagger} \right) \hat{F} \left(\hat{\mathbb{1}} - \hat{\rho} + \hat{\rho}^m \right) \quad (2.11)$$

$$\hat{\rho}^m = \sum_{i \in B_m} |\tilde{\psi}_{i^m}\rangle \langle \psi_{i^m}| = \sum_{l,\xi,\nu} R_{\xi\nu}^m |\phi_\xi^l\rangle \langle \phi_\nu^m| \quad (2.12)$$

where $\hat{\mathbb{1}}$ is the identity matrix, \hat{F} is the conventional Fock operator, \hat{F}^m is the Fock operator projected on block B_m , $\hat{\rho}$ is the density operator of the entire system, $\hat{\rho}^m$ is the non-Hermitian operator that represents the density of the block defined by using only ALMOs of block B_m , and R^m is the associated block density matrix. ν is a basis function of block B_m (see Eq. 2.8),

while ξ indexes the atom centered basis function of blocks B_l .

At each SCF iteration, these projected Fock equations are solved independently for each block, but coupled to each other between successive iterations due to the projection operators. In other words, each block B_m is optimized in the environment generated by all other blocks.

Since in mixed-state theory it is common to deal with orthonormal orbitals, we briefly discuss the connection between the use of reciprocal orbitals and a Lowdin orthonormalization:

$$|{}^L\psi_i\rangle = \sum_{\text{occ } j} |\psi_j\rangle \sigma_{ij}^{-\frac{1}{2}}, \quad {}^L\mathbf{T} = \mathbf{T}\sigma^{-\frac{1}{2}} \quad (2.13)$$

where i and j are general indexes, $\sigma^{-\frac{1}{2}}$ is the square root inverse of the overlap matrix and ${}^L\mathbf{T}$ is the Lowdin orbital coefficient matrix. The 1-electron density is equivalently expressed through the Lowdin orthonormalized molecular orbitals and the use of reciprocal orbitals, the latter avoiding the expensive computation of $\sigma^{-\frac{1}{2}}$:

$$\hat{\rho} = \sum_{\text{occ } i} |{}^L\psi_i\rangle\langle{}^L\psi_i| = \sum_{\text{occ } i} |\tilde{\psi}_i\rangle\langle\psi_i| = \sum_{\text{occ } i\nu, \xi} \tilde{T}_{\nu i} |\phi_\nu\rangle\langle\phi_\xi| T_{\xi i} = \sum_{\nu, \xi} \sum_{\text{occ } i} T_{\xi i} \tilde{T}_{i\nu}^\dagger |\phi_\nu\rangle\langle\phi_\xi| = \sum_{\nu, \xi} R_{\nu\xi} |\phi_\nu\rangle\langle\phi_\xi| \quad (2.14)$$

$$\mathbf{R} = {}^L\mathbf{T}{}^L\mathbf{T}^\dagger = \mathbf{T}\tilde{\mathbf{T}}^\dagger = \mathbf{T}(\mathbf{T}\sigma^{-1})^\dagger = \mathbf{T}\sigma^{-1}\mathbf{T}^\dagger \quad (2.15)$$

where \mathbf{R} is the conventional 1-electron density matrix. \mathbf{R} can then be fed to standard routines to determine the electron density in real space, compute gradients and so forth.

2.2.3 Ensemble BLW

Exact Ensemble Formulation

In order to adapt the BLW formalism to an ensemble formulation, we need to adapt the density matrices \mathbf{R} or, equivalently, the construction of the reciprocal orbitals (Eq. 2.10). In other words, we apply the general formula (Eq. 2.3) for the computation of an ensemble density matrix \mathbf{R}_Ω to non-orthogonal molecular orbitals:

$$\mathbf{R}_\Omega = \sum_{\chi \in \Omega} p_\chi \mathbf{T}_\chi = \sum_{\chi \in \Omega} p_\chi \mathbf{T}_\chi \sigma_\chi^{-1} \mathbf{R}_\chi^\dagger = \sum_{\chi \in \Omega} p_\chi \mathbf{T}_\chi (\mathbf{T}_\chi^\dagger \mathbf{S} \mathbf{T}_\chi)^{-1} \mathbf{T}_\chi^\dagger \quad (2.16)$$

where p_χ is the probability that the real system is in the quantum state χ , \mathbf{R}_χ is the density matrix of this quantum state, \mathbf{T}_χ is the occupied orbital coefficient matrix associated with the wavefunction Ψ_χ , σ_χ^{-1} is the overlap matrix of the quantum state χ , and \mathbf{S} is the basis set

Chapter 2. Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions

function overlap matrix which is common for all quantum states, since it is a property of the geometry and basis set.

Let us consider a general orbital coefficient matrix \mathbf{C} containing all the localized orbitals used to construct every \mathbf{T}_χ , the occupied orbital matrix of quantum state χ . Therefore, one can construct any \mathbf{T}_χ from \mathbf{C} :

$$\mathbf{T}_\chi = \mathbf{C}\Delta_\chi = \mathbf{C} \begin{pmatrix} \delta_{1\chi} & 0 & \cdots & 0 \\ 0 & \delta_{2\chi} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \delta_{n+k\chi} \end{pmatrix}, \quad \delta_{i\chi} = \begin{cases} 1 & \text{if } \psi_i \in \Psi_\chi, \\ 0 & \text{otherwise.} \end{cases} \quad (2.17)$$

where Δ_χ can be seen as a rescaling matrix with a dimension of $N \times N$.

Combining Eq. 2.17 with Eq. 2.16, we obtain the following reformulation:

$$\mathbf{R} = \sum_{\chi \in \Omega} p_\chi \mathbf{R}_\chi = \sum_{\chi \in \Omega} p_\chi \mathbf{T}_\chi \sigma_\chi^{-1} \mathbf{T}_\chi^\dagger = \sum_{\chi \in \Omega} p_\chi \mathbf{C} \Delta_\chi \sigma_\chi^{-1} \Delta_\chi \mathbf{C}^\dagger = \mathbf{C} \left(\sum_{\chi \in \Omega} p_\chi \Delta_\chi \sigma_\chi^{-1} \Delta_\chi \right) \mathbf{C}^\dagger = \mathbf{C}^\Omega \sigma^{-1} \mathbf{C}^\dagger \quad (2.18)$$

where ${}^\Omega \sigma^{-1} = \sum_{\chi \in \Omega} p_\chi \Delta_\chi \sigma_\chi^{-1} \Delta_\chi$.

Equation 2.18 requires σ_χ^{-1} to be computed for each state involved separately, i.e., it is not an efficient reformulation of Eq. 2.16. This contrasts with the canonical case, where Eq. 2.4 provides an efficient reformulation of Eq. 2.3, since the different quantum states involved are orthonormal among each other.

If we have k orbitals in addition to the n formally doubly occupied orbitals, then the maximum number $|\Omega|$ of quantum states to evaluate Eq. 2.17 is:

$$|\Omega| = \binom{n+k}{n} = \frac{(n+k)!}{n!k!} \quad (2.19)$$

In the worst case, $k = n$, $|\Omega| \approx \frac{2^{2n}}{\sqrt{\pi n}}$, i.e., the number of states to be considered is exponential in n . Hence, such a method is not applicable to sizable systems (e.g. a little more than 7×10^{20} states for the Pt(111) surface investigated in the later sections).

Why can there not be a simple reformulation of ensemble BLW, just like in the canonical case? Considering Eq. 2.14, we can understand that ALMOs are treated as if they would be orthonormalized when computing the density matrix. Therefore, scaling them by their occupation number is not an option, since this lack of normalization will simply be offset by the corresponding σ^{-1} . A second point of interference when aiming at a simplified ensemble

description comes from the interdependence of blocks. Let us consider a two block system. According to Eq. 2.10, the reciprocal orbitals of block 1 depend on the occupied orbitals of all blocks. Hence, in general (i.e., when the orbitals between blocks overlap) varying the occupations in block 2, enters as a varying environment for block 1 (and *vice versa*). In other words, the different quantum states in block 1 are subjected to different environments, depending on the quantum states of block 2. To put this in mathematical terms, let us consider two quantum states χ and χ' , with associated wave-functions Ψ_χ and $\Psi_{\chi'}$ and three different orbitals ψ_1 , ψ_2 and ψ_3 (considered doubly occupied) such that $\{\psi_1, \psi_2\} \subseteq \Psi_\chi$ and $\{\psi_1, \psi_3\} \subseteq \Psi_{\chi'}$, but $\psi_3 \notin \Psi_\chi$ and $\psi_2 \notin \Psi_{\chi'}$. If we would like to use the same set of ALMOs for describing Ψ_χ and $\Psi_{\chi'}$, then we would like $\tilde{\psi}_1$ to be equal to $\tilde{\psi}_1'$ (respectively the reciprocal orbital of ψ_1 computed in the quantum state χ and χ') so that Eq. 2.18 can be simplified in full analogy with the situation when using orthonormalized orbitals. However the reciprocal orbital $\tilde{\psi}_1$ in the quantum state χ does not change by replacing the orbital ψ_2 by ψ_3 , if and only if ψ_1 does not overlap with neither ψ_2 nor ψ_3 , and there exists no $\psi_i \notin \{\psi_1, \psi_2, \psi_3\}$ in any quantum state $\chi_j \in \Omega$ such that both ψ_1 and ψ_2 (or ψ_3) overlap with ψ_i :

$$|\tilde{\psi}_1\rangle = |\tilde{\psi}_1'\rangle \Leftrightarrow \begin{cases} \langle \psi_1 | \psi_2 \rangle = 0 \\ \wedge (\forall \chi \in \Omega, \nexists \psi_i \in \Psi_\chi \setminus \{\psi_1, \psi_2\}, \psi_1 \in \Psi_\chi \rightarrow (\langle \psi_1 | \psi_i \rangle \neq 0 \wedge \langle \psi_2 | \psi_i \rangle \neq 0)) \\ \wedge (\text{idem for } \psi_3 \text{ instead of } \psi_2) \end{cases} \quad (2.20)$$

Therefore, the occupation-state dependency of the orthonormalized orbitals is always present, except when the added or removed orbitals does not overlap with the rest of the system. Due to this occupation-state dependency of the orthonormalized orbitals, they cannot be used to construct an ensemble density matrix by rescaling them, in contrast to the canonical case.

As a conclusion, since the orbitals are non-orthogonal, when an orbital occupation is modified the whole system has to re-adapt. Note that this conclusion applies to the use of any non-orthogonal orbitals and not only to ALMOs.

Mean-field approximation to Ensemble BLW

The condition to formulate a computationally tractable approximation to ensemble BLW is that the contribution of each orbital to the ensemble density can be computed only once per SCF iteration and can then be weighted by the probability that the real system is in a quantum state containing this orbital (or equivalently, the probability that this orbital is occupied in the real system). In other words, we need to devise a scheme in which we have a common overlap matrix σ for all quantum states involved, i.e., we generate an average interaction field (over all quantum states) that is applied to every quantum state. Furthermore, we require that the scheme is equivalent to standard ensemble theory for a single block. This implies that Eq. 2.6

Chapter 2. Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions

is applied to each block separately (with the block specific chemical potential μ^m), so that the total number of electrons in each block remains an integer. With these two requirements, we arrive at a mean-field approximation to the overlap matrix (${}^\alpha\sigma$):

The interaction (i.e., overlap) between $|\psi_{im}\rangle$ and $|\psi_{jl}\rangle$ from any block, is rescaled by $\sqrt{n_i}\sqrt{n_j}$, while the self-overlap remains unmodified. This rescaling provides an approximate overlap matrix ${}^\alpha\sigma$:

$${}^\alpha\sigma_{ij} = \begin{cases} \langle\psi_{im}|\psi_{im}\rangle = 1 \\ \langle\psi_{im}|\psi_{jl}\rangle\sqrt{n_i}\sqrt{n_j} \end{cases} \quad (2.21)$$

Hence, the density matrix ${}^\alpha\mathbf{R}_\Omega$ can be written:

$${}^\alpha\mathbf{R}_\Omega = \mathbf{T}' {}^\alpha\sigma^{-1} \mathbf{T}'^\dagger \quad (2.22)$$

where \mathbf{T}' is the rescaled orbital coefficient matrix, in close analogy to Eq. 2.5, for the canonical case.

We use $|\psi_{im}^s\rangle$ to denote the presence of fractionally occupied ALMOs, i.e., the ones that necessitate the use of ${}^\alpha\sigma$. $|\psi_{im}^s\rangle$ have rescaled interactions (overlap), except with themselves. Therefore, we call them “selfish orbitals”. Unlike the canonical rescaled orbitals, selfish orbitals cannot be considered “shrunk”, they just interact less with their environment:

$$\forall j \neq i \quad \langle\psi_{im}^s|\psi_{jl}^s\rangle = \sqrt{n_i}\sqrt{n_j}\langle\psi_{im}|\psi_{jl}\rangle; \text{ but } \langle\psi_{im}^s|\psi_{im}^s\rangle = \langle\psi_{im}|\psi_{im}\rangle \quad (2.23)$$

where j goes over ALMOs of all blocks.

In practice, at each SCF iteration the orbital coefficients \mathbf{T} are scaled according to the (updated) occupation number, yielding \mathbf{T}' . With these scaled coefficients, the overlap matrix σ is computed. Then, the diagonal of σ is set to unity to obtain ${}^\alpha\sigma$. The density matrix is computed using Eq. 2.22 In brief, except for the use of \mathbf{T}' and ${}^\alpha\sigma$, the SCF procedure by Stoll is not modified.

Please note that for non-interacting systems (blocks that are far apart), the mean-field approximation reverts back to the canonical answer. Similarly, if the orbital occupations are either 0 or 2 (0 K limit of a system with a non-vanishing gap), our approximation provides the regular BLW result.

In our implementation in CP2K, we exploit the eigenvalue based optimization by Stoll^{15;19}, which has been implemented in CP2K^{51;53} and allows a variational optimization of the ensemble BLW. The computational cost of the extension to partially occupied orbitals is negligible thanks to the use of selfish orbitals and the additional storage used is limited to an array

containing the orbital energies.

As discussed in Appendix A, the error introduced by the mean-field approximation turns out to lead to losses of electrons in the Ensemble BLW. For example, for H₂O up to 0.005 electrons are lost when the molecule is closest to the metal surface. The energetic consequence on the polarization energy of this electron loss can be estimated to be up to 0.5 kcal mol⁻¹ assuming a metallic system with a typical workfunction around 5 eV. As a result, the charge-transfer energy, which does not suffer from this electron loss, is slightly overestimated compared to the polarization energy. Given that water, which is the least strongly adsorbed molecule studied herein, has an adsorption energy of -9.7 kcal mol⁻¹, such an error is deemed acceptable. In all cases no, or a small loss of electrons is encountered, never a gain in electrons. This can be rationalized in analogy to Hartree-Fock, where the average electron repulsion overestimates the actual electron repulsion. Hence, our mean-field approximation is an upper bound to exact ensemble BLW. Since ensemble DFT formulation is variational with respect to the electronic free energy (i.e., when accounting for the entropy related to the fractionally occupied orbitals), our Ensemble BLW-EDA defines the charge-transfer through a variational principle for the interaction of adsorbates with metallic surfaces.

2.3 Computational Details

The adsorbed structures were optimized with VASP 5.4.1^{54;55} using periodic boundary conditions applying the re-optimized Perdew, Burke and Ernzerhofer (PBE) functional to make it compatible with the non-local van der Waals (vdW) functional, i.e. the optPBE-vdW⁵⁶ functional. This functional has been found to be most accurate for adsorption energies on Pt(111)⁵⁷, together with PBE-dDsC⁵⁸, an other dispersion corrected density functional. Nonetheless, optPBE-vdW is available in CP2K, the code that is used for all Ensemble BLW-EDA computations. An energy cutoff of 400 eV is chosen for the expansion of the plane-wave basis set. The electron-ion interactions are described by the PAW formalism.^{59;60} The p(6×6) unit cell is built from bulk platinum (2.821 Å nearest neighbor distance) with four metallic layers. Additional tests regarding the need for K-point sampling reveal that the Γ -point optimized geometries are very close to the ones obtained with a 3 × 3 × 1 K-point mesh. The adsorption energy at the Γ -point is overestimated by 1-5 kcal mol⁻¹, corresponding to up to 10% at the equilibrium distance (see Table A.1). This accuracy is deemed acceptable for the current purpose, where the relative importance of different interaction energy components and their evolution as a function of the surface – adsorbate distance is analyzed. The out-of-plane vector of the unit cell was chosen to be ~ 23 Å to achieve a negligible interaction between periodic images.

The molecular orbitals were represented by a double- ζ Gaussian basis set with one set of polarization functions, called DZVP-MOLOPT-SR-GTH for both BLW-EDA and BSSE corrected

Chapter 2. Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions

SCF DFT simulations.⁶¹ A cutoff of 400 Ry was used to describe the electron density. The exchange-correlation (XC) energy was approximated with the optPBE-vdW⁵⁶ functional. The Brillouin zone was described at the Γ -point. Goedecker, Teter and Hutter (GTH) pseudo-potentials⁶² based on the PBE functional were used to describe the interactions between the valence electrons and the ionic cores, and the electronic smearing was approximated by a Fermi-Dirac distribution at 300 K, applied to all computations.

Figure A.1 provides a comparison between the total interaction energy (ΔE_{int} , vide infra, Eq. 2.24) as computed by standard KS-DFT with VASP and CP2K. In both codes, we compare a 10 and a 18 valence electron pseudopotential. The 18 valence electron potential of CP2K gives interaction energies very close to the one provided by VASP, independently if the latter uses 10 or a 18 valence electrons.

2.3.1 Energy Decomposition Analysis

The newly extended EDA scheme is now applied to charge transfer analysis of adsorption at a metallic surface. We select three prototypical couples of systems (see Figure 2.1), aiming at describing different types of bonding. The first couple is H_2O and H_2S , for which no strong bond formation is expected, although H_2S is interacting more strongly with Pt(111) than water. The second couple compares the di- σ and π adsorption modes of ethylene, while the third couple investigates the difference between CO adsorbed on fcc and top sites.

For all these systems, we have computed the interaction energy as a function of the distance between the surface and the adsorbate. All coordinates are fully optimized, except the z-coordinate of the heavy atoms closest to the surface and the two bottom layers that are kept fixed in their bulk position. We define the total interaction energy ΔE_{tot} as:

$$\Delta E_{tot} = E_{system} - E_{frag1,opt} - E_{frag2,opt} - \Delta E_{BSSE} \quad (2.24)$$

Where E_{system} is the standard KS-SCF energy of the full system, $E_{frag,opt}$ are the corresponding energies of the freely optimized fragments. Since the BLW is only defined in a localized basis set, we have to correct for the basis set superposition error (BSSE), which we do according to the counterpoise procedure of Boys and Bernardi⁶³, giving rise to the energy correction ΔE_{BSSE} . Note, however, that BSSE only affects the charge-transfer term, as all other terms are evaluated using exclusively the fragment basis set.

For the energy decomposition analysis, each system is divided into two blocks: a metallic block containing the metal surface, and an adsorbate block containing the adsorbed molecule. As common in BLW-EDA,^{22;23;64} we decompose the total interaction energy ΔE_{tot} into the following terms: preparation or deformation (ΔE_{deform}), the frozen energy term (ΔE_{frozen}) that describes the interaction of the two isolated fragment densities brought together and




Figure 2.1 – Top: H₂O (left) and H₂S (right) top adsorption on Pt(111). Middle: Ethylene adsorption in a di-σ (left) and π adsorption mode. Bottom: CO on fcc (left) and top (right) adsorption sites. All images refer to the equilibrium distance with respect to the surface.

Chapter 2. Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions

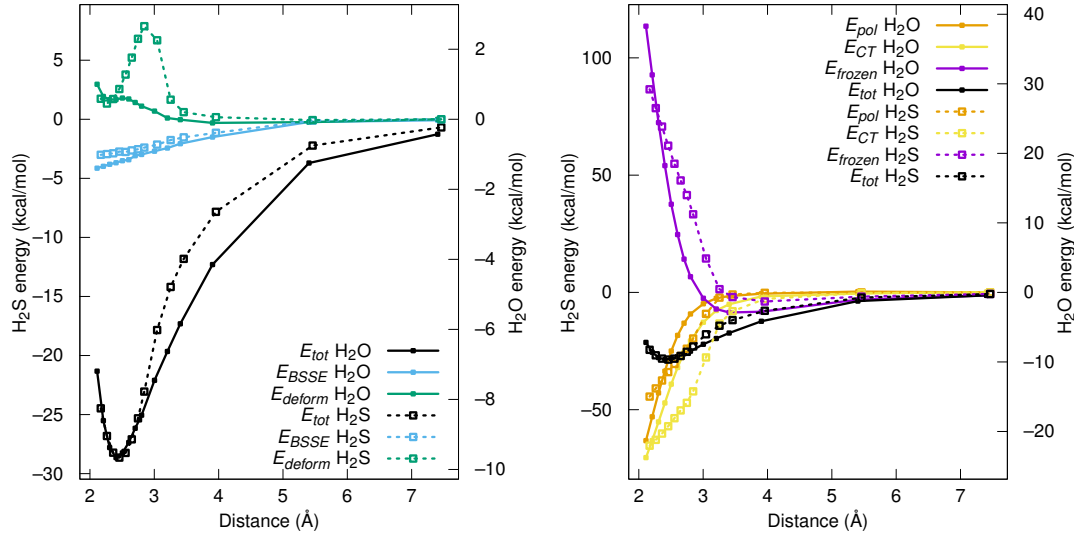


Figure 2.2 – Energy decomposition analysis for H₂O (full lines, right y axis) and H₂S (broken lines, left y-axis). The total interaction energy is compared to the BSSE and the deformation energy on the left, while the right quantifies the frozen monomers, polarization and charge transfer interaction energy. Note that the y-scales are aligned in such a way that the equilibrium interaction energy and the zero interaction energy are aligned for both systems, see Appendix A for separated Figures.

covers electrostatic interaction and Pauli repulsion²² as well as dispersion interactions⁶⁴, polarization (ΔE_{pol}) energy, obtained by variationally optimizing the BLW, and, finally, the charge transfer (ΔE_{CT}) interaction that includes covalent bond formation.

$$\Delta E_{tot} = \Delta E_{deform} + \Delta E_{frozen} + \Delta E_{pol} + \Delta E_{CT} \quad (2.25)$$

These terms can be expressed as:

$$\Delta E_{deform} = E_{frag1,sys} - E_{frag1,opt} + E_{frag2,sys} - E_{frag2,opt} \quad (2.26a)$$

$$\Delta E_{frozen} = E_{guess} - E_{frag1,sys} - E_{frag2,sys} \quad (2.26b)$$

$$\Delta E_{pol} = E_{BLW} - E_{guess} \quad (2.26c)$$

$$\Delta E_{CT} = E_{system} - E_{BLW} - \Delta E_{BSSE} \quad (2.26d)$$

where $E_{frag,sys}$ corresponds to the energy of a fragment in its final geometry adopted in the presence of the other fragment. E_{guess} is the systems energy obtained by a superposition of the fragment density matrices, which corresponds to the “frozen” density interaction energy approximation. E_{BLW} is the energy obtained by Ensemble BLW.

2.4 Results and Discussion

As a first application of Ensemble BLW-EDA to metal surfaces, we compare the adsorption of H_2O and H_2S on Pt(111). As shown Figure 2.1 for both molecules the most favorable adsorption configuration is atop, with the molecular plane approximately parallel to the surface. The total interaction energy as a function of the Pt–O or Pt–S distance shows that despite the nominally larger non-covalent radius of sulfur with respect to oxygen (1.8 and 1.5 Å according to Bondi⁶⁵), the minimum is found around 2.4 Å in both cases. This can be rationalized by the overall stronger interaction of Pt– H_2S (29 kcal mol⁻¹) compared to Pt– H_2O (9 kcal mol⁻¹), which allows S to approach the surface more closely than O with respect to its size. The BSSE contribution to the interaction energy is similar for H_2O and H_2S (roughly 10% of the total interaction energy). Therefore, the comparison between the two systems could be performed with a 10% accuracy without the BSSE correction, for instance for molecular dynamics simulations of H_2S dissolved in water. The deformation energy of the two systems show different patterns. While the energetic cost to deform water/Pt(111) to its optimal geometry is very small (at most 0.5 kcal mol⁻¹), H_2S /Pt(111) undergoes a deformation up to 7 kcal mol⁻¹ at a distance that is only slightly longer than the equilibrium distance, while it drops at shorter distances to about 1 kcal mol⁻¹. Analyzing the origin of the deformation energy, we identify the deformation of the metal surface as the major contributor. As shown in Figure A.2, the Pt atom on which the adsorbate is adsorbed, is “pulled out” of the surface at intermediate distances. After a certain elongation, it goes back to the original position, almost as if the spring had been overstretched and thus broken. Moving to the interaction energy components of the fragments after paying the preparation (deformation) energy (right hand graph of Figure 2.2), we first observe that the “frozen” term, i.e., the energy cost (or gain) of putting the fragments together without electronic density changes, follows the expected increase in atomic size when replacing oxygen by sulfur: The energy of assembly becomes positive for distances below 2.8 and 3.2 Å for H_2O and H_2S , respectively. At larger distances, the interaction energy is already negative without any further electronic relaxation. This is due to the subtle balance between Pauli repulsion and attraction by dispersion interactions. Note that electrostatic interactions are also included at this stage, but since the isolated metal surfaces do not have significant electrostatic moments (dipole, quadrupole etc.), they barely contribute in the present systems. Allowing the electron density to relax in the presence of the other fragment, but prohibiting any charge transfer (or direct orbital interaction), affords the stabilizing polarization energy E_{pol} . H_2S has a 2.7 times larger polarizability than H_2O (26.7 and 9.8 a.u.³, respectively,⁶⁶). Indeed, there is roughly a factor of three between E_{pol} for H_2S compared to H_2O at all relevant distances (> 2.3 Å), in close agreement with the corresponding factor for the total interaction energy. The charge transfer between the metal surface and the adsorbate, is, however, much more important for H_2S than for H_2O . This contrasting energy decomposition can be rationalized recalling that sulfides are softer bases than oxides and that platinum is a soft acid. Furthermore, interactions with sulfur can benefit from stabilization through hypervalence by increasing the ionic character of

bonds.⁶⁷

As a conclusion, the stabilizing interaction of a single water molecule with Pt(111) is due to equal amounts to polarization and to charge transfer. In other words, even in this case of rather weak physisorption, the use of a polarizable force field is unlikely to be enough to capture the physics of the interaction energy accurately. However, a good approximation can be obtained since the two components are well correlated, i.e., counting the polarization energy twice, BLW could capture a reasonable approximation of the total interaction energy. This insight might help the development of more accurate Pt/H₂O force fields.^{68;69} The situation for H₂S is more challenging to approximate: charge-transfer dominates the interaction energy, which also induces significant deformations of the surface. Hence, already for H₂S adsorption, explicit terms to mimic orbital/charge-transfer interactions are required. The situation can be simplified by excluding surface deformation. Given the scarcity of accurate force fields for the deformation of metal surfaces, this is the preferred setup anyway. Figure A.3 shows that on an ideal surface the deformation energy is negligible and the charge-transfer and polarization energies are now as smooth as for water. Furthermore, on an ideal surface, a very similar relative importance of the various contributions is obtained for H₂O and H₂S. Since E_{tot} is not very different when using an ideal Pt(111) surface, H₂S adsorption can be described without taking into account the surface deformation with only a small loss in accuracy. Given that E_{CT} is most strongly affected, this analysis also identifies the charge-transfer as the origin of the deformation.

On a more technical level, we have also investigated the influence of the electronic smearing temperature on the results. In Figure A.4 we show that increasing the electronic temperature from 300 K to 1000 K does not visibly affect the results. Even at 2000 K only changes of 3-4 kcal mol⁻¹ are observed, leaving the relative importance of different terms unchanged. As detailed in Appendix A, the small changes observed can be rationalized keeping in mind that higher electronic temperatures lead to the occupation of orbitals lying above the Fermi level.

The deformation energy observed for H₂S is small compared to the one obtained for ethylene adsorption (see Figure 2.3). This does not come as a surprise, given that the adsorbed geometry (Figure 2.1) suggests a re-hybridization of the sp² carbon atoms. This significant geometric rearrangement is best associated with a “state crossing”: sp² at long distances but ~ sp³ at short distances. This state-crossing is non-continuous, as best seen in the deformation energy, although other interaction energy components show a similar discontinuity. As mentioned above, the deformation is mainly due to the deformation of the metal surface, where the Pt atom is partially lifted out of the surface (see Figure A.2). The BSSE is again on the order of 10% of the total interaction energy around the minimum and the total interaction energy is, with 30 and 37 kcal mol⁻¹ for π and di- σ , respectively, similar. Furthermore, the larger geometric deformation (H-C-C-H dihedral angles of 132 and 156°) for the di- σ mode compared to the π mode is nicely captured by the deformation energy. Similar to H₂S, the maximum deformation




Figure 2.3 – Energy decomposition analysis for di- σ (broken lines, left y-axis) and π (full lines, right y axis) C_2H_4 . The total interaction energy is compared to the BSSE and the deformation energy on the left, while the right quantifies the frozen monomers, polarization and charge transfer interaction energy.

energy is obtained when the molecule is close enough to the surface to actually “feel” it, but not close enough to form strong covalent bonds. The deformation and the total interaction energy reflect quite well that the π mode approaches the surface less closely than the di- σ mode (Pt–C distance of 2.20 Å in π and 2.13 Å in di- σ). The right hand side of the graph provides a rationalization: the Pauli repulsion (contained in E_{frozen}) is longer-ranged for the π bond at the top of a Pt atom, than between the bridge site and carbon atoms that are not strictly on top of Pt atoms. The same observation also applies to the polarization and charge-transfer energy. However, the relative importance of repulsion, polarization and charge-transfer is surprisingly similar for both modes, demonstrating that the nature of the bonding interaction is the same. Nevertheless, the di- σ mode maximizes this bonding type which is characterized by $\sim 40\%$ polarization energy and $\sim 60\%$ charge-transfer, very similar to the interaction of H_2S with the same surface.

The last example we are discussing here is the classic case of CO adsorption on Pt(111). One of the challenging questions is whether CO is adsorbed on top or on bridge/hollow sites and how this evolves as a function of the surface charge.⁷⁰ The following analysis investigates exclusively the bonding patterns at the respective adsorption sites, without drawing any conclusions, neither on the reliability of the chosen density functional, nor on the actual site preference of CO on Pt(111).

In Figure 2.4 the data for CO adsorption on top and fcc sites is compared. At our level of theory, the fcc adsorption site is slightly preferred (-52 vs -49 kcal mol⁻¹). However, these graphs show very clearly the contrasting behavior of CO on these two sites: The top site

Chapter 2. Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions




Figure 2.4 – Energy decomposition analysis for top (full lines, right y axis) and fcc (broken lines, left y-axis) CO on Pt(111). The total interaction energy is compared to the BSSE and the deformation energy on the left, while the right quantifies the frozen monomers, polarization and charge transfer interaction energy.

experiences a significantly stronger Pauli repulsion, resulting in a longer equilibrium distance (2 Å compared to 1.5 Å). The deformation energy, on the other hand, is slightly lower at the equilibrium distance on the top site than on the fcc site, although the barrier for the state-crossing is roughly equal ($\sim 30 \text{ kcal mol}^{-1}$), shifted by 0.75 Å to longer distances for the top mode. Regarding the decomposition of the interaction energy we can first note that the Pauli repulsion at long distances is lower for the fcc mode compared to the top adsorption mode, but at shorter distances rises more quickly for the fcc site. Second, in contrast to the case of C_2H_4 , the ratio of polarization and charge-transfer energy is quite different for the two adsorption modes. Somewhat surprisingly, the fcc mode shows a similar pattern as H_2O adsorption, with E_{pol} and E_{CT} being roughly equal over the entire range of distances. In other words, polarization contributes more to the total binding than for the adsorption of C_2H_4 . While charge-transfer is relatively more important for the top adsorption mode, it is still far from reaching the importance it takes for ethylene adsorption. We suggest that this “lower than expected” importance of charge-transfer for the adsorption of CO on Pt(111) is due to the intricate donation and back-donation involved in CO adsorption, which can be seen as a particularly strong polarization effect. In other words, it is the polarization which is particularly strong for CO adsorption, not the charge-transfer that is small: this dative bond is the strongest bond investigated herein and orbital interactions are without any doubt key for its accurate description. However, our analysis demonstrates that the top site is more sensitive to charge-transfer than adsorption on the hollow site. This is in good agreement with previous reports,⁷⁰ but it is the first time that such a difference is clearly traced back to charge-transfer and disentangled from polarization effects.

2.5 Conclusions

In this work, we have presented the extension of the BLW formalism to systems that require fractionally occupied orbitals such as metals at finite temperature. Since the computational cost of the rigorous combinatorial formulation is unpractical, our extension relies on a mean-field approximation to ensemble BLW in the context of mixed-state theory. This approximation is based on a new concept: selfish orbitals. Selfish orbitals are normalized, but their interactions with other orbitals are scaled down according to their occupation number. This mean-field approximation is exact in the case of infinitesimal smearing (i.e. BLW with a gap) and/or 1-block systems (i.e. ensemble KS-DFT). Furthermore, in numerical examples studied, the estimated error in the polarization energy is roughly $0.5 \text{ kcal mol}^{-1}$.

The method has been implemented in CP2K and numerical applications have shown that H_2S is much more strongly bound to Pt(111) due to charge-transfer, while the contribution of polarization is on a similar level as for H_2O . The chemisorption of ethylene on Pt(111) is, however, dominated by the charge-transfer and the two modes (di- σ and π) show the same type of bonding, with proportions of different interaction energy components close to the ones of H_2S . CO adsorption, on the other hand, is shown to depend significantly on the adsorption site: although the bond is five times stronger than for H_2O , the hollow-site adsorption is characterized by a roughly equal contribution of polarization and charge-transfer, just like water adsorption. On the top site charge-transfer is somewhat more important, but even in this case its importance is relatively smaller than for H_2S or C_2H_4 . We ascribe this particularity to the donation/back-donation of CO, which can be seen as a particularly strong polarization contribution.

As demonstrated herein, BLW-EDA provides a convenient (i.e. easy to use) tool, providing polarization and charge-transfer energies at the metal interface, which is expected to provide valuable insight for catalysis and the understanding of the metal/gas and metal/liquid interface.

Bibliography

- [1] Heitler, W.; London, F. Wechselwirkung neutraler Atome und homo-polare Bindung nach der Quantenmechanik. *Z. Phys. A* **1927**, *44*, 455–472.
- [2] Ruedenberg, K. The Physical Nature of the Chemical Bond. *Rev. Mod. Phys.* **1962**, *34*, 326–376.
- [3] Bader, R. F. W. *Atoms in Molecules: A Quantum Theory*. **1990**,
- [4] Reed, A. E.; Curtiss, L. A.; Weinhold, F. Intermolecular interactions from a natural bond orbital, donor-acceptor viewpoint. *Chem. Rev.* **1988**, *88*, 899–926.

Chapter 2. Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions

- [5] Gonthier, J. F.; Steinmann, S. N.; Wodrich, M. D.; Corminboeuf, C. Quantification of “fuzzy” chemical concepts: a computational perspective. *Chem. Soc. Rev.* **2012**, *41*, 4671–4687.
- [6] Tsipis, A. C. DFT challenge of intermetallic interactions: From metallophilicity and metallaromaticity to sextuple bonding. *Coord. Chem. Rev.* **2017**, *345*, 229–262.
- [7] Pastorczak, E.; Corminboeuf, C. Perspective: Found in translation: Quantum chemical tools for grasping non-covalent interactions. *J. Chem. Phys.* **2017**, *146*, 120901.
- [8] Philipsen, P. H. T.; Baerends, E. J. Role of the Fermi Surface in Adsorbate-Metal Interactions: An Energy Decomposition Analysis. *J. Phys. Chem. B* **2006**, *110*, 12470–12479.
- [9] Cruz Hernandez, N.; Zicovich-Wilson, C. M.; Fdez. Sanz, J. The constrained space orbital variation analysis for periodic *ab initio* calculations. *J. Chem. Phys.* **2006**, *124*, 194105.
- [10] Raupach, M.; Tonner, R. A periodic energy decomposition analysis method for the investigation of chemical bonding in extended systems. *J. Chem. Phys.* **2015**, *142*, 194105.
- [11] Xin, H.; Linic, S. Analyzing relationships between surface perturbations and local chemical reactivity of metal sites: Alkali promotion of O₂ dissociation on Ag(111). *J. Chem. Phys.* **2016**, *144*, 234704.
- [12] Pecher, J.; Mette, G.; Durr, M.; Tonner, R. Site-Specific Reactivity of Ethylene at Distorted Dangling-Bond Configurations on Si(001). *ChemPhysChem* **2017**, *18*, 357–365.
- [13] Pecher, J.; Schober, C.; Tonner, R. Chemisorption of a Strained but Flexible Molecule: Cyclooctyne on Si(001). *Chem.-Eur. J.* **2017**, *23*, 5459–5466.
- [14] Mo, Y.; Song, L.; Wu, W.; Zhang, Q. Charge transfer in the electron donor-acceptor complex BH₃NH₃. *J. Am. Chem. Soc.* **2004**, *126*, 3974–3982.
- [15] Khaliullin, R. Z.; Head-Gordon, M.; Bell, A. T. An efficient self-consistent field method for large systems of weakly interacting components. *J. Chem. Phys.* **2006**, *124*, 204105.
- [16] Kaduk, B.; Kowalczyk, T.; Van Voorhis, T. Constrained Density Functional Theory. *Chem. Rev.* **2012**, *112*, 321–370.
- [17] Plaisance, C. P.; van Santen, R. A.; Reuter, K. Constrained-Orbital Density Functional Theory. Computational Method and Applications to Surface Chemical Processes. *J. Chem. Theory Comput.* **2017**, *13*, 3561–3574.
- [18] Mao, Y.; Ge, Q.; Horn, P. R.; Head-Gordon, M. On the Computational Characterization of Charge-Transfer Effects in Noncovalently Bound Molecular Complexes. *J. Chem. Theory Comput.* **2018**, *14*, 2401–2417.

- [19] Stoll, H.; Wagenblast, G.; Preuss, H. On the use of local basis sets for localized molecular orbitals. *Theor. Chim. Acta* **1980**, *57*, 169–178.
- [20] Gianinetti, E.; Raimondi, M.; Tornaghi, E. Modification of the Roothaan equations to exclude BSSE from molecular interaction calculations. *Int. J. Quantum Chem.* **1996**, *60*, 157–166.
- [21] Nagata, T.; Takahashi, O.; Saito, K.; Iwata, S. Basis set superposition error free self-consistent field method for molecular interaction in multi-component systems: Projection operator formalism. *J. Chem. Phys.* **2001**, *115*, 3553–3560.
- [22] Mo, Y.; Gao, J.; Peyerimhoff, S. D. Energy decomposition analysis of intermolecular interactions using a block-localized wave function approach. *J. Chem. Phys.* **2000**, *112*, 5530–5538.
- [23] Khaliullin, R. Z.; Cobar, E. A.; Lochan, R. C.; Bell, A. T.; Head-Gordon, M. Unravelling the Origin of Intermolecular Interactions Using Absolutely Localized Molecular Orbitals. *J. Phys. Chem. A* **2007**, *111*, 8753–8765.
- [24] Khaliullin, R. Z.; Bell, A. T.; Head-Gordon, M. Analysis of charge transfer effects in molecular complexes based on absolutely localized molecular orbitals. *J. Chem. Phys.* **2008**, *128*, 184112.
- [25] Azar, R. J.; Head-Gordon, M. An energy decomposition analysis for intermolecular interactions from an absolutely localized molecular orbital reference at the coupled-cluster singles and doubles level. *J. Chem. Phys.* **2012**, *136*, 024103.
- [26] Thirman, J.; Head-Gordon, M. An energy decomposition analysis for second-order Moller-Plesset perturbation theory based on absolutely localized molecular orbitals. *J. Chem. Phys.* **2015**, *143*, 084124.
- [27] Smits, G. F.; Krol, M. C.; Hart, W. J. v. d.; Altona, C. Theoretical investigations of the nature of intramolecular interactions. *Mol. Phys.* **1986**, *59*, 209–225.
- [28] Mo, Y.; Wu, W.; Song, L.; Lin, M.; Zhang, Q.; Gao, J. The Magnitude of Hyperconjugation in Ethane: A Perspective from Ab Initio Valence Bond Theory. *Angew. Chem., Int. Ed.* **2004**, *43*, 1986–1990.
- [29] Wodrich, M. D.; Gonthier, J. F.; Steinmann, S. N.; Corminboeuf, C. How Strained are Carbomeric-Cycloalkanes? *J. Phys. Chem. A* **2010**, *114*, 6705–6712.
- [30] Levine, D. S.; Horn, P. R.; Mao, Y.; Head-Gordon, M. Variational Energy Decomposition Analysis of Chemical Bonding. 1. Spin-Pure Analysis of Single Bonds. *J. Chem. Theory Comput.* **2016**, *12*, 4812–4820.

Chapter 2. Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions

- [31] Levine, D. S.; Head-Gordon, M. Energy decomposition analysis of single bonds within Kohn–Sham density functional theory. *Proc. Natl. Acad. Sci.* **2017**, *114*, 12649–12656.
- [32] Couty, M.; Bayse, C. A.; Hall, M. B. Extremely localized molecular orbitals (ELMO): a non-orthogonal Hartree-Fock method. *Theor. Chem. Acc.* **1997**, *97*, 96–109.
- [33] Fornili, A.; Sironi, M.; Raimondi, M. Determination of extremely localized molecular orbitals and their application to quantum mechanics/molecular mechanics methods and to the study of intramolecular hydrogen bonding. *J. Mol. Struct.* **2003**, *632*, 157–172.
- [34] Genoni, A.; Ghitti, M.; Pieraccini, S.; Sironi, M. A novel extremely localized molecular orbitals based technique for the one-electron density matrix computation. *Chem. Phys. Lett.* **2005**, *415*, 256–260.
- [35] Fornili, A.; Moreau, Y.; Sironi, M.; Assfeld, X. On the suitability of strictly localized orbitals for hybrid QM/MM calculations. *J. Comput. Chem.* **2006**, *27*, 515–523.
- [36] Genoni, A. X-ray Constrained Extremely Localized Molecular Orbitals: Theory and Critical Assessment of the New Technique. *J. Chem. Theory Comput.* **2013**, *9*, 3004–3019.
- [37] Steinmann, S. N.; Jana, D. E.; Wu, J. I.-C.; Schleyer, P. v. R.; Mo, Y.; Corminboeuf, C. Direct Assessment of Electron Delocalization Using NMR Chemical Shifts. *Angew. Chem., Int. Ed.* **2009**, *48*, 9828–9833.
- [38] Elgabarty, H.; Khaliullin, R. Z.; Kuhne, T. D. Covalency of hydrogen bonds in liquid water can be probed by proton nuclear magnetic resonance experiments. *Nat. Commun.* **2015**, *6*.
- [39] Steinmann, S. N.; Vogel, P.; Mo, Y.; Corminboeuf, C. The norbornene mystery revealed. *Chem. Commun.* **2011**, *47*, 227–229.
- [40] Steinmann, S. N.; Mo, Y.; Corminboeuf, C. How do electron localization functions describe π -electron delocalization? *Phys. Chem. Chem. Phys.* **2011**, *13*, 20584–20592.
- [41] Bomble, L.; Steinmann, S. N.; Perez-Peralta, N.; Merino, G.; Corminboeuf, C. Bonding analysis of planar hypercoordinate atoms via the generalized BLW-LOL. *J. Comput. Chem.* **2013**, *34*, 2242–2248.
- [42] Mo, Y.; Gao, J. Polarization and Charge-Transfer Effects in Aqueous Solution via Ab Initio QM/MM Simulations. *J. Phys. Chem. B* **2006**, *110*, 2976–2980.
- [43] Giese, T. J.; Chen, H.; Huang, M.; York, D. M. Parametrization of an Orbital-Based Linear-Scaling Quantum Force Field for Noncovalent Interactions. *J. Chem. Theory Comput.* **2014**, *10*, 1086–1098.

- [44] Demerdash, O.; Mao, Y.; Liu, T.; Head-Gordon, M.; Head-Gordon, T. Assessing many-body contributions to intermolecular interactions of the AMOEBA force field using energy decomposition analysis of electronic structure calculations. *J. Chem. Phys.* **2017**, *147*, 161721.
- [45] Bambagioni, V.; Bevilacqua, M.; Bianchini, C.; Filippi, J.; Lavacchi, A.; Marchionni, A.; Vizza, F.; Shen, P. K. Cover Picture: Self-Sustainable Production of Hydrogen, Chemicals, and Energy from Renewable Alcohols by Electrocatalysis (ChemSusChem 7/2010). *ChemSusChem* **2010**, *3*, 765–765.
- [46] Liu, P.; Norskov, J. K. Ligand and ensemble effects in adsorption on alloy surfaces. *Phys. Chem. Chem. Phys.* **2001**, *3*, 3814–3818.
- [47] Gianinetti, E.; Vandoni, I.; Famulari, A.; Raimondi, M. Extension of the SCF-MI method to the case of K fragments one of which is an open-shell system. *Adv. Quantum Chem.* **1999**, *31*, 251–266.
- [48] Horn, P. R.; Sundstrom, E. J.; Baker, T. A.; Head-Gordon, M. Unrestricted absolutely localized molecular orbitals for energy decomposition analysis: Theory and applications to intermolecular interactions involving radicals. *J. Chem. Phys.* **2013**, *138*, 134119.
- [49] Yang, W.; Mori-Sanchez, P.; Cohen, A. J. Extension of many-body theory and approximate density functionals to fractional charges and fractional spins. *J. Chem. Phys.* **2013**, *139*, 104114.
- [50] Steinmann, S. N.; Yang, W. Wave function methods for fractional electrons. *J. Chem. Phys.* **2013**, *139*, 074107.
- [51] Khaliullin, R. Z.; VandeVondele, J.; Hutter, J. Efficient Linear-Scaling Density Functional Theory for Molecular Systems. *J. Chem. Theory Comput.* **2013**, *9*, 4421–4427.
- [52] Mayer, I.; Surjan, P. R. Localization and delocalization. II. Role of overlap in interbond interactions. *J. Chem. Phys.* **1984**, *80*, 5649–5658.
- [53] CP2K Open Source Molecular Dynamics, available at <https://www.cp2k.org/>. <https://www.cp2k.org/>.
- [54] Kresse, G.; Hafner, J. Ab initio molecular dynamics for liquid metals. *Phys. Rev. B* **1993**, *47*, 558.
- [55] Kresse, G.; Furthmuller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **1996**, *54*, 11169.
- [56] Klimes, J.; Bowler, D. R.; Michaelides, A. Chemical accuracy for the van der Waals density functional. *J. Phys.: Condens. Matter* **2010**, *22*, 022201.

Chapter 2. Energy Decomposition Analysis for Metal Surface – Adsorbate Interactions by Block Localized Wave Functions

- [57] Gautier, S.; Steinmann, S. N.; Michel, C.; Fleurat-Lessard, P.; Sautet, P. Molecular adsorption at Pt(111). How accurate are DFT functionals? *Phys. Chem. Chem. Phys.* **2015**, *17*, 28921–28930.
- [58] Steinmann, S. N.; Corminboeuf, C. Comprehensive Benchmarking of a Density-Dependent Dispersion Correction. *J. Chem. Theory Comput.* **2011**, *7*, 3567–3577.
- [59] Blochl, P. E. Projector augmented-wave method. *Phys. Rev. B* **1994**, *50*, 17953.
- [60] Kresse, G.; Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B* **1999**, *59*, 1758.
- [61] VandeVondele, J.; Hutter, J. Gaussian basis sets for accurate calculations on molecular systems in gas and condensed phases. *J. Chem. Phys.* **2007**, *127*, 114105.
- [62] Goedecker, S.; Teter, M.; Hutter, J. Separable dual-space Gaussian pseudopotentials. *Phys. Rev. B* **1996**, *54*, 1703–1710.
- [63] Boys, S. F.; Bernardi, F. The calculation of small molecular interactions by the differences of separate total energies. Some procedures with reduced errors. *Mol. Phys.* **1970**, *19*, 553–566.
- [64] Steinmann, S. N.; Corminboeuf, C.; Wu, W.; Mo, Y. Dispersion-Corrected Energy Decomposition Analysis for Intermolecular Interactions Based on the BLW and dDXDM Methods. *J. Phys. Chem. A* **2011**, *115*, 5467–5477.
- [65] Bondi, A. van der Waals Volumes and Radii. *J. Phys. Chem.* **1964**, *68*, 441–451.
- [66] Miller, T. M. *CRC Handbook of Chemistry and Physics*; Taylor & Francis Group.
- [67] Reed, A. E.; Schleyer, P. v. R. Chemical bonding in hypervalent molecules. The dominance of ionic bonding and negative hyperconjugation over d-orbital participation. *J. Am. Chem. Soc.* **1990**, *112*, 1434–1445.
- [68] Steinmann, S. N.; Fleurat-Lessard, P.; Gotz, A. W.; Michel, C.; Ferreira de Morais, R.; Sautet, P. Molecular mechanics models for the image charge, a comment on “Including image charge effects in the molecular dynamics simulations of molecules on metal surfaces”. *J. Comput. Chem.* **2017**, *38*, 2127–2129.
- [69] Steinmann, S. N.; Ferreira De Morais, R.; Gotz, A. W.; Fleurat-Lessard, P.; Iannuzzi, M.; Sautet, P.; Michel, C. Force Field for Water over Pt(111): Development, Assessment, and Comparison. *J. Chem. Theory Comput.* **2018**, *14*, 3238–3251.
- [70] Kalhara Gunasooriya, G. T. K.; Saeys, M. CO Adsorption Site Preference on Platinum: Charge Is the Essence. *ACS Catal.* **2018**, *8*, 3770–3774.

3 Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis

This Chapter is based on the following article: P. Clabaut, R. Staub, J. Galiana, E. Antonetti, and S. N. Steinmann *J. Chem. Phys.* **2020**, *153*, 054703.

Water molecules adsorbed on noble metal surfaces are of fundamental interest in surface science, heterogeneous catalysis and as a model for the metal/water interface. Herein, we analyse 28 water structures adsorbed on five noble metal surfaces (Cu, Ag, Au, Pd, Pt) via density functional theory and energy decomposition analysis based on the block localized wave function technique. The structures, ranging from the monomers to ice adlayers, reveal that the charge-transfer from water to the surface is nearly independent from the charge-transfer between the water molecules, while the polarization energies are cooperative. Dense water-water networks with small surface dipoles, such as the $\sqrt{39} \times \sqrt{39}$ unit cell (experimentally observed on Pt(111)) are favored compared to the highly ordered and popular H^{up} and H^{down} phases. The second main result of our study is that the many-body interactions, which stabilize the water assemblies on the metal surfaces, are dominated by the polarization energies, with the charge-transfer scaling with the polarization energies. Hence, if an empirical model could be found that reproduces the polarization energies, the charge-transfer could be predicted as well, opening exciting perspectives for force field development.

3.1 Introduction

Ice-like water layers over noble metal surfaces are widely studied, both experimentally and theoretically.¹⁻⁴ Due to the sparsity of the characterization of the metal/water interface, they are sometimes considered model systems for the solid/liquid interface,⁵⁻⁹ even though the validity of this extrapolation is far from obvious.^{10;11} Furthermore, the ice adlayers are regularly used to model the metal/liquid interface in (electro-)catalysis.¹²⁻¹⁶ The alternatives for approximate treatments of the solvent are implicit solvents,¹⁷ which do not compete with adsorbates for surface sites,¹⁸ microsolvation^{19;20} which solvates adsorbates only locally, and ab initio molecular dynamics, which is computationally very expensive²¹.

The most commonly reported and applied ice adlayers over closed-packed noble-metal surfaces are the $\sqrt{3} \times \sqrt{3}$ H^{up} and H^{down} models, going back to the seminal STM work of Doering on Ru.¹ However, larger unit cells have been observed for Pt(111)² and explained in terms of more disordered ice-like layers featuring ring-structures of various sizes.³

Previous theoretical studies have focused on the bonding mechanism of individual monomers on metal surfaces^{22;23} or on the possibility of water dissociation.²⁴ Herein, we focus on non-dissociated water layers, fully covering the noble metal surfaces. The purpose of this study is, on the one hand, to elucidate the relative stability of these ice-like structures on five noble metal surfaces (Cu, Pd, Ag, Pt, Au) and, on the other hand, to identify the driving force of their formation via energy decomposition analysis (EDA). We rely on dispersion corrected DFT to achieve a balanced description between water–water and water–metal interactions.^{25;26}

Energy decomposition analysis is a powerful tool which is mostly applied in molecular chemistry,²⁷⁻²⁹ but also increasingly in condensed phase³⁰ and at surfaces.³¹⁻³³ EDAs, like most concepts in chemistry, make reference to quantities that are neither observable, nor uniquely defined, just like the definition of an atom in a molecule. Such noumenons are, nevertheless, widely accepted to be useful.³⁴ The lack of unique definition spurs debate in the community,^{35;36} which we interpret as a sign of the importance of the concept and not of its futility. Hence, we herein exploit EDA to gain insight into the role of polarization and charge-transfer for the interaction of water with noble metal surfaces.

We have recently extended the block localized wave function (BLW) technique³⁷⁻³⁹ to metallic surfaces.³³ The BLW based EDA now allows to decompose the adsorption energy into four terms: deformation, frozen, polarization and charge-transfer, which encompasses electron sharing. This energy decomposition not only provides deep insight into the bonding, but also allows to gain information for force field development.⁴⁰ The charge-transfer (chemisorption) is the term that is the most difficult to reproduce, as it is intrinsically a many-body term with no generally applicable analytical expression known for it. The polarization is, on the other hand, a better understood many-body term, which can be modelled via induced dipoles, themselves

modelled according to different techniques. BLW, which includes polarization at the DFT level, also defines the limit of the precision that can be expected from a polarizable force field in the absence of error cancellations between different interaction energy components.

To achieve this detailed insight, the remaining of the work is structured as follows: After describing the computational details, we analyse the stability of the various ice-layers on the five investigated metal surfaces. Next, we perform an EDA for the ice-layers, but also for 23 smaller (monomer to heptamer) clusters (see Fig. B.4 and B.5). According to these computations, the polarization interaction is strongly correlated to the charge-transfer energy, so that the total interaction can be estimated based on the (linear-scaling⁴¹) BLW energy. Furthermore, we quantify the cooperativity between water–water and water–metal polarization interactions and evidence a competition between the water–water and water–metal charge-transfer interactions. Note, that hydrogen bonds are, as attested by various previous studies, not of pure electrostatic origin, but show characteristics of weak covalent bonds,^{30;42;43} which is indistinguishable from charge-transfer in BLW. It is this partial covalent character that is responsible for the synergy or competition with other interactions.

3.2 Methods

We start by defining the total adsorption energy of a given system:

$$\Delta E_{ads} = E_{SCF} - E_{surf^{opt}} - n \cdot E_{W^{opt}} + \Delta E_{BSSE} \quad (3.1)$$

where E_{SCF} is the standard KS-SCF energy of the full system, $E_{surf^{opt}}$ and $E_{W^{opt}}$ are the corresponding energies of the freely optimized surface and water molecule, respectively. n is the number of water molecules in a given system. Since the BLW is only defined in a localized basis set, we have to correct for the basis set superposition error (BSSE), which we do according to the counterpoise procedure of Boys and Bernardi⁴⁴, giving rise to the (by definition positive) energy correction ΔE_{BSSE} .

As common in BLW-EDA,^{27;45;46} we decompose the total adsorption energy ΔE_{ads} into:

$$\Delta E_{ads} = \Delta E_{deform} + \Delta E_{frozen} + \Delta E_{pol} + \Delta E_{CT} \quad (3.2)$$

where ΔE_{deform} is the preparation or deformation energy, ΔE_{frozen} is the frozen energy term that describes the interaction of the isolated fragment densities brought together and covers electrostatic interaction and Pauli repulsion²⁷ as well as dispersion interactions⁴⁶. ΔE_{pol} is the polarization energy which is obtained by variationally optimizing the BLW. ΔE_{CT} is, finally, the charge transfer interaction that includes the covalent bond formation. Note, that the BSSE only affects the charge-transfer term, as all other terms are evaluated using the

Chapter 3. Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis

same fragment-decomposed basis set. It is important to note that the distinction between ΔE_{pol} and ΔE_{CT} depends on the choice of the basis set, as in the complete basis set limit the variational optimization at the origin of ΔE_{pol} retrieves the full interaction energy. However, in several previous studies the basis set influence has been found to be rather limited when using “standard” basis sets.^{46;47} Nevertheless, an extension has been proposed to fully overcome this issue if it would severely affect the results.⁴⁸

The following equation summarizes the scheme and different terms. Further details on the computation of these terms are given in the corresponding equations as indicated:

$$E_{isolated} \xrightarrow[\text{Eq. 3.4a}]{\Delta E_{deform}} E_{fragments} \xrightarrow[\text{Eq. 3.4b}]{\Delta E_{frozen}} E_{SFD} \xrightarrow[\text{Eq. 3.4c}]{\Delta E_{pol}} E_{BLW} \xrightarrow[\text{Eq. 3.4d}]{\Delta E_{CT}} E_{cor} \xrightarrow{-\Delta E_{BSSE}} E_{SCF} \quad (3.3)$$

$\underbrace{\hspace{15em}}_{\Delta E_{ads}(\text{Eq.3.2})}$
 $\underbrace{\hspace{10em}}_{\Delta E_{int}(\text{Eq.3.5})}$
 $\underbrace{\hspace{10em}}_{\Delta E_{int}^{BLW}(\text{Eq.3.6})}$

where $E_{isolated} = E_{surf^{opt}} + n \cdot E_{W_i^{opt}}$ is the sum of the electronic energy of each fragment optimized separately, $E_{fragments} = E_{surf^{sys}} + \sum_i^n E_{W_i^{sys}}$ is the sum of the energy of each fragment evaluated in its final geometry. The superscript “sys” corresponds to the energy of a fragment in the geometry adopted in the presence of the other fragment *s*. E_{SFD} is the total energy after Superposition of the Fragment Densities, E_{BLW} is the total energy obtained by the Block Localized Wavefunction³³ and E_{cor} corresponds to the final energy of the complete system, corrected for the BSSE, while E_{SCF} is the energy obtained by a standard SCF computation.

This leads to the following definitions for the four terms of the adsorption energy as decomposed in Eq. 3.2:

$$\Delta E_{deform} = E_{surf^{sys}} - E_{surf^{opt}} + \sum_i^n E_{W_i^{sys}} - n \cdot E_{W^{opt}} \quad (3.4a)$$

$$\Delta E_{frozen} = E_{SFD} - E_{surf^{sys}} - \sum_i^n E_{W_i^{sys}} \quad (3.4b)$$

$$\Delta E_{pol} = E_{BLW} - E_{SFD} \quad (3.4c)$$

$$\Delta E_{CT} = E_{SCF} - E_{BLW} + \Delta E_{BSSE} \quad (3.4d)$$

Furthermore, we define the interaction energy, ΔE_{int} as the adsorption energy excluding the deformation energy, i.e.,

$$\Delta E_{int} = \Delta E_{frozen} + \Delta E_{pol} + \Delta E_{CT} \quad (3.5)$$

Similarly, we define the BLW interaction energy as the interaction energy that excludes the charge-transfer interaction:

$$\Delta E_{BLW} = \Delta E_{frozen} + \Delta E_{pol} \quad (3.6)$$

Since ΔE_{BLW} include all polarization contributions at the DFT level but excludes any charge-transfer, it can be understood as the interaction energy of an “ideal” polarizable force field. Furthermore, computationally its evaluation can be performed with near linear scaling with respect to the number of fragments,⁴¹ which contrasts with the cubic scaling for the computations including the charge-transfer interactions.

For the energy decomposition analysis and its interpretation, each system is either divided into two blocks (one for the metallic surface, one for all the n water molecules together) or into $n + 1$ blocks.

Taking the frozen interaction as an example, we denote the standard decomposition:

$$\Delta E_{frozen} = \Delta E_{frozen}(W_1, \dots, W_n, surf) \quad (3.7)$$

as the situation where every water molecule W_i is treated as a separate subsystem. This contrasts with decomposition into two blocks, the surface and the adlayer:

$$\Delta E_{frozen}^{surf-layer} = \Delta E_{frozen}(\bigcup_i W_i, surf) \quad (3.8)$$

where all the water molecules are treated together as a single block and the surface is a second block.

Finally, in order to assess many-body effects, we also determine the “additive” frozen interaction:

$$\Delta E_{frozen}^{add} = \sum_i^n \Delta E_{frozen}(W_i, surf) \quad (3.9)$$

where we perform n separate computations, one for each water molecule, and then sum the corresponding contributions.

The standard decomposition leads to the most complete interaction while Eq. 3.8 excludes the water–water interaction components and Eq. 3.9 is free of any many-body interactions. It is, therefore, possible to define the missing part of the interaction component

$$\Delta E_{frozen}^{NonAdd} = \Delta E_{frozen} - \Delta E_{frozen}^{add} \quad (3.10)$$

which represent the non-additive part of the interaction.

Chapter 3. Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis

Analogous equations to Eq. 3.7-3.10 can be written for the polarization and charge-transfer energy.

For the purpose of comparison with experimental estimates, we also compute an approximate surface energy, defined as:

$$\Gamma = \frac{\Delta E_{ads}}{A} \quad (3.11)$$

where ΔE_{ads} is the adsorption energy of the adlayer and A the area of the corresponding surface, i.e., the unit-cell.

3.3 Computational Details

In order to avoid BSSE during geometry optimizations, the adsorbed structures were optimized with the plane-wave code VASP 5.4.1^{49;50} using periodic boundary conditions applying the re-optimized Perdew, Burke and Ernzerhofer functional to make it compatible with the non-local van der Waals (vdW) functional, in short optPBE-vdW⁵¹ functional. An energy cutoff of 400 eV is chosen for the expansion of the plane-wave basis set, which yields converged adsorption energies in agreement with our previous studies^{18;52;53}. The electron-ion interactions are described by the PAW formalism.^{54;55} The unit cells are built from bulk metals (2.821, 2.580, 2.955, 2.943 and 2.797 Å nearest neighbor distance for Pt, Cu, Au, Ag, and Pd, respectively) with four metallic layers, the top two of which are allowed to relax. The out-of-plane vector of the unit cell was chosen to be ~ 20 Å to achieve a negligible interaction between periodic images. Geometries are considered converged when the maximum gradient on all atoms is below 0.05 eV/Å, with the electronic structure being converged to a precision of 10^{-6} eV. VASP was also used to determine the surface dipole moment and the workfunction of the various systems. Following our previous studies,^{18;56} the dipole moment was determined using the self-consistent dipole correction. The latter not only computes the dipole moment, but also decouples the periodic images, which avoids spurious polarization of the system. The workfunction was, however, computed for a centro-symmetric slab of seven layers, which was found sufficient in our previous study on formic acid decomposition over Pd(111).⁵⁶

In CP2K,^{57;58} which uses atom-centered basis functions as required for the BLW-EDA, the molecular orbitals were represented by a double- ζ Gaussian basis set with one set of polarization functions, called DZVP-MOLOPT-SR-GTH for both BLW-EDA and BSSE corrected SCF DFT simulations.⁵⁹ A cutoff of 400 Ry was used to describe the electron density in agreement with our previous study³³. The exchange-correlation (XC) energy was approximated with the optPBE-vdW⁵¹ functional. Like in VASP, the Brillouin zone was described at the Γ -point. Goedecker, Teter and Hutter (GTH) pseudo-potentials⁶⁰ based on the PBE functional were used to describe the interactions between the valence electrons and the ionic cores, and the




Figure 3.1 – Structures of the ice adlayers on Pt(111) together with their short-hand notation as used herein. The unit cell is indicated in green. For the small unit cells, a supercell is chosen to have comparable sizes for all systems and allow the use of the Γ point only in the DFT computations. The in-plane unit-cell vectors are of following lengths (in Å) for Pt (and accordingly rescaled for the other metals, see sec. 3.3): 14.66×14.66 ; 14.66×14.66 ; 19.54×14.66 ; 17.16×17.16 ; and 17.62×17.62 .

electronic smearing was approximated by a Fermi-Dirac distribution at 300 K. As discussed in our previous publication³³ the 18 valence electron potential is necessary for Pt to obtain similar results between CP2K and VASP. For Cu and Au the 11 valence electron potential is applied. For the adopted choice, Fig. B.1 in Appendix B provides the comparison between CP2K and VASP, showing a satisfactory correlation for our purposes ($R^2 > 0.99$), indicating good numerical convergence of the results for both codes. We have performed additional tests with the larger TZVP-MOLOPT-(SR)-GTH basis set. As shown in Fig. B.6, ΔE_{pol} increases by about 10% at the expense of ΔE_{CT} , which is similar to the earlier reports.^{46;47}

In order to identify the water molecules in the ice-layers, where the atoms are ordered by elements rather than molecule, we have used our in-house code *imecs*, which is provided in the supplementary information.

3.4 Results and Discussion

3.4.1 Relative Stability of Ice-like Layers

We are comparing five previously reported ice-like layers (depicted in Fig. 3.1) on five metals, which we will denote H^{up} , H^{down} , $chain-H^{down}$, $\sqrt{37}$ and $\sqrt{39}$. The nominal coverage of these structures is 0.67 ML for the first three, 0.70 and 0.72 ML for the last two, where ML stands for monolayer with respect to the surface metal atoms. Figure 3.2 shows that the $\sqrt{39}$ structure leads to the lowest surface energy Γ (see Eq. 3.11) for almost all metals, closely followed by $\sqrt{37}$. Cu(111) is the exception in the sense that it is the only metal investigated herein for which $\Gamma(\sqrt{37}) < \Gamma(\sqrt{39})$. Concomitantly, Cu(111) has the lowest interatomic distance of 2.58 Å followed by 2.80 Å for Pd, the second smallest metal investigated here.

In terms of absolute values, the surface energy of Pt(111) for the $\sqrt{37}$ structure is 1.73 kcal/(mol·Å²), which compares to 0.46 kcal/(mol·Å²) for the adhesion of solid water at ~ 100 K⁶¹ and 3.45 kcal/(mol·Å²) for the Pt/liquid water tension.⁶² For a broader comparison to experiment, Fig.

Chapter 3. Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis




Figure 3.2 – Surface energies Γ in kcal/mol/Å² relative to the most stable ice adlayer, i.e., $\sqrt{39}$ for all metals except for Cu, where $\sqrt{37}$ is slightly more stable. The higher the bar, the less stable is the corresponding structure.




Figure 3.3 – The change of workfunction ($\Delta\Phi$) when an ice adlayer is adsorbed on a (111) noble metal surface.

B.2 shows the correlation between the lowest surface energy of each metal and the reported experimental value. Except for Cu, the trends are nicely reproduced. Since, furthermore, the interatomic distance of Cu (2.58 Å) is less compatible with the ideal H-bond length of ~ 2.8 Å, this might indicate that liquid water behaves differently on Cu(111) compared to the other noble metal surfaces.

3.4.2 Electronic Analysis

Before moving to the energy decomposition analysis, we here investigate the electronic nature of the various interfaces by computing the surface dipole moment and the workfunction Φ . The workfunction is intimately connected to the electrochemical potential and it has been argued that the H^{up} and H^{down} phases should co-exist over large potential ranges.^{9;10} However at that time the three other surfaces investigated here have not been assessed.

All ice adlayers taken alone, except H^{up} for all metals and chain- H^{down} on top of Cu(111), feature a positive dipole moment, meaning that there is a positive charge accumulation on the “bottom” and a negative one on the “top”. The maximum (1.4 eÅ) is obtained for $\sqrt{39}$

Chapter 3. Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis

over Pd(111), while the minimum is found for H^{up} over Cu(111) (-2.7 eÅ). The water layers, when optimized on different metallic surfaces (and hence, on different lattice sizes), undergo noticeable geometry distortions. A specific layer, evaluated without metallic slab but in the perturbed geometry corresponding to different metals, can exhibit a range of dipole moment up to 1.3 eÅ. This maximum is obtained for the H^{up} layer optimized on Au(111) (-1.4 eÅ) compared to the one from Cu(111) (-2.7 eÅ). Since water is adsorbed only slightly stronger on Cu than on Au(111), this shows that it is mostly the lattice mismatch, and not so much the interaction strength with the metal, that affects the electronic structures via geometrical constraints.

Similarly, the change in workfunction upon adsorption of an ice adlayer depends significantly on the metal (larger changes in absolute values for Pt, Pd and Cu than for Ag and Au) and on the ice layer (see Fig. 3.3). In particular, the workfunction is lowered by almost 3 eV when adsorbing the H^{up} layer on Pt, Pd and Cu, but “only” 1.6 V on Au and Ag. Given the very reductive nature of the H^{up} structure,⁹ its stability is doubtful in itself. Even though one could have expected that the H^{down} layer has the opposite effect, this is not the case and the workfunction still drops for Pt and Pd (-0.8 eV), but remains unchanged for Cu and increases slightly (0.2 eV) for Ag and Au. This not only shows that a purely geometric analysis of the structure is not enough to retrieve the trends on the electronic structure, but also that the metal-dependant interaction plays a major role.

As expected based on basic physical principles,⁶³ the surface dipole moments of the hydrated metal surfaces are correlated with the change in workfunction ($\Delta\Phi$), with an intercept of zero (see Fig. B.3). In other words, a positive surface dipole moment is associated with a positive change in workfunction and vice versa. Given the generally positive dipole moments for the isolated layers as discussed above, the dominance of negative changes in Φ , and, thus, the negative dipole moments for hydrated surfaces, require additional explanations. Indeed, the change in dipole moment upon adsorption is generally negative (the one exception being H^{up} on Cu(111)), with an average of -0.8 eÅ and a minimum of -1.6 eÅ (chain-H^{down} on Pt(111)). This nicely demonstrates a “universal” interaction between water and noble metal surfaces featuring a net polarization (or charge-transfer) from water to the surface, i.e., the surface becomes more negatively charged and behaves as a more reductive system compared to vacuum. This conclusion is fully confirmed by the analysis of the density reorganization upon adsorption, as shown in Fig. B.9-B.11, which represent the density difference between the fully relaxed density and the superposition of the density of the surface and the adlayer. The averaged profiles as a function of the out-of-plane distance demonstrate the density accumulation in the region on top of the last metallic layer for all five adstructures. The negative contributions, indicating the origin of the density accumulation, are, however, distinct from one adstructure to the other. In the case of H^{up}, the density comes from the water adlayer. For the other structures, it is a combination of charge transfer from the water-layer and a polarization of the metallic system, down to the second metal layer. The top and side




Figure 3.4 – Average (per water molecule) contributions of frozen, polarization, and charge-transfer to the total interaction energy between water structures (oligomer of adlayer) and metal surfaces. The error bar gives the standard deviation among all the 28 considered systems. Red dots give the specific values for the H^{up} layers and horizontal dashes, those for the $\sqrt{39}$ layers.

views of the isosurfaces of the density reorganization (Fig. B.9) nicely illustrates the spatial heterogeneity of the $\sqrt{37}$ and $\sqrt{39}$ structures, which can be seen as a low-energy realization of the proposed mixtures of H^{up} and H^{down} structures proposed by Filhol and Doublet.⁹ Note, however, that the arrangements with the lowest surface energy ($\sqrt{39}$ and $\sqrt{37}$) feature a $\Delta\Phi$ close to zero and thus also the smallest surface dipole moments.

In summary, both the energetic and the electronic structure analysis support the idea that the lowest energy arrangement of water on noble metal surfaces might resemble the $\sqrt{39}$ structure, i.e., densely packed, but containing various relative orientations of the water molecules.

3.4.3 Energy Decomposition Analysis

Water–Metal Interaction

The first, fundamental, question addressed herein is how the interaction of water with a given metal surface depends on the arrangement of the water molecules and on the nature of the

Chapter 3. Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis

metal surface. This question is, furthermore, of importance when aiming at the development of a second generation force field, improving over the existing ones that are fitted to monomer interaction energies, i.e., missing all many body terms. Therefore, we start by analyzing the interaction of the preformed adlayers with the metallic surface, i.e. $\Delta E_{int}^{surf-layer}$, and each of its components, defined in analogy to Eq. 3.8. This means that the deformation energy is excluded, while the water molecules interact with each other freely, i.e., the water–water charge-transfer associated with the hydrogen bonds is present at all stages of the analysis. As a consequence, the water–water CT does not directly contribute to the studied energy difference: the “frozen” term, $\Delta E_{frozen}^{surf-layer}$, solely accounts for the electrostatic, steric and dispersion interaction between the adlayer and the metal surface. The polarization term $\Delta E_{pol}^{surf-layer}$ is mainly composed of the polarization of the metallic surface and the adlayer, but also contains a response of the water–water interaction due to this polarization. Finally, the charge-transfer contribution $\Delta E_{CT}^{surf-layer}$ captures the charge-transfer between the metal surface and the adlayers and its repercussions on the water–water interaction (see Fig. B.9 for a visualization). To simplify the discussion, we will only discuss the case of three metals: Pt, Cu and Au. Indeed, the corresponding values for Pd resemble Pt very closely and the same is true for the couple Ag and Au. On the other hand, we enrich the discussion by including oligomeric clusters on the surface (see Fig B.4 and B.5), in order to deduce more general trends than just observation of the five ice adlayers. The oligomers offer a larger diversity of structural motifs than the five adlayers. Furthermore, compared with the periodic adlayers that need to be stretched or compressed to fit into the unit cell, the oligomers can relax and accommodate more easily the various lattice constants.

Fig. 3.4 reports average energies per water molecule for the interaction energy and its components. As expected based on the single molecule adsorption²⁵, the average interaction energy is largest for Pt (-9.5 kcal/mol) and smallest for Au (-6.2 kcal/mol). Note, that this is less than the single molecule adsorption (-10.6 and -7.5 kcal/mol), indicating that the net effect of high coverage is slightly repulsive.

When moving to the components, we can first note a general trend for all components to be, in absolute value, more important for Pt than for Cu, than for Au. For instance, the steric repulsion, at the origin of the positive sign for $\Delta E_{frozen}^{surf-layer}$, is highest for Pt and almost zero for Au. This can be traced back to “geometrical” reasons, with a mean distance Au–O of 3.20 Å vs. Pt–O of 2.98 Å, which is a consequence of the overall stronger adsorption on Pt, which leads to shorter internuclear distances. The origin of this strong difference in geometry, and thus steric repulsion, is mostly found in $\Delta E_{CT}^{surf-layer}$, which is more than twice for Pt compared to Au (-10.1 vs -4.2 kcal/mol). The same proportion applies to the polarization energy, but overall, $\Delta E_{pol}^{surf-layer}$ is less stabilizing than $\Delta E_{CT}^{surf-layer}$, except for Cu, where they average to -5.9 and -5.4 kcal/mol, respectively.

To give an illustration of the spread of the individual components and the profound difference

between H^{up} compared to the other ice adlayers, let us discuss them at the example of Pt(111), even though the observations and conclusions for the other metal surfaces would barely differ. First, many water molecules in the ice adlayers are not adsorbed in the optimal single molecule geometry. This contributes to a lowering of the repulsion ($\Delta E_{\text{frozen}}^{\text{surf-layer}}$) for the ice adlayers compared to the average (4.6 vs. 8.9 kcal/mol). For H^{up} this repulsion is even only 0.6 kcal/mol, illustrating the little steric hindrance between the ice-layer and the metal surface. The polarization energy is similarly small for H^{up} (-1.6 kcal/mol) while the two H^{down} adlayers feature $\Delta E_{\text{pol}}^{\text{surf-layer}} \approx -5$ kcal/mol. $\Delta E_{\text{pol}}^{\text{surf-layer}}$ reaches even ~ -8 kcal/mol for the more complex $\sqrt{37}$ and $\sqrt{39}$ structures, a value that compares well to an average of -8.3 kcal/mol for all 28 systems considered. The situation for the charge-transfer, $\Delta E_{\text{CT}}^{\text{surf-layer}}$, is close to the observations for $\Delta E_{\text{pol}}^{\text{surf-layer}}$, i.e., H^{up} only marginally benefits from CT ($\Delta E_{\text{CT}}^{\text{surf-layer}} = -3.3$ kcal/mol), while the other structures are, with -7.8 kcal/mol somewhat shy of the average $\Delta E_{\text{CT}}^{\text{surf-layer}}$ of -10.1 kcal/mol. These observations demonstrate that H^{up} behaves differently compared to the other adlayers. However, most of the other ice adlayers, and in particular the $\sqrt{37}$ and $\sqrt{39}$ structures which are the most stable ones, are closely related to adsorption patterns that can be mimicked via oligomers. Indeed, especially on Au and Cu metallic surfaces, the specific value of the energetic components of the surface-layer interaction for the $\sqrt{39}$ structure (for example) is found within the standard deviation among all considered systems. On Pt, however, the $\sqrt{39}$ layer exhibits a lower than average $\Delta E_{\text{frozen}}^{\text{surf-layer}}$ (5.3 kcal/mol against 8.9 kcal/mol). A similar behavior, but with opposite sign, is observed for $\Delta E_{\text{CT}}^{\text{surf-layer}}$, so that $\Delta E_{\text{int}}^{\text{surf-layer}}$ is found within the standard deviation. Consequently, the $\sqrt{39}$ structure is a very stable adlayer that is well represented by oligomers.

The comparison of $\sqrt{37}$ and H^{up} based on energetic quantities could not have been deduced from the analysis of the flow of electron density as represented in Figs. B.9-B.11. This is in full agreement with our previous study on molecular complexes, where we demonstrated that the electron flow and the associated interaction strength are not directly related.⁶⁴ In the case of $\sqrt{39}$, the explanation is particularly simple: Even though the overall workfunction is barely affected by the nearly vanishing surface dipole moment, this global property hides the complexity of the local charge rearrangements, associated with the polarization and charge-transfer (see Fig. B.9).

From the perspective of designing a force field, the most important question at this point is if the charge-transfer energy between the ice-like layer and the metallic surface is indeed required. Hence, Fig. 3.5 reports the interaction energy of the water subsystem with the metal surface when charge-transfer is neglected ($\Delta E_{\text{BLW}}^{\text{surf-layer}}$) as a function of the total interaction energy ($\Delta E_{\text{int}}^{\text{surf-layer}}$). To better distinguish the behavior of the oligomers (o) and the complete ice adlayers (l), the two groups are depicted with different symbols, but using the same color. For the oligomers (o) Fig. 3.5 it is evident that for Pt (green) there is no relation between the BLW (polarization-only) and the total interaction energy. However, for Cu

Chapter 3. Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis




Figure 3.5 – The interaction of the water subsystem with the metal surface at the BLW level, $\Delta E_{BLW}^{surf-layer}$, is plotted against the corresponding total interaction energy, $\Delta E_{int}^{surf-layer}$. The oligomers (o) and adlayers (l) are given by separated symbols. The (l) point most left and right corresponds to $\sqrt{39}$ and H^{up} , respectively.




Figure 3.6 – Correlation of $\Delta E_{CT}^{surf-layer}$ with $\Delta E_{pol}^{surf-layer}$. The empty symbols correspond to the three outlier adlayers (H^{up} , H^{down} and chain- H^{down}) which are excluded from the correlation.

(orange) and Au (blue), where the role of CT is less important, there is a reasonable correlation between the two quantities, suggesting that relative adsorption energies could already be estimated at the BLW level. The BLW computations could benefit from a significant speedup due to their (near) linear scaling, dramatically reducing the computational cost of sampling phase space at the metal/liquid interface. When considering the ice adlayers (l) we first see a rough correlation for all three metals which supports the suggestion that CT might not be necessary for relative energies at the metal/liquid interface. At a closer look, the value for H^{up} point (rightmost points of the (l) series) does not fit in the correlation for any of the metals. Apparently, H^{up} has a non-typical behavior, meaning its properties are significantly different from other water arrangements on noble metal surfaces. We, therefore, advise against its use in practical applications as a model for the water/metal interface.

Even if in the absence of CT a good correlation with $\Delta E_{int}^{surf-layer}$ can be obtained, in absolute terms it cannot be neglected: $\Delta E_{int}^{surf-layer}$ in the absence of CT, i.e., $\Delta E_{BLW}^{surf-layer}$, is not even stabilizing for Pt and only mildly so for Cu and Au (see Fig. 3.5). This demonstrates that $\Delta E_{CT}^{surf-layer}$ is a significant term over all the metals and most important on Pt. Platinum is known to be more oxophilic than Au and Cu, which is also seen in the water monomer binding

Chapter 3. Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis

energy, which is -10.6 for Pt(111) vs -8.5 and -7.5 kcal/mol for Cu and Au, respectively. This oxophilicity can explain the importance of CT over Pt(111). Strikingly, with the exception of the H^{up} , H^{down} and chain- H^{down} adlayers, $\Delta E_{CT}^{\text{surf-layer}}$ can be estimated from $\Delta E_{pol}^{\text{surf-layer}}$ (see Fig. 3.6). The slope of $\Delta E_{CT}^{\text{surf-layer}}$ vs $\Delta E_{pol}^{\text{surf-layer}}$ only slightly depends on the metal when excluding the “outliers”, which are H^{up} , H^{down} and chain- H^{down} . The slope is close to unity for Pt and Au, whereas it is only 0.7 for Cu. As shown in Fig. B.6, the basis set dependence of ΔE_{pol} and ΔE_{CT} of about $\sim 10\%$ is identical for all metals considered. The slope for Cu rises to 0.9 when excluding all ice adlayers, revealing once again the impact of the lattice mismatch. Hence, if an accurate prediction of the metal/water polarization energy could be found via an empirical force field, the corresponding charge-transfer term could be estimated without a detailed physical model. This possibility opens encouraging perspectives for the next generation of water/metal force fields. Indeed, polarizable force fields for metals have been developed in the past,^{65–67} but rarely coupled to polarizable water models,⁶⁸ so that their full potential might not have been reached so far.^{21;69} Furthermore, only our current work quantifies the polarization energy that should be aimed at, an important quantity when fitting an empirical force field.

Adsorption energy of water at noble metal surfaces

Having established that the interaction energy between an ice adlayer, or just a water oligomer, and a noble metal surface can be expressed in terms of the frozen energy and a scaled polarization energy, we now tackle the more general question of the total adsorption energy on metal surface. ΔE_{ads} , (Eq. 3.2) accounts for all the many-body interaction terms, i.e., the water–water many-body interactions that are already present in the absence of a metal surface,⁷⁰ the water–metal many-body interactions at the interface and, moreover, the change of the water–water interaction due to the presence of the metal surface.

To set the stage, Fig. 3.7 represents the same kind of analysis for ΔE_{ads} , as Fig. 3.4 does for $\Delta E_{int}^{\text{surf-layer}}$, i.e., the different interaction energy components per water molecule for each metal. The first, general, comment is that the two Figures look quite similar, with the same increase in absolute value of all terms when going from Au to Cu and then to Pt. The additional energy contribution, ΔE_{deform} , turns out to be of minor importance overall (<1 kcal/mol). Even for Cu(111) the deformation energy is not larger than for Pt, despite the more important lattice mismatch. This can be traced back to the relative rigidity of the water molecules compared to the softer hydrogen bond interactions between them. Hence, while the monomer geometry does not respond much to the unit cell, it is the assembly into an adlayer that has to adapt upon adsorption.

On average, ΔE_{ads} only differs by ~ 2 kcal/mol per water molecule between Au (weakest adsorption) and Pt (strongest adsorption), even though the magnitude of the major adsorption




Figure 3.7 – Average (per water molecule) contributions of deformation, frozen, polarization, and charge-transfer to the total adsorption energy of water structures (oligomer of adlayer) on metal surfaces. The error bar gives the standard deviation among all the 28 considered systems.

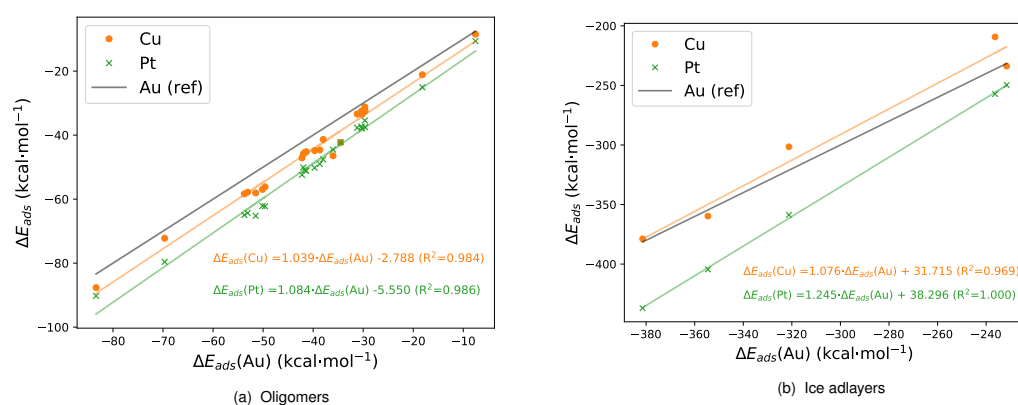


Figure 3.8 – The adsorption energy for oligomers (a) and ice adlayers (b) on Cu(111) and Pt(111) is plotted against the adsorption energy on Au(111).

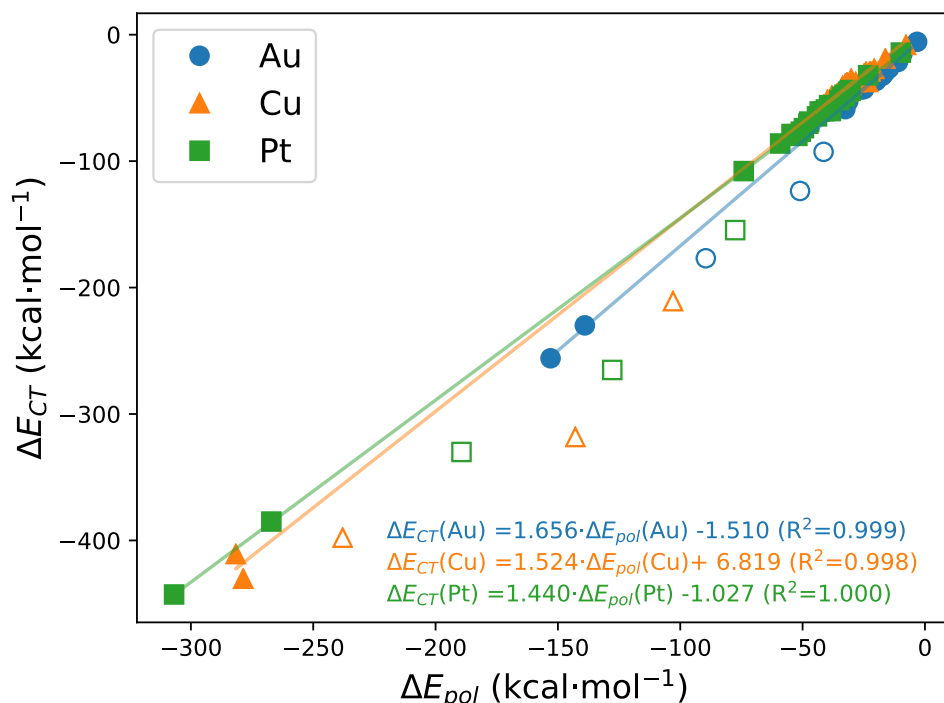


Figure 3.9 – The charge-transfer ΔE_{CT} is correlated to ΔE_{pol} for the adsorption energy. The empty symbols correspond to the three outlier adlayers (H^{up} , H^{down} and chain- H^{down}) which are excluded from the correlation.

energy components differ by at least a factor of two. In order to uncover if this similarity is only true on average or if it is a “universality” of the interaction of water with any of the noble metal surfaces, Fig. 3.8 reports the correlation of adsorption energies on Pt(111) and Cu(111) with the more physisorption-like adsorption on Au(111). Due to the large absolute difference between the adsorption energies of oligomers (up to -90 kcal/mol) and ice adlayers (up to -380 kcal/mol), the two families of systems are separated. The oligomers (Fig. 3.8a) have slope close to unity and the intercept reflects the stronger adsorption of a single water molecule on Cu and Pt compared to Au(111). For the ice adlayers (Fig. 3.8b), Cu(111) is nearly indistinguishable from Au(111). The combination of the two figures clearly shows that even though water oligomers are more strongly bound on Cu(111) than on Au(111), the lattice-mismatch affects the water adlayers significantly. ΔE_{ads} for ice adlayers on Pt(111) is, with a slope of 1.24 against Au(111), stronger and indicative of additional stabilization on Pt(111) compared to the other noble metals and compared to the oligomers. This stabilization is presumably due to a combination of stronger chemisorption and a well-matching metal lattice.

Despite this seemingly simple distinction between adlayers and oligomers when analyzing

the differences between metals, the individual components offer a complementary insight. In Fig. 3.9 we trace ΔE_{CT} as a function of ΔE_{pol} for the three metal surfaces. When excluding the three “exceptional” and energetically less stable ice adlayers (H^{up} , H^{down} and chain- H^{down}), a correlation with $R^2 > 0.99$ is obtained with slopes of about 1.5 for all three metals. This slope is higher compared to the near unity slope from Fig. 3.6, where only the interaction between the adlayer and the metal surface was analyzed. The origin of the difference is two-fold: first and foremost, the water–water interaction, which is directly present in the scaling of Fig. 3.9, features a comparably stronger CT component with respect to the polarization energy. Second, Fig. 3.9 also contains the full cooperativity between water–water and water–metal interactions, which are quantified in more details in the next subsection. The variation of slope between Fig. 3.9 and Fig. 3.6 nicely illustrates that even though charge-transfer and polarization are related (with the former vanishing in the complete basis set limit), the precise relationship between the two components depends on the nature of the probed interaction. The linear correlation of Fig. 3.9 means that ΔE_{pol} is sufficient to retrieve the complex many-body physics of ΔE_{CT} , even including the water–water interaction. Therefore, modelling the polarization energy in the absence of charge-transfer should be enough to capture the essential features for the full adsorption energies ΔE_{ads} .

Non-Additivity and Cooperativity of Water–Water–Metal Many-Body Interactions

Operationally, modification of the water–water interaction at the metal interface cannot be distinguished from the modification of the water–metal interaction due to the presence of co-adsorbed water molecules. We first quantify and compare the non-additivity of the interaction energy ($\Delta E_{int}^{NonAdd} = \Delta E_{frozen}^{NonAdd} + \Delta E_{pol}^{NonAdd} + \Delta E_{CT}^{NonAdd}$) for oligomers on Pt(111) and Au(111). The non-additivity (Eq. 3.10) measures the difference in the interaction energy between the sum of single water molecules interacting with the surface and the assembly of all water molecules interacting with the metal surface.

For the oligomers, ΔE_{int}^{NonAdd} contributes to more than 30% to the total interaction energy. In other words, the non-additivity is significant for a quantitative understanding of the interactions at the metal/water interface. Fig. B.7 demonstrates, however, that ΔE_{int}^{NonAdd} correlates ($R^2 = 0.93$ and slope of almost unity) between the two extreme metals, Au (weak adsorption) and Pt (strong adsorption), suggesting that it is a “universal” quantity. The components of ΔE_{int}^{NonAdd} do not all behave the same: The repulsive frozen term is very weakly correlated ($R^2 \approx 0.5$), but noticeably smaller for Au than for Pt (roughly one third). This is to be expected since the water molecules are further away from the surface on Au than on Pt. It is the non-additive charge-transfer and polarization energies that are correlated between Pt and Au and thus bring about the correlation between the metals.

We now unravel the origin of the non-additivity in terms of contributing components. Overall,




Figure 3.10 – Representation of the cartwheel (θ) and propeller (ϕ) angles used to describe the orientation of water molecules toward the surface. μ is the dipole of the water molecule. $\theta = 0^\circ$ when μ is aligned with the z axis, $\phi = 0^\circ$ when the two hydrogens are at an equal distance to the surface.

$\Delta E_{frozen}^{NonAdd}$ is destabilizing while both the polarization and charge-transfer energies are stabilizing the adlayer and are responsible for $\sim 40\%$ and 60% of $\Delta E_{int}^{NonAdd} - \Delta E_{frozen}^{NonAdd}$ (i.e. the stabilizing component), respectively and are correlated with each other ($R^2 = 0.98$). The major exception to this trend is, again, the H^{up} layer for which the frozen interaction is attractive, but the polarization and charge transfer provide less additional stability compared to the average.

Having established the “universal” character of the non-additivity interaction, we now focus on the case of Pt(111) to obtain a geometric understanding of its origin. Since the structures are essentially two-dimensional, we do not simply determine the coordination number,⁷¹ but perform a directional analysis: in each structure, the H-bond ($H \cdots O$ distance below 2.5 \AA) acceptors are identified. Then, they are classified according to the Pt–O distance ($< 3.0 \text{ \AA}$ for chemisorbed water molecules, $> 3.0 \text{ \AA}$ for physisorbed molecules). Two angles are additionally introduced to describe the water orientation with respect to the surface, namely the cartwheel angle θ and the propeller angle ϕ (see Fig 3.10 for a graphical definition). These angles describe, respectively, the angle between the dipole of water and the normal to the surface, and the rotation of the hydrogen pair around the dipole of water, following our previous studies on the development of water–metal surface force fields.^{21;53} The value of θ is zero when the dipole is pointing away from the surface and rises when the molecule is bending toward the surface. Also, a value of ϕ of 0° indicates that the two hydrogens are equally close to the surface, while a value of 180° indicate that the difference in distance to the surface between the two hydrogens




Figure 3.11 – Representation of the non-additive energy contribution of physisorbed H-bond acceptors as a function of their θ and ϕ characteristics. Dots indicate the observed points. The data is also available in Table B.1

Chapter 3. Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis

is maximal.

In order to simplify the analysis, the two-dimensional space spanned by θ and ϕ is divided into 9 rectangles for the physisorbed water molecules. The limits of these rectangles are optimized to find the optimal linear model reproducing the non-additivity for all 27 structures, i.e., including the ice adlayers. For the chemisorbed molecules, only two combinations are necessary (see Table B.1). The root mean square error of this linear regression amounts to 1.37 kcal/mol (see Fig. B.8), demonstrating the good predictive power of this simple model. The advantage of the linear model is that we also identify the geometrical arrangements that are responsible for the non-additivity. The corresponding energy coefficients for chemisorbed molecules (see Table B.1) indicate that the typical adsorption minimum of a single water molecule ($\theta \approx 80^\circ$, $\phi \approx 0^\circ$, Pt–O $\approx 2.5\text{\AA}$) is the worst H-bond acceptor ($\Delta E_{int}^{acceptor, NonAdd} = -3.0$ kcal/mol), i.e., does not contribute significantly to the non-additivity. This is compatible with the observation that the oxygen atom is already interacting with the metal surface via its lone-pairs. Therefore, its electrons are less available to interact with a third hydrogen atom. In contrast, the typical building block of the ice-like layers, where the chemisorbed water molecule is tilted so that the hydrogens are pointing away from the surface ($\theta \approx 50^\circ$, $\phi \approx 0$, Pt–O $\approx 2.5\text{\AA}$) is a better H-bond acceptor (-5.4 kcal/mol). The best H-bond acceptors are, however, not chemisorbed to the surface but physisorbed (Pt–O > 3.0), and present all small ϕ values ($< 70^\circ$), meaning that the two hydrogens are at somewhat similar distances from the surface. The very best region (-10.4 kcal/mol) is found for $\theta > 120^\circ$, which corresponds to two hydrogens pointing toward the surface, but with a lone pair of the oxygen atom pointing in the direction of the surface, and thus, towards a potential H-bond donor. A more in-depth study on model systems would be necessary in order to deduce clearer trends and adapted functional forms to reproduce these trends in an empirical force field. Nevertheless, based on Fig. 3.11 it is clear that a strategy based on correction maps (CMAPs) as introduced for protein backbone angles,⁷² could successfully retrieve the non-additivity contributions. The only extension to the CMAP approach would be the introduction of a distance dependence, in analogy to the θ dependence introduced in our GAL forcefields.^{21;53}

The synergy between water–water charge transfer and water–metal charge transfer

The synergistic energy can be defined as

$$\Delta E_{CT}^{syn} = \Delta E_{CT}^{NonAdd} - E_{CT}(\bigcup_i W_i) = \Delta E_{CT} - \sum_i^n \Delta E_{CT}(W_i, surf) - E_{CT}(\bigcup_i W_i), \quad (3.12)$$

which represents the CT-associated energetic difference between, on one hand, the overall CT and, on the other hand, the sum of each individual charge transfer between a single water molecule and the surface plus the charge transfer within the isolated water layer. This synergy is therefore positive (destabilizing) if there is a competition between these charge transfers

and negative for (stabilizing) cooperativity. Overall, cooperativity is found with $\Delta E_{CT}^{syn} = -1.37$ kcal/mol per water molecule on average for Pt and -0.57 kcal/mol for Au. For H^{up} , however, ΔE_{CT}^{syn} is positive for both Pt and Au (1.1 and 0.5 kcal/mol per water molecule, respectively). On the opposite, the most stable layers ($\sqrt{37}$, $\sqrt{39}$) are the ones where this cooperation is higher, with also small net dipole moments and large contributions due to polarization.

This observation is in full agreement with our discussion in sec. 3.4.3, where we have highlighted the important density reorganizations of $\sqrt{39}$ that, nevertheless, lead to a small surface dipole moment.

3.5 Conclusion

The detailed analysis of the electronic and geometrical characteristics of 28 diverse water arrangements (from monomer to ice adlayers) over the (111) surface of five noble metals (Cu, Ag, Au, Pd and Pt) has allowed to identify trends and key factors for the stability of water arrangements on metallic surfaces. We identified dense-packed layers to be the most stable structures, like the $\sqrt{39} \times \sqrt{39}$ unit cell for Ag, Au, Pd and Pt and the $\sqrt{37} \times \sqrt{37}$ unit-cell for Cu. This stability was found to correspond to the smallest change of workfunction upon adsorption. The H^{down} and H^{up} structures, which are often cited as model structure for water/metal interfaces, are less stable and lead to workfunction changes up to 3 eV.

The energy decomposition analysis, relying on the block localized wave function (BLW), shows that the charge-transfer from water to the surface is overall nearly independent from the charge-transfer between the water molecules, the latter being key for the H-bonding. ΔE_{CT} , which is the computationally most costly term, is found to be linearly correlated to ΔE_{pol} . Hence, it can be predicted at minimal cost. As a consequence, the polarization energy can be seen as the most important contribution to the adsorption energy.

Remarkably, the polarization energy provides about 40% additional stability at the interface compared to the single water molecule adsorption, and displays a strong correlation with its charge-transfer counterpart ($R^2 = 0.99$ on Pt(111)). This is good news for force field developments, where models for the polarization energy could be included to capture the many-body effects. Alternatively, this cooperativity can be largely reproduced by an additive model based on the geometric parameters of the H-bond acceptor molecule. In summary, our investigation highlights the closely related physics that governs the various noble-metal – water interaction and suggests that polarization energies should be enough to retrieve most of the complex many-body interactions at the metal/water interface.

Bibliography

- [1] Doering, D. L.; Madey, T. E. The adsorption of water on clean and oxygen-dosed Ru(011). *Surface Science* **1982**, *123*, 305–337.
- [2] Glebov, A.; Graham, A. P.; Menzel, A.; Toennies, J. P. Orientational ordering of two-dimensional ice on Pt(111). *The Journal of Chemical Physics* **1997**, *106*, 9382–9385.
- [3] Nie, S.; Feibelman, P. J.; Bartelt, N. C.; Thürmer, K. Pentagons and Heptagons in the First Water Layer on Pt(111). *Phys. Rev. Lett.* **2010**, *105*, 026102.
- [4] Carrasco, J.; Hodgson, A.; Michaelides, A. A molecular perspective of water at metal interfaces. *Nature Materials* **2012**, *11*, 667–674.
- [5] Ogasawara, H.; Brena, B.; Nordlund, D.; Nyberg, M.; Pelmenchikov, A.; Pettersson, L. G. M.; Nilsson, A. Structure and Bonding of Water on Pt(111). *Phys. Rev. Lett.* **2002**, *89*, 276102.
- [6] Rossmeisl, J.; Norskov, J. K.; Taylor, C. D.; Janik, M. J.; Neurock, M. Calculated Phase Diagrams for the Electrochemical Oxidation and Reduction of Water over Pt(111). *The Journal of Physical Chemistry B* **2006**, *110*, 21833–21839.
- [7] Schnur, S.; Groß, A. Properties of metal–water interfaces studied from first principles. *New J. Phys.* **2009**, *11*, 125003.
- [8] Schiros, T.; Andersson, K. J.; Pettersson, L. G. M.; Nilsson, A.; Ogasawara, H. Chemical bonding of water to metal surfaces studied with core-level spectroscopies. *J. Electron Spectrosc. Relat. Phenom.* **2010**, *177*, 85–98.
- [9] Filhol, J.-S.; Doublet, M.-L. An ab initio study of surface electrochemical disproportionation: The case of a water monolayer adsorbed on a Pd(111) surface. *Electrocatalysis* **2013**, *202*, 87–97.
- [10] Vassilev, P.; van Santen, R. A.; Koper, M. T. M. Ab initio studies of a water layer at transition metal surfaces. *J. Chem. Phys.* **2005**, *122*, 054701, Publisher: American Institute of Physics.
- [11] Bjorneholm, O.; Hansen, M. H.; Hodgson, A.; Liu, L.-M.; Limmer, D. T.; Michaelides, A.; Pedevilla, P.; Rossmeisl, J.; Shen, H.; Tocci, G.; Tyrode, E.; Walz, M.-M.; Werner, J.; Bluhm, H. Water at Interfaces. *Chem. Rev.* **2016**, *116*, 7698–7726.
- [12] Desai, S. K.; Pallassana, V.; Neurock, M. A Periodic Density Functional Theory Analysis of the Effect of Water Molecules on Deprotonation of Acetic Acid over Pd(111). *J. Phys. Chem. B* **2001**, *105*, 9171–9182.
- [13] Filhol, J.-S.; Neurock, M. Elucidation of the Electrochemical Activation of Water over Pd by First Principles. *Angew. Chem., Int. Ed.* **2006**, *45*, 402–406.

- [14] Skúlason, E.; Karlberg, G. S.; Rossmeisl, J.; Bligaard, T.; Greeley, J.; Jonsson, H.; Norskov, J. K. Density functional theory calculations for the hydrogen evolution reaction in an electrochemical double layer on the Pt(111) electrode. *Phys. Chem. Chem. Phys.* **2007**, *9*, 3241–3250.
- [15] Tian, F.; Jinnouchi, R.; Anderson, A. B. How Potentials of Zero Charge and Potentials for Water Oxidation to OH(ads) on Pt(111) Electrodes Vary With Coverage. *J. Phys. Chem. C* **2009**, *113*, 17484–17492, Publisher: American Chemical Society.
- [16] Hussain, J.; Jónsson, H.; Skúlason, E. Calculations of Product Selectivity in Electrochemical CO₂ Reduction. *ACS Catal.* **2018**, *8*, 5240–5249.
- [17] Mathew, K.; Kolluru, V. S. C.; Mula, S.; Steinmann, S. N.; Hennig, R. G. Implicit self-consistent electrolyte model in plane-wave density-functional theory. *J. Chem. Phys.* **2019**, *151*, 234101.
- [18] Steinmann, S. N.; Sautet, P.; Michel, C. Solvation free energies for periodic surfaces: comparison of implicit and explicit solvation models. *Phys. Chem. Chem. Phys.* **2016**, *18*, 31850–31861.
- [19] Hibbitts, D. D.; Loveless, B. T.; Neurock, M.; Iglesia, E. Mechanistic Role of Water on the Rate and Selectivity of Fischer-Tropsch Synthesis on Ruthenium Catalysts. *Angew. Chem. Int. Ed.* **2013**, *52*, 12273–12278.
- [20] Schweitzer, B.; Steinmann, S. N.; Michel, C. Can microsolvation effects be estimated from vacuum computations? A case-study of alcohol decomposition at the H₂O/Pt(111) interface. *Phys. Chem. Chem. Phys.* **2019**, *21*, 5368–5377.
- [21] Steinmann, S. N.; Ferreira De Morais, R.; Gotz, A. W.; Fleurat-Lessard, P.; Iannuzzi, M.; Sautet, P.; Michel, C. Force Field for Water over Pt(111): Development, Assessment, and Comparison. *J. Chem. Theory Comput.* **2018**, *14*, 3238–3251.
- [22] Carrasco, J.; Michaelides, A.; Scheffler, M. Insight from first principles into the nature of the bonding between water molecules and 4d metal surfaces. *J. Chem. Phys.* **2009**, *130*, 184707.
- [23] Michaelides, A.; Ranea, V. A.; de Andres, P. L.; King, D. A. General Model for Water Monomer Adsorption on Close-Packed Transition and Noble Metal Surfaces. *Phys. Rev. Lett.* **2003**, *90*, 216102.
- [24] Schiros, T.; Ogasawara, H.; Näslund, L.-Å.; Andersson, K. J.; Ren, J.; Meng, S.; Karlberg, G. S.; Odelius, M.; Nilsson, A.; Pettersson, L. G. M. Cooperativity in Surface Bonding and Hydrogen Bonding of Water and Hydroxyl at Metal Surfaces. *J. Phys. Chem. C* **2010**, *114*, 10240–10248.

Chapter 3. Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis

- [25] Carrasco, J.; Klimeš, J.; Michaelides, A. The role of van der Waals forces in water adsorption on metals. *J. Chem. Phys.* **2013**, *138*, 024708, Publisher: American Institute of Physics.
- [26] Carrasco, J.; Santra, B.; Klimeš, J.; Michaelides, A. To Wet or Not to Wet? Dispersion Forces Tip the Balance for Water Ice on Metals. *Phys. Rev. Lett.* **2011**, *106*, 026101.
- [27] Mo, Y.; Gao, J.; Peyerimhoff, S. D. Energy decomposition analysis of intermolecular interactions using a block-localized wave function approach. *J. Chem. Phys.* **2000**, *112*, 5530–5538.
- [28] Khaliullin, R. Z.; Bell, A. T.; Head-Gordon, M. Analysis of charge transfer effects in molecular complexes based on absolutely localized molecular orbitals. *J. Chem. Phys.* **2008**, *128*, 184112.
- [29] Steinmann, S. N.; Vogel, P.; Mo, Y.; Corminboeuf, C. The norbornene mystery revealed. *Chem. Commun.* **2011**, *47*, 227–229.
- [30] Elgabarty, H.; Khaliullin, R. Z.; Kuhne, T. D. Covalency of hydrogen bonds in liquid water can be probed by proton nuclear magnetic resonance experiments. *Nat. Commun.* **2015**, *6*.
- [31] Philipsen, P. H. T.; Baerends, E. J. Role of the Fermi Surface in Adsorbate-Metal Interactions: An Energy Decomposition Analysis. *J. Phys. Chem. B* **2006**, *110*, 12470–12479.
- [32] Pecher, L.; Laref, S.; Raupach, M.; Tonner, R. Ethers on Si(001): A Prime Example for the Common Ground between Surface Science and Molecular Organic Chemistry. *Angew. Chem., Int. Ed.* **2017**, *56*, 15150–15154.
- [33] Staub, R.; Iannuzzi, M.; Khaliullin, R. Z.; Steinmann, S. N. Energy Decomposition Analysis for Metal Surface-Adsorbate Interactions by Block Localized Wave Functions. *Journal of Chemical Theory and Computation* **2019**, *15*, 265–275.
- [34] Parr, R. G.; Ayers, P. W.; Nalewajski, R. F. What Is an Atom in a Molecule? *J. Phys. Chem. A* **2005**, *109*, 3957–3959.
- [35] Bader, R. F. W. Letter to the editor: Quantum mechanics, or orbitals? *International Journal of Quantum Chemistry* **2003**, *94*, 173–177.
- [36] Zhao, L.; Hermann, M.; Schwarz, W. H. E.; Frenking, G. The Lewis electron-pair bonding model: modern energy decomposition analysis. *Nat Rev Chem* **2019**, *3*, 48–63.
- [37] Stoll, H.; Wagenblast, G.; Preuss, H. On the use of local basis sets for localized molecular orbitals. *Theor. Chim. Acta* **1980**, *57*, 169–178.
- [38] Mo, Y.; Peyerimhoff, S. D. Theoretical analysis of electronic delocalization. *J. Chem. Phys.* **1998**, *109*, 1687–1697.

- [39] Steinmann, S. N.; Jana, D. F.; Wu, J. I.-C.; Schleyer, P. v. R.; Mo, Y.; Corminboeuf, C. Direct Assessment of Electron Delocalization Using NMR Chemical Shifts. *Angew. Chem., Int. Ed.* **2009**, *48*, 9828–9833.
- [40] Demerdash, O.; Mao, Y.; Liu, T.; Head-Gordon, M.; Head-Gordon, T. Assessing many-body contributions to intermolecular interactions of the AMOEBA force field using energy decomposition analysis of electronic structure calculations. *J. Chem. Phys.* **2017**, *147*, 161721.
- [41] Khaliullin, R. Z.; VandeVondele, J.; Hutter, J. Efficient Linear-Scaling Density Functional Theory for Molecular Systems. *Journal of Chemical Theory and Computation* **2013**, *9*, 4421–4427.
- [42] Martín Pendás, A.; Blanco, M. A.; Francisco, E. The nature of the hydrogen bond: A synthesis from the interacting quantum atoms picture. *J. Chem. Phys.* **2006**, *125*, 184112, Publisher: American Institute of Physics.
- [43] Arunan, E.; Desiraju, G. R.; Klein, R. A.; Sadlej, J.; Scheiner, S.; Alkorta, I.; Clary, D. C.; Crabtree, R. H.; Dannenberg, J. J.; Hobza, P.; Kjaergaard, H. G.; Legon, A. C.; Mennucci, B.; Nesbitt, D. J. Defining the hydrogen bond: An account (IUPAC Technical Report). *Pure and Applied Chemistry* **2011**, *83*, 1619–1636, Publisher: De Gruyter Section: Pure and Applied Chemistry.
- [44] Boys, S. F.; Bernardi, F. The calculation of small molecular interactions by the differences of separate total energies. Some procedures with reduced errors. *Mol. Phys.* **1970**, *19*, 553 – 566.
- [45] Khaliullin, R. Z.; Cobar, E. A.; Lochan, R. C.; Bell, A. T.; Head-Gordon, M. Unravelling the Origin of Intermolecular Interactions Using Absolutely Localized Molecular Orbitals. *J. Phys. Chem. A* **2007**, *111*, 8753–8765.
- [46] Steinmann, S. N.; Corminboeuf, C.; Wu, W.; Mo, Y. Dispersion-Corrected Energy Decomposition Analysis for Intermolecular Interactions Based on the BLW and dDXDM Methods. *J. Phys. Chem. A* **2011**, *115*, 5467–5477.
- [47] Khaliullin, R. Z.; Bell, A. T.; Head-Gordon, M. Electron Donation in the Water-Water Hydrogen Bond. *Chem. Eur. J.* **2009**, *15*, 851–855.
- [48] Horn, P. R.; Mao, Y.; Head-Gordon, M. Probing non-covalent interactions with a second generation energy decomposition analysis using absolutely localized molecular orbitals. *Phys. Chem. Chem. Phys.* **2016**, *18*, 23067–23079, Publisher: The Royal Society of Chemistry.
- [49] Kresse, G.; Hafner, J. Ab initio molecular dynamics for liquid metals. *Phys. Rev. B* **1993**, *47*, 558.

Chapter 3. Water Adlayers on Noble Metal Surfaces: Insights from Energy Decomposition Analysis

- [50] Kresse, G.; Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **1996**, *54*, 11169.
- [51] Klimes, J.; Bowler, D. R.; Michaelides, A. Chemical accuracy for the van der Waals density functional. *J. Phys.: Condens. Matter* **2010**, *22*, 022201.
- [52] Gautier, S.; Steinmann, S. N.; Michel, C.; Fleurat-Lessard, P.; Sautet, P. Molecular adsorption at Pt(111). How accurate are DFT functionals? *Phys. Chem. Chem. Phys.* **2015**, *17*, 28921–28930.
- [53] Clabaut, P.; Fleurat-Lessard, P.; Michel, C.; Steinmann, S. N. Ten Facets, One Force Field: The GAL19 Force Field for Water–Noble Metal Interfaces. *Journal of Chemical Theory and Computation* **2020**,
- [54] Blochl, P. E. Projector augmented-wave method. *Phys. Rev. B* **1994**, *50*, 17953.
- [55] Kresse, G.; Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B* **1999**, *59*, 1758.
- [56] Wang, P.; Steinmann, S. N.; Fu, G.; Michel, C.; Sautet, P. Key Role of Anionic Doping for H₂ Production from Formic Acid on Pd(111). *ACS Catal.* **2017**, *7*, 1955–1959.
- [57] 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. *Comput. Phys. Commun.* **2005**, *167*, 103–128.
- [58] Hutter, J.; Iannuzzi, M.; Schiffmann, F.; VandeVondele, J. cp2k: atomistic simulations of condensed matter systems. *WIREs Comput Mol Sci* **2014**, *4*, 15–25.
- [59] VandeVondele, J.; Hutter, J. Gaussian basis sets for accurate calculations on molecular systems in gas and condensed phases. *J. Chem. Phys.* **2007**, *127*, 114105.
- [60] Goedecker, S.; Teter, M.; Hutter, J. Separable dual-space Gaussian pseudopotentials. *Phys. Rev. B* **1996**, *54*, 1703–1710.
- [61] Rumpitz, J. R.; Campbell, C. T. Adhesion Energies of Solvent Films to Pt(111) and Ni(111) Surfaces by Adsorption Calorimetry. *ACS Catal.* **2019**, 11819–11825.
- [62] Tyson, W.; Miller, W. Surface free energies of solid metals: Estimation from liquid surface tension measurements. *Surface Science* **1977**, *62*, 267–276.
- [63] Mortensen, J. J.; Hammer, B.; Nørskov, J. K. Alkali Promotion of N₂ Dissociation over Ru(0001). *Phys. Rev. Lett.* **1998**, *80*, 4333–4336.
- [64] Steinmann, S. N.; Piemontesi, C.; Delachat, A.; Corminboeuf, C. Why are the Interaction Energies of Charge-Transfer Complexes Challenging for DFT? *J. Chem. Theory Comput.* **2012**, *8*, 1629–1640.

- [65] Siepmann, J. I.; Sprik, M. Influence of surface topology and electrostatic potential on water/electrode systems. *J. Chem. Phys.* **1995**, *102*, 511–524, Publisher: American Institute of Physics.
- [66] Iori, F.; Corni, S. Including image charge effects in the molecular dynamics simulations of molecules on metal surfaces. *J. Comput. Chem.* **2008**, *29*, 1656–1666.
- [67] Pensado, A. S.; Padua, A. A. H. Solvation and Stabilization of Metallic Nanoparticles in Ionic Liquids. *Angew. Chem. Int. Ed.* **2011**, *50*, 8683–8687.
- [68] Golze, D.; Iannuzzi, M.; Nguyen, M.-T.; Passerone, D.; Hutter, J. Simulation of Adsorption Processes at Metallic Interfaces: An Image Charge Augmented QM/MM Approach. *J. Chem. Theory Comput.* **2013**, *9*, 5086–5097.
- [69] Steinmann, S. N.; Fleurat-Lessard, P.; Götz, A. W.; Michel, C.; Ferreira de Morais, R.; Sautet, P. Molecular mechanics models for the image charge, a comment on “including image charge effects in the molecular dynamics simulations of molecules on metal surfaces”. *J. Comput. Chem.* **2017**, *38*, 2127–2129.
- [70] Reddy, S. K.; Straight, S. C.; Bajaj, P.; Huy Pham, C.; Riera, M.; Moberg, D. R.; Morales, M. A.; Knight, C.; Götz, A. W.; Paesani, F. On the accuracy of the MB-pol many-body potential for water: Interaction energies, vibrational frequencies, and classical thermodynamic and dynamical properties from clusters to liquid water and ice. *J. Chem. Phys.* **2016**, *145*, 194504, Publisher: American Institute of Physics.
- [71] Staub, R.; Steinmann, S. N. Parameter-free coordination numbers for solutions and interfaces. *J. Chem. Phys.* **2020**, *152*, 024124.
- [72] Mackerell, A. D.; Feig, M.; Brooks, C. L. Extending the treatment of backbone energetics in protein force fields: Limitations of gas-phase quantum mechanics in reproducing protein conformational distributions in molecular dynamics simulations. *Journal of Computational Chemistry* **2004**, *25*, 1400–1415.



4 Parameter-Free Coordination Numbers for Solutions and Interfaces

This Chapter is based on the following article: R. Staub, and S. N. Steinmann *J. Chem. Phys.* **2020**, *152*, 024124.

Coordination numbers are among the central quantities to describe the local environment of atoms and are thus used in various applications such as structure analysis, fingerprints and parameters. Yet, there is no consensus regarding a practical algorithm for calculating coordination numbers, and many proposed methods are designed for specific systems. In this work, we propose a scale-free and parameter-free algorithm for nearest neighbor identification. This algorithm extends the powerful Solid-Angle based Nearest-Neighbor (SANN) framework to explicitly include local anisotropy. As such, our Anisotropically corrected Solid-Angle based Nearest-Neighbor (ASANN) algorithm provides with a fast, robust and adaptive method for computing coordination numbers. The ASANN algorithm is applied to flat and corrugated metallic surfaces to demonstrate that the expected coordination numbers are retrieved without the need for any system-specific adjustments. The same applies to the description of the coordination numbers of metal atoms in AuCu nano-particles and we show that ASANN based coordination numbers are well adapted for automatically counting neighbors and the establishment of cluster expansions. Analysis of classical molecular dynamics simulations of an electrified graphite electrode reveals a strong link between the coordination number of Cs⁺ ions and their position within the double layer, a relation that is absent for Na⁺, which keeps its first solvation shell even close to the electrode.

4.1 Introduction

Detailed analysis of chemical systems allows to gain insight into their working mechanisms, to rationalize their properties and sometimes even to predict the impact of chemical modifications such as functionalization of molecules or alloying of surfaces. Roughly two big categories can be distinguished: energetic and structural characterizations. While energy decomposition analysis provides detailed energetic insight in molecular chemistry, solid interfaces and in liquid phase,¹⁻⁴ it is the coordination environment that commonly defines the building block for local structural interpretations. The main applications in chemistry involve hydration numbers of ions^{5;6}, which are of particular interest in the context of biochemistry,⁷ the link between structure and activity in heterogeneous catalysis^{8;9} and the characterization of alloys.^{10;11}

Intuitively, the coordination number of a particle is defined as the number of nearest neighboring particles. However, this does not provide a rigorous definition: what is neighboring after all? Chemists tend to link neighboring to the idea of bonding: two items are connected if there is a strong interaction between them. But again, “strong” is not defined and most of the time the strength of the interaction is not easily accessible. Therefore, it is far more common and convenient to think in terms of spatial proximity, echoing the initial terminology of “neighbors”.

In chemistry and physics, many different algorithms are used to define the connectivity between atoms, using more or less additional parameters (e.g., nature of the atoms or reference bond lengths). We propose to classify the algorithms based on their input, using only two criteria: (i) Usage of parameters (e.g., particles’ nature, reference bond lengths). Any algorithm that depends on more than a set of point coordinates falls into this category and is denoted with a P^1 , while parameter-free algorithms are denoted by P^0 . (ii) Local adaptivity. By this we mean that the connectivity between two particles depends on the presence of other surrounding particles and denote it by A^1 , while the absence of local adaptivity corresponds to A^0 . Since these two criteria are non-exclusive, Figure 4.1 provides relevant examples of algorithms for the four possible combinations.

The simplest coordination number algorithm is based on tabulated “typical” bond lengths: only if the atoms are closer, they are connected. Hence, it falls into the class $P^1 A^0$. This simple and rapid algorithm is most popular in visualization softwares,¹² where bond length cut-offs are typically constructed from predefined covalent radii¹³ of considered elements.

Instead of using predetermined typical/average bond lengths that might not be of relevance for the considered sample, coordination cut-offs can be estimated from the sample itself. This led to the development of multiple parameter-free methods^{7;10;11;14-19} using a cut-off based on the radial distribution function $g(r)$. However, such algorithms require enough data to perform

Criteria	Locally-adaptive (A^1)	Locally-non-adaptive (A^0)
Parameter-free (P^0)	Voronoi SANN ASANN	Radial distribution analysis
Parameter-dependent (P^1)	RAD	Reference cut-off

Figure 4.1 – Overview of available coordination number algorithms. We report only popular algorithms of relevance for the context of the ASANN algorithm development, with the references provided in the text.

meaningful radial distribution analysis. Finally, these algorithms have an inherently mean-field description of coordination since they use averaged cut-offs. Hence, local fluctuations are not necessarily well handled as these algorithms do not take the remaining neighboring particles into account, i.e., it is a $P^0 A^0$ algorithm.

Constructing locally adaptive coordination numbers requires an analysis of the surroundings of each possible connection, which necessarily increases the computational cost. The best known algorithm to extract context-dependent coordination numbers uses a Voronoi tessellation,^{18;20} where space is decomposed into regions (polyhedra for 3D structures), based on which particle is the closest. Two particles are connected if their corresponding regions share a face. This algorithm is parameter-free and thus falls into the $P^0 A^1$ category, which means it can be applied without any additional knowledge on the set of particles considered. However, it has been shown that Voronoi-based coordination numbers have a tendency to be overestimated,¹⁸⁻²¹ which can be alleviated by re-weighting.¹¹ Furthermore, 3D Voronoi tessellations are computationally expensive²². In the same $P^0 A^1$ category, the Solid-Angle based Nearest-Neighbor algorithm (SANN) algorithm²¹ was developed as an efficient alternative to Voronoi tessellations. SANN attributes a solid angle to each possible neighbor and iteratively increases the cut-off radius to reach a sum of 4π for the solid angles, corresponding to a fully surrounded coordination environment. This elegant method is faster than the Voronoi tessellation and, furthermore, reduces the overestimation of the coordination numbers, without the introduction of any parameters²¹. However, SANN suffers from coordination number overestimation at interfaces²³.

Recently, the Relative Angular Distance (RAD) algorithm²³ was developed to amend the SANN overestimation at interfaces. The reduction of the coordination number is achieved by accounting for blocking: the coordination of two particles can be obstructed by particles between the considered pair. Compared to other available blocking-based methods^{24;25} RAD

does not require predefined cut-off distances, but still requires additional parameters, so that the algorithm falls in the $P^1 A^1$ class. An earlier $P^0 A^1$ algorithm²⁶ accounted for blocking by considering only bonds that are not intersecting a triangle formed by shortest bonds. However, this idea was not further explored as its computational cost is expected to be large, and we found the method not suitable to describe surfaces.

Herein we propose to augment the SANN algorithm with a correction term that takes the local anisotropy into account, while preserving all properties of SANN, i.e., low computational cost, no parameters, being locally adaptive ($P^0 A^1$ class) and being scale free, i.e., an isotropic scaling of all coordinates does not change the coordination number. This scheme is called Anisotropically corrected Solid-Angle based Nearest-Neighbors (ASANN).

To set the stage, we introduce the SANN algorithm in section 4.2.1. Then, in section 4.2.2, we identify the SANN isotropic description as source of coordination number overestimation and propose a natural solution to tackle local anisotropy through ASANN. After the computational details in section 4.3, section 4.4 contains tests and applications of ASANN for characterizing metal surfaces, alkali metal ion coordination at an electrified interface and for the construction of model Hamiltonians for alloy nanoparticles.

4.2 Algorithms

4.2.1 The SANN algorithm

The reader is strongly encouraged to read the original SANN article²¹ for details. Nevertheless, the main points are outlined below to make the ASANN algorithm understandable.

SANN defines a coordination number $CN_{\text{SANN}}(i)$ of particle i to be the number of the m nearest neighbors found within the individual coordination sphere of radius $R_i^{(m)}$ (also called shell radius). SANN elegantly provides a self-consistent definition of m and the associated coordination sphere radius $R_i^{(m)}$ as detailed below.

Let $r_{i,j}$ be the distance between particle i and its j -th nearest neighbor (denoted simply as j in the following). By definition of $R_i^{(m)}$ and m , one can already write the inequality:

$$r_{i,m} \leq R_i^{(m)} < r_{i,m+1} \quad (4.1)$$

Within SANN, the m nearest neighbors are defined such that the sum of their solid angles (with respect to i) is equal to 4π . This requirement is related to the intuitive idea that the nearest neighbors of i are the smallest set of particles fully covering its field of view. Indeed, each nearest neighbor j of particle i is associated with a solid angle $\Omega_{i,j}$, which is defined by $\Omega_{i,j} = 2\pi(1 - \cos(\theta_{i,j}))$, where the angle $\theta_{i,j}$ is depicted on Figure 4.2.




Figure 4.2 – Schematic representation of the solid angle definition $\Omega_{i,j} = 2\pi(1 - \cos(\theta_{i,j}))$ associated with neighbor j at distance $r_{i,j}$ of central particle i , and using a coordination radius of $R_i^{(m)}$.




Figure 4.3 – Illustration of the SANN iterative algorithm in 2D. Blue dots represents neighbors of the considered particle (green dot), while other particles are displayed in purple. For each increasing number of neighbors m (starting from 3), a coordination sphere radius is computed such that the SANN equality is respected (sum of neighboring solid angle requirement). The algorithm continues until the coordination radius is well-defined (i.e. surrounding particles are considered neighbors if and only if they are within the coordination radius).

The above definitions are summarized in the following equation:

$$4\pi = \sum_{j=1}^m \Omega_{i,j} = \sum_{j=1}^m 2\pi(1 - \cos(\theta_{i,j})) = \sum_{j=1}^m 2\pi \left(1 - \frac{r_{i,j}}{R_i^{(m)}} \right) \quad (4.2)$$

One can rewrite the two fundamental Eqs 4.1 and 4.2 into a much simpler form:

$$R_i^{(m)} = \frac{\sum_{j=1}^m r_{i,j}}{m-2} < r_{i,m+1} \quad (4.3)$$

The inequality of Eq. 4.3 allows to easily compute the $\text{CN}_{\text{SANN}}(i) = m$ by iteratively increasing $R_i^{(m)}$ from $m = 3$ until the inequality is respected (see Figure 4.3). The convergence of the algorithm was mathematically proven²¹.

SANN is scale free, i.e., a uniform scaling of all coordinates will not affect the coordination

number. This property can have counter-intuitive consequences when the analyzed phase features significant density fluctuations. This is particularly problematic when applied to the gas-liquid interface, where thermal fluctuations might create nearly isolated particles. The scale-free algorithm will assign a significantly larger coordination sphere to such an isolated particle, which potentially leads to a higher coordination number compared to particles in dense regions. However, such problematic cases can be spotted by monitoring the coordination sphere radius.

Furthermore, in situations where the environment of j is more dense than the one of i , it is possible that j is considered a neighbor of i , even though i is not considered a neighbor of j . Further along this line, we note that the solid-angle based approach behind SANN is inherently linked to the assumption of locally identical coordination radii. This hypothesis is the only meaningful approach within a parameter-free approach and represents a first order expansion for the coordination sphere. As a consequence, directional bonding (e.g., in organic molecules) which leads to locally more compact geometries is can lead to an overestimation of the coordination number compared to conventional views.

4.2.2 Anisotropically corrected Solid-Angle based Nearest-Neighbors, ASANN algorithm

The SANN algorithm implicitly uses a locally isotropic description of the coordination shell, i.e., the angular distribution of neighboring particles is not taken into account. At surfaces or tips there are, however, simply no neighbors in some directions, leading to empty sections in the coordination sphere. Therefore, the SANN radius must expand to reach more neighbors in the non-empty directions in order to compensate and reach a total sum of 4π . Hence, one observes an over-estimation of coordination numbers when neighboring particles are not isotropically distributed.²³ This issue was already reported in the literature, and led to the development of the relative angular distance (RAD) algorithm. However, RAD requires the particles sizes as input parameters, and is fundamentally different from SANN. Here, we propose an Anisotropically corrected Solid-Angle based Nearest neighbors (ASANN) algorithm, which is based on the SANN framework, but includes a parameter-free correction term to correctly describe intrinsically anisotropic coordination shells.

The basic idea of the ASANN algorithm is that some directions are more relevant for coordination than others. Therefore, instead of requiring the neighbors to fill up the whole field of view, ASANN only requires to fill up the relevant neighboring region of space (see Figure 4.4).

An estimator of the first order local anisotropy is obtained based on the barycentre. Specifically, we define the SANN-based barycentre $\vec{G}_i^{(m)}$ as the position barycentre over all SANN-coordinated neighbors of i , weighted by their SANN-based solid angle $\Omega_{i,j}$ (with respect to i). A similar solid-angle based neighbors weighting has been proposed in the context of the




Figure 4.4 – Illustration of the SANN coordination number overestimation, and founding idea of the ASANN algorithm: tackle local anisotropy by focusing on relevant regions for coordination. This leads to the consideration of a coordination spherical cap instead of a sphere. The chosen spherical cap reproduces the local anisotropy, through our first order local anisotropy estimator (using a SANN-based barycentre with solid-angle weighting).

Voronoi tessellation for its canonicity.¹¹ Mathematically, this can be written:

$$\vec{G}_i^{(m)} = \frac{\sum_{j=1}^m \Omega_{i,j} \vec{r}_{i,j}}{\sum_{j=1}^m \Omega_{i,j}} \quad (4.4)$$

This SANN-based barycentre is in agreement with the properties expected for retrieving the first order local anisotropy. Indeed, it identifies a global privileged direction in the distribution of coordinating neighbors and, more importantly, it provides a quantitative estimation of the first order local anisotropy magnitude $\alpha_i^{(m)}$ defined by:

$$\alpha_i^{(m)} = \frac{\|\vec{G}_i^{(m)}\|}{R_i^{(m)}} \quad (4.5)$$

Instead of considering a coordination sphere, the ASANN algorithm uses a coordination spherical cap \mathcal{C} which focuses on the locally relevant region of space for coordination. \mathcal{C}_i is the spherical cap centered on i having the same anisotropic properties as the neighboring particles through our first order local anisotropy estimator. This is achieved by setting the radius $R_{\mathcal{C}_i} = R_i^{(m)}$ and requiring that its barycentre $\vec{G}_{\mathcal{C}_i}$ is confounded with the SANN-based barycentre $\vec{G}_i^{(m)}$:

$$\begin{aligned} \mathcal{C}_i &\text{ is centred on } i \\ R_{\mathcal{C}_i} &= R_i^{(m)} \\ \vec{G}_{\mathcal{C}_i} &= \vec{G}_i^{(m)} \end{aligned} \quad (4.6)$$

The solid angle $\Omega_{\mathcal{C}_i}$ associated with this spherical cap (with respect to i) can be written as:

$$\Omega_{\mathcal{C}_i} = 4\pi(1 - \gamma_{\mathcal{C}_i}) \quad (4.7)$$

Chapter 4. Parameter-Free Coordination Numbers for Solutions and Interfaces

where $\gamma_{\mathcal{C}_i} \in [0, 1[$ corresponds to the angular correction coefficient, which only depends on the first order local anisotropy magnitude $\alpha_i^{(m)}$ (see section C.1 in Appendix C for a derivation):

$$\gamma_{\mathcal{C}_i} = \frac{\alpha_i^{(m)} + \sqrt{\alpha_i^{(m)2} + 3\alpha_i^{(m)}}}{3} \quad (4.8)$$

Therefore, instead of filling up the whole field of view ($\sum_{j=1}^m \Omega_{i,j} = 4\pi$), ASANN neighbors only cover the spherical cap. This leads to the fundamental equation of ASANN:

$$\sum_{j=1}^{m'} \Omega'_{i,j} = \sum_{j=1}^{m'} 2\pi \left(1 - \frac{r_{i,j}}{R_i^{(m')}} \right) = \Omega_{\mathcal{C}_i} = 4\pi(1 - \gamma_{\mathcal{C}_i}) \quad (4.9)$$

where the ASANN coordination number m' and the coordination cap radius $R_i^{(m')}$ are defined by extending the SANN coordination radius definition to include the angular correction coefficient $\gamma_{\mathcal{C}_i}$:

$$R_i^{(m')} = \frac{\sum_{j=1}^{m'} r_{i,j}}{m' - 2 \times (1 - \gamma_{\mathcal{C}_i})} < r_{i,m'+1} \quad (4.10)$$

Note that for isotropic distributions ($\gamma_{\mathcal{C}_i} = 0$), Eq. 4.10 goes back to Eq. 4.3, which means that SANN and ASANN coincide for this case.

Using the inequality of Eq. 4.10, the ASANN coordination number m' can be easily computed by iteratively computing $R_i^{(m')}$, starting with $m' = \lfloor 2(1 - \gamma_{\mathcal{C}_i}) \rfloor + 1$, and increasing m' until the inequality is respected.

In principle, this angular correction could be computed self-consistently, i.e., the angular correction term is recomputed at each iteration along with the ASANN coordination radius. This self-consistency breaks the analytic, direct link between the fundamental equation (Eq. 4.9) and the coordination radius definition (Eq. 4.10). Therefore, we only explore the ‘‘perturbative’’ angular correction, where the SANN-based barycentre $\vec{G}_i^{(m)}$ and the associated angular correction term $\gamma_{\mathcal{C}_i}$ are computed only once, using the SANN-based coordination radius and coordinating neighbors. Then, the iterative ASANN algorithm (Eq. 4.10) is performed with a fixed $\gamma_{\mathcal{C}_i}$. As demonstrated in the supporting information, this algorithm converges to a unique solution. The overall workflow for ASANN is depicted in Figure 4.5, which also provides information regarding the computational cost of ASANN compared to SANN: Essentially, there is a factor of two between them, since the iterations have to be performed twice. Knowing that ASANN coordination number is lower or equal to the one from SANN (see Appendix C for a

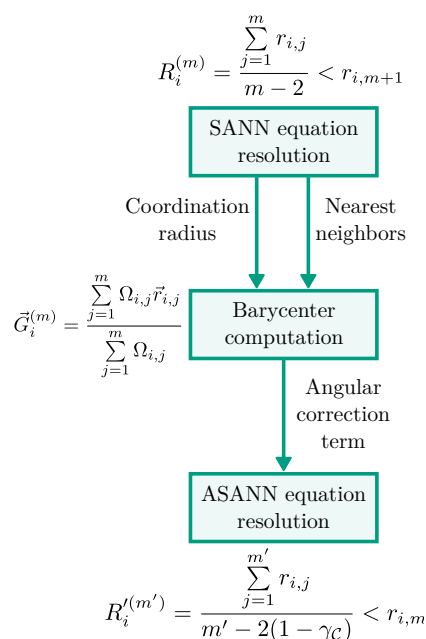


Figure 4.5 – Schematic representation of the ASANN workflow: Based on the SANN coordination sphere, the barycenter and the corresponding angular correction is computed, leading to the ASANN coordination spherical cap.

proof), there are, in general, less ASANN than SANN iterations.

Based on general properties of ASANN, partially inherited from SANN, we show the expected applicability domain of the two algorithms in Figure 4.6. The first major limitation is related to low coordination numbers (e.g., four) in combination with isotropic distributions (e.g., a tetrahedron). In this case, the coordination radius $R_i^{(4)}$ of Eq. 4.3 is twice the distance between atom i and its nearest neighbors j . If there are atoms closer than $2r_{ij}$, the inequality is not satisfied and the coordination shell is expanded to include next nearest neighbors. This is the case for most organic molecules (see Figs. C.2-C.5 and the associated discussion), but also in ionic compounds where small cations are located in tetrahedral sites as shown in Appendix C (Tables C.4 to C.15), where we have analyzed a variety of salts, oxides and perovskites. The particularity of these systems is most clearly seen for the case of a graphene sheet (Fig. C.5): the barycenter of neighboring atoms is confounded with the central atom, so there is no first order anisotropy and ASANN is equivalent to SANN. There are three nearest neighbor atoms at a distance r_{C-C} and the next nearest neighbors are already at $\sqrt{3}r_{C-C}$, again due to the directional bonding, but which is well below the $R_i^{(3)} = 3r_{C-C}$ (Eq. 4.3 for the perfect symmetry of this system). If there would not be an anisotropic preference due to bond formation characteristic for an sp^2 carbon atom, there would also be “space” for at least two more particles one above and one below the plane within such a coordination sphere. But this is not the case and, therefore, for geometrical reasons, ASANN significantly overestimates




Figure 4.6 – Venn diagram depicting the applicability domains of both SANN and ASANN. By design, ASANN extends the applicability domain of SANN to systems presenting local anisotropy.

the conventional coordination number of 3, by adding the 6 next nearest neighbors, reaching a $CN_{ASANN}=9$ for carbon in graphene. The same geometrical reasons and deviations from homogeneous distributions apply to the ionic compounds referred to above. Since ASANN is built on a first order anisotropy, the application to nanopores or slits is also at the border of its applicability: if the SANN coordination sphere spans two surfaces, the first order anisotropy vanishes and ASANN does not improve upon SANN. Higher order anisotropy corrections might correct SANN in this situation as well as for the isotropic, low-coordination environments mentioned above. Second, density fluctuations are problematic for the scale free algorithms. Therefore, cross-checks with the coordination radii are required for biphasic systems and the liquid-gas interface. As we will show by numerical results below, ASANN extends, however, the applicability of SANN to (inhomogeneous) solutions, the solid-liquid interface and rough surfaces including nano-particles.

The time complexity of (A)SANN is in the worst-case $O(nk)$ for pre-treated input, where n is the number of particles, and k is the maximum coordination number. In fact, the most time consuming step is due to the input pre-processing, required for both SANN and ASANN algorithms. This pre-processing comprises the computation of pair-wise distances (and vectors for ASANN) in $O(n^2)$, and sorting those distances in $O(n^2 \log(n))$. It is, therefore, this step that is also responsible for the space complexity of (A)SANN in $O(n^2)$. In practice, this pre-processing quickly becomes the computational bottleneck for large systems. As a consequence, particular care was taken for optimizing these steps in our Python 3 implementation, heavily relying on the Numpy library^{27;28}. We thus combine low-level language performance for large systems, with convenient Python 3 syntax and modularity. As an example, the computation of the Cs coordination numbers in a 742 atom system (vide infra), took around 0.14 seconds and the computation of all coordination numbers in a 2942 atom system took about 2 seconds

(93% is spent on input pre-processing) on a personal laptop (Core i5-6300HQ CPU).

4.3 Computational Details

The presented algorithms have been implemented using Python3²⁹ and are available in the supporting information and also on the webpage of the authors. The implementation relies on the Numpy library^{27;28} whenever possible and ASE¹² for reading common coordinate formats. A particular effort was made to provide with extensively commented code, including documentation and illustrative examples.

Figures have been generated with a custom ASE¹² fork to visualize the coordination numbers.

All DFT computations are performed using VASP^{30;31}. The PBE exchange-correlation functional³² was applied and a 400 eV energy cut-off was chosen for the plane-wave basis set. The PAW pseudo-potentials were used for describing the electron-ion interactions.^{33;34}

Some of the optimized metallic surfaces and metal clusters have been previously published by our group.³⁵⁻³⁷ The graphite electrode was set up as follows: A symmetric six layer p(2×2) graphite slab has been built based on the optimized bulk geometry using the PBE-dDsC³⁸ level of theory. A 9×9×1 Monkhorst-Pack K-point grid was used to integrate the Brillouin zone. To assess the charge distribution at negative potentials, the number of electrons was varied and the linearized Poisson-Boltzmann equation was utilized to neutralize the unit cell as implemented in the VASPsol package.^{39;40} Similarly to our previous study on electroreduction in aprotic medium,⁴¹ the isodensity value defining the cavity has been lowered to 5 10⁻⁴ e/Å³. In the case of graphite, this value ensures an absence of implicit solvent between the dispersion bound graphene layers.

The potential of zero charge of graphite is, according to our setup, -0.2 V vs SHE, which compares reasonably well with experiments (0.1 to 0.3 V depending on the nature of the carbon electrode⁴²⁻⁴⁴) considering the idealized morphology of our electrode and the expected systematic error due to the simplified electrolyte model as discussed in the context of the potential dependent adsorption of pyridine on a gold single crystal surface⁴⁵. A negative charge of -0.25 electrons in the unit cell (-0.125 electrons per surface) led to a potential drop of 2.4 V compared to the potential of zero charge and thus corresponds to a significant negative polarization. Based on the DFT computations, the unit cell was replicated four times in each in-plane direction, resulting in 128 carbon atoms per layer and a total charge of 4 e-. Following a setup very similar to ref 46, the Lennard-Jones (LJ) parameters for graphite are taken from UFF.⁴⁷ The initial solvent distribution is obtained from the predefined TIP3P⁴⁸ box with about 35 Å of water surrounding the system, resulting in a box of 716 water molecules. To simulate a 1 M electrolyte, 15 Na⁺ (or Cs⁺) and 11 Cl⁻ ions, described by their standard AMBER force field parameters, are added using tleap of the AmberTools. MM computations are performed

with NAMD 2.9⁴⁹ for 5 ns (time step of 2 fs), applying a Langevin thermostat (300 K) with a damping coefficient of 1 ps^{-1} . The Langevin barostat (1 bar) is applied with a piston period of 100 fs and a decay of 150 fs.

4.4 Results and Discussion

4.4.1 Metallic surfaces

As a first test, we apply ASANN to various metallic surfaces. The selection of surfaces is intended to span a wide range of typical cases to assess the applicability domain of SANN and ASANN.

For face centered cubic (fcc) metallic surfaces, the metal-metal contact distance in the bulk is a reliable measure for identifying nearest neighbors. Reassuringly, closed packed, flat surfaces are well described by SANN and ASANN, giving identical results that are in perfect agreement with cutoff based methods (see Appendix C). In other words, flat surfaces are not anisotropic enough to induce an expansion of the SANN coordination sphere to the next nearest neighbors.

An increased local anisotropy can be obtained for nano-rods. The large rod, constructed from the corresponding bulk without any geometry optimization (Figure 4.7), is still well and identically described by SANN and ASANN.

However, stepped surfaces (see Figures 4.8 and 4.9) are more challenging for SANN, while ASANN consistently retrieves the cutoff based coordination number. In particular, SANN is very sensitive to small local surface induced contractions, while ASANN is more robust, avoiding coordination number overestimations seen in Figure 4.9 for SANN. Stepped surfaces can therefore be considered at the limit of the applicability domain of SANN, due to their local convexity.

To simulate low-coordinated atoms,⁹ add-atoms have also been analyzed. Based on visual inspection (and confirmed by cutoff based methods), the square of addatoms on a (100) surface is characterized by two nearest neighbors from the square (the third is not “touching”, i.e., it is a next-nearest neighbor) and four from the (100) plane below, yielding a coordination number of six. Similarly, an addatom of the triangle on a (111) surface possesses two nearest neighbors from the triangle and three from the supporting (111) plane, resulting in a coordination number of five. Recalling that the flat (100) surface of fcc closed packed metals has a coordination number of 8 for its surface atoms, the $\text{CN}_{\text{SANN}}=8$ for surface adatoms reveal a clear overestimation (see Figures 4.10 and 4.11), qualitatively failing to identify the low-coordination environment. This contrasts with CN_{ASANN} that is in full agreement with visual inspection and cutoff based methods.




Figure 4.7 – Transversal cut of a periodic nano-rod of an fcc metal, with superimposed coordination numbers. SANN and ASANN coordination numbers are identical.




Figure 4.8 – Side view of a stepped gold surface (664), with superimposed coordination numbers. SANN and ASANN coordination numbers are identical.

Chapter 4. Parameter-Free Coordination Numbers for Solutions and Interfaces

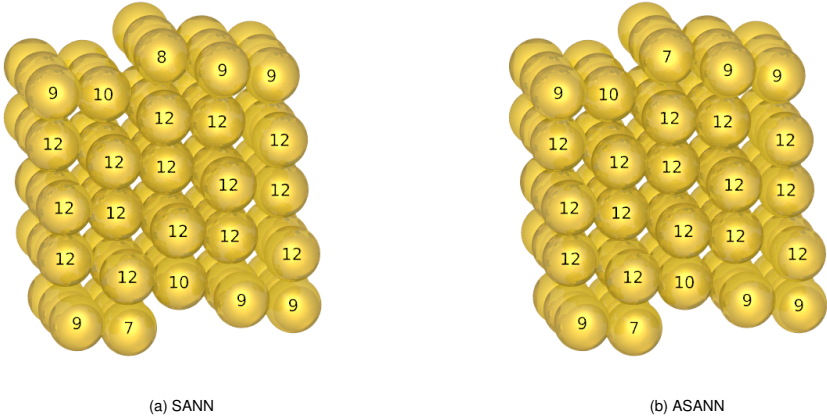


Figure 4.9 – Side view of a stepped gold surface (644), with superimposed SANN (a) and ASANN (b) coordination numbers.




Figure 4.10 – Representation of a p(3x3) Cu(100) surface with a square of 4 Cu addatoms. Coordination numbers from SANN (a) and ASANN (b) are superimposed.




Figure 4.12 – Low energy, amorphous Au_{38} nanoparticle. Coordination numbers from SANN (a) and ASANN (b) are superimposed. SANN overestimation is particularly visible at the tips.

graphite electrode. On the one hand, the experimental capacitance of a graphitic electrode is not strongly dependent on the alkali cation in the electrolyte,⁵⁹ which might suggest that the double layer structure is similar for all alkali metal cations. On the other hand, Cs^+ ions are, in contrast to Na^+ ions, known to be soft and large, implying a low charge density at their surface. As a consequence, Cs^+ ions are expected to partially lose their solvation shell to approach the negatively charged electrode surface more closely, which increases both the VdW interactions and the Coulomb attraction.

We probe the difference between Cs^+ and Na^+ based on classical molecular dynamics simulations of a graphite carbon electrode at a potential of about -2.5 V (see section 4.3) in ~ 1 M aqueous CsCl and NaCl solutions. To analyze the double layer, we investigate the existence of a link between coordination number of ions in water (determined by ASANN, SANN and a $g(r)$ based cut-off algorithm) and their adsorption state.

Figure 4.13 displays the structuring of the electrolyte at the electrode surface. As commonly observed⁶⁰, for both electrolyte solutions the distribution curve of the excess cation density presents two well-defined peaks, corresponding to two distinct layers of preferred presence near the anode. The strong first cation excess peak is followed by a nearly charge compensated zone before a second, although lower, cation excess peak occurs which then decays to zero towards the bulk solution. Besides, the anion (Cl^-) distribution is barely structured and independent on the nature of the cation. The details of the charge-excess profiles are quite different: Cs^+ excess density has the first peak at about 3.5 Å, while the corresponding one for Na^+ is at 4.8 Å. This difference indicates that Cs^+ can indeed approach the electrode more closely. The second peak is, accordingly, shifted from 6 Å to 7.5 Å when going from caesium to




Figure 4.13 – Ionic density of ~ 1 M CsCl (top panel) and NaCl (bottom panel) aqueous solutions near a graphite anode (-2.5 V vs SHE). The excess counter-ions (compared with the co-ions, equivalent here to the charge density) and the co-ions density are represented for both systems. Three distinct zones are identified, corresponding to different adsorption/trapping states for the counter-ions. The average ASANN coordination number for the counter-ions (with water molecules and co-ions) is superimposed for each zone.

Average coordination numbers of Cs ⁺ ions			
Method	Desolvated trapping layer $h < 4.7 \text{ \AA}$	Solvated trapping layer $4.7 < h < 7 \text{ \AA}$	Bulk-like layer $h > 7 \text{ \AA}$
ASANN	7.4 (1.1)	9.6 (1.2)	9.4 (1.2)
SANN	9.1 (1.1)	10.6 (1.2)	10.5 (1.2)
fixed cutoff	7.1 (1.1)	9.2 (1.3)	9.0 (1.3)

Table 4.1 – Comparison of average Cs⁺ ions coordination numbers determined by ASANN, SANN and a fixed cutoff (4.0 \AA) for each defined region. The data is displayed in the form: mean (standard deviation).

sodium.

Table 4.1 and 4.2 report the ASANN and SANN coordination numbers of Cs⁺ and Na⁺, respectively. For reference purposes, the radial distribution function of the O-Cs⁺ was analysed and a first coordination sphere radius cutoff of 4.0 \AA was identified. From this radius, cutoff-based coordination numbers have been computed. All three algorithms agree that on average at least one water molecule is lost when Cs⁺ is located in the first layer, while the second layer is indistinguishable from the bulk solution in terms of the coordination numbers. This can be rationalized based on Van der Waals (VdW) radii: a fully solvated Cs ion cannot approach the graphitic electrode closer than roughly 7 \AA ($r_C + r_{Cs^+} + 3 \text{ \AA}$, where 3 \AA is a rough estimate for the size of a water molecule).

Chapter 4. Parameter-Free Coordination Numbers for Solutions and Interfaces

Average coordination numbers of Na ⁺ ions			
Method	First trapping layer $h < 6.5 \text{ \AA}$	Second trapping layer $6.5 < h < 9 \text{ \AA}$	Bulk-like layer $h > 9 \text{ \AA}$
ASANN	6.1 (0.6)	6.1 (0.6)	6.1 (0.5)
SANN	6.5 (0.7)	6.6 (0.8)	6.5 (0.8)
fixed-cutoff	5.8 (0.5)	5.8 (0.5)	5.8 (0.5)

Table 4.2 – Comparison of average Na⁺ ions coordination numbers determined by ASANN, SANN and a fixed cutoff (3.2 Å) for each defined region. The data is displayed in the form: mean (standard deviation).

For Na⁺ ions, however, the coordination numbers are the same in the two excess regions as in the bulk, despite the qualitative similarity between the two cation excess profiles. Even the nearest Na⁺ layer ($\sim 4.8 \text{ \AA}$) is not close enough to the surface to induce a partial desolvation (see Table 4.2). In other words, in contrast to Cs⁺, the solvation shell of Na⁺ is strong and stable across the double layer and the three methods agree on an approximate average coordination number of six, indicating minor fluctuations and a rather isotropic distribution, in agreement with spectroscopic data.⁶

To quantitatively determine if ASANN provides with “good” coordination numbers, we assess their ability to discriminate between “close” and “bulk” Cs ions. Specifically, we want to quantify if CN_{ASANN} is statistically lower (i.e., below CN_{thresh}) at the interface than in the bulk. For the sake of comparison, we apply the same analysis also to SANN and to the cutoff based algorithm.

To have a quantifiable measure whether or not CN_{ASANN} is lower at the interface than in the bulk, we explore the relation between closeness and loss of coordination via a Receiver Operating Characteristic (ROC) curve analysis. The ROC curve is one of the most rigorous and well-established tools to evaluate the discriminating power for a binary classifier (here close/not close) depending on a threshold value (here CN_{ASANN}).^{61;62} The ROC curve of a class of binary classifiers is constructed by plotting the True Positive Rate (TPR) against the False Positive Rate (FPR) computed on a data set, for different threshold values.

Here, the TPR is the ratio of $CN^i < CN_{\text{thresh}}$ for the i Cs ions “close” to the electrode to the number of Cs ions “close” to the electrode overall, while the False Positive Rate is the ratio of $CN^i < CN_{\text{thresh}}$ for Cs ions *not* “close” to the electrode to the number of Cs ions *not* “close” to the electrode overall. The threshold value can be varied between the minimum and maximum observed with a given algorithm. Here, we define “close” as the Cs – surface distance $\leq 4.7 \text{ \AA}$.

Such an analysis is reported on Figure 4.14. Let us first look at the ASANN results to explain how to read a ROC curve. Since the CNs are reasonable descriptors, for low CN_{thresh} (e.g., 7), only “close” Cs ions are identified (FPR=0%), but not all of them (TPR \approx 20%), as shown in Table C.16a. Similarly, at high CN (e.g., 10), the TPR reaches \approx 100%, i.e., all “close” Cs ions have CNs




Figure 4.14 – ROC curve for classifying Cs ions for being “close” to the carbon anode based on the coordination number being below a given threshold. When moving from TPR=20% to TPR=55% with ASANN (blue curve), CN_{thresh} is increased from 7 to 8 (see Table C.16a in Appendix C for the raw data). This exploits the geometrical constraint responsible for a partial desolvation when close enough to the surface. The distance criterium for being close enough to the surface is chosen to be 4.7 \AA , since it corresponds to the empirical frontier between the partially desolvated adsorption zone and the fully solvated adsorption zone for Cs ions near the graphite anode (see Fig. 4.13).

below this value, but about 50% (FPR) of the ions *not* close to the surface have CNs below this value. Hence, increasing the threshold further does not change the TPR, but the FPR reaches 100%.

A major ROC curve analysis is performed through the Area Under the Curve (AUC), as it provides an aggregate measure of classification performance⁶³. In this particular example, the AUC is associated with the probability that a random interfacial Cs ion has a lower coordination number than a random bulk Cs ion.

With an AUC of 0.90, we find a strong correlation between CN_{ASANN} and the proximity of Cs^+ and the surface. This finding is coherent with ASANN providing reasonable coordination numbers in this system. The same analysis performed with SANN only yields an AUC of 0.81. The cutoff-based algorithm yields an AUC of 0.86, outperforming SANN, but not reaching the performance of ASANN. This is rather remarkable, given that the cutoff algorithm is tuned to the system by analyzing the radial distribution function, while ASANN does not know anything about the specificities of the system. On a physical level, the major difference between the cutoff algorithm and ASANN is that the ASANN coordination sphere adapts itself to the local environment, while the cutoff does not. This means that the cutoff based CNs are more

sensitive to thermal fluctuations. Indeed, we found cutoff-based coordination numbers to be slightly noisier ($\approx +8\%$) than their ASANN/SANN counterparts within each region of the studied system (see Table 4.1), rationalizing the poorer AUC performance.

From this comparison it can be concluded that fixed-cutoff based algorithms are somewhat challenged when dealing with liquids, while SANN is not suitable for interfaces. In contrast, ASANN combines the best features of these algorithms and performs well. Therefore, ASANN is most relevant for describing solid/liquid interfaces.

4.4.3 Application to Au NP energy fitting

Exploring the energy landscape of alloy nano-particles from first principles is computationally very expensive. Therefore, cheap energy evaluations based on simplified Hamiltonians are often applied.^{64–66} Furthermore, if the model Hamiltonian captures the essential physics, it can also be used to gain insight into the driving force of the alloy formation. Inspired from cluster expansions, we here test the use of a simple model Hamiltonian H_m for bimetallic materials. H_m depends only on the number of atoms of each element and the number of bonds between each possible elements pairs:

$$H_m = \sum_{\chi} \Gamma_{\chi} N_{\chi} + \sum_{\chi, \chi' \geq \chi} \Gamma_{\chi, \chi'} N_{\chi, \chi'} \quad (4.11)$$

In other words, it is a two-body model Hamiltonian. For alloys of elements with quite different atomic radii, the alloy nano-particles are not necessarily very regular and selecting characteristic distances is not straightforward. Therefore, an automatic, parameter free algorithm to identify bonds is particularly beneficial for these systems. In the following, we apply this simplified model to bimetallic Au-Cu materials, where H_m becomes

$$\begin{aligned} H_m &= \Gamma_{Au} N_{Au} + \Gamma_{Cu} N_{Cu} + \Gamma_{Au, Au} N_{Au, Au} + \Gamma_{Au, Cu} N_{Au, Cu} + \Gamma_{Cu, Cu} N_{Cu, Cu} \\ &= \Gamma \cdot \mathbf{N} \end{aligned} \quad (4.12)$$

with Γ the set of parameters to fit. \mathbf{N} is constructed by enumeration of atoms and bonds of each type. The automatic bond enumeration was performed using ASANN and SANN for comparison purposes. In order to have consistent bond description, a bond between atoms A and B is only counted if and only if B is found as nearest neighbor of A and reciprocally. Such filtering is particularly beneficial to SANN, where buried atoms are mostly spared from coordination overestimation. Hence, the bidirectional bond definition improves the soundness of bonds between bulk and surface atoms.

For fitting the model Hamiltonian, we use a structurally diverse set of 27 Au-Cu bimetallic nanoparticles composed of 38 atoms with various compositions and morphologies and 6

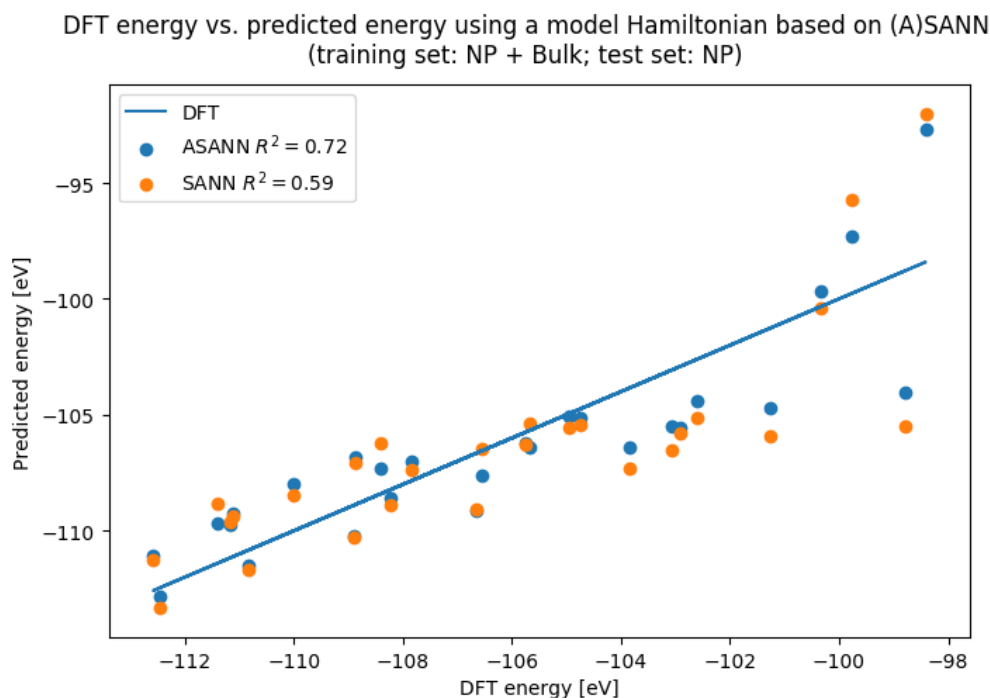


Figure 4.15 – Parity plot for the total energy predicted by a 2-body model Hamiltonian fitted using the SANN vs. ASANN coordination algorithm. The whole set of 27 nanoparticles + 6 Bulk structures was used as training set, and the test set was composed of the 27 Au-Cu bimetallic nanoparticles only for graphical considerations.

Au-Cu bulk structured alloys with various Au-Cu ratios. The coordinates are provided in the supplementary information.

When fitting only the nanoparticle set, the algorithms lead to model Hamiltonians with mean absolute deviation (MAD) of 0.8 eV for ASANN and 0.7 eV for SANN, see Figure C.1. In other words, both approaches are very similar, with SANN performing slightly better than ASANN. However, the ASANN-derived parameters are found to be more robust than their SANN-based counterparts when adding the bulk structured alloys to the training set (i.e. diversifying the training set). The fitted parameters are not only slightly less impacted (see Tables C.1 and C.2 in Appendix C), but the overall model performance is better with ASANN-based bonds (MAD near 1.5 eV for ASANN, versus 1.8 eV for SANN). This is especially true when using only the nanoparticles as test set. The predicted energy for the nanoparticles with respect to the DFT reference energies yields a coefficient of determination $R^2 = 0.72$ and $R^2 = 0.59$ with ASANN and SANN based 2-body patterns, respectively (see Figure 4.15). Similarly, the pure bulks are better described by ASANN (MAD=0.37 eV), than SANN (MAD=0.60 eV). Hence, the better performance of SANN for exclusively fitting the nanoparticles set can be traced back to error cancellation between coordination number overestimation and missing higher-order many-body terms at the surface. In summary, we conclude that ASANN provides a balanced and

physical sound description between bulk and surface-rich nanoparticles. Therefore, ASANN can also be applied to alloys of elements with quite different atomic radii. Similarly, ASANN could be used to provide a simple estimation of the surface of nanoparticles, by considering the union (or surface sum) of the planar edge of each coordination cap. A more rigorous approach using ASANN to define the surface of nanoparticles would be through the definition of a locally adaptative coordination radius, defining the local coarseness of the surface.

4.5 Conclusion

In this work we have extended the solid-angle based nearest-neighbor (SANN) algorithm to account for the local anisotropy, defining the anisotropically corrected solid-angle based nearest-neighbor (ASANN) method. ASANN allows to efficiently compute parameter-free coordination numbers and significantly improves over SANN at interfaces, where the coordination environment is highly anisotropic, such as nano-particles and partially desolvated ions at solid/liquid interfaces, where SANN leads to overestimations of the coordination number. In all situations tested, the ASANN coordination numbers are either identical or physically more sound than the ones provided by SANN and on par with system-specific, cut-off based coordination numbers. Therefore, ASANN provides a robust method to compute coordination numbers for describing alloy nano-particles, extended surfaces and (electrified) solid/liquid interfaces.

Bibliography

- [1] Reed, A. E.; Curtiss, L. A.; Weinhold, F. Intermolecular interactions from a natural bond orbital, donor-acceptor viewpoint. *Chem. Rev.* **1988**, *88*, 899–926.
- [2] Elgabarty, H.; Khaliullin, R. Z.; Kuhne, T. D. Covalency of hydrogen bonds in liquid water can be probed by proton nuclear magnetic resonance experiments. *Nat. Commun.* **2015**, *6*.
- [3] Pecher, L.; Laref, S.; Raupach, M.; Tonner, R. Ethers on Si(001): A Prime Example for the Common Ground between Surface Science and Molecular Organic Chemistry. *Angew. Chem., Int. Ed.* **2017**, *56*, 15150–15154.
- [4] Staub, R.; Iannuzzi, M.; Khaliullin, R. Z.; Steinmann, S. N. Energy Decomposition Analysis for Metal Surface–Adsorbate Interactions by Block Localized Wave Functions. *Journal of Chemical Theory and Computation* **2019**, *15*, 265–275.
- [5] Marcalo, J.; De Matos, A. P. A new definition of coordination number and its use in lanthanide and actinide coordination and organometallic chemistry. *Polyhedron* **1989**, *8*, 2431–2437.

- [6] Mahler, J.; Persson, I. A Study of the Hydration of the Alkali Metal Ions in Aqueous Solution. *Inorg. Chem.* **2012**, *51*, 425–438.
- [7] Varma, S.; Rempe, S. B. Coordination numbers of alkali metal ions in aqueous solutions. *Biophysical Chemistry* **2006**, *124*, 192 – 199.
- [8] Falicov, L. M.; Somorjai, G. A. Correlation between catalytic activity and bonding and coordination number of atoms and molecules on transition metal surfaces: Theory and experimental evidence. *Proceedings of the National Academy of Sciences* **1985**, *82*, 2207–2211.
- [9] Calle-Vallejo, F.; Martínez, J. I.; García-Lastra, J. M.; Sautet, P.; Loffreda, D. Fast Prediction of Adsorption Properties for Platinum Nanocatalysts with Generalized Coordination Numbers. *Angewandte Chemie International Edition* **2014**, *53*, 8316–8319.
- [10] Frank, F. C.; Kasper, J. S. Complex alloy structures regarded as sphere packings. I. Definitions and basic principles. *Acta Cryst* **1958**, *11*, 184–190.
- [11] O’Keeffe, M. A proposed rigorous definition of coordination number. *Acta Crystallogr A Cryst Phys Diffr Theor Gen Crystallogr* **1979**, *35*, 772–775.
- [12] Larsen, A. H. et al. The atomic simulation environment—a Python library for working with atoms. *Journal of Physics: Condensed Matter* **2017**, *29*, 273002.
- [13] Cordero, B.; Gómez, V.; Platero-Prats, A. E.; Revés, M.; Echeverría, J.; Cremades, E.; Barragán, F.; Alvarez, S. Covalent radii revisited. *Dalton Trans.* **2008**, 2832–2838.
- [14] Coulson, C. A.; Rushbrooke, G. S. On the Interpretation of Atomic Distribution Curves for Liquids. *Phys. Rev.* **1939**, *56*, 1216–1223.
- [15] Kirkwood, J. G.; Boggs, E. M. The Radial Distribution Function in Liquids. *The Journal of Chemical Physics* **1942**, *10*, 394–402.
- [16] Brostow, W. Radial distribution function peaks and coordination numbers in liquids and in amorphous solids. *Chemical Physics Letters* **1977**, *49*, 285 – 288.
- [17] Mikolaj, P. G.; Pings, C. J. The Use of the Coordination Number in the Interpretation of Fluid Structure. *Physics and Chemistry of Liquids* **1968**, *1*, 93–108.
- [18] Bernal, J. D. An attempt at a molecular theory of liquid structure. *Trans. Faraday Soc.* **1937**, *33*, 27–40.
- [19] Rahman, A. Liquid Structure and Self-Diffusion. *The Journal of Chemical Physics* **1966**, *45*, 2585–2592.

Chapter 4. Parameter-Free Coordination Numbers for Solutions and Interfaces

- [20] Finney, J. L.; Bernal, J. D. Random packings and the structure of simple liquids. I. The geometry of random close packing. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* **1970**, *319*, 479–493.
- [21] van Meel, J. A.; Fillion, L.; Valeriani, C.; Frenkel, D. A parameter-free, solid-angle based, nearest-neighbor algorithm. *The Journal of Chemical Physics* **2012**, *136*, 234107.
- [22] Ledoux, H. Computing the 3D Voronoi Diagram Robustly: An Easy Explanation. 4th International Symposium on Voronoi Diagrams in Science and Engineering (ISVD 2007). 2007; pp 117–129.
- [23] Higham, J.; Henchman, R. H. Locally adaptive method to define coordination shell. *The Journal of Chemical Physics* **2016**, *145*, 084108.
- [24] Noel, J. K.; Whitford, P. C.; Onuchic, J. N. The Shadow Map: A General Contact Definition for Capturing the Dynamics of Biomolecular Folding and Function. *The Journal of Physical Chemistry B* **2012**, *116*, 8692–8702.
- [25] Taylor, R. Which intermolecular interactions have a significant influence on crystal packing? *CrystEngComm* **2014**, *16*, 6852–6865.
- [26] Collins, R. On the entropy of the liquid state. *Proceedings of the Physical Society* **1965**, *86*, 199–212.
- [27] Oliphant, T. E. *Guide to NumPy*, 2nd ed.; CreateSpace Independent Publishing Platform: USA, 2015.
- [28] Walt, S. v. d.; Colbert, S. C.; Varoquaux, G. The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science & Engineering* **2011**, *13*, 22–30.
- [29] Van Rossum, G.; Drake, F. L. *Python 3 Reference Manual*; CreateSpace: Paramount, CA, 2009.
- [30] Kresse, G.; Hafner, J. Ab initio molecular dynamics for liquid metals. *Phys Rev B* **1993**, *47*, 558.
- [31] Kresse, G.; Furthmuller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys Rev B* **1996**, *54*, 11169.
- [32] Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Phys Rev Lett* **1996**, *77*, 3865.
- [33] Blochl, P. E. Projector augmented-wave method. *Phys Rev B* **1994**, *50*, 17953.
- [34] Kresse, G.; Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys Rev B* **1999**, *59*, 1758.

- [35] Viola, A.; Peron, J.; Kazmierczak, K.; Giraud, M.; Michel, C.; Sicard, L.; Perret, N.; Beau-nier, P.; Sicard, M.; Besson, M.; Piquemal, J.-Y. Unsupported shaped cobalt nanoparticles as efficient and recyclable catalysts for the solvent-free acceptorless dehydrogenation of alcohols. *Catal. Sci. Technol.* **2018**, *8*, 562–572.
- [36] Dhifallah, M.; Iachella, M.; Dhouib, A.; Di Renzo, F.; Loffreda, D.; Guesmi, H. Support Effects Examined by a Comparative Theoretical Study of Au, Cu, and CuAu Nanoclusters on Rutile and Anatase Surfaces. *J. Phys. Chem. C* **2019**, *123*, 4892–4902.
- [37] Mom, R. V.; Melissen, S. T. A. G.; Sautet, P.; Frenken, J. W. M.; Steinmann, S. N.; Groot, I. M. N. The Pressure Gap for Thiols: Methanethiol Self-Assembly on Au(111) from Vacuum to 1 bar. *J. Phys. Chem. C* **2019**, *123*, 12382–12389.
- [38] Steinmann, S. N.; Corminboeuf, C. Comprehensive Benchmarking of a Density-Dependent Dispersion Correction. *J Chem Theory Comput* **2011**, *7*, 3567.
- [39] Mathew, K.; Sundararaman, R.; Letchworth-Weaver, K.; Arias, T. A.; Hennig, R. G. Implicit solvation model for density-functional study of nanocrystal surfaces and reaction pathways. *J Chem Phys* **2014**, *140*, 084106.
- [40] Mathew, K.; Kolluru, V. S. C.; Mula, S.; Steinmann, S. N.; Hennig, R. G. Implicit self-consistent description of electrolyte in plane-wave density-functional theory. *J. Phys. Chem.* **2019**, DOI: 10.1063/1.5132354.
- [41] Steinmann, S. N.; Michel, C.; Schwiedernoch, R.; Sautet, P. Impacts of electrode potentials and solvents on the electroreduction of CO₂: a comparison of theoretical approaches. *Phys. Chem. Chem. Phys.* **2015**, *17*, 13949–13963.
- [42] Kastening, B.; Spinzig, S. Electrochemical polarization of activated carbon and graphite powder suspensions: Part II. Exchange of ions between electrolyte and pores. *Journal of Electroanalytical Chemistry and Interfacial Electrochemistry* **1986**, *214*, 295–302.
- [43] Zebardast, H. R.; Rogak, S.; Asselin, E. Potential of zero charge of glassy carbon at elevated temperatures. *Journal of Electroanalytical Chemistry* **2014**, *724*, 36–42.
- [44] Poon, J.; Batchelor-McAuley, C.; Tschulik, K.; Compton, R. G. Single graphene nanoplatelets: capacitance, potential of zero charge and diffusion coefficient. *Chemical Science* **2015**,
- [45] Steinmann, S. N.; Sautet, P. Assessing a First-Principles Model of an Electrochemical Interface by Comparison with Experiment. *J. Phys. Chem. C* **2016**, *120*, 5619–5623.
- [46] Steinmann, S. N.; Sautet, P.; Michel, C. Solvation free energies for periodic surfaces: comparison of implicit and explicit solvation models. *Phys. Chem. Chem. Phys.* **2016**, *18*, 31850–31861.

Chapter 4. Parameter-Free Coordination Numbers for Solutions and Interfaces

- [47] Rappe, A. K.; Casewit, C. J.; Colwell, K. S.; Goddard, W. A.; Skiff, W. M. UFF, a full periodic table force field for molecular mechanics and molecular dynamics simulations. *J Am Chem Soc* **1992**, *114*, 10024.
- [48] Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of simple potential functions for simulating liquid water. *J Chem Phys* **1983**, *79*, 926.
- [49] Phillips, J. C.; Braun, R.; Wang, W.; Gumbart, J.; Tajkhorshid, E.; Villa, E.; Chipot, C.; Skeel, R. D.; Kalé, L.; Schulten, K. Scalable molecular dynamics with NAMD. *J Comput Chem* **2005**, *26*, 1781.
- [50] Resasco, J.; Chen, L. D.; Clark, E.; Tsai, C.; Hahn, C.; Jaramillo, T. F.; Chan, K.; Bell, A. T. Promoter Effects of Alkali Metal Cations on the Electrochemical Reduction of Carbon Dioxide. *J. Am. Chem. Soc.* **2017**, *139*, 11277–11287.
- [51] Pérez-Gallent, E.; Marcandalli, G.; Figueiredo, M. C.; Calle-Vallejo, F.; Koper, M. T. M. Structure- and Potential-Dependent Cation Effects on CO Reduction at Copper Single-Crystal Electrodes. *J. Am. Chem. Soc.* **2017**, *139*, 16412–16419.
- [52] Ringe, S.; Clark, E. L.; Resasco, J.; Walton, A.; Seger, B.; Bell, A. T.; Chan, K. Understanding cation effects in electrochemical CO₂ reduction. *Energy Environ. Sci.* **2019**, 10.1039.C9EE01341E.
- [53] Ho, T. A.; Argyris, D.; Cole, D. R.; Striolo, A. Aqueous NaCl and CsCl Solutions Confined in Crystalline Slit-Shaped Silica Nanopores of Varying Degree of Protonation. *Langmuir* **2012**, *28*, 1256–1266.
- [54] Bose, S.; Kuila, T.; Mishra, A. K.; Rajasekar, R.; Kim, N. H.; Lee, J. H. Carbon-based nanostructured materials and their composites as supercapacitor electrodes. *J. Mater. Chem.* **2011**, *22*, 767–784.
- [55] Chen, X.; Paul, R.; Dai, L. Carbon-based supercapacitors for efficient energy storage. *Natl Sci Rev* **2017**, *4*, 453–489.
- [56] Merlet, C.; Péan, C.; Rotenberg, B.; Madden, P. A.; Simon, P.; Salanne, M. Simulating Supercapacitors: Can We Model Electrodes As Constant Charge Surfaces? *J. Phys. Chem. Lett.* **2013**, *4*, 264–268.
- [57] Forse, A. C.; Merlet, C.; Griffin, J. M.; Grey, C. P. New Perspectives on the Charging Mechanisms of Supercapacitors. *J. Am. Chem. Soc.* **2016**, *138*, 5731–5744.
- [58] Petrović, M. et al. The mechanism of caesium intercalation of graphene. *Nat Commun* **2013**, *4*, 2772.

- [59] Golub, D.; Soffer, A.; Oren, Y. The electrical double layer of carbon and graphite electrodes: Part V. Specific interactions with simple ions. *Journal of Electroanalytical Chemistry and Interfacial Electrochemistry* **1989**, *260*, 383–392.
- [60] Welch, D. A.; Mehdi, B. L.; Hatchell, H. J.; Faller, R.; Evans, J. E.; Browning, N. D. Using molecular dynamics to quantify the electrical double layer and examine the potential for its direct observation in the in-situ TEM. *Advanced Structural and Chemical Imaging* **2015**, *1*, 1.
- [61] Spackman, K. A. Signal Detection Theory: Valuable Tools for Evaluating Inductive Learning. Proceedings of the Sixth International Workshop on Machine Learning. San Francisco (CA), 1989; pp 160–163.
- [62] Fawcett, T. An introduction to ROC analysis. *Pattern Recognition Letters* **2006**, *27*, 861–874.
- [63] Ferri, C.; Hernández-Orallo, J.; Flach, P. A. A coherent interpretation of AUC as a measure of aggregated classification performance. Proceedings of the 28th International Conference on Machine Learning (ICML-11). 2011; pp 657–664.
- [64] Chen, W.; Schmidt, D.; Schneider, W. F.; Wolverton, C. Ordering and Oxygen Adsorption in Au–Pt/Pt(111) Surface Alloys. *J. Phys. Chem. C* **2011**, *115*, 17915–17924.
- [65] Zhu, B.; Creuze, J.; Mottet, C.; Legrand, B.; Guesmi, H. CO Adsorption-Induced Surface Segregation and Formation of Pd Chains on AuPd(100) Alloy: Density Functional Theory Based Ising Model and Monte Carlo Simulations. *J. Phys. Chem. C* **2016**, *120*, 350–359.
- [66] Vignola, E.; Steinmann, S. N.; Le Mapihan, K.; Vandegehuchte, B. D.; Curulla, D.; Sautet, P. Acetylene Adsorption on Pd–Ag Alloys: Evidence for Limited Island Formation and Strong Reverse Segregation from Monte Carlo Simulations. *J. Phys. Chem. C* **2018**, *122*, 15456–15463.



5 Efficient Recursive Least Squares Solver for Rank-Deficient Matrices

This Chapter is under review: R. Staub, and S. N. Steinmann *Appl. Math. Comput.*

Updating a linear least squares solution can be critical for near real-time signal-processing applications. The Greville algorithm proposes a simple formula for updating the pseudoinverse of a matrix $A \in \mathbb{R}^{n \times m}$ with rank r . In this paper, we explicitly derive a similar formula by maintaining a general rank factorization, which we call rank-Greville. Based on this formula, we implemented a recursive least squares algorithm exploiting the rank-deficiency of A , achieving the update of the minimum-norm least-squares solution in $O(mr)$ operations and, therefore, solving the linear least-squares problem from scratch in $O(nmr)$ operations. We empirically confirmed that this algorithm displays a better asymptotic time complexity than LAPACK solvers for rank-deficient matrices. The numerical stability of rank-Greville was found to be comparable to Cholesky-based solvers. Nonetheless, our implementation supports exact numerical representations of rationals, due to its remarkable algebraic simplicity.

5.1 Introduction

In this chapter, we are interested in computationally efficient algorithms for solving the recursive least-squares problem on rank-deficient matrices.

Let $\Gamma_i \in \mathbb{R}^m$ represent an observation of m variables, called regressors, associated with measurement y_i and let X_n be the unknown parameters of the linear relation:

$$A_n X_n + \epsilon_n = Y_n \quad (5.1)$$

where A_n is an $n \times m$ matrix representing n such observations, Y_n contains associated measurements and ϵ_n is the random disturbance:

$$A_n = \begin{pmatrix} \Gamma_1^\top \\ \Gamma_2^\top \\ \vdots \\ \Gamma_n^\top \end{pmatrix}, \quad Y_n = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} \quad (5.2)$$

The solution X_n to the general linear least-squares problem of equation 5.1 is defined as:

$$X_n = \underset{x \in S}{\operatorname{argmin}}(\|x\|_2), \quad S = \underset{x}{\operatorname{argmin}}(\|A_n x - Y_n\|_2) \quad (5.3)$$

This solution is unique¹, and sometimes called minimum-norm least-squares solution. Because of its uniqueness, it is sometimes simply referred to as the least-squares solution.

As demonstrated in the seminal paper by Penrose, the least-squares solution X_n can also be written²:

$$X_n = A_n^+ Y_n \quad (5.4)$$

where A_n^+ is the pseudoinverse of A_n , also called generalized inverse, or Moore-Penrose inverse. Due to its practical importance, the numerical determination of the generalized inverse remains an active topic of research.^{3;4}

The pseudoinverse $A^+ \in \mathbb{C}^{m \times n}$ of any matrix $A \in \mathbb{C}^{n \times m}$ is uniquely⁵ characterized by the four Penrose equations⁶:

$$A A^+ A = A \quad (5.5)$$

$$A^+ A A^+ = A^+ \quad (5.6)$$

$$(A A^+)^T = A A^+ \quad (5.7)$$

$$(A^+ A)^T = A^+ A \quad (5.8)$$

Here, we are interested in a particular problem, i.e., updating a least-squares solution (or the generalized inverse) when a new observation (Γ_{n+1}, y_{n+1}) is added. This is typically called the recursive least-squares (RLS) problem for which the updated solution X_{n+1} is usually written^{1;7}:

$$X_{n+1} = X_n + K \times (y_{n+1} - \Gamma_{n+1}^\top X_n) \quad (5.9)$$

where $y_{n+1} - \Gamma_{n+1}^\top X_n$ is the predicted residual (or *a priori* error), and K is called the Kalman gain vector.

Algorithms that allow to update an already known previous solution can be of critical importance for embedded-systems signal processing, for example, as near real-time solutions might be required and new observations are added continuously⁷. Therefore, recursive least-squares algorithms significantly benefit from the computational efficiency introduced by updating the least-squares solution instead of recomputing it from scratch.

If A_n has full column rank, the recursive least-squares solution X_{n+1} of equation (5.9) can be straightforwardly computed using normal equations¹. This RLS algorithm has a time complexity of $O(m^2)$ for each update. Therefore, computing the solution for n successive observations, using equation (5.9), lead to a total time complexity of $O(nm^2)$.

However in the general case, A_n can be rank deficient, i.e. neither full column rank nor full row rank. This is indeed the case if we want to sample a large variable space (column deficiency), while accumulating data redundancy on the subspace of observed variables (row deficiency). Handling rank deficiency is, for example, desirable in neurocomputational learning applications⁸. Several algorithms have been developed to solve the rank-deficient recursive least-squares problem. In particular, the Greville algorithm⁹ was designed for the recursive least-squares problem specifically, whereas most of the other algorithms are common least-squares solvers (based on Cholesky decomposition, QR factorization, SVD, ...) adapted to support updates of A_n (without the need to recompute the whole solution)¹.

The Greville algorithm provides an updated least-squares solution, and additionally, an updated pseudoinverse at the same cost. This update step still has computational complexity in $O(m^2)$, independently of the rank deficiency of A_n . This leads to a $O(nm^2)$ time complexity for the computation of the full solution.ⁱ Variants of this algorithm were developed¹⁰ based on the implicit decomposition of A_n , but still with an $O(m^2)$ update complexityⁱⁱ.

In this chapter, we write and implement a recursive least-squares algorithm that has single-update time complexity in $O(mr)$ (i.e. $O(nmr)$ total time complexity), where r is the rank of

ⁱNote that one can easily reach $O(\max(n, m) \min(n, m)^2)$ using the property $(A_n^\top)^+ = (A_n^+)^T$.

ⁱⁱUsing the notations defined below, Albert and Sittler maintain the $m \times m$ matrices $(1 - C_n^\top \tilde{C}_n)$ and $(\tilde{C}_n^\top P_n^{-1} \tilde{C}_n)$, whereas we maintain C_n , \tilde{C}_n and P_n^{-1} (whose sizes are reduced).

the matrix A_n . The underlying idea is to maintain a general rank decomposition into a full row rank matrix (purely non-overdetermined system) and a full column rank matrix (purely non-underdetermined system), which are much easier to treat in a recursive least-squares procedure. Indeed, due to the rank deficiency of A_n these matrices have reduced sizes, leading to more efficient updates.

The remarkable simplicity of this approach makes it compatible with exact numerical representations in practice, without the need to use expensive symbolic computing. We also explore slightly more sophisticated rank decompositions, effectively bridging the gap between the Greville algorithm and QR-based recursive least-squares solvers.

5.2 Derivation

5.2.1 General non-orthogonal rank factorization

Let A_n be a $n \times m$ matrix of rank r . A_n can be expressed as the product⁹:

$$A_n = B_n C_n \tag{5.10}$$

where B_n is a $n \times r$ matrix, and C_n is a $r \times m$ matrix, both of rank r , which corresponds to a full-rank factorization.

Let us consider a maximal free family $\mathcal{B}_{\mathcal{A}}$ of observations among observations $\Gamma_i \in \mathbb{R}^m$ in A_n . Note that $\mathcal{B}_{\mathcal{A}}$ is a basis of $\text{Im}(A_n^\top)$. Such basis is represented by the rows of C_n . Interpreting rows of C_n as observations $\Gamma_{C_i} \in \mathbb{R}^m$, we find that each observation Γ_{C_i} in C_n is linearly independent of the others. Hence C_n can then be seen as a purely non-overdetermined system. This system can be thought as linking the fitted value of each observation in $\mathcal{B}_{\mathcal{A}}$, to the fitted values for the m variables themselves.

Interpreting rows of B_n as observations $\gamma_i \in \mathbb{R}^r$, we find that each observation γ_i in B_n is the observation Γ_i of A_n expressed in the $\mathcal{B}_{\mathcal{A}}$ basis. Therefore, B_n can be seen as a purely non-underdetermined system, since each observation in $\mathcal{B}_{\mathcal{A}}$ is observed at least once. One can consider that this system links the value of each observation Γ_i in A_n , to the fitted values for each observation in $\mathcal{B}_{\mathcal{A}}$.

Theorem 1. *The pseudoinverse A_n^+ of A_n can be computed in $O(nmr)$ if matrices B_n and C_n verifying equation 5.10 are known. An explicit formula is then given by:*

$$A_n^+ = C_n^\top (C_n C_n^\top)^{-1} (B_n^\top B_n)^{-1} B_n^\top \tag{5.11}$$

Proof. By definition of B_n and C_n , $B_n^\top B_n$ and $C_n C_n^\top$ are both $r \times r$ matrices of rank r , and therefore non-singular. As a consequence, the explicit formula given is well defined as long as

B_n and C_n correspond to a full-rank factorization.

It is straightforward to check that Eq. 5.11 satisfies all four Penrose equations (Eq. 5.5 to 5.8), and therefore, represents an acceptable pseudoinverse of A_n ⁹. By the unicity of the Moore-Penrose inverse, we conclude that $A_n^+ = C_n^\top (C_n C_n^\top)^{-1} (B_n^\top B_n)^{-1} B_n^\top$. Computing the pseudoinverse can, therefore, be reduced to computing the inverse of two $r \times r$ matrices and three matrix multiplications giving rise (in any order) to a total $O(nmr)$ time complexity, that could even be reduced by using faster algorithms¹¹. \square

We are now interested in the update of the pseudoinverse of A_n when adding a new observation Γ_{n+1} . Let us define A_{n+1} as:

$$A_{n+1} = \begin{pmatrix} A_n \\ \Gamma_{n+1}^\top \end{pmatrix} = B_{n+1} C_{n+1} \quad (5.12)$$

We distinguish two cases depending on the linear dependency of Γ_{n+1} with respect to previous observations $\Gamma_1, \dots, \Gamma_n$. Note that we can equally well only consider the observations in $\mathcal{B}_{\mathcal{A}}$, since $\mathcal{B}_{\mathcal{A}}$ is a basis for $\text{Vect}(\Gamma_1, \dots, \Gamma_n)$.

Let $P_{\mathcal{B}_{\mathcal{A}}}$ be the projector into $\text{Vect}(\mathcal{B}_{\mathcal{A}}) = \text{Im}(C^\top)$. We define $\Gamma_{n+1}^p \in \mathbb{R}^m$ the projection of Γ_{n+1} into $\text{Vect}(\mathcal{B}_{\mathcal{A}})$. We also define $\gamma_{n+1}^p \in \mathbb{R}^r$ as Γ_{n+1}^p expressed in the $\mathcal{B}_{\mathcal{A}}$ basis. If $\mathcal{B}_{\mathcal{A}}$ was an orthonormal basis, the decomposition γ_{n+1}^p could be easily computed by inner products with $\gamma_{n+1}^p = C_n \Gamma_{n+1}$. However, in the general (non-orthonormal) case, the decomposition γ_{n+1}^p of $P_{\mathcal{B}_{\mathcal{A}}} \Gamma_{n+1}$ can be obtained using the dual $\tilde{\mathcal{B}}_{\mathcal{A}}$ of $\mathcal{B}_{\mathcal{A}}$, represented by the rows of \tilde{C}_n defined as:

$$\begin{aligned} \tilde{C}_n &= (C_n C_n^\top)^{-1} C_n \\ &= \sigma_n^{-1} C_n \end{aligned} \quad (5.13)$$

where $\sigma_n = C_n C_n^\top$ is the Gram matrix of observations in $\mathcal{B}_{\mathcal{A}}$. γ_{n+1}^p can then be expressed by:

$$\begin{aligned} \gamma_{n+1}^p &= \tilde{C}_n P_{\mathcal{B}_{\mathcal{A}}} \Gamma_{n+1} \\ &= \tilde{C}_n \Gamma_{n+1} \end{aligned} \quad (5.14)$$

and Γ_{n+1}^p can then be expressed by:

$$\begin{aligned} \Gamma_{n+1}^p &= P_{\mathcal{B}_{\mathcal{A}}} \Gamma_{n+1} \\ &= C_n^\top \tilde{C}_n \Gamma_{n+1} \\ &= C_n^\top \gamma_{n+1}^p \end{aligned} \quad (5.15)$$

We define the rejection vector $\Gamma_{n+1}^{rej} \in \mathbb{R}^m$ associated with the projection of Γ_{n+1} into $\mathcal{B}_{\mathcal{A}}$:

$$\begin{aligned}
 \Gamma_{n+1} &= P_{\mathcal{B}_{\mathcal{A}}} \Gamma_{n+1} + \Gamma_{n+1}^{rej} \\
 &= \Gamma_{n+1}^p + \Gamma_{n+1}^{rej} \\
 &= C_n^\top \gamma_{n+1}^p + \Gamma_{n+1}^{rej} \\
 &= C_n^\top \tilde{C}_n \Gamma_{n+1} + \Gamma_{n+1}^{rej}
 \end{aligned} \tag{5.16}$$

It becomes clear that Γ_{n+1} is linearly dependent from the previous observations $\Gamma_1, \dots, \Gamma_n$, if and only if, Γ_{n+1}^{rej} is null. Note that γ_{n+1}^p , Γ_{n+1}^p and Γ_{n+1}^{rej} can be computed in $O(mr)$ if \tilde{C}_n and C_n are already known.

The pseudoinverse A_n^+ can then be rewritten:

$$A_n^+ = \tilde{C}_n^\top P_n^{-1} B_n^\top \tag{5.17}$$

where

$$P_n = B_n^\top B_n = \sum_{i=1}^n \gamma_i^p \cdot \gamma_i^{p\top} \tag{5.18}$$

We finally define $\zeta_{n+1} \in \mathbb{R}^r$ and $\beta_{n+1} \in \mathbb{R}^n$ as:

$$\zeta_{n+1} = P_n^{-1} \gamma_{n+1}^p, \quad \beta_{n+1} = B_n \zeta_{n+1} \tag{5.19}$$

Note that if \tilde{C}_n , C_n and P_n^{-1} are already known, ζ_{n+1} can be computed in $O(mr)$, but β_{n+1} can only be computed in $O(\max(m, n)r)$ if B_n is also known.

Theorem 2. *If $\Gamma_{n+1} \neq 0$ is linearly independent from previous observations, the pseudoinverse A_{n+1}^+ of A_{n+1} can be updated in $O(mn)$ if A_n^+ , B_n , P_n^{-1} , C_n and \tilde{C}_n are known. An explicit formula is then given by:*

$$A_{n+1}^+ = \begin{pmatrix} A_n^+ & 0 \end{pmatrix} + \frac{\Gamma_{n+1}^{rej}}{\|\Gamma_{n+1}^{rej}\|_2^2} \begin{pmatrix} -\beta_{n+1}^\top & 1 \end{pmatrix} \tag{5.20}$$

Proof. First, let us observe how the full-rank factorization $A_{n+1} = B_{n+1} C_{n+1}$ is impacted by adding an observation Γ_{n+1} linearly independent from previous observations. Concerning C_{n+1} , we have:

$$\begin{aligned}
 \text{Im}(C_{n+1}^\top) &= \text{Im}(A_{n+1}^\top) = \text{Vect}(\mathcal{B}_{\mathcal{A}} \cup \{\Gamma_{n+1}\}) \neq \text{Vect}(\mathcal{B}_{\mathcal{A}}) = \text{Im}(A_n^\top) = \text{Im}(C_n^\top) \\
 &\Rightarrow C_{n+1} \neq C_n
 \end{aligned} \tag{5.21}$$

Adding Γ_{n+1} to the rows of C_n leads to:

$$C_{n+1} = \begin{pmatrix} C_n \\ \Gamma_{n+1}^\top \end{pmatrix}, \quad B_{n+1} = \begin{pmatrix} B_n & 0 \\ 0 & 1 \end{pmatrix} \quad (5.22)$$

It becomes clear from this definition that B_{n+1} has full column rank since B_n has full column rank, and also that C_{n+1} has full row rank. Therefore, B_{n+1} and C_{n+1} represent an acceptable full-rank decomposition of A_{n+1} , since we have:

$$B_{n+1}C_{n+1} = \begin{pmatrix} B_n & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} C_n \\ \Gamma_{n+1}^\top \end{pmatrix} = \begin{pmatrix} B_n C_n \\ \Gamma_{n+1}^\top \end{pmatrix} = \begin{pmatrix} A_n \\ \Gamma_{n+1}^\top \end{pmatrix} = A_{n+1} \quad (5.23)$$

Second, we apply Theorem 1 to A_{n+1} :

$$\begin{aligned} A_{n+1}^+ &= C_{n+1}^\top (C_{n+1} C_{n+1}^\top)^{-1} (B_{n+1}^\top B_{n+1})^{-1} B_{n+1}^\top \\ &= \begin{pmatrix} C_n^\top & \Gamma_{n+1} \end{pmatrix} \begin{pmatrix} C_n C_n^\top & C_n \Gamma_{n+1} \\ (C_n \Gamma_{n+1})^\top & \Gamma_{n+1}^\top \Gamma_{n+1} \end{pmatrix}^{-1} \begin{pmatrix} B_n^\top B_n & 0 \\ 0 & 1 \end{pmatrix}^{-1} \begin{pmatrix} B_n^\top & 0 \\ 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} C_n^\top & \Gamma_{n+1} \end{pmatrix} \begin{pmatrix} C_n C_n^\top & C_n \Gamma_{n+1} \\ (C_n \Gamma_{n+1})^\top & \Gamma_{n+1}^\top \Gamma_{n+1} \end{pmatrix}^{-1} \begin{pmatrix} (B_n^\top B_n)^{-1} & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} B_n^\top & 0 \\ 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} C_n^\top & \Gamma_{n+1} \end{pmatrix} \begin{pmatrix} \sigma_n & C_n \Gamma_{n+1} \\ (C_n \Gamma_{n+1})^\top & \Gamma_{n+1}^\top \Gamma_{n+1} \end{pmatrix}^{-1} \begin{pmatrix} P_n^{-1} B_n^\top & 0 \\ 0 & 1 \end{pmatrix} \end{aligned} \quad (5.24)$$

Finally, we apply a generic block-wise inversion scheme:

$$\begin{aligned} \begin{pmatrix} D & E \\ E^\top & F \end{pmatrix}^{-1} &= \begin{pmatrix} D^{-1} + D^{-1} E (F - E^\top D^{-1} E)^{-1} E^\top D^{-1} & -D^{-1} E (F - E^\top D^{-1} E)^{-1} \\ -(F - E^\top D^{-1} E)^{-1} E^\top D^{-1} & (F - E^\top D^{-1} E)^{-1} \end{pmatrix} \\ &= \begin{pmatrix} D^{-1} + D^{-1} E S^{-1} E^\top D^{-1} & -D^{-1} E S^{-1} \\ -S^{-1} E^\top D^{-1} & S^{-1} \end{pmatrix} \end{aligned} \quad (5.25)$$

where $S = F - E^\top D^{-1} E$, $D = \sigma_n$, $E = C_n \Gamma_{n+1}$ and $F = \Gamma_{n+1}^\top \Gamma_{n+1}$, since σ_n is non-singular and

$\Gamma_{n+1}^\top \Gamma_{n+1}$ is not null since $\Gamma_{n+1} \neq 0$. This leads to the pseudoinverse formula:

$$\begin{aligned}
 A_{n+1}^+ &= \begin{pmatrix} C_n^\top & \Gamma_{n+1} \end{pmatrix} \begin{pmatrix} \sigma_n & C_n \Gamma_{n+1} \\ (C_n \Gamma_{n+1})^\top & \Gamma_{n+1}^\top \Gamma_{n+1} \end{pmatrix}^{-1} \begin{pmatrix} P_n^{-1} B_n^\top & 0 \\ 0 & 1 \end{pmatrix} \\
 &= \begin{pmatrix} C_n^\top & \Gamma_{n+1} \end{pmatrix} \begin{pmatrix} \sigma_n^{-1} + \sigma_n^{-1} C_n \Gamma_{n+1} S^{-1} (C_n \Gamma_{n+1})^\top \sigma_n^{-1} & -\sigma_n^{-1} C_n \Gamma_{n+1} S^{-1} \\ -S^{-1} (C_n \Gamma_{n+1})^\top \sigma_n^{-1} & S^{-1} \end{pmatrix} \begin{pmatrix} P_n^{-1} B_n^\top & 0 \\ 0 & 1 \end{pmatrix} \\
 &= \begin{pmatrix} C_n^\top & \Gamma_{n+1} \end{pmatrix} \begin{pmatrix} \sigma_n^{-1} + \gamma_{n+1}^p S^{-1} \gamma_{n+1}^{p\top} & -S^{-1} \gamma_{n+1}^p \\ -S^{-1} \gamma_{n+1}^p & S^{-1} \end{pmatrix} \begin{pmatrix} P_n^{-1} B_n^\top & 0 \\ 0 & 1 \end{pmatrix} \\
 &= \begin{pmatrix} C_n^\top & \Gamma_{n+1} \end{pmatrix} \left(\begin{pmatrix} \sigma_n^{-1} & 0 \\ 0 & 0 \end{pmatrix} + S^{-1} \begin{pmatrix} \gamma_{n+1}^p \gamma_{n+1}^{p\top} & -\gamma_{n+1}^p \\ -\gamma_{n+1}^p & 1 \end{pmatrix} \right) \begin{pmatrix} P_n^{-1} B_n^\top & 0 \\ 0 & 1 \end{pmatrix} \\
 &= \begin{pmatrix} A_n^+ & 0 \end{pmatrix} + S^{-1} \begin{pmatrix} C_n^\top & \Gamma_{n+1} \end{pmatrix} \begin{pmatrix} \gamma_{n+1}^p \gamma_{n+1}^{p\top} & -\gamma_{n+1}^p \\ -\gamma_{n+1}^p & 1 \end{pmatrix} \begin{pmatrix} P_n^{-1} B_n^\top & 0 \\ 0 & 1 \end{pmatrix} \\
 &= \begin{pmatrix} A_n^+ & 0 \end{pmatrix} + S^{-1} \begin{pmatrix} C_n^\top \gamma_{n+1}^p \gamma_{n+1}^{p\top} P_n^{-1} B_n^\top - \Gamma_{n+1} \gamma_{n+1}^{p\top} P_n^{-1} B_n^\top & -C_n^\top \gamma_{n+1}^p + \Gamma_{n+1} \end{pmatrix} \\
 &= \begin{pmatrix} A_n^+ & 0 \end{pmatrix} + S^{-1} \begin{pmatrix} -\Gamma_{n+1}^{rej} \beta_{n+1}^\top & \Gamma_{n+1}^{rej} \end{pmatrix}
 \end{aligned} \tag{5.26}$$

where the Schur complement S of σ_n is written as:

$$\begin{aligned}
 S &= \Gamma_{n+1}^\top \Gamma_{n+1} - \Gamma_{n+1}^\top C_n^\top \sigma_n^{-1} C_n \Gamma_{n+1} \\
 &= \Gamma_{n+1}^\top \Gamma_{n+1} - \Gamma_{n+1}^\top \Gamma_{n+1}^p \\
 &= \Gamma_{n+1}^\top \Gamma_{n+1}^{rej} \\
 &= \Gamma_{n+1}^{rej\top} \Gamma_{n+1}^{rej} = \left\| \Gamma_{n+1}^{rej} \right\|_2^2
 \end{aligned} \tag{5.27}$$

S is, therefore, the square of the norm of the component of Γ_{n+1} along the orthogonal complement of $\text{Im}(A_n^\top)$. S is invertible since $\Gamma_{n+1}^{rej} \neq 0$.

Γ_{n+1}^{rej} and β_{n+1} can be computed in $O(\max(m, n)r)$ if C_n, \tilde{C}_n, B_n and P_n^{-1} are already known. Therefore, the time complexity bottleneck is the outer product $\Gamma_{n+1}^{rej} \beta_{n+1}^\top$, leading to a total update in $O(mn)$. \square

Theorem 3. *If Γ_{n+1} is a linear combination of previous observations, the pseudoinverse A_{n+1}^+ of A_{n+1} can be updated in $O(mn)$ if $A_n^+, B_n, P_n^{-1}, C_n$ and \tilde{C}_n are known. An explicit formula is then given by:*

$$A_{n+1}^+ = \begin{pmatrix} A_n^+ & 0 \end{pmatrix} + \frac{\tilde{C}_n^\top \zeta_{n+1}}{1 + \gamma_{n+1}^p \zeta_{n+1}^\top \zeta_{n+1}} \begin{pmatrix} -\beta_{n+1}^\top & 1 \end{pmatrix} \tag{5.28}$$

Proof. First, let us observe how the full-rank factorization $A_{n+1} = B_{n+1} C_{n+1}$ is impacted by

adding an observation Γ_{n+1} that is a linear combination of previous observations. Notice that:

$$\begin{aligned}\Gamma_{n+1} &= P_{\mathcal{B}_{\mathcal{A}}} \Gamma_{n+1} + \Gamma_{n+1}^{rej} = P_{\mathcal{B}_{\mathcal{A}}} \Gamma_{n+1} = \Gamma_{n+1}^p \\ &= C_n^\top \tilde{C}_n \Gamma_{n+1} \\ &= C_n^\top \gamma_{n+1}^p\end{aligned}\tag{5.29}$$

Since $\Gamma_{n+1} \in \text{Vect}(\Gamma_1, \dots, \Gamma_n) = \text{Vect}(\mathcal{B}_{\mathcal{A}})$, $\mathcal{B}_{\mathcal{A}}$ is still a basis for $\text{Vect}(\Gamma_1, \dots, \Gamma_n, \Gamma_{n+1})$. As a consequence, we can take $C_{n+1} = C_n$, leading to:

$$C_{n+1} = C_n, \quad B_{n+1} = \begin{pmatrix} B_n \\ \gamma_{n+1}^p{}^\top \end{pmatrix}\tag{5.30}$$

From this definition follows that C_{n+1} still has full row rank, and also that B_{n+1} has full column rank since B_n has full column rank. Therefore, B_{n+1} and C_{n+1} represent an acceptable full-rank decomposition of A_{n+1} , since we have:

$$B_{n+1} C_{n+1} = \begin{pmatrix} B_n \\ \gamma_{n+1}^p{}^\top \end{pmatrix} C_n = \begin{pmatrix} B_n C_n \\ \gamma_{n+1}^p{}^\top C_n \end{pmatrix} = \begin{pmatrix} A_n \\ \Gamma_{n+1}{}^\top \end{pmatrix} = A_{n+1}\tag{5.31}$$

Second, we apply Theorem 1 to A_{n+1} :

$$\begin{aligned}A_{n+1}^+ &= C_{n+1}^\top (C_{n+1} C_{n+1}^\top)^{-1} (B_{n+1}^\top B_{n+1})^{-1} B_{n+1}^\top \\ &= C_n^\top (C_n C_n^\top)^{-1} \left(\begin{pmatrix} B_n^\top & \gamma_{n+1}^p \end{pmatrix} \begin{pmatrix} B_n \\ \gamma_{n+1}^p{}^\top \end{pmatrix} \right)^{-1} \begin{pmatrix} B_n^\top & \gamma_{n+1}^p \end{pmatrix} \\ &= C_n^\top (C_n C_n^\top)^{-1} \left(B_n^\top B_n + \gamma_{n+1}^p \gamma_{n+1}^p{}^\top \right)^{-1} \begin{pmatrix} B_n^\top & \gamma_{n+1}^p \end{pmatrix} \\ &= \tilde{C}_n^\top \left(P_n + \gamma_{n+1}^p \gamma_{n+1}^p{}^\top \right)^{-1} \begin{pmatrix} B_n^\top & \gamma_{n+1}^p \end{pmatrix}\end{aligned}\tag{5.32}$$

Finally, we apply the Sherman-Morrison formula stating that for any non-singular matrix $G \in \mathbb{R}^{n \times n}$ and any vector $v \in \mathbb{R}^n$, if $G + vv^\top$ is non-singular, then:

$$G + vv^\top = G^{-1} - \frac{G^{-1} v v^\top G^{-1}}{1 + v^\top G^{-1} v}\tag{5.33}$$

with $G = P_n$ and $v = \gamma_{n+1}^p$, since P_n and $P_{n+1} = B_{n+1} B_{n+1}^\top$ are non-singularⁱⁱⁱ. This leads to

ⁱⁱⁱ $P_{n+1} = B_{n+1} B_{n+1}^\top$ is non-singular, since P_{n+1} is square with full rank. Indeed, $\text{rank}(P_{n+1}) = \text{rank}(B_{n+1}) = r$, using theorem 5.5.4 of¹²

the pseudoinverse formula:

$$\begin{aligned}
 A_{n+1}^+ &= \tilde{C}_n^\top \left(P_n + \gamma_{n+1}^p \gamma_{n+1}^{p\top} \right)^{-1} \begin{pmatrix} B_n^\top & \gamma_{n+1}^p \end{pmatrix} \\
 &= \tilde{C}_n^\top \left(P_n^{-1} - \frac{P_n^{-1} \gamma_{n+1}^p \gamma_{n+1}^{p\top} P_n^{-1}}{1 + \gamma_{n+1}^{p\top} P_n^{-1} \gamma_{n+1}^p} \right) \begin{pmatrix} B_n^\top & \gamma_{n+1}^p \end{pmatrix} \\
 &= \tilde{C}_n^\top \left(P_n^{-1} - \frac{\zeta_{n+1} \zeta_{n+1}^\top}{1 + \gamma_{n+1}^{p\top} \zeta_{n+1}} \right) \begin{pmatrix} B_n^\top & \gamma_{n+1}^p \end{pmatrix} \\
 &= \begin{pmatrix} A_n^+ & \tilde{C}_n^\top \zeta_{n+1} \end{pmatrix} - \tilde{C}_n^\top \frac{\zeta_{n+1} \zeta_{n+1}^\top}{1 + \gamma_{n+1}^{p\top} \zeta_{n+1}} \begin{pmatrix} B_n^\top & \gamma_{n+1}^p \end{pmatrix} \\
 &= \begin{pmatrix} A_n^+ & \tilde{C}_n^\top \zeta_{n+1} \end{pmatrix} - \frac{\tilde{C}_n^\top \zeta_{n+1}}{1 + \gamma_{n+1}^{p\top} \zeta_{n+1}} \begin{pmatrix} \beta_{n+1}^\top & \zeta_{n+1}^\top \gamma_{n+1}^p \end{pmatrix} \\
 &= \begin{pmatrix} A_n^+ & 0 \end{pmatrix} + \frac{\tilde{C}_n^\top \zeta_{n+1}}{1 + \gamma_{n+1}^{p\top} \zeta_{n+1}} \begin{pmatrix} -\beta_{n+1}^\top & 1 + \gamma_{n+1}^{p\top} \zeta_{n+1} - \zeta_{n+1}^\top \gamma_{n+1}^p \end{pmatrix} \\
 &= \begin{pmatrix} A_n^+ & 0 \end{pmatrix} + \frac{\tilde{C}_n^\top \zeta_{n+1}}{1 + \gamma_{n+1}^{p\top} \zeta_{n+1}} \begin{pmatrix} -\beta_{n+1}^\top & 1 \end{pmatrix}
 \end{aligned} \tag{5.34}$$

γ_{n+1}^p , ζ_{n+1} and β_{n+1} can be computed in $O(\max(m, n)r)$ if \tilde{C}_n , B_n and P_n^{-1} are already known. Therefore, the time complexity bottleneck is the outer product $(\tilde{C}_n^\top \zeta_{n+1}) \beta_{n+1}^\top$, leading to a total update in $O(mn)$. \square

Corollary 3.1. *For any observation $\Gamma_{n+1} \in \mathbb{R}^m$, the pseudoinverse A_{n+1}^+ of A_{n+1} can be updated in $\Theta(mn)$ if A_n^+ , B_n , P_n^{-1} , C_n and \tilde{C}_n are known.*

Indeed, at least $n \times m$ terms of the pseudoinverse need to be updated when adding a new observation, in the general case^{iv}. Therefore, the pseudoinverse update has a fundamental cost component that cannot be improved, hence the $\Theta(mn)$ complexity. This limitation is not present in the recursive least-squares problem. In this problem, we are only interested in updating the least-squares solution X_{n+1} when adding a new observation Γ_{n+1} with associated target y_{n+1} :

$$X_{n+1} = \begin{pmatrix} A_n & \\ & \Gamma_{n+1}^\top \end{pmatrix}^+ \begin{pmatrix} Y_n \\ y_{n+1} \end{pmatrix} = A_{n+1}^+ \begin{pmatrix} Y_n \\ y_{n+1} \end{pmatrix} \tag{5.35}$$

Theorem 4. *If $\Gamma_{n+1} \neq 0$ is linearly independent from previous observations, the least-squares solution X_{n+1} can be updated in $O(mr)$ if X_n , C_n and \tilde{C}_n are known. An explicit formula (in*

^{iv}To be convinced, consider $A_n = I_n$ the identity matrix and $\Gamma_{n+1} = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}$.

the form of equation 5.9) is then given by:

$$X_{n+1} = X_n + \frac{\Gamma_{n+1}^{rej}}{\|\Gamma_{n+1}^{rej}\|_2^2} (y_{n+1} - \Gamma_{n+1}^\top X_n) \quad (5.36)$$

Proof. First, let us inject theorem 2 into the definition of X_{n+1} :

$$\begin{aligned} X_{n+1} &= A_{n+1}^+ \begin{pmatrix} Y_n \\ y_{n+1} \end{pmatrix} \\ &= A_n^+ Y_n + \frac{\Gamma_{n+1}^{rej}}{\|\Gamma_{n+1}^{rej}\|_2^2} \begin{pmatrix} -\beta_{n+1}^\top & 1 \end{pmatrix} \begin{pmatrix} Y_n \\ y_{n+1} \end{pmatrix} \\ &= X_n + \frac{\Gamma_{n+1}^{rej}}{\|\Gamma_{n+1}^{rej}\|_2^2} (y_{n+1} - \beta_{n+1}^\top Y_n) \end{aligned} \quad (5.37)$$

Let us simplify further this equation by recognizing $\beta_{n+1}^\top Y_n$ as the fitted target associated with Γ_{n+1} :

$$\begin{aligned} X_{n+1} &= X_n + \frac{\Gamma_{n+1}^{rej}}{\|\Gamma_{n+1}^{rej}\|_2^2} (y_{n+1} - \Gamma_{n+1}^\top \tilde{C}_n^\top P_n^{-1} B_n^\top Y_n) \\ &= X_n + \frac{\Gamma_{n+1}^{rej}}{\|\Gamma_{n+1}^{rej}\|_2^2} (y_{n+1} - \Gamma_{n+1}^\top A_n^+ Y_n) \\ &= X_n + \frac{\Gamma_{n+1}^{rej}}{\|\Gamma_{n+1}^{rej}\|_2^2} (y_{n+1} - \Gamma_{n+1}^\top X_n) \\ &= X_n + \frac{\Gamma_{n+1}^{rej}}{\|\Gamma_{n+1}^{rej}\|_2^2} \Delta y_{n+1} \end{aligned} \quad (5.38)$$

where $\Delta y_{n+1} = y_{n+1} - \Gamma_{n+1}^\top X_n$ is the difference between the expected/fitted target (i.e. $\Gamma_{n+1}^\top X_n$) and the real target y_{n+1} associated with the new observation Γ_{n+1} (i.e. the predicted residual, or *a priori* error). We identify $\frac{\Gamma_{n+1}^{rej}}{\|\Gamma_{n+1}^{rej}\|_2^2}$ to be the associated Kalman gain vector in this case¹.

Γ_{n+1}^{rej} can be computed in $O(mr)$ if C_n and \tilde{C}_n are already known, which is the time complexity bottleneck of the whole update step. \square

Theorem 5. *If Γ_{n+1} is a linear combination of previous observations, the least-squares solution X_{n+1} can be updated in $O(mr)$ if X_n , C_n , \tilde{C}_n and P_n^{-1} are known. An explicit formula (in the*

form of equation 5.9) is then given by:

$$X_{n+1} = X_n + \frac{\tilde{C}_n^\top \zeta_{n+1}}{1 + \gamma_{n+1}^p \zeta_{n+1}^\top \zeta_{n+1}} (y_{n+1} - \Gamma_{n+1}^\top X_n) \quad (5.39)$$

Proof. Let us proceed similarly to theorem 4, by injecting theorem 3 into the definition of X_{n+1} :

$$\begin{aligned} X_{n+1} &= A_{n+1} + \begin{pmatrix} Y_n \\ y_{n+1} \end{pmatrix} \\ &= A_n + Y_n + \frac{\tilde{C}_n^\top \zeta_{n+1}}{1 + \gamma_{n+1}^p \zeta_{n+1}^\top \zeta_{n+1}} \begin{pmatrix} -\beta_{n+1}^\top & 1 \end{pmatrix} \begin{pmatrix} Y_n \\ y_{n+1} \end{pmatrix} \\ &= X_n + \frac{\tilde{C}_n^\top \zeta_{n+1}}{1 + \gamma_{n+1}^p \zeta_{n+1}^\top \zeta_{n+1}} (y_{n+1} - \beta_{n+1}^\top Y_n) \\ &= X_n + \frac{\tilde{C}_n^\top \zeta_{n+1}}{1 + \gamma_{n+1}^p \zeta_{n+1}^\top \zeta_{n+1}} (y_{n+1} - \Gamma_{n+1}^\top X_n) \\ &= X_n + \frac{\tilde{C}_n^\top \zeta_{n+1}}{1 + \gamma_{n+1}^p \zeta_{n+1}^\top \zeta_{n+1}} \Delta y_{n+1} \end{aligned} \quad (5.40)$$

where $\Delta y_{n+1} = y_{n+1} - \Gamma_{n+1} X_n^\top$ (i.e. the predicted residual, or *a priori* error). We identify $\frac{\tilde{C}_n^\top \zeta_{n+1}}{1 + \gamma_{n+1}^p \zeta_{n+1}^\top \zeta_{n+1}}$ to be the associated Kalman gain vector in this case.

γ_{n+1}^p and ζ_{n+1} can be computed in $O(mr)$ if C_n , \tilde{C}_n and P_n^{-1} are already known. The whole update step can then be performed in $O(mr)$ operations. \square

Theorem 6. For any new observation $\Gamma_{n+1} \in \mathbb{R}^m$, the matrices C_{n+1} , \tilde{C}_{n+1} and P_{n+1}^{-1} can be updated in $O(mr)$ if C_n , \tilde{C}_n and P_n^{-1} are already known.

Proof. The updating formula naturally depends on the linear dependency of Γ_{n+1} from previous observations (i.e. whether Γ_{n+1}^{rej} is non-null). Let us note that Γ_{n+1}^{rej} itself can be computed in $O(mr)$ operations if \tilde{C}_n and C_n are known.

If Γ_{n+1} is linearly independent from previous observations (i.e. $\Gamma_{n+1}^{rej} \neq 0$), equation 5.22 is valid, leading to:

$$\begin{aligned} P_{n+1}^{-1} &= (B_{n+1}^\top B_{n+1})^{-1} = \left(\begin{pmatrix} B_n^\top & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} B_n & 0 \\ 0 & 1 \end{pmatrix} \right)^{-1} = \begin{pmatrix} B_n^\top B_n & 0 \\ 0 & 1 \end{pmatrix}^{-1} \\ &= \begin{pmatrix} P_n^{-1} & 0 \\ 0 & 1 \end{pmatrix} \end{aligned} \quad (5.41)$$

$$C_{n+1} = \begin{pmatrix} C_n \\ \Gamma_{n+1}^\top \end{pmatrix} \quad (5.42)$$

Using equation 5.26, we can write:

$$\begin{aligned} \tilde{C}_{n+1} &= (C_{n+1} C_{n+1}^\top)^{-1} C_{n+1} = \begin{pmatrix} \sigma_n^{-1} + \frac{\gamma_{n+1}^p \gamma_{n+1}^{p\top}}{\|\Gamma_{n+1}^{rej}\|_2^2} & -\frac{\gamma_{n+1}^p}{\|\Gamma_{n+1}^{rej}\|_2^2} \\ -\frac{\gamma_{n+1}^{p\top}}{\|\Gamma_{n+1}^{rej}\|_2^2} & \frac{1}{\|\Gamma_{n+1}^{rej}\|_2^2} \end{pmatrix} \begin{pmatrix} C_n \\ \Gamma_{n+1} \end{pmatrix} \\ &= \begin{pmatrix} \tilde{C}_n - \gamma_{n+1}^p \frac{\Gamma_{n+1}^{rej\top}}{\|\Gamma_{n+1}^{rej}\|_2^2} \\ \frac{\Gamma_{n+1}^{rej\top}}{\|\Gamma_{n+1}^{rej}\|_2^2} \end{pmatrix} \end{aligned} \quad (5.43)$$

These formulae can be applied in $O(mr)$ operations, since γ_{n+1}^p and Γ_{n+1}^{rej} can themselves be computed in $O(mr)$ operations if \tilde{C}_n and C_n are already known.

If Γ_{n+1} is a linear combination of previous observations (i.e. $\Gamma_{n+1}^{rej} = 0$), equation 5.30 is valid, leading to:

$$C_{n+1} = C_n, \quad \tilde{C}_{n+1} = \tilde{C}_n \quad (5.44)$$

Using equation 5.34, we can write:

$$P_{n+1}^{-1} = P_n^{-1} - \frac{\zeta_{n+1} \zeta_{n+1}^\top}{1 + \gamma_{n+1}^p \zeta_{n+1}^\top} \quad (5.45)$$

This formula can be applied in $O(mr)$ operations^v, since γ_{n+1}^p and ζ_{n+1} can themselves be computed in $O(mr)$ operations if P_n^{-1} , \tilde{C}_n and C_n are already known.

□

Corollary 6.1. For any $n \times m$ matrix A_n of rank r and any vector $Y_n \in \mathbb{R}^n$, the least-squares solution $X_n = A_n^+ Y_n$ can be computed in $O(mnr)$ operations.

5.2.2 Orthogonal rank factorization

The theorems regarding the update of the pseudo-inverse A_{n+1}^+ (theorems 2 and 3) and least-squares solution X_{n+1} (theorems 4 and 5) are valid for any decomposition satisfying

^vNote that this complexity reduces to $O(r^2)$ if γ_{n+1}^p is already known.

equation 5.10. Therefore, the rows of C_n can be required to form an orthogonal basis. This can be easily achieved by storing only the rejection vectors Γ_{n+1}^{rej} into C_{n+1} . More generally, one can store rescaled rejection vectors $\alpha_{n+1}\Gamma_{n+1}^{rej}$ instead, with $0 \neq \alpha_{n+1} \in \mathbb{R}$. Theorem 6 remains valid, with equations 5.41, 5.42 and 5.43 becoming:

$$C_{n+1} = \begin{pmatrix} C_n \\ \alpha_{n+1}\Gamma_{n+1}^{rej \top} \end{pmatrix} \quad (5.46)$$

$$\tilde{C}_{n+1} = \begin{pmatrix} \tilde{C}_n \\ \frac{\Gamma_{n+1}^{rej \top}}{\alpha_{n+1}\|\Gamma_{n+1}^{rej}\|_2} \end{pmatrix} \quad (5.47)$$

$$\begin{aligned} P_{n+1}^{-1} &= (B_{n+1}^\top B_{n+1})^{-1} \\ &= \left(\begin{pmatrix} B_n^\top & \gamma_{n+1}^p \\ 0 & \frac{1}{\alpha_{n+1}} \end{pmatrix} \begin{pmatrix} B_n & 0 \\ \gamma_{n+1}^{p \top} & \frac{1}{\alpha_{n+1}} \end{pmatrix} \right)^{-1} \\ &= \begin{pmatrix} P_n^{-1} & -\alpha_{n+1}\zeta_{n+1} \\ -\alpha_{n+1}\zeta_{n+1}^\top & \alpha_{n+1}^2(1 + \gamma_{n+1}^{p \top} \zeta_{n+1}) \end{pmatrix} \end{aligned} \quad (5.48)$$

In particular, one can consider $\alpha_{n+1} = \|\Gamma_{n+1}^{rej}\|_2^{-1}$. In this case the rows of C_n form an orthonormal basis, i.e., C_n is an orthogonal matrix (i.e. $\tilde{C}_n = C_n$), and equation 5.10 becomes a Gram-Schmidt based thin LQ decomposition. This variant offers slightly reduced storage and update computational time.

5.3 Implementation

Based on the theorems above, one can devise a simple algorithm satisfying corollary 6.1, the pseudocode of which is shown in Algorithm 1.

These formulae have been implemented in Python3 using the Numpy library. In addition to the least-squares update algorithm (in $O(mr)$ operations), this implementation also supports pseudo-inverse (in $O(mn)$ operations) and covariance matrix¹³ updates (in $O(m^2)$ operations).

The orthogonal and orthonormal basis variants described in section 5.2.2 have also been implemented for least-squares update (in $O(mr)$ operations) and pseudo-inverse update (in $O(mn)$ operations).

In practice, checking if $\Gamma_{n+1}^{rej} \neq 0$ is ill-defined with floating-point arithmetic. Yet, it is crucial in

Algorithm 1 Rank-deficient RLS, also called rank-Greville

```

1: procedure UPDATELEASTSQUARES( $\Gamma, y, X, C, \tilde{C}, P^{-1}$ )
2:    $\gamma \leftarrow \tilde{C}\Gamma$ 
3:    $\Gamma_r \leftarrow \Gamma - C^\top \gamma$ 
4:   if  $\Gamma_r \neq 0$  then
5:      $K \leftarrow \frac{\Gamma_r}{\|\Gamma_r\|_2^2}$ 
6:      $C \leftarrow \begin{pmatrix} C \\ \Gamma^\top \end{pmatrix}$ 
7:      $\tilde{C} \leftarrow \begin{pmatrix} \tilde{C} - \gamma K^\top \\ K^\top \end{pmatrix}$ 
8:      $P^{-1} \leftarrow \begin{pmatrix} P^{-1} & 0 \\ 0 & 1 \end{pmatrix}$ 
9:   else
10:     $\zeta \leftarrow P^{-1}\gamma$ 
11:     $K \leftarrow \frac{\tilde{C}^\top \zeta}{1 + \gamma^\top \zeta}$ 
12:     $P^{-1} \leftarrow P^{-1} - \frac{\zeta \zeta^\top}{1 + \gamma^\top \zeta}$ 
13:   end if
14:    $X \leftarrow X + K \times (y - \Gamma^\top X)$ 
15: end procedure

```

this algorithm as it determines the effective rank of the linear system. Therefore, we define a threshold eps so that Γ_{n+1} is considered a linear combination of previous observations if and only if:

$$\|\Gamma_{n+1}^{rej}\|_2 < eps \quad \text{or} \quad \|\Gamma_{n+1}^{rej}\|_2 < eps \times \|\Gamma_{n+1}\|_2 \quad (5.49)$$

By default, eps is set to $(m^2 r + m r + m) \times \epsilon_M$ in order to account for rounding error propagation, with ϵ_M being the machine precision.

It is important to note that the general non-orthogonal basis implementation does support exact representations such as those defined in Python's *fractions* module. Indeed, this scheme (as the Greville algorithm) uses only operations well defined on rational numbers. One should note that the orthogonal basis variant scheme is also compatible with exact numerical representations as long as the rescaling factors α_{n+1} can themselves be represented exactly.

5.4 Numerical Tests

All computations were performed using Python 3.6.9^{14;15} with Scipy version 1.4.1¹⁶ and Numpy version 1.18.3^{17;18} linked with OpenBLAS on an Intel Xeon W-2123 CPU with DDR4-2666 RAM. The code used for the numerical tests is available along with our Python3 imple-

mentation in the supporting information and will soon become available on <https://github.com/RubenStaub/rank-greville>.

5.4.1 Computational efficiency

In this section we empirically evaluate the computational efficiency achieved by the rank factorization Greville algorithm described in this chapter (Algorithm 1), referred to as “rank-Greville”. Note that for comparison purposes, its total complexity (i.e. computing the full least-squares solution from scratch in $O(nmr)$ operations) is studied here, even though this algorithm was not specifically designed as a least-squares solver from scratch^{vi}. Rather, rank-Greville was designed for the recursive update in applications requiring a constantly up-to-date solution.

For this analysis to be meaningful, the rank-Greville algorithm time efficiency is compared against standard algorithms from the LAPACK library:

- “gelsy” refers to the DGELSY least-squares solver from the LAPACK library based on a complete orthogonal factorization (using QR factorization with column pivoting)^{19;20}.
- “gelss” refers to the DGELSS least-squares solver from the LAPACK library based on SVD¹⁹.
- “gelsd” refers to the DGELSD least-squares solver from the LAPACK library also based on SVD, using a diagonal form reduction¹⁹.

These routines were accessed through the `scipy.linalg.lstsq` wrapper function from the Scipy library.

For these tests, we consider the computation from scratch of the least-squares solution X_n verifying:

$$R_{\mathcal{N}}(n, m, r)X_n = R_{\mathcal{N}}(n, 1, n) \quad (5.50)$$

where the random matrix $R_{\mathcal{N}}(n, m, r) \in \mathbb{R}^{n \times m}$ has rank r and its elements are sampled from the normal $\mathcal{N}(0, 1)$ distribution. For reproducibility purposes the pseudo-random number generator was reset before each $R_{\mathcal{N}}(n, m, r)$ computation.

In order to assess the scaling properties of these algorithms, we evaluate (using Python’s `timeit` module) the time elapsed for solving equation 5.50 using various ranges for the parameters n , m , and r :

^{vi}Indeed, in this case rank-Greville must perform much more memory allocation and copy than a conventional solver, explaining, in part, the large prefactor

- Full-rank square matrices:

$$r = n = m \tag{5.51}$$

Figure 5.1 highlights the $O(n^3)$ asymptotic complexity of all algorithms in this case.

- Full-rank rectangular matrices:

$$r = n \leq m \tag{5.52}$$

By fixing the number of rows n , Figure 5.2 highlights the $O(mn^2)$ asymptotic complexity of all algorithms in this case^{vii}.

- Rank-deficient square matrices:

$$r \leq n = m \tag{5.53}$$

At fixed rank $r = 100$, Figure 5.3 nicely highlights the exploitation of rank-deficiency from rank-Greville to reach an $O(mnr)$ asymptotic complexity.

- Optimal (time-efficiency wise) applicability domains:

$$r \leq n \leq m \tag{5.54}$$

Figure 5.4 illustrates which algorithm is the fastest (and by which margin) for a range of parameters ratios. Even though not specifically designed for solving the linear least-squares problem from scratch, the rank-Greville algorithm appears to be more efficient than other LAPACK solvers, but only when the observations matrix has particularly low rank $r \lesssim 0.15 \times \min(n, m)$.

These tests confirm the lowest $O(mnr)$ asymptotic complexity of the rank-Greville algorithm compared with other LAPACK solvers ($O(m^2n)$ or $O(mn^2)$) for solving the least-squares problem from scratch. We note, nonetheless, that rank-Greville has a larger pre-factor, in part due to the additional work of maintaining a constantly up-to-date minimum-norm least-squares solution, typical of RLS solvers.

^{vii}Additional tests at $m > 4 \times 10^5$ seem to confirm the asymptotic linear dependency on m for the DGELSS solver.

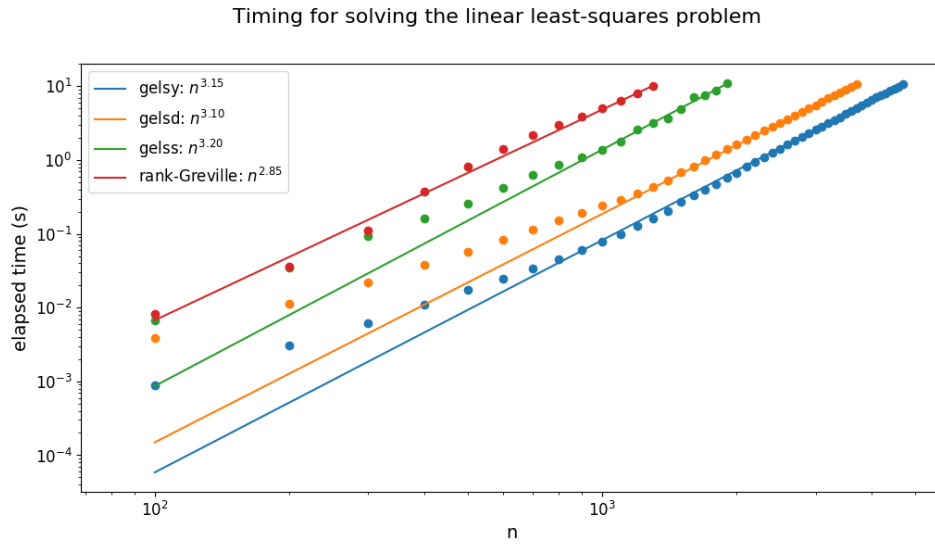


Figure 5.1 – Timings for solving the linear least-squares problem on a random full-rank square observations matrix $R_{\mathcal{N}}(n, n, n)$. The asymptotic dependency with respect to n is fitted on the last points and reported in the legend.




Figure 5.2 – Timings for solving the linear least-squares problem on a random full-rank observations matrix $R_{\mathcal{N}}(n, m, n)$ with a fixed number of observations $n = 100$. The asymptotic dependency with respect to m is fitted on the last points and reported in the legend.




Figure 5.3 – Timings for solving the linear least-squares problem on a random rank-deficient square observations matrix $R_{\mathcal{N}}(n, n, r)$ with a fixed rank $r = 100$. The asymptotic dependency with respect to n is fitted on the last points and reported in the legend.




Figure 5.4 – The fastest algorithm is represented for various n/m and r/n ratios, with $m = 4000$. The contour plot represents the interpolated relative margin by which an algorithm is the fastest (e.g. $\times 1.25$ means that the execution time for second fastest algorithm was 1.25 times larger than for the fastest one).

5.4.2 Numerical stability

In this section we evaluate the numerical stability achieved by the approach described in this chapter for computing the pseudoinverse.

For a more meaningful analysis, we compare our rank-Greville algorithm with standard algorithms from the LAPACK library described in 5.4.1, and other related algorithms:

- “Cholesky” refers to the *scipy.linalg.cho_solve* solver applied to the Cholesky factorization of $M = A^\top A$.
- “Greville” refers to a basic implementation of the original Greville algorithm^{9;21}.
- “orthogonal” refers to the orthogonal variant described in section 5.2.2 with rescaling factors $\alpha_{n+1} = 1$.
- “orthonormal” refers to the orthonormal variant described in section 5.2.2 (i.e. orthogonal variant with rescaling factors $\alpha_{n+1} = \|\Gamma_{n+1}^{rej}\|_2^{-1}$).

In order to assess the numerical stability of the algorithms described in this chapter, we rely on the measures defined in^{20;22}:

- The stability factor of an algorithm with respect to the computation of the pseudoinverse A^+ of A is given by:

$$e_{\text{algo}} = \frac{\|A_{\text{algo}}^+ - A^+\|_2}{\epsilon_M \|A^+\|_2 \kappa_2(A)} \quad (5.55)$$

where A_{algo}^+ is the pseudoinverse of A computed by the algorithm, A^+ is the exact pseudoinverse, $\|\cdot\|_2$ is the 2-norm (e.g. for a matrix A , $\|A\|_2 = \max(\sigma(A))$ is the largest singular value of A), $\kappa_2(A) = \frac{\max(\sigma(A))}{\min(\sigma(A))}$ is the condition number of A and ϵ_M is the machine precision.

This estimator is related to the forward stability of such algorithm, and should be bounded by a small constant depending on m and n .

- Similarly, we refer to the residual error as:

$$res_{\text{algo}} = \frac{\|A_{\text{algo}}^+ A - I\|_2}{\|A\|_2 \|A_{\text{algo}}^+\|_2} \quad (5.56)$$

where I is the identity.

This estimator is related to the mixed forward-backward stability of such algorithm, and should be of the order of the machine precision ϵ_M .

During the test, the machine precision used corresponds to $\epsilon_M \approx 2.22 \times 10^{-16}$.

This evaluation was performed empirically, using the matrices defined in²²:

- Pascal matrices $P(n) \in \mathbb{Z}^{n \times n}$.

These are full-rank square matrices whose inverses $P(n)^{-1} \in \mathbb{Z}^{n \times n}$ can be computed exactly since their elements are also integers.

Empirical results are reported in Tables 5.1 and 5.2. We found Greville-like algorithms, as well as the Cholesky decomposition to be orders of magnitudes less robust than the standard LAPACK solvers, with respect to both numerical stability indicators. Nonetheless, one should note that, as expected, when using *fractions.Fraction* based numerical representation, the rank Greville algorithm and its orthogonal variant were able to compute the exact inverse.

- Random matrices $R_{\mathcal{N}} \in \mathbb{R}^{3n \times n}$ whose elements are sampled from the normal $\mathcal{N}(0, 1)$ distribution.

Pseudoinverses generated by the DGELSY solver were used as reference for computing the stability factor since they display the lowest residual error and the empirical results are reported in Tables 5.3 and 5.4. We also found that the numerical stability indicators for the Greville-like algorithms are significantly dependent on the seed used by the pseudorandom number generator, unlike other algorithms tested.

- Random ill-conditioned matrices $R_{\mathcal{N}}^4 \in \mathbb{R}^{n \times n}$, taken as the fourth power of random square matrices $R_{\mathcal{N}} \in \mathbb{R}^{n \times n}$ whose elements are sampled from the $\mathcal{N}(0, 1)$ distribution.

Similarly as above, we used for pseudoinverse reference $(R_{\mathcal{N}}^4)^+ = (R_{\mathcal{N}}^+)^4$, the fourth power of the pseudoinverse generated by the DGELSY solver.

Empirical results are reported in Tables 5.5 and 5.6, which show that for these ill-conditioned matrices, the Cholesky-based solver is, overall, the least stable algorithm tested herein.

- Random matrices USV^T , where $U \in \mathbb{R}^{5n \times n}$ and $V \in \mathbb{R}^{n \times n}$ are random column orthogonal matrices and $S = \text{diag}(1, 2^{\frac{1}{2}}, \dots, 2^{\frac{n-1}{2}})$.

In this case, $(USV^T)^+ = VS^{-1}U^T$ was used for pseudoinverse reference.

Empirical results are reported in Tables 5.7 and 5.8. For these tests, the *rcond/eps* parameter was set to 10^{-8} . This was required by Greville-like algorithms to reach a reasonable solution.

Chapter 5. Efficient Recursive Least Squares Solver for Rank-Deficient Matrices

Table 5.1 – Empiric stability factors associated with the pseudoinverse computation of Pascal matrices $P(n)$.

n	$\kappa_2(A)$	$e_{\text{orthonormal}}$	$e_{\text{orthogonal}}$	$e_{\text{rank-Greville}}$	e_{Greville}	e_{Cholesky}	e_{gelsy}	e_{gelsd}	e_{gelss}
4	6.92e+2	8.80e-2	1.11e-1	1.67e+0	4.73e-2	5.88e+0	3.86e-2	6.67e-2	6.76e-2
6	1.11e+5	7.18e+2	1.06e+2	2.12e+2	6.58e+2	2.44e+3	5.04e-3	4.74e-2	4.74e-2
8	2.06e+7	3.42e+5	1.93e+3	2.18e+4	2.08e+5	1.59e+5	3.86e-3	8.55e-3	8.55e-3
10	4.16e+9	1.08e+6	1.08e+6	1.08e+6	1.08e+6	7.34e+5	3.00e-4	1.86e-3	1.86e-3

Table 5.2 – Empiric residual errors associated with the pseudoinverse computation of Pascal matrices $P(n)$.

n	$\kappa_2(A)$	$res_{\text{orthonormal}}$	$res_{\text{orthogonal}}$	$res_{\text{rank-Greville}}$	res_{Greville}	res_{Cholesky}	res_{gelsy}	res_{gelsd}	res_{gelss}
4	6.92e+2	1.83e-15	5.99e-16	3.85e-16	1.66e-17	4.01e-15	5.29e-17	5.39e-17	5.42e-17
6	1.11e+5	1.81e-13	4.16e-14	4.71e-14	1.46e-13	1.04e-12	2.49e-17	4.77e-17	4.47e-17
8	2.06e+7	7.60e-11	6.17e-13	4.84e-12	4.61e-11	9.24e-11	6.06e-17	3.16e-17	1.58e-17
10	4.16e+9	1.42e-9	1.61e-9	1.37e-9	1.54e-9	2.35e-8	4.78e-17	3.26e-17	1.54e-17

- Kahan matrices²³ $K(c, s) \in \mathbb{R}^{n \times n}$ with $c^2 + s^2 = 1$ and $n = 100$.

An explicit formula is available for the inverse²³, and was used as pseudoinverse reference.

Empirical results are reported in Table 5.9. Unlike what was reported in²², we did not find the pure SVD-based solver to perform significantly worse than other LAPACK solvers. Furthermore, after setting the *rcond/eps* parameter low enough to tackle the extreme ill-conditionality of the Kahan matrices (i.e. $rcond < \kappa_2(A)^{-1}$), all algorithms (except Cholesky) performed relatively well, with the QR-based LAPACK solver performing the best.

The algorithms described in this chapter (rank-Greville and variants), including the original Greville algorithm, perform roughly equivalently in terms of numerical stability. Furthermore, the stability of these Greville-like algorithms seems much less dependent on $\kappa_2(A)$ than the Cholesky based algorithm. As expected, we found these algorithms to be neither mixed forward-backward nor forward stable. As a consequence, the much more robust LAPACK least-squares solvers are to be recommended when numerical stability is crucial. However, the stability of Greville-like algorithms are competitive compared to the Cholesky-based LS solvers.

A compromise between update efficiency and numerical stability can be searched among the

Table 5.3 – Empiric stability factors associated with the pseudoinverse computation of random matrices $R_{\mathcal{N}} \in \mathbb{R}^{3n \times n}$ with elements distributed from $\mathcal{N}(0, 1)$.

n	$\kappa_2(A)$	$e_{\text{orthonormal}}$	$e_{\text{orthogonal}}$	$e_{\text{rank-Greville}}$	e_{Greville}	e_{Cholesky}	e_{gelsd}	e_{gelss}
4	2.44e+0	1.06e+2	7.46e+1	2.02e+1	2.84e+0	7.57e-1	2.08e+0	1.67e+0
6	2.25e+0	4.00e+0	4.98e+0	2.56e+0	1.05e+0	9.20e-1	1.35e+0	1.86e+0
8	2.74e+0	3.98e+0	4.43e+0	2.86e+0	2.23e+0	1.39e+0	1.81e+0	2.63e+0
10	2.84e+0	1.14e+3	4.60e+2	3.18e+2	2.24e+1	1.54e+0	1.99e+0	2.21e+0

5.4. Numerical Tests

Table 5.4 – Empiric residual errors associated with the pseudoinverse computation of random matrices $R_{\mathcal{N}} \in \mathbb{R}^{3n \times n}$ with elements distributed from $\mathcal{N}(0, 1)$.

n	$\kappa_2(A)$	$res_{\text{orthonormal}}$	$res_{\text{orthogonal}}$	$res_{\text{rank-Greville}}$	res_{Greville}	res_{Cholesky}	res_{gelsy}	res_{gelsd}	res_{gelss}
4	2.44e+0	3.45e-14	2.55e-14	6.36e-15	7.20e-16	1.25e-16	2.24e-16	4.38e-16	4.68e-16
6	2.25e+0	8.53e-16	1.69e-15	7.61e-16	2.09e-16	2.32e-16	2.24e-16	4.10e-16	6.77e-16
8	2.74e+0	1.24e-15	1.05e-15	3.88e-16	3.39e-16	2.77e-16	2.48e-16	4.96e-16	6.16e-16
10	2.84e+0	4.75e-13	1.73e-13	3.73e-14	3.22e-15	4.36e-16	2.64e-16	6.33e-16	8.02e-16

Table 5.5 – Empiric stability factors associated with the pseudoinverse computation of random ill-conditioned matrices $R_{\mathcal{N}}^4 \in \mathbb{R}^{n \times n}$.

n	$\kappa_2(A)$	$e_{\text{orthonormal}}$	$e_{\text{orthogonal}}$	$e_{\text{rank-Greville}}$	e_{Greville}	e_{Cholesky}	e_{gelsy}	e_{gelsd}	e_{gelss}
6	3.39e+2	5.52e+0	3.04e+0	1.39e+1	8.87e+0	3.72e+1	5.59e-2	6.86e-2	7.36e-2
8	1.08e+2	1.50e+0	3.07e+0	4.16e+0	6.85e+0	2.40e+0	1.66e-1	1.58e-1	1.72e-1
10	2.92e+7	3.53e+0	2.62e+0	8.25e+0	5.62e+0	6.60e+5	2.07e-1	1.49e-2	1.49e-2
12	5.78e+4	2.47e+1	1.64e+1	2.41e+2	3.06e+1	1.62e+3	2.80e-2	7.38e-2	7.38e-2

Table 5.6 – Empiric residual errors associated with the pseudoinverse computation of random ill-conditioned matrices $R_{\mathcal{N}}^4 \in \mathbb{R}^{n \times n}$.

n	$\kappa_2(A)$	$res_{\text{orthonormal}}$	$res_{\text{orthogonal}}$	$res_{\text{rank-Greville}}$	res_{Greville}	res_{Cholesky}	res_{gelsy}	res_{gelsd}	res_{gelss}
6	3.39e+2	1.40e-15	7.32e-16	3.38e-15	2.03e-15	1.01e-14	5.06e-17	8.23e-17	4.55e-17
8	1.08e+2	8.97e-16	1.55e-15	1.05e-15	1.63e-15	2.90e-15	1.02e-16	3.00e-16	2.79e-16
10	2.92e+7	1.57e-15	8.01e-16	1.83e-15	1.25e-15	9.58e-10	6.37e-17	1.71e-16	1.58e-16
12	5.78e+4	6.09e-15	5.01e-15	5.36e-14	6.90e-15	1.55e-12	4.71e-17	9.97e-17	1.36e-16

Table 5.7 – Empiric stability factors associated with the pseudoinverse computation of random matrices $USV^T \in \mathbb{R}^{5n \times n}$, where U and V are random column orthogonal matrices and $S = \text{diag}(1, 2^{\frac{1}{2}}, \dots, 2^{\frac{n-1}{2}})$.

n	$\kappa_2(A)$	$e_{\text{orthonormal}}$	$e_{\text{orthogonal}}$	$e_{\text{rank-Greville}}$	e_{Greville}	e_{Cholesky}	e_{gelsy}	e_{gelsd}	e_{gelss}
10	2.26e+1	6.72e+0	4.60e+0	4.19e+1	1.09e+1	1.85e+0	2.77e-1	2.94e-1	2.89e-1
15	1.28e+2	1.63e+2	1.60e+2	2.40e+3	2.99e+3	1.16e+1	1.70e-1	2.85e-1	2.93e-1
20	7.24e+2	1.57e+2	2.41e+2	1.52e+4	1.07e+4	4.24e+1	1.60e-1	1.85e-1	1.85e-1

Table 5.8 – Empiric residual errors associated with the pseudoinverse computation of random matrices $USV^T \in \mathbb{R}^{5n \times n}$, where U and V are random column orthogonal matrices and $S = \text{diag}(1, 2^{\frac{1}{2}}, \dots, 2^{\frac{n-1}{2}})$.

n	$\kappa_2(A)$	$res_{\text{orthonormal}}$	$res_{\text{orthogonal}}$	$res_{\text{rank-Greville}}$	res_{Greville}	res_{Cholesky}	res_{gelsy}	res_{gelsd}	res_{gelss}
10	2.26e+1	2.68e-15	4.25e-15	2.52e-15	5.92e-16	1.70e-15	1.01e-16	1.62e-16	1.26e-16
15	1.28e+2	1.67e-14	9.69e-15	9.20e-14	2.59e-13	7.37e-15	7.58e-17	1.19e-16	1.30e-16
20	7.24e+2	1.65e-14	6.11e-14	6.04e-13	4.93e-13	3.53e-14	4.31e-17	9.96e-17	5.61e-17

Table 5.9 – Empiric residual errors associated with the pseudoinverse computation of Kahan matrices $K(c, s) \in \mathbb{R}^{100 \times 100}$, with $c^2 + s^2 = 1$.

c	$\kappa_2(A)$	$res_{\text{orthonormal}}$	$res_{\text{orthogonal}}$	$res_{\text{rank-Greville}}$	res_{Greville}	res_{Cholesky}	res_{gelsy}	res_{gelsd}	res_{gelss}
0.10	5.42e+4	5.57e-17	3.00e-17	3.50e-17	2.64e-17	3.40e-13	5.24e-17	7.61e-17	1.33e-16
0.15	1.13e+7	1.11e-17	2.16e-17	1.03e-17	1.19e-17	1.60e-10	1.52e-17	5.61e-17	7.97e-17
0.20	2.18e+9	8.10e-18	7.96e-18	2.42e-18	2.41e-18	failure	7.40e-18	6.30e-17	5.69e-17
0.25	4.37e+11	2.07e-18	1.06e-18	1.14e-18	1.17e-18	1.61e-6	8.74e-18	4.85e-17	4.39e-17
0.30	9.77e+13	3.01e-19	4.31e-19	1.92e-19	1.94e-19	7.98e-4	1.82e-19	1.33e-16	2.43e-17
0.35	2.57e+16	3.57e-20	4.27e-20	2.73e-20	2.18e-20	failure	5.05e-20	9.43e-18	2.03e-18
0.40	8.36e+18	3.46e-21	3.51e-21	3.09e-21	2.55e-21	1.02e-4	1.08e-20	3.17e-19	3.04e-19

full QR decomposition updating algorithms¹, or even faster, stable updating algorithms for Gram-Schmidt QR factorization^{1;24}.

5.5 Conclusion

In this chapter, we first derive a simple explicit formula for the recursive least-squares problem using a general rank decomposition update scheme. Based on this, we devise a transparent, Greville-like algorithm. We also introduce two variants bridging the gap between Greville's algorithm and QR-based least-squares solvers. In contrast to Greville's algorithm, we maintain a rank decomposition at each update step. This allows us to exploit rank-deficiency to reach an asymptotic computational complexity of $O(mr)$ for updating a least-squares solution when adding an observation, leading to a total complexity of $O(mnr)$ for computing the full least-squares solution. This complexity is lower than Greville's algorithm or any commonly available solver tested, even though a truncated QR factorization based solver can achieve such a $O(mnr)$ bound for computing the full least-squares solution²⁵. Nonetheless, a $O(mr)$ bound for the least-squares solution update step is, to our knowledge, lower than those achieved by the more sophisticated updating algorithms explicitly reported in the literature. We have implemented these algorithms in Python3, using Numpy. This publicly available implementation offers a recursive least-squares solver ($O(mr)$ operations per update), with optional pseudoinverse ($O(mn)$) and covariance support ($O(m^2)$). The numerical stability of these Greville-like algorithms were empirically found to be significantly inferior compared to common LAPACK solvers. However, it is important to note that the algebraic simplicity of some of these Greville-like methods make them compatible with exact numerical representation, without the need to use symbolic computing.

Bibliography

- [1] Bjorck, A. *Numerical Methods for Least Squares Problems*; Other Titles in Applied Mathematics; Society for Industrial and Applied Mathematics, 1996.
- [2] Penrose, R. On best approximate solutions of linear matrix equations. *Mathematical Proceedings of the Cambridge Philosophical Society* **1956**, 52, 17–19.
- [3] Rakha, M. A. On the Moore–Penrose generalized inverse matrix. *Applied Mathematics and Computation* **2004**, 158, 185–200.
- [4] Toutounian, F; Ataei, A. A new method for computing Moore–Penrose inverse matrices. *Journal of Computational and Applied Mathematics* **2009**, 228, 412–417.
- [5] Moore, E. On the reciprocal of the general algebraic matrix. *Bull. Amer. Math. Soc.* **1920**, 26, 394–395.

-
- [6] Penrose, R. A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society* **1955**, 51, 406–413.
- [7] Björck, Å. In *Encyclopedia of Optimization*; Floudas, C. A., Pardalos, P. M., Eds.; Springer US: Boston, MA, 2009; pp 1856–1866.
- [8] Courrieu, P. Fast Computation of Moore-Penrose Inverse Matrices. *ArXiv* **2005**, *abs/0804.4809*.
- [9] Greville, T. N. E. Some Applications of the Pseudoinverse of a Matrix. *SIAM Review* **1960**, 2, 15–22.
- [10] Albert, A.; Sittler, R. W. A Method for Computing Least Squares Estimators that Keep Up with the Data. *Journal of the Society for Industrial and Applied Mathematics Series A Control* **1965**, 3, 384–417.
- [11] Le Gall, F. Faster Algorithms for Rectangular Matrix Multiplication. 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science. 2012; pp 514–523.
- [12] Mirsky, L. *An Introduction to Linear Algebra*; Dover Books on Mathematics; Dover, 1990.
- [13] Greene, W. *Econometric Analysis*; Prentice Hall, 2003.
- [14] Van Rossum, G.; Drake, F. L. *Python 3 Reference Manual*; CreateSpace: Paramount, CA, 2009.
- [15] Pérez, F.; Granger, B. E. IPython: a System for Interactive Scientific Computing. *Computing in Science and Engineering* **2007**, 9, 21–29.
- [16] Virtanen, P. et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **2020**, 17, 261–272.
- [17] Oliphant, T. E. *Guide to NumPy*, 2nd ed.; CreateSpace Independent Publishing Platform: USA, 2015.
- [18] Walt, S. v. d.; Colbert, S. C.; Varoquaux, G. The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science & Engineering* **2011**, 13, 22–30.
- [19] Anderson, E.; Bai, Z.; Bischof, C.; Blackford, L. S.; Demmel, J.; Dongarra, J.; Du Croz, J.; Greenbaum, A.; Hammarling, S.; McKenney, A.; Sorensen, D. *LAPACK Users' Guide*, 3rd ed.; Society for Industrial and Applied Mathematics, 1999.
- [20] Higham, N. J. *Accuracy and Stability of Numerical Algorithms*, 2nd ed.; Society for Industrial and Applied Mathematics, 2002.
- [21] Ben-Israel, A.; Greville, T. N. *Generalized inverses: theory and applications*; Springer Science & Business Media, 2003; Vol. 15.

Chapter 5. Efficient Recursive Least Squares Solver for Rank-Deficient Matrices

- [22] Smoktunowicz, A.; Wrobel, I. Numerical aspects of computing the Moore-Penrose inverse of full column rank matrices. *BIT Numerical Mathematics* **2012**, *52*, 503–524.
- [23] Kahan, W. Numerical Linear Algebra. *Canadian Mathematical Bulletin* **1966**, *9*, 757–801.
- [24] Daniel, J. W.; Gragg, W. B.; Kaufman, L.; Stewart, G. W. Reorthogonalization and stable algorithms for updating the Gram-Schmidt QR factorization. *Mathematics of Computation* **1976**, *30*, 772–795.
- [25] Businger, P.; Golub, G. H. Linear Least Squares Solutions by Householder Transformations. *Numer. Math.* **1965**, *7*, 269–276.



6 UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

Model Hamiltonians based on the so called cluster expansion (CE), which consist of a linear fit of parameters corresponding to geometric patterns, provide an efficient and rigorous means to quickly evaluate the energy of diverse arrangements of adsorbate mixtures on reactive surfaces. The model Hamiltonians are particularly well suited to explore the phase-space of adsorbates on catalytic surfaces and the kinetics of processes of key importance in heterogeneous catalysis, where high coverages, leading to significant lateral interactions, are omnipresent. However, establishing the model Hamiltonian is a tedious task, requiring the construction and optimization of many geometries. Today, most of these geometries are constructed by hand, based on chemical intuition or random choices. Hence, the quality of the training set is unlikely to be optimal and its construction is not reproducible. Herein, we propose a reformulation of the construction of the training set as a strategy-based game, aiming at an efficient exploration of the relevant patterns constituting the model Hamiltonian. Based on this reformulation, we exploit a typical active learning solution for machine-learning such a strategy game: an upper confidence tree (UCT) based framework. However, in contrast to standard games, evaluating the true score is computationally expensive, as it requires a costly geometry optimization. Hence, we augment the UCT with a pre-exploration step, inspired by the variance-based Design of Experiments methods. The pre-exploration step uses the constantly updated model Hamiltonian to determine which additional structure is most likely to improve the model Hamiltonian. This novel framework, called Reinforcement Sampling allows to automatically construct a well adapted training set, minimizing computational cost and user-intervention. As a proof of principle, we apply Reinforcement Sampling on the CO oxidation reaction on Pt(111), for which a reliable model Hamiltonian has been established previously. The results demonstrate the effectiveness of the custom built UCT and the significant benefit of the pre-exploration.




Figure 6.1 – Synoptic illustration of some of the many interdependent processes that can occur, in practice, during an heterogeneous catalysis.

6.1 Introduction

Heterogeneous catalytic processes play a major role in chemistry, with numerous real-world applications¹. Heterogeneous catalysis aims at optimizing reaction rates (i.e. the kinetics of a reaction) through reactive surfaces, where the reactants are adsorbed, migrate, react and the desired products are desorbed. In practice, however, many interdependent processes should be taken into consideration² (e.g., competition, poisoning or side-reactions, see figure 6.1), leading to the need for powerful modelling tools for studying heterogeneous reactions, in order to identify the key steps involved and eventually optimize them through rational design³. The traditional kinetic models, such as the Langmuir-Hinshelwood model^{4;5}, allow for computing analytical expressions for the total reaction rate. However, these models do not explicitly capture the influence of other adsorbates present on the surface (i.e. the adsorbates do not explicitly “feel” each other), since these models rely on mean-field descriptions. Nonetheless, these interactions between adsorbates (called lateral interactions) can have significant effects^{6;7}, especially at high coverage. Therefore, their description and inclusion can be essential to properly describe catalytic processes^{8;9}.

Explicitly accounting for all possible surface states \mathcal{S}_i would lead to an intractable set of rate equations, and the associated master equations^{10;11} can no longer be, reasonably, solved analytically:

$$\frac{dP(\mathcal{S}_i, t)}{dt} = \sum_{j \neq i} (R(\mathcal{S}_j \rightarrow \mathcal{S}_i)P(\mathcal{S}_j, t) - R(\mathcal{S}_i \rightarrow \mathcal{S}_j)P(\mathcal{S}_i, t)), \quad (6.1)$$

where $P(\mathcal{S}_i, t)$ is the probability of the system being in the state \mathcal{S}_i at time t and $R(\mathcal{S}_i \rightarrow \mathcal{S}_j)$ is the transition rate associated with the transition $\mathcal{S}_i \rightarrow \mathcal{S}_j$ (i.e. the conditional probability of hopping to state \mathcal{S}_j from state \mathcal{S}_i).

Instead, this set of differential equations ought to be solved numerically. In a Kinetic Monte-

Carlo^{10;11} (KMC) approach, the total reaction rate is estimated by a Monte-Carlo method, using the transition rates $R(\mathcal{S}_i \rightarrow \mathcal{S}_j)$ predicted by Transition State Theory from the corresponding activation energies E^\ddagger . To account for lateral interactions, a Brønsted-Evans-Polanyi¹²⁻¹⁴ (BEP) relation is commonly used to estimate the corrected activation energies from the corrected reaction energies, due to lateral interactions.

As a consequence, a KMC approach including lateral interactions requires to estimate the energy of each surface state \mathcal{S}_i considered, including the energy correction due to lateral interactions present in that state \mathcal{S}_i . Directly computing the energy of an occupied surface from first-principles calculations (such as DFT) would require too much computational time to efficiently sample the accessible states and provide meaningful kinetic predictions. Instead, an approximate model Hamiltonian is typically used^{10;11}, especially for describing lattice-based systems such as metallic surfaces, where geometries corresponding to energy minima are well defined and separated (i.e. the adsorbates are preferentially located at active sites) and the relevant geometry space can therefore be discretized.

A model Hamiltonian (or effective Hamiltonian) is a model linking the energy of a system to its geometrical features¹⁵⁻¹⁷. The most simple yet rigorous model Hamiltonian for lattice-based active surfaces (and accounting for lateral interactions) is probably a linear model based on a Cluster Expansion¹⁸⁻²¹ (CE). In this CE-based model Hamiltonian, a linear relation is considered between the total adsorption energy of an occupied surface and the number of each type of lateral interactions present among the adsorbates at the surface. Using a lattice-based representation of the surface, the types of lateral interactions become enumerable (since the positions are discretized) but infinitely many distinct types of lateral interactions can be defined. Therefore, a truncated CE is used in practice²², by considering only the lateral interactions that display significant energy contributions. In other words, a CE-based model Hamiltonian is a simple linear model whose parameters are these significant energy contributions of the corresponding relevant lateral interactions. Therefore, identifying and estimating these energy contributions from first-principles calculations (usually based on DFT) is essential.

Indeed, like any empirical model, a CE-based model Hamiltonian must be trained on a set of empirical data. In this case, the training set is composed of optimized geometries representing different adsorption states of the active surface of interest and their associated energy, usually computed at a DFT level. Therefore, the training set of a CE-based model Hamiltonian must be composed of reaction-specific geometries (with the active surface and adsorbates involved in the catalytic process being modelled), so that a universal training set cannot be designed for such a model (or would require a combinatorially large database). As a consequence,

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

the training set is an essential component of a KMC simulation, as it must be generated for each catalytic process studied. Yet this training set is usually generated by the chemist²³, selecting around a few hundreds hand-picked geometries using their own intuition, and/or trying to sample geometries at pseudo-random/exhaustively. Therefore, this essential task is particularly time consuming, hardly reproducible, and leads to a possibly suboptimal training set.

In the recent years, most efforts were made toward optimizing the selection of relevant lateral interactions to include in the CE-based model Hamiltonian (i.e. optimizing the choice of parameters)²⁴⁻²⁸. While some efforts were also made for optimizing the training set, these were basic variance-based Design of Experiments integrated in the dual optimization of both the terms and the geometries in a greedy approach^{23;29}, and were mainly designed to be updated on-the-fly for applications to ground state predictions³⁰. To our knowledge, the *a priori* selection of geometries to evaluate (at a DFT level) with the aim of generating an adequate training set for KMC applications has not yet been properly discussed, especially since the development of recent reinforcement learning based active learning schemes³¹⁻³³.

In this contribution, we address these challenges by proposing an active learning framework for the construction of an adequate training set for KMC applications, by optimizing the selection of geometries to include *before* evaluating their energy by first-principles calculations, without imposing the choice of a CE-based Hamiltonian. This active learning algorithm combines Reinforcement Learning techniques (i.e. a Upper Confidence Tree) with classical variance-based Design of Experiments³¹ into a tool called Reinforcement Sampling, aimed at automating the construction of a CE-based model Hamiltonian training set from scratch, while reducing the computational time required for its creation and achieving a final training set that is consistently better designed than from a pseudo-random sampling.

6.2 Cluster Expansion Hamiltonian

The Cluster Expansion (CE) Hamiltonian¹⁸⁻²¹ is a linear model Hamiltonian based on lateral interactions, enumerated by their corresponding correction order, for predicting the adsorption energy of a set of adsorbates at an active surface. A linear model Hamiltonian (or effective Hamiltonian) is a linear model whose aim is to reproduce the energy $E(g_i)$ (here, the target function, or regressand, is the DFT-based adsorption energy) of a geometry g_i , as a linear combination of its input features Γ_i (called regressors):

$$E(g_i) = \Gamma_i^\top \beta + \epsilon_i, \tag{6.2}$$

where ϵ_i is the residual error. Therefore, a model Hamiltonian provides a simple and fast energy predictor, as the regressors are typically features that can be computationally inexpensive to obtain.

In a CE Hamiltonian, the regressors Γ are the multiplicities of each possible lateral interaction (i.e. the number of times each possible lateral interaction is found in the geometry). The parameters β of the model are, therefore, the energy contribution of each lateral interaction. These interactions and their contributions are seen as corrective terms of increasing order to the total energy:

$$E(g_i) = \underbrace{E^{(1)}(g_i)}_{\text{1-body model}} + \underbrace{\sum_{n=2}^{N_{\text{ads}}} E^{(n)}(g_i)}_{\text{many-body correction terms}}, \quad (6.3)$$

where N_{ads} is the number of adsorbates in g_i and $E^{(n)}(g_i)$ is the adsorption energy contributions of the n -body effects only:

$$E^{(n)}(g_i) = \sum_{\substack{\{a_1, \dots, a_n\} \subseteq g_i \\ a_1 \neq \dots \neq a_n}} \left(E(\{a_1, \dots, a_n\} @ \mathcal{S}) - \sum_{k=1}^{n-1} E^{(k)}(\{a_1, \dots, a_n\} @ \mathcal{S}) \right) \quad (6.4)$$

where $\{a_1, \dots, a_n\} @ \mathcal{S}$ is an altered version of the geometry g_i where only the n adsorbates a_1, \dots, a_n of g_i are considered to be present (and adsorbed at the surface).

A full explicit expansion, where the lateral interaction contributions appear, would be:

$$\begin{aligned} E(g_i) &= \underbrace{\sum_{a \in g_i} E(a @ \mathcal{S})}_{\text{1-body terms}} + \underbrace{\sum_{\substack{\{a, b\} \subseteq g_i \\ a \neq b}} E(\{a, b\} @ \mathcal{S}) - E(a @ \mathcal{S}) - E(b @ \mathcal{S})}_{\text{2-body terms}} \\ &+ \underbrace{\sum_{\substack{\{a, b, c\} \subseteq g_i \\ a \neq b \neq c}} E(\{a, b, c\} @ \mathcal{S}) - E(\{a, b\} @ \mathcal{S}) - E(\{a, c\} @ \mathcal{S}) - E(\{b, c\} @ \mathcal{S}) + E(a @ \mathcal{S}) + E(b @ \mathcal{S}) + E(c @ \mathcal{S})}_{\text{3-body terms}} \\ &+ \dots \end{aligned} \quad (6.5)$$

In practice, only a truncated version of this model is considered²², typically up to the 2 or 3-body terms⁹. Indeed, higher order corrections tend to fade out quickly. Besides, a cutoff is usually used to discard lateral interactions where the adsorbates are too far away from to significantly “feel” each other. As a consequence, the number of lateral interactions to consider become tractable, since only a few contributions are non negligible.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

Of course, the sums might contain redundant contributions if the same corresponding lateral interaction is found multiple times, leading ultimately to the linear model described equation 6.2. It now becomes clear that, in the cluster expansion model, the residual error ϵ_i comes from the discarded terms and the DFT-related error.

Assuming that the CE Hamiltonian satisfies the assumptions of the linear regression modelⁱ, an ordinary least squares method can be applied to retrieve an unbiased, consistent and efficient estimator³⁴ $\hat{\beta}$ for the energy contributions of the considered lateral interactions. One should note that a guess β_{guess} can be given, so that instead of $\hat{\beta}$ being the minimum-norm least-squares solution, it is searched as the least-squares solution with minimum Euclidean distance to β_{guess} ⁱⁱ. Such guess would therefore only have an impact where no empirical data is available (i.e. when the CE Hamiltonian is underdetermined).

Futhermore, the uncertainty on the parameters $\hat{\beta}$ is given by the covariance matrix $\text{Var}(\hat{\beta})$. On a purely overdetermined system of linear equationsⁱⁱⁱ, this covariance can be estimated with³⁴:

$$\text{Var}(\hat{\beta}) \approx s^2 (\Gamma^\top \Gamma)^{-1}, \quad (6.6)$$

where $s^2 = \frac{\mathbf{E} - \Gamma^\top \hat{\beta}}{n-r}$ is an estimator of the variance of the noise σ^2 , with $\mathbf{E} = (E_1 \dots E_n)^\top$, n the number of observations and r the rank of Γ (i.e. $n - r$ is the number of linearly redundant observations).

However, in the general case and especially during the training of a non-truncated CE Hamiltonian, the number of observations is likely to be less than the number of regressors, so that one must deal with a underdetermined system. Nonetheless, any linear system of equation can be decomposed into a purely non-overdetermined system and a purely non-underdetermined system (see chapter 5). Therefore, it is always possible to define the unique maximum subspace^{iv} $\text{Vect}(\Gamma_{over})$ onto which a model Hamiltonian is overdetermined.

As a consequence, it is possible to estimate a covariance matrix embedded unto that subspace $\text{Var}(\hat{\beta})_{empirical}$, using equation 6.6. Concerning the rest of the regressors space, the user-defined variance of the guessed parameters $\text{Var}(\beta_{guess})$ can be used onto that space.

ⁱsee section D.1 in appendix D

ⁱⁱThis can be done by considering a simple minimum-norm least-square solution $\widehat{(\beta - \beta_{guess})}$ on the linear equation $E_j - \Gamma_i^\top \beta_{guess} = \Gamma_i^\top (\beta - \beta_{guess}) + \epsilon_i$, and choosing $\widehat{(\beta - \beta_{guess})} + \beta_{guess}$ as the final estimator.

ⁱⁱⁱTo be precise: the covariance formula is valid if the observation matrix $\Gamma^\top = (\Gamma_1 \dots \Gamma_n)^\top$ has full column-rank. In other words, if the whole regressors space is spanned by the observations.

^{iv}This subspace is eventually null (if all observations are linearly independent), otherwise it is the subspace spanned the all the observations $\text{Vect}(\Gamma_{over}) = \text{Vect}(\Gamma)$.

6.3. Presentation of the problem and the strategy

Combining these two terms (assuming that guessed parameters are uncorrelated with empirically estimated parameters) provides a covariance matrix estimator $\text{Var}(\hat{\beta}) = \text{Var}(\hat{\beta})_{\text{empirical}} + \text{Var}(\beta_{\text{guess}})$ that can be defined even on a non purely overdetermined model Hamiltonian.

Using this estimator, it is possible to estimate the uncertainty associated with the prediction of the energy E_j on a given input Γ_j (called prediction error):

$$\begin{aligned}\text{Var}(E_j - \Gamma_j^\top \hat{\beta}) &= \text{Var}(\Gamma_j^\top \beta + \epsilon_j - \Gamma_j^\top \hat{\beta}) \\ &= \sigma^2 + \Gamma_j^\top \text{Var}(\hat{\beta}) \Gamma_j \\ &\approx s^2 + \Gamma_j^\top \text{Var}(\hat{\beta}) \Gamma_j\end{aligned}\tag{6.7}$$

As a consequence, $\text{Var}(\hat{\beta})$ is an important tool for estimating the quality of a prediction, and more globally, of the full CE Hamiltonian in terms of convergence^v.

6.3 Presentation of the problem and the strategy

Building a model Hamiltonian requires two main components: a functional form for the effective Hamiltonian and a training dataset to feed to the model. Regarding the first component, we place our study in the context of a cluster expansion (CE) Hamiltonian for describing the many-body effects. The basic notions of this energetic model are presented in section 6.2. The main focus of our study is, however, related to the second ingredient, i.e., building the most appropriate training set which we here call “adequate training set” from scratch in an automated and reproducible way. The developed framework is general and is, thus, not strongly linked to the first component. In other words, instead of training a CE Hamiltonian, one could construct a training set for semi-empirical methods or neural networks in an analogous way.

6.3.1 Definition of an “adequate” training set

Before anything, we define an adequate training set as a set of geometries satisfying three criteria:

- **Diverse:** We wish our model Hamiltonian to be able to describe most likely events and even some rare events. Therefore, an adequate training set should contain a broad and diverse set of geometries. Ideally, such sampling should roughly cover the whole chemical space available. Mathematically, we require: any possible geometry should have a non-zero probability to be found in an adequate training set.

^vIndeed, since $\hat{\beta}$ is an efficient estimator, $\text{Var}(\hat{\beta})$ represents the inverse Fisher information matrix.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

- **Chemically relevant:** Just as a chemist would use its chemical intuition to choose geometries displaying large many-body effects, we wish our model Hamiltonian to better describe strong lateral interactions than negligible ones. This is translated mathematically as: the higher the absolute energy deviation from the 1-body model associated with a geometry, the more likely an adequate training set should contain this geometry. For a discussion of this particular choice to measure the relevance of a geometry, see section D.3 of appendix D.
- **Reasonably-sized:** In a typical setting, all possible geometry combinations cannot even be enumerated in a reasonable time. Furthermore, including a new geometry in the training set is a computationally expensive task: in a typical case, it requires to run a geometry optimization or at least a single point energy computation at a DFT level. Therefore, an adequate training set should not contain geometries that can be removed without a major impact on the quality of the resulting model Hamiltonian. In other words, the training should stop as soon as the model would reach sufficient convergence.

In summary, building an adequate training set requires a balance between focusing on groups of geometries that seem chemically relevant (i.e. sets of similar geometries that were found to contain geometries with large associated many-body energy effects) and exploring novel groups of geometries.

6.3.2 Upper Confidence Trees for solving the adequate training set problem

In this section, we describe our training set selection problem using a related problem with known solution. Indeed, the balance between exploration and exploitation is a well-known dilemma that has been most studied with the notorious multi-armed bandit problem.

Multi-armed bandit problem

The multi-armed bandit problem can be formulated as follows: at a casino, there are multiple armed bandit machines (B_1, B_2, \dots, B_k). Each has its own average payout $\mathbb{E}(W_i)$ that is *a priori* unknown to the player. The aim is to find a playing strategy that has minimal cumulative regret (i.e. the total number of times a sub-optimal bandit has been played).

Of course, if the average payouts $\mathbb{E}(W_i)$ were known, the optimal strategy would be obvious: play repeatedly the armed bandit with highest payout on average. But since the average payouts are *a priori* unknown, they must be estimated during the play. It then becomes clear that an optimal strategy should maintain a balance between estimating the payout of bandits (i.e. exploration) and focusing on the apparently most profitable bandits (i.e. exploitation), as the latter evolve as a function of the exploration.

6.3. Presentation of the problem and the strategy

This is reminiscent of the exploration vs. exploitation balance required for building an adequate training set. Therefore, we propose in this chapter to identify the building of an adequate training set as a multi-armed bandit problem. In such embedding, the bandit machines B_i are identified as possible groups G_i of similar geometries to be included in the training set, playing a bandit B_i means adding to the training set an arbitrarily chosen geometry $g_j \in G_i$ within the related group G_i and the associated payout W_i of playing B_i is taken as the absolute deviation from the 1-body model $W_i = |E_j - E_j^{alone}|$ associated with the chosen geometry g_j , where $E_j^{alone} = \sum_k E_k$ is the sum of the individual adsorption energies E_k for each adsorbate \mathcal{M}_k present in g_j as if adsorbed alone at the surface.^{vi} Let us illustrate these notations with a more chemical example by considering $g_j = (\text{H}_2\text{O} + \text{NH}_3)@\text{Pt}(111)$ a geometry where one water molecule and one ammonia molecule are adsorbed on top of a Pt(111) surface: then $E_j = E((\text{H}_2\text{O}-\text{NH}_3)@\text{Pt}(111)) - E(\text{Pt}(111)) - E(\text{H}_2\text{O}) - E(\text{NH}_3)$ is the DFT-based adsorption energy of one molecule of water and ammonia together, and $E_j^{alone} = E_{\text{H}_2\text{O}} + E_{\text{NH}_3} = \left(E(\text{H}_2\text{O}@\text{Pt}(111)) - E(\text{Pt}(111)) - E(\text{H}_2\text{O}) \right) + \left(E(\text{NH}_3@\text{Pt}(111)) - E(\text{Pt}(111)) - E(\text{NH}_3) \right)$ is the sum of $E_{\text{H}_2\text{O}}$ and E_{NH_3} the adsorption energies of a single water/ammonia molecule, individually. So $W_i = |E_j - E_j^{alone}|$ is the absolute value of the energy correction from the one-body model, due to the lateral interaction $\text{H}_2\text{O}-\text{NH}_3$

The Upper Confidence Bound approach

The multi-armed bandit problem has been much studied in the literature, and the most efficient strategies are commonly based on Upper Confidence Bound (UCB) approaches.

The general idea behind UCB-like approaches can seem intuitive: each bandit is given a score $S_i = \mathbb{E}(W_i)_{emp} + b_i(n_i, T)$, where n_i is the number of times bandit B_i was played, T is the total number of plays (i.e. $\sum_i n_i = T$), $\mathbb{E}(W_i)_{emp}$ is the empirical estimator of the average payout and $b_i(n_i, T)$ is called the exploration bias. The UCB strategy is then: repeatedly play the bandit B_i with the highest UCB score S_i .

Such UCB approach is inherently optimistic. Indeed, the exploration bias acts as an upper confidence interval for the empirical estimate of the payout. From this perspective, it becomes clear that the UCB score provides an upper confidence bound for the true payout $\mathbb{E}(W_i)$. Of course, the main challenge lies in determining the exploration bias. Finding an efficient functional form for the exploration bias is still an active field of research, even though many functions have already been designed³⁵⁻³⁹. The most widespread formula probably comes

^{vi}Note that, as in a typical multi-armed bandit problem, the payout W_i can be seen as a random variable since g_j is chosen randomly in G_i .

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

from the UCB1 algorithm³⁵, with $b_i(n_i, T) = C\sqrt{\frac{\log(n_i)}{T}}$ where C is an input parameter acting as a confidence level.

However, a UCB approach alone is not enough to efficiently solve our training set building problem. Indeed, the organization of geometries into groups is still not well-defined. In particular, optimising their number/sizes is a major problem which is exacerbated by the sheer number of possible geometries: too large groups would lead to a loss of group specificity, while too many small groups would require too much exploration to properly sample them (i.e. there would be a loss in interpolation/generalisation power, in the sense that the information gained from each play would only be applicable to a small number of geometries)^{vii}. This specificity vs generalisation power trade-off is usually solved by combining the UCB approach with a Monte-Carlo Tree Search (MCTS) scheme, effectively taking the best of both worlds between coarse-grain and fine-grain sampling.

Combining a Monte-Carlo Tree Search with the Upper Confidence Bound, leading to the Upper Confidence Tree construction of the training set

Instead of using a single arbitrary group size, a MCTS framework uses a recursive grouping. In our case, all geometries are organized in a tree structure: each single geometry is represented by a leaf, and each node \mathcal{N}_i (a branch) represents therefore a group of geometries G_i composed of all the geometries $g_j \in G_i$ associated with the leaves descending from that node. The tree starts out with a single root (or the trunk of the tree to push the analogy to biology). The more similar two geometries are, the closer their corresponding leaves in the tree structure are (i.e. the more nodes are shared in both paths from the root node)^{viii}.

Selecting a leaf in a MCTS framework is done by a best-first search composed of the following steps:

- Selection/Expansion: Starting from the root node, the best child node is selected until an unexplored node (or a leaf corresponding to a single geometry) is reached. In a pure MCTS framework, each node \mathcal{N}_i is evaluated using a score S_i taken as the empirical average payout of its associated group of geometries $\mathbb{E}(W_i)_{emp} = \frac{\sum_j |E_j - E_j^{alone}|}{n_i}$ where n_i is the number of times node \mathcal{N}_i was explored, E_j is the adsorption energy of geometry g_j , E_j^{alone} is the sum of all individual adsorption energies for each adsorbate in geometry g_j , and the sum \sum_j is over all n_i corresponding geometries $g_j \in G_i$ already explored/chosen.

^{vii}Note that this can be seen as a typical coarse-grain versus fine-grain sampling trade-off.

^{viii}This is a direct generalisation of the definition of groups of geometries as sets of similar geometries.

6.3. Presentation of the problem and the strategy

- **Playout/rollout:** Once an unexplored node \mathcal{N}_i is selected, a geometry g_j is arbitrarily chosen from the associated group of geometries G_i .
- **Backpropagation:** The scores of all ancestors of \mathcal{N}_i (i.e. all nodes selected during the selection/expansion step) are updated with the lateral interaction energy $E_j - E_j^{alone}$ of the chosen geometry $g_j \in G_i$.

During the selection step, choosing the best child node can be seen as a typical multi-armed bandit problem (as described in section 6.3.2, since each child node represents a distinct group of geometries). Therefore, this framework is expected to benefit from an UCB-like approach by adding an upper confidence bound to the empirical average payout used as a score for choosing a best child node. An MCTS using such score $S_i = \mathbb{E}(W_i)_{emp} + b_i(n_i, T)$ (instead of an empirical average payout alone) is referred as an Upper Confidence Tree (UCT).

In an MCTS scheme, the selection step can be seen as a successive selection over increasingly specific groups. In this manner, each selection step benefits in its early stages from coarse-grain sampling effects where large sub-groups of apparently sub-optimal geometries are excluded (i.e. at low depth into the tree the choice is made among large groups of geometries), while deeper in the tree the selection is performed over increasingly specific groups, where the exponentially reduced number of reachable geometries at this point makes such finely tuned sampling achievable.

The sampling associated with such framework displays an interesting property: while the whole geometry space is sampled, the regions with apparently larger many-body energy effects associated are more sampled. As a consequence, such MCTS trees are in practice explored/evaluated in a largely unbalanced manner (i.e. some branches are much more explored than others).

In summary, the UCT framework is expected to be particularly relevant for solving our adequate training set building problem: the diverse property is expected to be fulfilled since the whole geometry space is sampled, while the chemically relevant property is expected to be fulfilled by the enhanced sampling of regions associated with apparently larger many-body energy effects. Finally, the reasonably-sized property depends on the convergence rate of the chosen UCT framework. This last property is closely related to the learning rate of the UCT framework, as it can be seen more clearly by reformulating our adequate training set building problem into a strategy learning task.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

Indeed, UCTs are a common tool of reinforcement learning^{ix} used for designing optimal strategies. An easy to grasp real-world application can be found in strategy-based game playing, where UCT-based approaches demonstrated spectacular performance^{32;33;40}. Consequently, on the one hand, a UCT approach is expected to be relevant for solving our adequate training set building problem, while on the other hand, UCTs are known to perform remarkably well for playing strategy-based games. This suggests that our adequate training set building problem can be reformulated in terms of a strategy-based game. Such analogy not only provides with an alternative point of view for better visualization purposes, but is also relevant in developing an intuition of the mechanisms involved in our novel reinforcement sampling approach described in this chapter.

6.3.3 Reformulation of the training set construction as a game

Let us describe the construction of a single geometry as a strategy-based game entitled “Build a configuration”, whose goal is to create a configuration (i.e. an assembly of molecular adsorbates occupying specific active sites of a catalytic surface). A configuration can be seen as an abstract description of a geometry. “Build a configuration” is actually a simple game:

- This is a turn-based 1-player game^x.
- This game is played on a predefined board that represents a surface or any lattice-based space. A board is composed of an arbitrary number of cells C_i representing active sites \mathcal{S}_i on that surface. Just as active sites are positioned next to each other, cells can be directly connected and considered adjacent. In addition to relative positioning, the relative orientation of active sites can be taken into account by considering oriented cells. Mathematically, a board is a node-colored and edge-colored graph whose nodes are cells^{xi}.
- This game requires a predefined set of pawns. Each pawn P_i represents a molecular adsorbate \mathcal{M}_i . Pawns can be placed on the board, in cells. A pawn P_i placed in a board cell C_j represents a molecule \mathcal{M}_i adsorbed at the active site \mathcal{S}_j on the surface. However, the chemical nature of the adsorbate is not enough: orientation should also be taken into consideration. Therefore, pawns can come with predefined orientations. In fact, pawns contain all necessary parameters for adsorbing a given adsorbate (offset from the surface, angles, ...). Finally, just as a multidentate adsorbate can occupy multiple active sites, pawns can spawn across multiple cells

^{ix}Reinforcement learning is a sub-field of machine learning, along with supervised and unsupervised learning

^xJust as a chemist would naturally place adsorbates on the surface, one at a time.

^{xi}Note that while this definition allows for the description of finite surfaces and periodic surfaces (spheres or torus), in practice, a path connecting an active site with one of its periodic images cannot be properly represented with such graph. As a consequence, a large enough unit cell should be considered when building an associated board.

6.3. Presentation of the problem and the strategy

- This game is played in N_{ads} turns, where N_{ads} is the target number of adsorbed molecular entities desired in the final configuration. At each turn, the player selects a pawn and places it on the board. Therefore, playing an instance of this game is equivalent to sequentially building an input geometry by adding adsorbates on a surface, with specific parameters (positions, offsets, orientations), one at a time.
- The board is initially empty. At the end of the N_{ads} turns, the final configuration created (by placing pawns on the board) is given a score. This score represents the absolute total lateral interaction energy (i.e. the absolute many-body energy effects, or deviation from the 1-body model) associated with the corresponding geometry. Indeed, in our Cluster Expansion based model Hamiltonian framework, we are only interested in learning the many-body correction terms to the adsorption energy. Let us distinguish which components are included/excluded in the final score:
 - Excluded components: Each pawn P_i placed on the board is worth points in an amount that depends on the occupied board cell C_j . These points represent the stabilization energy associated with the adsorption of a single molecule \mathcal{M}_i at the \mathcal{S}_j active site of the surface considered, alone (i.e. no other molecules are adsorbed on the surface). These contributions are *a priori* known to the player beforehand, as preliminary studies of competitiveness among active sites are usually performed before conducting kinetic Monte-Carlo simulations. Therefore, these points are not taken into account for the score, since they do not contain multi-body effects.
 - Included components: Bonuses and maluses are granted for each many-body pattern present in the final configuration. These points represent the energy contributions of lateral interactions between adsorbed molecular entities. Lateral interactions can be stabilizing or destabilizing and are *a priori* unknown to the player, so that they can be seen as hidden score contributions whose sum is only revealed at the end of the game. These points are the only ones considered for computing the final score, and can result in a positive or negative total sum corresponding to the total many-body effects (that can be globally stabilizing or destabilizing).
 - Final score computation: the final score is taken as the absolute value of the number of points earned for the corresponding configuration.

This scoring is coherent with a description relying on a cluster expansion based model Hamiltonian, where the total adsorption energy of an occupied surface is taken as the sum of the individual adsorption energies of each adsorbate as if taken alone plus correction terms corresponding to many-body lateral interactions. Here, we are only interested in the correction terms to the adsorption energy (associated with many-body lateral interactions), and the final score correspond to the total strength of the cumulative many-body effects (either globally stabilizing or not).

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

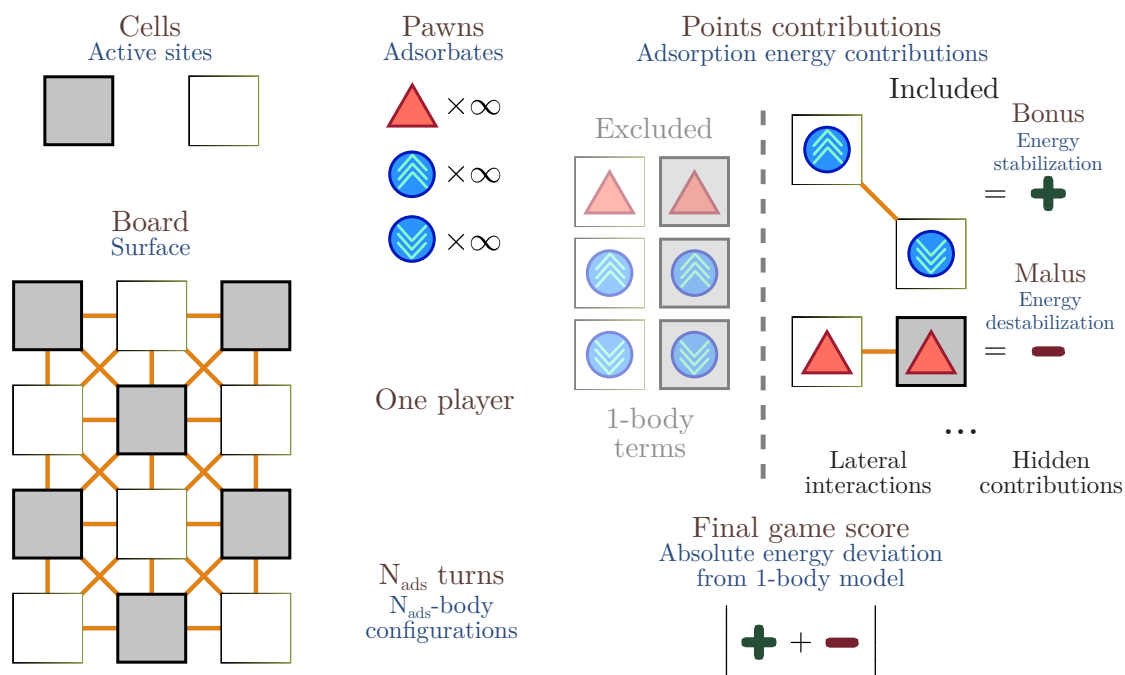


Figure 6.2 – Synoptic illustration of the “Build a configuration” game. The parallel with hetero-catalysis is highlighted by explicitly linking game features to surface science concepts. At each turn, the player chooses a pawn and places it on a free available cell. This is equivalent to the chemist placing adsorbates at non-occupied active sites on a surface. After N_{ads} turns, the occupied board (called configuration) represents an occupied surface with N_{ads} adsorbates. Each many-body pattern found on the final configuration earns points, in an amount that is based on the strength of the associated lateral interaction. These points can be positive or negative, but the individual details are not known to the player. The final score is the absolute value of the sum of these points (one could imagine some sort of game where the aim would be to be the best mythical creature, either “angel” by accumulating the most “good karma” points, or “demon” by accumulating the most “bad karma” points). This score of a final configuration is chosen to match the absolute adsorption energy deviation, from the 1-body model, of the corresponding geometry.

Figure 6.2 represents a synoptic illustration of the “Build a configuration” game features.

Despite its apparent simplicity, “Build a configuration” is a strategy-based game due to the additional points contributions (i.e. pattern-related bonuses and penalties), even though these contributions are hidden (i.e. *a priori* unknown to the player). This is better illustrated with an example: in chess, early queen attacks would seem advantageous to the novice player, but latter in the game, these moves would surely prove disappointing and the player pay the hidden cost associated with such risky moves. The quality of a move is learned by experience and simulation, so that an experimented player will eventually develop an intuition helping them devise strategies, just as a chemist uses their own intuition to build chemically relevant structures (e.g. creating a H-bond network, ...).

6.4. Designing a UCT for model Hamiltonians: Reinforcement Sampling

UCTs provide a general framework for efficiently learning (i.e. with as few simulations as possible) how to play strategy-based games. Therefore, we expect a UCT-based approach to be relevant for learning how to play this “Build a configuration” in a minimal amount of simulations (and associated computational time for evaluating the DFT-based adsorption energies corresponding to final configurations). This learning is based on exploring game strategies leading to high-score final configurations and extracting expertise/intuition for devising strategies to explore novel configurations with even higher scores. Therefore, a UCT-based learning is designed to generate a diverse set of high-quality configurations during its training, using the smallest amount of simulations (i.e. reducing the number of sub-optimal configurations generated). With this game reformulation, it becomes clearer why the set of configurations explored by a UCT framework is expected to provide an adequate training set.

By all means “Build a configuration” is an unusual game where, in particular, simulations are expensive (since evaluating the score requires a full DFT-based geometry optimization). Therefore, the Reinforcement Sampling framework for model Hamiltonians introduces tweaks from the pure UCT framework, in practice, to better fit the specific properties of the “Build a configuration” game.

6.4 Designing a UCT for model Hamiltonians: Reinforcement Sampling

Since Reinforcement Sampling for model Hamiltonians is not simply a pure UCT framework, its design is not straightforward, as it encompasses further optimizations and approximations to tackle the technical difficulties and specifics relative to the adequate training set problem.

6.4.1 Patterns

Patterns have a crucial role when dealing with CE-based model Hamiltonians (where there are also sometimes called figures), as they represent lateral interactions. Their definition and identification are of utmost importance for building meaningful model Hamiltonians. In particular, careful considerations must be applied to recognize equivalent patterns, since lateral interactions are defined up to global translation, rotation and chirality (i.e. the energy of an isolated sub-system is invariant by translation, rotation and reflection).

Let us introduce notations to rigorously describe patterns from a mathematical point of view:

- Abstract binding sites (ABS): Let an ABS be an abstract representation of a binding site.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

An ABS needs to contain only 3 pieces of information: chemical nature of the binding site ; orientation relative to a reference axis (if applicable, e.g. bridge sites) ; chirality. Note that abstract binding sites are defined for both active sites on the surface (e.g. Pt top-site) and binding sites on adsorbates (e.g. $[\text{SCN}]^-$ S-site). These ABS can, therefore, be thought of as anchor points (for both adsorbates and surface) and are assumed to be enough for describing unambiguously ligand-surface connections, since additional adsorption parameters (offset, bond angle, ...) can be encoded in the chemical type. Note that ABS do not contain information about absolute positioning, and therefore are unchanged by global translation.

- Abstract surface: Similarly to board (see 6.3.3), surfaces are abstractly represented by a node-colored and edge-colored directed graph whose nodes are ABS representing surface active sites. Adjacent sites are connected by directed edges containing only essential geometrical information: distance between sites and orientation relative to the reference axis.
- Abstract path: Starting with paths on the surface, the associated abstract description can be generic enough to no longer require an arbitrary reference axis, making abstract paths invariant by global rotation. A directed path (i.e. an ordered sequence of pairwise adjacent active sites) on the surface is abstractly represented by 4 components that are invariant by translation and rotation:
 - Ordered sequence of ABS, corresponding to surface active sites along the path, in a generic version where the absolute orientation is discarded (i.e. only chemical type and chirality are considered). Note that the first and last ABS are not explicitly included, as an abstract path focuses on the connection between endpoints.
 - Ordered sequence of distance between adjacent sites considered along the path.
 - Ordered sequence of edge-in–edge-out angles along the path (i.e. relative orientation of next ABS with respect to previous ABS, for each ABS considered along the abstract path).
 - Ordered sequence of edge-in–active site angles along the path (i.e. intrinsic orientation of the current ABS relative to the orientation of previous ABS, for each ABS considered along the abstract path).
- Abstract occupied surfaces (AOS): Here, we are interested in modelling the occupied active sites (and corresponding adsorbates) and their connections, with translational and rotational invariance. Theoretically, a complete description of the connection between two active sites on the surface should include all possible paths connecting them, so that the full environment (even far from both sites) can be taken into account. However, in practice, local effects are significantly stronger, so that we consider only shortest

6.4. Designing a UCT for model Hamiltonians: Reinforcement Sampling

paths (in terms of numbers of sites along the path)^{xii}. An occupied surface is therefore abstractly represented by a node-colored and edge-colored directed multigraph whose nodes are occupied surface ABS (each coupled with the associated adsorbate ABS) and directed edges are abstract paths representing shortest paths between them. Note that, just as two occupied active sites can have multiple shortest paths, a multigraph can be defined with multiple edges connecting the same nodes.

Figure 6.3 shows a synoptic illustration of these notions and their relation to pattern definition.

Theoretically, patterns are rigorously defined by vertex-induced subgraphs of AOS (i.e. a subset of nodes and all edges whose both endpoints are within that subset). However, such definition is rather impractical in terms of both manipulation/storage and comparison. Therefore, in practice, we consider two main approximations:

- Simple subgraphs: when two occupied active sites are connected by multiple shortest paths, only one is considered per pattern. This approximation helps reduce the possible number of patterns and extend the extrapolative power of the model Hamiltonian being trained (since two rigorous patterns sharing the same simple subgraph, with maximum size, will be considered with correlated energy effects).
- Simple paths: instead of considering the whole connectivity of a subgraph, a pattern is defined as a simple path (i.e. no node seen twice on the path) on an AOS. This major approximation is only impacting patterns with at least 3 bodies where discarded connections (i.e. shortest paths) contain ABS that are not already included (otherwise, the encoded geometrical information is enough to ensure that a single pattern can be represented by a single path). This approximation also helps reduce the possible number of patterns definable and extend the extrapolative power of the model Hamiltonian being trained. Finally, paths are much more convenient for storage and comparison purposes^{xiii}.

The major benefit of these approximations lies in the reduction of the maximum number of learnable patterns, since the number of possible paths is much lower than the number of possible graphs of same size, especially for patterns with many bodies involved. This reduction in the number of regressors naturally induces energy correlations between full patterns. Consequently, such representation helps the extrapolation of unseen full patterns from shared paths.

^{xii}The significance of local effects compared to long distance ones can also be seen in the relevance of a truncated cluster expansion model.

^{xiii}Comparing paths is straightforward, unlike graphs requiring isomorphism detection (even though such comparison can be done in linear time for planar graphs).

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

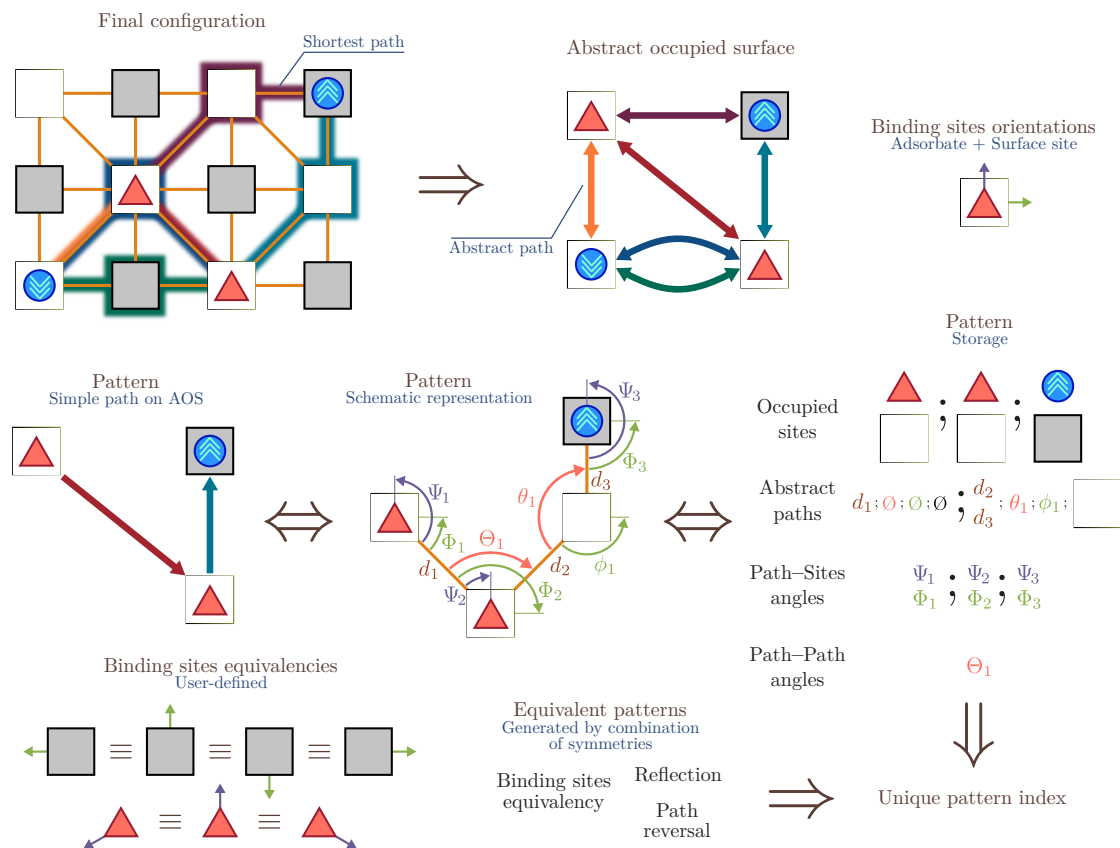


Figure 6.3 – Synoptic illustration of the pattern definition. From a final configuration, all shortest paths between any two occupied sites are computed, and represented here with colored shadows. Note that a cutoff of maximum 2 edges was applied here (so that occupied sites are not considered connected in the abstract occupied surface (AOS) if they are more than two edges away). These paths only are considered in the AOS. Note that the AOS is a directed multigraph, since abstract paths are directed and two occupied sites can have multiple shortest paths between them. A pattern is then defined as a simple path (no same edge/node explored twice, maximum one edge between two nodes) on the AOS, instead of the usual connected simple sub-graph (no multiple edges between two nodes). Using the properties of sites seen along the path (including chemical properties, distances and orientations of both the adsorbates and surface sites, if any), a translational and rotational-invariant description is formed, while avoiding redundancies. Each time a new pattern is detected, all equivalent patterns are generated by the combination of all applicable symmetries (user-defined ABS equivalencies, reflection and reversal). All these equivalent patterns are assigned a unique pattern index.

6.4. Designing a UCT for model Hamiltonians: Reinforcement Sampling

More precisely, in this framework, patterns are taken as simple paths on an AOS, and are composed of 4 types of information that are invariant by global translation and rotation:

- Ordered sequence of couples of ABS (in their generic version with chemical type and chirality only) corresponding to the occupied active sites and related adsorbates.
- Ordered sequence of single shortest abstract paths along the pattern path.
- Ordered sequence of in-path–out-path angles.
- Ordered sequence of in-path–ABS angles (for both the occupied surface ABS and adsorbate ABS, if applicable).

Figure 6.3 illustrates this definition on an example.

This representation of patterns is therefore invariant by global translation and rotation. But patterns (paths) should also be invariant by reflection and by reversing the path direction. These symmetries are actually not enough to ensure that patterns are uniquely identified since each ABS can display additional symmetries (e.g. fcc sites on a 111 surface have C_3 symmetry, and the N-site of NH_3 belongs to the C_{3v} point group). These symmetries are taken into account by generating all equivalent patterns constructed from the equivalents of all ABS involved. For each such equivalent pattern (path), its reversed and/or reflected versions are also generated to cover all identified sources of invariance that can be applied to a path on the surface.

However, patterns defined as simple paths on an AOS correspond to walks^{xiv} on the surface (and not only to paths), allowing for the description of embedded graphs on the surface (by representing a graph traversal, for example). Therefore complex patterns might require additional symmetries considerations due to the presence of:

- Automorphisms (i.e. multiple identical traversals of the same graph) when an embedded graph on the surface can be represented by multiple identical patterns.
- Equivalencies (i.e. multiple different traversals of the same graph) when multiple different patterns represent the same embedded graph on the surface.

We suggest to detect these additional symmetries on-the-fly by merging the detected patterns sharing the same connections on the surface. Finally, on top of these symmetries, systematic

^{xiv}Paths where nodes and edges can be repeated.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

correlations can occur. These correlations might give rise to additional pattern multiplicities that ought to be defined by the user (i.e. expert-based knowledge), since an automatic and accurate detection of such correlations would likely require the exploration of all possible final configurations.

A classical way to uniquely define equivalent patterns would be to consider only the minimum equivalent pattern, using a lexicographic order for example⁴¹, each time a pattern is identified. Such approach becomes quickly computationally expensive, especially since pattern detection is constantly performed in a Reinforcement Sampling run.

Therefore, we suggest to use a time-memory trade-off by defining a lookup table. Since the same patterns are constantly being identified, it is most beneficial to perform the exhaustive generation of equivalents only once per pattern. All of these equivalent patterns are stored under the same unique identifier, leading to subsequent equivalency evaluations in $O(1)$ operations.

6.4.2 Basic UCT design

The UCT approach is highly generic, so applying it to a specific case is not straightforward, as one needs to choose and tune implementation details.

Tree structure (modular organization)

One of the most crucial elements of an MCTS instance is undoubtedly the organization of the tree used. In other words, how the final configurations are organized in groups. Groups are meant to gather similar configurations, and the whole tree can be seen as a metric: the more similar two configurations are, the closer they should be on the tree (see section 6.3.2). In a MCTS framework, groups are used to extrapolate the score of close configurations^{xv}, so that the convergence rate is dependent on the metric used.

In the case of final configurations, the optimal *a priori* metric would be based on the sharing of similar lateral interactions. However, the partitioning of final configurations with respect to lateral interactions similarity is not a trivial task, since one would need to enumerate all possible configurations and identify all patterns beforehand, or privilege some lateral interactions to be more important^{xvi}. We implemented one of the latter approaches by considering a partitioning inspired by game playing, where final configurations are partitioned by the first

^{xv}This is reminiscent of the fundamental principle “similar structures have similar properties” of QSPR models (Quantitative Structure-Property Relationship) in Chemoinformatics, where the metric is an essential component.

^{xvi}One can think of this as a multi-scale clustering task with unknown data...

6.4. Designing a UCT for model Hamiltonians: Reinforcement Sampling

move (i.e. the first turn of a “Build a configuration” game), then sub-partitioned based on the second move, and so forth.

However, the number of possible moves per turn is quite large, so partitioning only on complete moves would lead to a large branching factor (i.e. the number of children per node) and a small maximal depth that would reduce the relevance of an MCTS approach. Inspired by move grouping^{42;43}, we subdivide each move into the following sub-partitions, in that order:

- Chemical nature of the adsorbate (e.g. H₂O O-site, NH₃ N-site, ...).
- Chemical nature of the active site (e.g. Pt, Au, ...).
- Type of the active site (e.g. top, bridge, fcc, hcp, ...).
- Site number (i.e. position on the surface).
- Orientation of the adsorbate, with respect to the arbitrary axis, if applicable.
- Chirality state of the adsorbate.

These sub-partitions and this specific order were subjectively chosen based on chemical intuition in order to produce intermediate groups with a relative similarity in terms of lateral interactions. Figure 6.4 illustrate this partitioning on a simple example.

This game-playing inspired partitioning provides an acceptable tree structure that is, furthermore, coherent with a metric based on the sharing of similar lateral interactions. Indeed, let us consider a node \mathcal{N}_i that is located at a depth corresponding to k turns of choices/partitioning. In other words, \mathcal{N}_i corresponds to an incomplete configuration (with k adsorbates) generated after k turns. Therefore, all final configurations (with N_{ads} adsorbates) descending from that node \mathcal{N}_i will result from the addition of $N_{\text{ads}} - k$ new adsorbates, corresponding to the remaining $N_{\text{ads}} - k$ turns. Since all adsorbates placed on the surface are not modified by subsequent moves, any final configuration or group/node descending from \mathcal{N}_i is guaranteed to contain the incomplete k adsorbates based configuration represented by \mathcal{N}_i , along with all patterns present in this incomplete configuration. As a consequence, this partitioning ensures that, for any branch of the corresponding tree, the further down a node/group is in that branch, the more lateral interactions are shared among that group, since patterns are cumulative^{xvii}.

^{xvii}Note that some branches are likely to converge faster toward homogeneous groups than others. Or maybe another partitioning would have provided better group homogeneity convergence. Actually, no guarantees are made concerning group homogeneity, since we only discuss the number of shared patterns (which are, nonetheless, a pre-requisite for group homogeneity in terms of lateral interactions). Still, this basic turn-based partitioning guarantees that the number of shared patterns within groups is bound to increase along any path from such tree.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

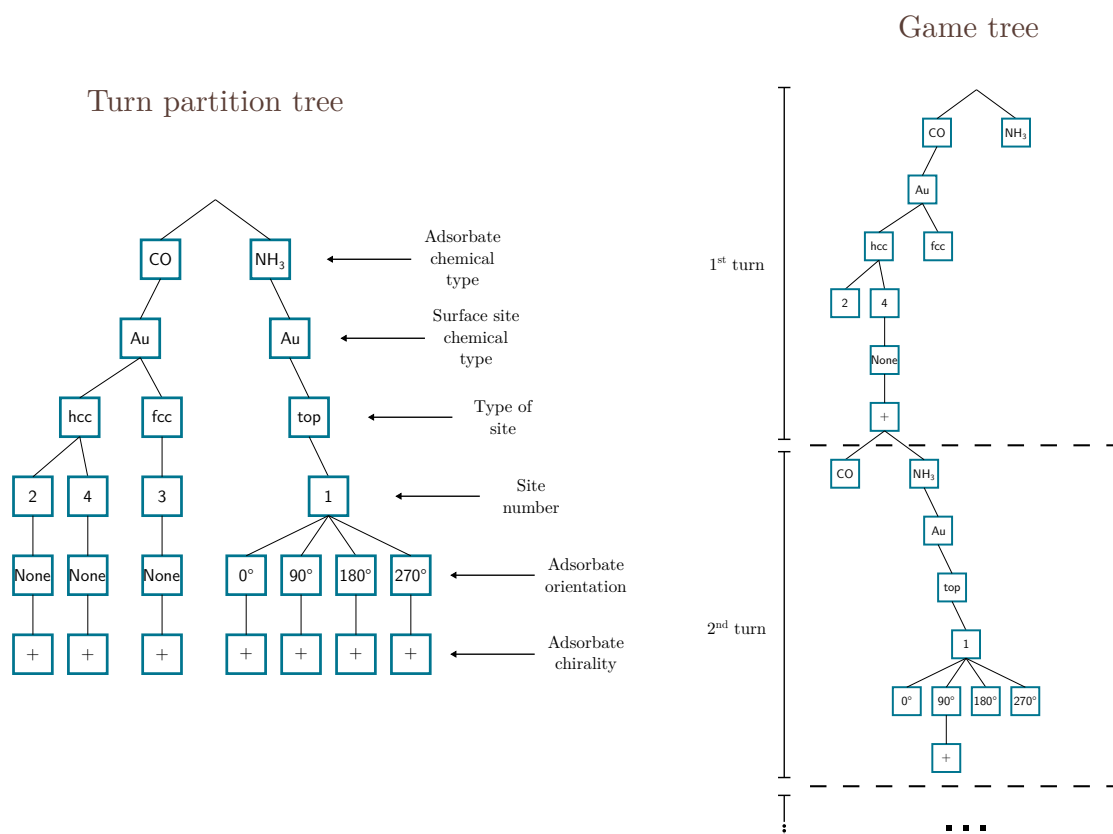


Figure 6.4 – Illustration of the turn-based partitioning for constructing the full game tree. All possible moves during a turn are partitioned (see the Turn partition tree) according to properties of the new site(s) being occupied (at the end of that turn). This partitioning is fixed and based on chemical intuition (i.e. expert selection). The leaves (i.e. terminal nodes) of this turn partition tree represent acceptable moves for the current turn. In this illustration, the turn partition tree represents the first turn. Following turns are represented by their own altered partition tree (some moves might become unavailable, while others might finally become available). The full game tree is globally organized/partitioned in a turn-based manner. Indeed, each node representing a game state after a single turn is expanded into all 2-turns game states that can be played from that single turn game state. As a consequence, all nodes representing a k -turns game states are at the same depth in that tree, and the more turns are required to reach a certain game state, the deeper its corresponding node is located in the tree. Each turn is sub-partitioned according to the turn partition tree, leading to the full structure depicted in this figure.

6.4. Designing a UCT for model Hamiltonians: Reinforcement Sampling

This turn-based partitioning also provides highly modular trees that can easily be extended/truncated for generating final configurations with different numbers of adsorbates (i.e. changing the number of total turns, see figure 6.4). To that purpose, one might want to re-use previous data that was generated from configurations with a different number of adsorbates, instead of restarting from scratch. In order to help generating data that can be relatively generalized to any final number of adsorbates, we store properties per adsorbate, so that all scores used are at least defined for any number of adsorbates^{xviii}.

Finally, this partitioning produces trees that are convenient to visualize, as it illustrates the possible evolution of a “Build a configuration” game. Each end of turn corresponds to a fixed depth, and adding a configuration (i.e. exploring the tree from node \mathcal{N}_i) is made with the full knowledge of the incomplete configuration corresponding to \mathcal{N}_i , in a convenient way (i.e. the position, orientation, chirality and nature of previously placed adsorbates are fully known). This structure makes, therefore, the definition of custom symmetries (for avoiding the redundancy of configurations within the tree) more user-friendly (e.g. fixing the position of the first adsorbate, restraining the orientation of the second adsorbate with respect to the first one, ...).

Concerning symmetries, a major source of redundancy comes from the ordering in which adsorbates are placed, since adsorbates can be placed in any order to generate the same final configuration. As a consequence, we force an arbitrary lexicographic order on the sequence of adsorbates placed. This lexicographic order ensures that for each branch in the tree, no other permutation of the associated sequence of moves can be found in any other branch of the tree. In practice, this is done by imposing each move to be strictly greater than the previous one, using the usual lexicographic order \leq_{move} on the tuple associated with a move (adsorbate chemical nature, surface site nature, site type, site number, orientation, chirality).

Normalized exploration rate

Due to the lexicographic order, the children of a node are likely unbalanced in terms of number of accessible final configurations (e.g. all the other moves can be explored after the \leq_{move} -smallest move whereas no move is accessible after the \leq_{move} -largest one). This contributes to create unbalanced branches, even when fully explored, leading to the over-representation of \leq_{move} -small moves as first, or so, moves.

Such imbalance can also result from the unequal representation of some properties within the 1-body terms. For example, if multiple distinct orientations (0° , 90° , 180° , 270°) of an NH_3

^{xviii}Note that this trick remains a somewhat crude extrapolation from dissimilar data.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

molecule are considered whereas CO does not have alternative orientations, then the number of moves placing an NH_3 molecule is roughly 4 times more than for placing a CO molecule.

By default, a UCT approach will uniformly sample children that seem equally promising, independently of the number of geometries that they represent (i.e. independently of the number of accessible final configurations from them). For example, a UCT-based sampling of the 1-body configurations with the NH_3 and CO adsorbates considered above being equally relevant would lead to sampled configurations having NH_3 and CO equally represented (i.e. NH_3 configurations and CO configurations in equal proportions). Indeed, a classic UCT framework has a tendency to uniformly sample equally promising moves.

However, we highlighted that tree imbalance leads to some moves being over-represented in the set of accessible final configurations. Therefore, when such imbalance is present, uniformly sampling moves leads to a nonuniform sampling of configurations. Indeed, with equally promising branches, a move leading to more accessible final configurations should be sampled more than a move leading to just a few configurations, if we are interested in representing all similarly promising configurations equally. This is well illustrated using our equally promising NH_3/CO example, where each orientation of NH_3 is sampled roughly 4 times less than the single orientation of CO if the chemical type related moves (i.e. choosing either NH_3 or CO) are equally sampled, despite NH_3 configurations displaying a greater diversity due to additional sub-partitioning (based on orientation).

A normalized exploration rate, inspired by Bosc *et al.*⁴⁴, can be used to take into account the number of accessible geometries and consider a uniform effective exploration ratio between unbalanced branches. The general idea behind this is intuitive: each exploration is weighted by the number of explorations that would be required on the most populated sibling to sample the same fraction of descendants. Back to our NH_3/CO example: each exploration of the CO node shall be counted as 4 explorations (i.e. the number of explorations required to equivalently sample the NH_3 node)^{xix}. Mathematically, this reads as:

$$n_i^{\text{norm}} = n_i \times \frac{\max_{j \in \text{Sbl}(i)} (d_j)}{d_i}, \quad (6.8)$$

where n_i the number of times node \mathcal{N}_i was explored, n_i^{norm} is the normalized number of explorations of \mathcal{N}_i , d_j is the number of descendants of node \mathcal{N}_j and $\text{Sbl}(i)$ are the siblings of node \mathcal{N}_i (i.e. the children of the parent of node \mathcal{N}_i). Note that with this definition, the

^{xix}Another way to see it is: counting in terms of NH_3 -equivalent explorations.

6.4. Designing a UCT for model Hamiltonians: Reinforcement Sampling

normalized number of explorations of the parent node T^{norm} (required for computing the UCB pre-factor) is taken as the sum of the normalized number of explorations of all its children.

Using these normalized number of explorations for computing the UCB score (through the exploration bias term, see section 6.3.2) should contribute to the reduction of the sampling bias introduced by unbalanced partitioning in the tree structure.

KL-UCB score

Many functional forms are available for the exploration bias term $b_i(n_i^{\text{norm}}, T^{\text{norm}})$ defined in section 6.3.2. While variants of the original UCB1 policy have been developed, displaying better empirical performance on specific classes of the multi-armed bandit problem, KL-UCB was found to be the only algorithm performing consistently better than the UCB1 policy³⁹.

In our case, since DFT computations are expensive, we can typically explore up to only a few hundred geometries. For this reason, we rely on a KL-UCB framework, as this algorithm is expected to be efficient even with a small number of observations³⁹. Finally, KL-UCB requires only the empirical mean $\mathbb{E}(W_i)_{emp}$ and the bounds of the target function W_i of node \mathcal{N}_i , in comparison to other variants requiring additional information (e.g. the empirical variance³⁶). Besides, KL-UCB is parameter-free in practice.

The KL-UCB score is based on the Kullback-Leibler divergence, a metric for comparing probability distributions, and its functional form is a bit different from the classical $S_i = \mathbb{E}(W_i)_{emp} + b_i(n_i, T)$ UCB form. The KL-UCB score of a node \mathcal{N}_i is defined as:

$$S_i = \mu_i^{-1} \left(\max_q \left\{ q \in [0, 1] \mid KL_{div}(\mu_i(\mathbb{E}(W_i)_{emp}), q) \leq \frac{\log(T^{\text{norm}})}{n_i^{\text{norm}}} \right\} \right), \quad (6.9)$$

where μ_i is a normalization function for the target variable W_i , n_i^{norm} is the number of normalized explorations of node \mathcal{N}_i , T^{norm} is the number of normalized explorations of its parent node, and KL_{div} is the Bernoulli Kullback-Leibler divergence:

$$KL_{div}(p, q) = p \log\left(\frac{p}{q}\right) + (1-p) \log\left(\frac{1-p}{1-q}\right) \quad (6.10)$$

Note that the target function W_i needs to be bounded in $[0, 1]$ for computing the KL-UCB score. However, the exact bounds of W_i are *a priori* unknown: even though the absolute energy deviation is always positive, its effective lower bound is dependent on the geometries

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

within G_i and its effective higher bound can only be known exactly by evaluating all the geometries in G_i . Therefore an estimate of the effective bounds $[\min(W_i), \max(W_i)]$ of W_i can be computed from its empirical minimum $\min(W_i)_{emp}$ and maximum $\max(W_i)_{emp}$ assuming, rather crudely, a uniform distribution^{xx}. Using these bounds, the normalization function μ_i is designed to rescale the W_i variable to $[0, 1]$.

6.4.3 Pre-exploration

Despite the efficiency of KL-UCB to deal with even a small number of observations, a DFT-based Reinforcement Sampling is far from the usual application of a UCT framework. The reason behind this lies in the most relevant difference between “Build a configuration” and classical board games: the computational cost associated with evaluating the score of a final game state. Indeed, conventional games rely on the application of a simple set of rules for quickly evaluating a final game state, whereas “Build a configuration” requires a DFT-based geometry optimization which is orders of magnitude more computationally expensive. As a consequence, additional efforts should be made to favour the selection of adequate configurations, especially during the playout/rollout phase where the choice is made randomly in “vanilla” MCTS.

Indeed, instead of a random sampling of unexplored nodes, a rollout policy can be applied through heuristics^{45–47} or other exploration techniques⁴⁸, in order to filter out configurations that are expected to be sub-optimal. Since including a final configuration requires computationally expensive calculations, such policy is therefore especially important here to avoid spending a large amount of computational time on sub-optimal configurations.

After a pre-exploration step composed of P successive random pre-explorations, P novel^{xxi} final configurations are selected. These configurations are expected to display large absolute energy deviation, according to the UCT-based policy that selected them. As a consequence, the pre-exploration step can be seen as a broad sampling of configurations, with predicted strong many-body effects, from which one should choose the next configuration to be included in the model Hamiltonian training set.

In fact, a random rollout would avoid such major convergence issues by eventually selecting an unexpectedly good final configuration. Fortunately, instead of relying on random discovery, it is possible to actively select the configuration that would benefit the most to the model

^{xx}This is known as the German tank problem, see section 6.4.4

^{xxi}In the sense that none of them was already included in the training set (i.e. chosen for DFT-based evaluation), but a same configuration could be pre-explored multiple times.

6.4. Designing a UCT for model Hamiltonians: Reinforcement Sampling

Hamiltonian in terms of diversity. In fact, the diversity requirement is only there to prevent the resulting model Hamiltonian from systematically focusing on lateral interactions whose contributions are already well defined, instead of exploring lateral interactions whose contributions are ill-defined. Mathematically, the uncertainty associated with a linear model is properly defined by the variance of its parameters $\text{Var}(\hat{\beta})$. Each new observation reduces the variance matrix $\text{Var}(\hat{\beta})$ by a symmetric positive semi-definite^{xxii} matrix $\Delta \text{Var}(\hat{\beta})$ that roughly depends on how redundant is the new observation with respect to the previous ones (globally, the more a configuration has already been seen, the less its inclusion impacts the variance matrix). By estimating the changes in the variance matrix $\Delta \text{Var}(\hat{\beta})$ ^{xxiii}, one can choose the configuration that benefits the most the model Hamiltonian. One should note that this is a typical variance-based DoE approach³¹. A convenient quality estimator for the variance reduction can be seen in the trace $\text{Tr}(\Delta \text{Var}(\hat{\beta}))$, considering that all lateral interactions are equivalently important^{xxiv}.

As a conclusion, we propose a rollout policy that consists in selecting, among the pre-explored configurations, the configuration expected to benefit the most the model Hamiltonian being learned. This policy can be seen as an additional compromise between exploration (choosing the configuration with *a priori* least redundant information) and exploitation (the UCT-based selection policy acts as a filter for predicted high-score configurations). Finally, even if the whole UCT-based sampling went to fail, Reinforcement Sampling would still be similar to an A-optimal variance-based active learning algorithm³¹ due to the $\text{Tr}(\Delta \text{Var}(\hat{\beta}))$ -based final selection during the rollout step. One should, therefore, expect Reinforcement Sampling to be more efficient than random sampling as an active learning algorithm.

6.4.4 Implementation

The Reinforcement Sampling framework described in this chapter has been implemented in Python3, with a modular design and a relatively few number of required parameters. Our implementation heavily relies on the NetworkX library for all graph-related algorithms, and Numpy whenever applicable. The code is available on GitHub, along with a rich documentation and examples.

^{xxii}A symmetric matrix M is positive semi-definite if $\forall x, \quad x^\top M x \geq 0$.

^{xxiii}Assuming a constant variance of the noise.

^{xxiv}This A-optimal design (i.e. minimizing the trace of the inverse Fisher information matrix of the whole model) is a popular choice for active learning algorithms³¹

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

This module was designed as a versatile tool, with numerous flags/options for addressing most cases. For example, our implementation can easily be turned into a purely pseudorandom sampler, or even using a pre-exploration step on a pseudorandom sampling, leading to a typical variance-based active learning algorithm.

Let us review some major implementation details that are not straightforward from the Reinforcement Sampling features described in this chapter, for application to CE-based model Hamiltonians.

Node storage

As mentioned in section 6.3.2, UCT trees are explored in a largely unbalanced manner. This feature is taken advantage of by storing only the evaluated nodes and expanding the tree on-the-fly, enabling the effective storage in memory of an incomplete combinatorially large space that would otherwise be impossible to store in full^{xxv}.

In practice, the expansion of the tree is made by maintaining a turn partition tree containing all available sub-moves until the end of the current turn (see Figure 6.4). All move filters are applied to this turn partition tree by removing the sub-moves filtered out along with their descendants, then pruning the above branches by removing those with no complete turn-move leaf remaining (using an inverse topological ordering).

By default, two filters are applied, at the beginning of a new turn, when the turn partition tree is reset:

- Lexicographic ordering: all turn-moves that are \leq_{move} -smaller than the previous turn-move are removed.
- No doubly occupied active sites: all turn-moves that would occupy an already occupied active site are removed.

Additional user-defined filters (also called symmetries) can be applied, leading to hardly predictable move filtering from one turn to the next turns. Therefore, during the expansion, all subsequent moves could be found unsuitable, leading to a failed exploration.

Note that the lexicographic ordering is applied by deleting incompatible sub-trees while exploring the turn partition tree in a BFS^{xxvi} order to save computational time.

^{xxv}Using coarse approximations, the number of nodes evaluated after T exploration steps is linear in T , since it is loosely bound by $O(Td \max_i(B_i))$ where d is the tree depth and $\max_i(B_i)$ is the maximal number of children.

^{xxvi}Breadth First Search, exploring the siblings of a node before its children

Biased exploration

As mentioned in section 6.4.2, the lexicographic order filtering induces a sampling bias in the space of accessible final configuration (if the moves sampling is unbiased), that ought to be corrected by a bias in the sampling of moves. To do so, the number of final configurations that can be accessed from a node should be estimated. From a lexicographic order filtering only, estimating the number of final configurations from a complete turn-move, that is the $(N + 1)$ -th \leq_{move} -largest turn-move and after n additional turns, is equivalent to the number of strictly increasing sequences of size n with N possible values (i.e. $\binom{N}{n} = \frac{N!}{n!(N-n)!}$). Therefore, the number of final configurations accessible from any sub-move can be computed from those of its children, and are therefore updated in an inverse topological order, in practice.

However, the lexicographic order filtering is not the only filter applied. In particular, another important filter is applied to avoid doubly occupied active sites. Yet this filter, simply “removes” an active site, reducing the number of final configurations roughly equivalently between node siblings. In fact, the only imbalance effects from such filter comes from the correlation with other filters (e.g. if CO can only be adsorbed at a fcc site while NH₃ can be adsorbed at fcc or hcp site, the occupation of a hcp site will impact the remaining possible adsorptions of NH₃ but not CO). However, such effects are hardly predictable, and in real-life applications, should not be significant compared to the lexicographic order filtering. Therefore, only the lexicographic order filtering is considered for estimating the number of final configurations and computing the normalized exploration rate.

Finally, custom (i.e. user-defined) symmetry functions (defined to avoid redundancies specific to the modeled system) may also affect unequally the number of descendants, and are not considered in this normalized exploration rate approach (as there is no general rule for quickly predicting their effects beforehand).

Pre-exploration bias

First of all, the final selection from the configurations sampled during the pre-exploration step optimizes the reduction of the trace of the covariance matrix for the model parameters $\text{Tr}(\Delta \text{Var}(\hat{\beta}))$ (see section 6.4.3). However, when using a large uniform guess for the parameters variance, any new observation that is linearly independent from the previous observations roughly induces the same reduction of the covariance matrix trace^{xxvii}, regardless of the number of new patterns discovered. This behavior is undesired because it reduces the likelihood of the new observation to explore an unseen subspace of relevant patterns (the fewer the

^{xxvii}Indeed, any non-redundant new observation is equivalent in terms of increasing the rank of Γ (i.e. reducing the null space of Γ).

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

number of new variables in a non-redundant observation, the fewer the number of variable subspaces, spanned by a subsets of variables, are impacted by this observation). Furthermore, the reduction in extrapolation is not expected to benefit much the framework since it has been partially designed to deal with extrapolated models. Therefore a bonus score is added for each new pattern discovered.

Finally, during the pre-exploration step, a balance between exploration and exploitation is achieved by the UCT framework based on the predicted absolute energy deviation from the current model. However, this model is likely partially trained and meant to be updated, leading to the predicted energies being modified. This would not be very problematic if the rollout was purely random, but in our case, the rollout policy is expected to be biased toward apparently relevant configurations. As a consequence, the model quality directly impacts the rollout exploration bias. Therefore, nodes explored at different times might have been explored using different policies, so that their reward cannot be compared. For example, a node \mathcal{N}_1 that has been explored in the early stages of Reinforcement Sampling should contain rewards from geometries that were expected to be chemically relevant (i.e. large absolute energy deviation) from the early and incomplete form of the model, but are now not so chemically relevant with respect to the final model, whereas a node \mathcal{N}_2 explored in the final stages have been explored with a much better model, leading to effectively better rewards $\mathbb{E}(W_2^{pred})_{emp} \stackrel{?}{>} \mathbb{E}(W_1^{pred})_{emp}$ with the current model. Therefore, to accelerate the re-exploration of nodes explored with a poorly trained model, we in principle allow the weighting of exploration count, so that each node exploration can be weighted according to the current model quality for being considered in the UCB score (i.e. an exploration with a robust model can count as multiple explorations with a poorly trained model). However, since we currently lack the justification for a particular weighting function, we do not exploit this capability in practice.

Estimation of score statistics (mean, min/max)

First of all, the KL-UCB score ought to be computed on normalized rewards. Therefore, one should estimate the distribution maximum $\max(W_i)$ and minimum $\min(W_i)$ from the empirical maximum $\max(W_i)_{emp}$ and minimum $\min(W_i)_{emp}$. As mentioned in section 6.4.2, assuming a uniform distribution, this is a typical German tank problem that lead to the estimates:

$$\begin{aligned} \min(W_i) &\approx \frac{n \times \min(W_i)_{emp} - \max(W_i)_{emp}}{n - 1} \\ \max(W_i) &\approx \frac{n \times \max(W_i)_{emp} - \min(W_i)_{emp}}{n - 1} \end{aligned} \tag{6.11}$$

6.4. Designing a UCT for model Hamiltonians: Reinforcement Sampling

where n is the number of records seen. The details of the derivation can be found in the section D.2 of appendix D.

For the “real” DFT-based energy deviations, the rewards are directly independent from the model, and can therefore be stored safely. However, the predicted absolute energy deviations do depend on the model, and therefore, cannot be permanently stored in the nodes (since these predicted deviations would become outdated as soon as the model is updated), see section 6.4.3. So instead, the inputs (for the model Hamiltonian) are stored, each in the form of a hash table linking for each pattern index found, the number of corresponding patterns found in the final configuration considered. These pattern multiplicities (i.e. the model Hamiltonian input) corresponding to the geometry g_i will be referred as Γ_i in the following. Therefore, each node \mathcal{N} should theoretically store Γ_i for each geometry g_i explored from that node \mathcal{N} . However, such storage would require large amounts of memory. So, in order to reduce the memory consumption, all the explored inputs are not explicitly stored in the nodes. Instead, our implementation relies on approximations to drastically reduce the storage in each node.

Let us now consider the approximation concerning the average predicted absolute energy deviation. First of all, updating the mean of absolute deviation for a node would require the storage of all seen configurations for that node. Instead of greatly increasing the memory consumption, we use an approximation that requires only to store $\mathbb{E}(\Gamma_i)_{emp}$, which is equivalent to a single configuration per node in terms of memory:

$$\mathbb{E}(W_i^{pred})_{emp} = \mathbb{E}(|\Gamma_i^\top \hat{\beta}'|)_{emp} \approx \mathbb{E}(\Gamma_i^\top |\hat{\beta}'|)_{emp} = \mathbb{E}(\Gamma_i)_{emp}^\top |\hat{\beta}'|, \quad (6.12)$$

where $\hat{\beta}' = \hat{\beta} - \hat{\beta}_{1\text{-body}}$ are the energy contributions (of the CE-based model Hamiltonian) of many-body lateral interactions only (i.e. the contributions of the single-body model are subtracted to those of the full model).

This rough approximation is undoubtedly prone to errors, in general, if both stabilizing and destabilizing lateral interactions are defined (this approximation is exact if only stabilizing/destabilizing lateral interactions are found). In fact, the most significant errors introduced by this approximation come from configurations where there is a coexistence of apparently largely stabilizing and destabilizing lateral interactions. Fortunately, these errors are expected to be relatively mitigated within a linear model dominated by extrapolation^{xxviii}.

^{xxviii} Indeed, apparent lateral interactions are extrapolated largely stabilizing (or destabilizing) when they are not found canceled out in configurations. Therefore, if the already selected configurations are representative enough for the pre-explored configurations, it is relatively unlikely that apparently largely stabilizing and destabilizing lateral interactions be found together in pre-explored configurations.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

For the maximum/minimum predicted absolute energy deviations, additional approximations are made to limit the memory consumption to $O(m)$ per node (i.e. roughly the storage required for a fixed number of configurations). The approximation chosen uses $\min(\Gamma_i)_{emp}$ and $\max(\Gamma_i)_{emp}$, corresponding to the minimum/maximum number of corresponding patterns seen in a single configuration for each pattern index^{xxix}. This approximation uses therefore the equivalent of two configurations per node in terms of storage, and defines $\min(W_i^{pred})_{emp}$ and $\max(W_i^{pred})_{emp}$ as:

$$\begin{aligned} \min(W_i^{pred})_{emp} &\geq 0 \\ \max(W_i^{pred})_{emp} &\leq \max \left\{ \begin{array}{l} |\hat{\beta}'_{>0} \min(\Gamma_i)_{emp} + \hat{\beta}'_{<0} \max(\Gamma_i)_{emp}| \\ |\hat{\beta}'_{>0} \max(\Gamma_i)_{emp} + \hat{\beta}'_{<0} \min(\Gamma_i)_{emp}| \end{array} \right. , \end{aligned} \quad (6.13)$$

where $\hat{\beta}'_{>0}$ are the model parameters (for the energy deviation) for which only positive energy contributions are kept (i.e. negative terms are set to zero), and reciprocally for $\hat{\beta}'_{<0}$. Therefore, $\hat{\beta}'_{>0} \max(W_i^{pred})_{emp} + \hat{\beta}'_{<0} \min(W_i^{pred})_{emp}$ is an upper bound on the maximal energy deviation observed, and $\hat{\beta}'_{>0} \min(W_i^{pred})_{emp} + \hat{\beta}'_{<0} \max(W_i^{pred})_{emp}$ is a lower bound on the minimal energy deviation observed. Consequently, this approximation gives lower/upper bounds for the true empirical minimum/maximum of the absolute energy deviation, respectively. As stated previously, this approximation is rather crude, but can be partially justified in the case of a linear model dominated by extrapolation.

Finally, each energy deviation prediction comes with an uncertainty that can be estimated through confidence intervals. In order to account for these confidence intervals, additional approximations are considered to keep a storage in $O(m)$ per node. This additional approximated correction is explained in details in section D.2 of appendix D, along with a reformulation and explanations of most approximations above, applied to the predicted energies (as it is simpler to grasp than with absolute deviation energies, without changing much of the reasoning).

Solvers

For computing the KL-UCB score, our implementation uses the `scipy.optimize.brentq` solver.

For updating the linear model associated with the internal model Hamiltonian, and estimating the covariance change, we use the Recursive Least Squares algorithm rank-Greville, described in details in chapter 5.

^{xxix}Once again, this is stored in a hash table to save space in case a small fraction of distinct pattern indexes were found.

6.5 Results

6.5.1 Protocol presentation

As a proof of principle, the Reinforcement Sampling approach is applied to a real-life application of CO oxidation on Pd(111), by learning the CE-based model Hamiltonian described by S. Piccinin and M. Stamatakis,⁹ and used as reference.

Pattern definition

The model of ref⁹ defines 52 relevant “figures” (which roughly correspond to what we call patterns) composed of up to 3 adsorbates/sites and exploring up to the second nearest neighbors; four of these figures correspond to steric clashes. Therefore, in our framework, we are only interested in patterns involving up to 3 adsorbates, 3 sites (occupied or not), and occupied sites are directly connected in the AOS representation if there are not further away than second nearest-neighbors. The adsorbates (CO and O) were defined with no characteristic orientation (i.e. treated with D_∞ symmetry group on the surface). The active sites (fcc and hcp) rigorously have D_3 for symmetry group, but their symmetry is not considered in the reference model Hamiltonian, so for coherence purposes, we treat them with the D_∞ symmetry group.

Using these restrictions, we find 56 distinct patterns^{xxx} corresponding to the relevant figures defined in the reference model Hamiltonian and 138 possible patterns. This difference of 4 additional relevant patterns are attributed to triangle shaped figures, since multiple 3-sites paths can define the same triangle. We identify these patterns and set them to be equivalent, with adequate multiplicities so that there is a bijection between relevant patterns and relevant figures, leading to a final number of 52 distinct relevant patterns and a total of 134 possible patterns. However, we convert the 4 relevant patterns corresponding to steric clashes to a user-defined filter that prevents any configuration from displaying these clashes. Accounting for these additional constraints, only 70 distinct possible patterns remain from the 138 initially detected patterns, 48 of which are distinct relevant patterns, i.e., corresponding to a relevant figure. While the four single-body terms are taken to be known from the beginning, the remaining 24 2-body patterns and 20 3-body patterns are to be learned via a relevant training set.

Model quality evaluation

In this test, we are interested in evaluating the performance of the Reinforcement Sampling framework for generating a training set that is optimal for learning the relevant many-body

^{xxx}These patterns are defined as paths, see section 6.4.1

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

energy contributions (i.e. training a relevant model Hamiltonian), using only a limited number of first-principles calculations.

In order to assess this performance, we use the Reinforcement Sampling framework to learn the reference model Hamiltonian described by S. Piccinin and M. Stamatakis, using only a limited number of access to the predictions of this reference Hamiltonian. Indeed, this model Hamiltonian already defines which patterns are relevant, so that the relevant many-body energy contributions are already given. Therefore, we wish to evaluate the final training set on its performance for estimating, not all the possible many-body energy contributions, but only the relevant ones, as identified in the model Hamiltonian.

Of course, these contributions are not known beforehand by the framework which has only access to a limited number of predictions from the reference model Hamiltonian, similarly as if the predictions were computed from first-principles calculations. One should also note that the selection of relevant parameters is not addressed here (actually, the internal model Hamiltonian defines all possible contributions): only the training set is evaluated, not the eventual dimensionality reduction (in fact, that is why we use a reference model Hamiltonian, so that the relevant contributions are already identified, and therefore, the performance of our framework is not impacted by the dimensionality reduction performance).

As a consequence, we evaluate the performance of a training set $\Gamma_1, \Gamma_2, \dots, \Gamma_n$ by:

- First, training a linear model (i.e. a model Hamiltonian) for estimating the 44 relevant multi-body energy contributions $\hat{\beta}_{rel}$ from these configurations and their associated energy (predicted from the reference model Hamiltonian).
- Then, the covariance matrix of the estimated contributions is computed $\text{Var}(\hat{\beta}_{rel})$.
- Finally, we use the trace of the covariance matrix $\text{Tr}(\text{Var}(\hat{\beta}_{rel}))$ as an estimator of the robustness of these parameters, and therefore as an estimator of the quality of the training set $\Gamma_1, \Gamma_2, \dots, \Gamma_n$.

In fact, the covariance matrix is independent from the predicted energies, and can be computed^{xxxi} as:

$$\text{Var}(\hat{\beta}_{rel}) = (X_{rel}^\top X_{rel})^{-1}, \quad (6.14)$$

where X is the observation matrix associated with the training set $\Gamma_1, \Gamma_2, \dots, \Gamma_n$, restricted to the 44 relevant many-body patterns.

^{xxxi} Assuming a non-underdetermined linear system.

In the unfortunate case where the training set does not sample the whole 44-patterns subspace, $X_{rel}^\top X_{rel}$ is singular. We can still evaluate the model quality on the subspace effectively sampled by using the eigenvalues^{xxxii} of $X_{rel}^\top X_{rel}$.

Modelling parameters

For this test, we limit the number of predictions from the reference model Hamiltonian to 100. This leads to the final training having 100 configurations. This compares to the 92 DFT configurations used in ref⁹.

The Reinforcement Sampling framework aims at reducing the computational cost spent of first-principles calculations for creating a training set. Therefore, it is natural to consider the smallest surface for the geometries in the training set. The smallest supercell where all patterns can be defined (due to periodic boundary conditions) is the 4×3 cell, therefore composed of 12 unit cells each containing a single fcc and a single hcp site, leading to 24 active sites in total.

Random sampling

First of all, we estimate the quality of a purely random sampling (still using a normalized exploration rate, to ensure a relatively unbiased sampling in the configuration space), and study the impact of coverage on the final training set quality.

Figure 6.5 represents the final training set quality (defined in section 6.5.1) obtained by simply sampling from random. From these results, it is clear that the coverage has a strong influence on the quality of the final training set.

At low coverage (i.e. small N_{ads}), each configuration displays only a small number of patterns, leading to poor quality training sets. Indeed, a simple argument is that if one doubles the number of each pattern found in each observations, then the covariance matrix of the parameters becomes $(2X_{rel}^\top 2X_{rel})^{-1} = \frac{1}{4} \text{Var}(\hat{\beta}_{rel})$. Another way to grasp this idea is to consider that the residual error ϵ_i associated with a configuration Γ_i is “split”/“diluted” into all the patterns seen, so that the more patterns the fewer residual “per pattern”.

Whereas at high coverage, the diversity of patterns seen is reduced from one observation to the other. Indeed, the patterns are being correlated, leading to less well-defined contributions. For example, on a fully covered surface, any 2-body pattern is never seen isolated, but always

^{xxxii}Indeed, since $X_{rel}^\top X_{rel}$ is symmetric, its eigenvalues are well-defined and the sum of their inverse is equal to the trace of the inverse matrix. Therefore, we can use the sum of the inverse eigenvalues that are strictly positive.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling




Figure 6.5 – Evaluation of the quality of the random sampling with respect to the coverage. The quality of a sampling is defined by the ability of the corresponding training set (obtained after such sampling) to properly determine the energy contributions $\hat{\beta}_{rel}$ of patterns defined as relevant. This ability is measured in terms of $\text{Tr}(\text{Var}(\hat{\beta}_{rel}))$, the trace of the covariance matrix of $\hat{\beta}_{rel}$ obtained after fitting a CE-based model Hamiltonian (with $\hat{\beta}_{rel}$ as parameters) on this training set. Therefore, the better (i.e. the most adapted) a sampling, the lower the $\text{Tr}(\text{Var}(\hat{\beta}_{rel}))$. A range of coverage are explored, by varying the number of adsorbates (i.e. occupied sites) on a 4×3 supercell composed of 24 sites. For each coverage, the random seed is reset to 0 and 10 consecutive tests are run. The error bars are estimated from the corresponding Student's t -distribution.

found with the same group of 3-body patterns, leading to its contribution being eventually indiscernible from those of the 3-body patterns group.

In the absence of correlation, we expect this performance to be roughly correlated to the total possible number of distinct configurations, and here, the optimal coverage corresponds to a third the sites being occupied $N_{\text{ads}} = 8$.

Nonetheless, the full subspace of relevant patterns is not systematically explored, even with $N_{\text{ads}} = 8$, leading to the final relevant model Hamiltonian being underdetermined and therefore a subset of relevant contributions being indiscernible. From a general point of view, random sampling does not seem well-adapted to describe systems subjected to constraints or strong correlations.

One should also note that below $N_{\text{ads}} \leq 7$ the random sampling does not consistently succeed in exploring all 44 relevant many-body patterns. Finally, configurations with $N_{\text{ads}} > 12$ are undefined due to the steric clashes constraints.

Since we search a method that is consistently better than random, we consider in the following the most favorable coverage for a random sampling: $N_{\text{ads}} = 8$.

Pure UCT sampling

Using a pure UCT-based sampling (i.e. using a single pre-exploration step) leads to an insufficient exploration after 100 steps. Indeed, after choosing the first configuration at random, the pure UCT-based framework failed to explore all 44 many-body patterns in almost all cases ($\approx 80\%$ of failure) with $N_{\text{ads}} = 8$.

Therefore, the exploration rate of the pure UCT framework is insufficient for this application.

Pre-explored random sampling

Here we consider a full exploration approach, by applying the variance-based selection of the rollout policy to the pre-explored configurations sampled at random.

Indeed, at the end of the rollout policy, the selected configuration is chosen among the pre-explored configurations as the one that improves the most the internal model Hamiltonian.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling




Figure 6.6 – Evaluation of the effects of pre-exploration (i.e. performing a variance-based DoE selection on a pool of 10 random configurations) on the quality of the random sampling. The quality of a sampling is defined by the ability of the corresponding training set (obtained after such sampling) to properly determine the energy contributions $\hat{\beta}_{rel}$ of patterns defined as relevant. This ability is measured in terms of $\text{Tr}(\text{Var}(\hat{\beta}_{rel}))$, the trace of the covariance matrix of $\hat{\beta}_{rel}$ obtained after fitting a CE-based model Hamiltonian (with $\hat{\beta}_{rel}$ as parameters) on this training set. Therefore, the better (i.e. the most adapted) a sampling, the lower the $\text{Tr}(\text{Var}(\hat{\beta}_{rel}))$. A range of coverage are explored, by varying the number of adsorbates (i.e. occupied sites) on a 4×3 supercell composed of 24 sites. For each coverage, the random seed is reset to 0 and 10 consecutive tests are run. The error bars are estimated from the corresponding Student's t -distribution.

So that adding a few pre-exploration steps should increase the exploration rate and the diversity of the training set. It is crucial to note that the variance-based final selection aims at reducing the full internal model Hamiltonian, i.e. the trace of the covariance $\text{Tr}(\text{Var}(\hat{\beta}_{full}))$ of all 70 patterns contributions, not only the relevant ones $\hat{\beta}_{rel}$ identified by the reference model Hamiltonian (since they are *a priori* unknown to the Reinforcement Sampling). Yet by reducing $\text{Tr}(\text{Var}(\hat{\beta}_{full}))$, we expect to also approximately reduce the variance for any subset of contributions (therefore also for $\hat{\beta}_{rel}$).

Using 10 pre-explored configuration sampled at random (instead of selected by the UCT framework) per step and applying the DoE-inspired variance based selection leads to a significant and systematic improvement over the random sampling for generating training sets that efficiently discriminate the relevant contributions, see figure 6.6.

Pre-explored UCT sampling (pure Reinforcement Sampling)

In this approach, at each step, the variance-based selection is applied to 10 configurations that are expected to be chemically relevant (strong many-body effects) as they are pre-explored using the UCT framework. Therefore, the subspace of relevant pattern is expected to be sampled.

Indeed, by filtering the pre-explored nodes based on their predicted chemical relevance, we hope to sample more thoroughly the space of relevant patterns, and therefore better discriminating the relevant contributions $\hat{\beta}_{rel}$, by pre-exploring preferentially configurations displaying relevant patterns. Indeed, for example, if all pre-explored configurations contain only relevant patterns, then optimizing $\text{Tr}(\text{Var}(\hat{\beta}_{full}))$ is equivalent to optimizing $\text{Tr}(\text{Var}(\hat{\beta}_{full}))$ among these configurations.

Using a 10 step pre-exploration from the UCT approach (i.e. pure Reinforcement Sampling approach) leads indeed to an improved exploration rate compared to the pure UCT approach, since all 44 relevant multi-body patterns are now found in every run. Besides, the final model is systematically found non-underdetermined, compared to the pure random sampling.

Nonetheless, compared with the pure pre-explored random sampling, the quality of the training set is systematically reduced for estimating the relevant contributions defined in the reference model Hamiltonian (see figure 6.7).

This is characteristic of a biased sampling, and this, despite the variance-based selection at the end of the rollout policy. Therefore, this sampling bias (even within the subspace of relevant configurations) is also strongly present in the pre-explored configurations (this is coherent with the poor exploration rate of the pure UCT-based framework), and cannot be removed during the final variance-based selection. Let us investigate this issue.

First of all, let us exclude the eventual UCT being completely confused by the extrapolated linear model and systematically sampling patterns that are not relevant. Indeed, the proportion of relevant patterns per configuration is equivalent to the proportion obtained by random sampling. Furthermore, the final training set still provides more robust parameters for the relevant contributions than by random sampling, which is a consequence of the subspace of relevant patterns being better sampled.

Therefore, this bias could be explained by two main mechanisms:

- The UCT framework is stuck in a local minimum. So, this would typically be a convergence issue. Indeed, the average absolute energy deviation is not found significantly

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling




Figure 6.7 – Comparison of the Reinforcement Sampling (10 pre-exploration steps from the UCT) with pre-exploration from random (10 pre-exploration steps from random) and random only. The quality of a sampling is defined by the ability of the corresponding training set (obtained after such sampling) to properly determine the energy contributions $\hat{\beta}_{rel}$ of patterns defined as relevant. This ability is measured in terms of $\text{Tr}(\text{Var}(\hat{\beta}_{rel}))$, the trace of the covariance matrix of $\hat{\beta}_{rel}$ obtained after fitting a CE-based model Hamiltonian (with $\hat{\beta}_{rel}$ as parameters) on this training set. Therefore, the better (i.e. the most adapted) a sampling, the lower the $\text{Tr}(\text{Var}(\hat{\beta}_{rel}))$. A range of coverage are explored, by varying the number of adsorbates (i.e. occupied sites) on a 4×3 supercell composed of 24 sites. For each coverage, the random seed is reset to 0 and 10 consecutive tests are run. The error bars are estimated from the corresponding Student's t -distribution.

larger than with a random sampling (4.30 eV instead of 4.31 eV).

- The relevant patterns have highly dispersed contributions (i.e. some contributions are negligible compared to others, but are still labelled “relevant”), leading the UCT framework to disregard these negligible “relevant” patterns. Indeed, the strongest relevant contributions is 0.2849 eV, while the weakest is 0.0016 eV. Besides, if the 20 strongest contributions only are considered relevant, then the Reinforcement Sampling is globally performing better than the pre-explored random sampling and significantly better than the random sampling, see figure 6.8.

Nonetheless, both mechanisms are likely to be involved. Indeed, to test the convergence issue hypothesis, 1000 steps of the Reinforcement Sampling framework are run with $N_{\text{ads}} = 8$ and compared with a random sampling. The average absolute deviation is now larger than with a random sampling (4.83 eV per configuration instead of 4.31 eV), suggesting a convergence issue (since such difference was not observed after only 100 steps). However, the training set quality was now lower than with a random sampling (0.47 eV^2 instead of 0.36 eV^2)^{xxxiii}, suggesting that this convergence issue is not particularly relevant.

As a consequence, the UCT framework induces a strong bias toward the sampling of the most significant interactions, which can lead to the full subspace of relevant patterns being sampled in an heterogeneous manner, leading to some directions being poorly sampled, and therefore leading to a poorly trained model (defined on arbitrary patterns or with highly dispersed contributions) compared to a pre-explored random sampling that is prone to sampling homogeneously the full 70 pattern space. In such a case, the gain from biasing toward more chemically relevant configurations is not as essential as ensuring an homogeneously sampled pattern space.

Impact of the number of pre-exploration steps

While the benefit of the DoE-inspired variance-based pre-exploration (i.e. choosing the configuration minimizing $\text{Tr}(\text{Var}(\hat{\beta}_{full}))$ from $N_{\text{pre-exp}}$ pre-explored configurations) are clear from these results, the importance of the number of pre-exploration steps $N_{\text{pre-exp}}$ is studied in figure 6.9. Globally, increasing the number of pre-exploration steps $N_{\text{pre-exp}}$ leads to overall better performance, but comes at the cost of additional computational time (scaling linearly with $N_{\text{pre-exp}}$).

^{xxxiii}If only the 20 strongest contributions are considered: pure Reinforcement Sampling is better than random sampling 0.09 eV^2 instead of 0.13 eV^2 , but slightly worse than pre-explored random sampling (0.08 eV^2). But Reinforcement Sampling is now better than pre-explored random sampling for determining the up-to 8 strongest contributions $7.8e-3 \text{ eV}^2$ instead of $1.3e-2 \text{ eV}^2$, which is coherent with an increased bias toward the strongest contributions. Otherwise, the decreased performance can simply arise from synergistic correlations that lead to higher absolute deviation energy.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling




Figure 6.8 – Comparison of the Reinforcement Sampling (10 pre-exploration steps from the UCT) with pre-exploration from random (10 pre-exploration steps from random) and random only. The quality of a sampling is defined by the ability of the corresponding training set (obtained after such sampling) to properly determine the strongest energy contributions $\hat{\beta}_{best}$. This ability is measured in terms of $\text{Tr}(\text{Var}(\hat{\beta}_{best}))$, the trace of the covariance matrix of $\hat{\beta}_{best}$ obtained after fitting a CE-based model Hamiltonian (with $\hat{\beta}_{best}$ as parameters) on this training set. Therefore, the better (i.e. the most adapted) a sampling, the lower the $\text{Tr}(\text{Var}(\hat{\beta}_{best}))$. A range of coverage are explored, by varying the number of adsorbates (i.e. occupied sites) on a 4×3 supercell composed of 24 sites. For each coverage, the random seed is reset to 0 and 10 consecutive tests are run. The error bars are estimated from the corresponding Student's t -distribution.




Figure 6.9 – Influence of the number of pre-exploration steps on the sampling quality. The quality of a sampling is defined by the ability of the corresponding training set (obtained after such sampling) to properly determine the energy contributions $\hat{\beta}_{rel}$ of patterns defined as relevant. This ability is measured in terms of $\text{Tr}(\text{Var}(\hat{\beta}_{rel}))$, the trace of the covariance matrix of $\hat{\beta}_{rel}$ obtained after fitting a CE-based model Hamiltonian (with $\hat{\beta}_{rel}$ as parameters) on this training set. Therefore, the better (i.e. the most adapted) a sampling, the lower the $\text{Tr}(\text{Var}(\hat{\beta}_{rel}))$. A range of coverage are explored, by varying the number of adsorbates (i.e. occupied sites) on a 4×3 supercell composed of 24 sites. For each coverage, the random seed is reset to 0 and 10 consecutive tests are run. The error bars are estimated from the corresponding Student's t -distribution.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

The apparent mitigated improvement for when configurations are pre-explored from UCT instead of random sampling might come from the strong bias imposed by the UCT framework, making it harder to find significantly different configurations during the pre-exploration. Besides, this pre-exploration step might come in direct competition with the exploitation optimization approach of the UCT framework.

6.5.2 Conclusion

In this chapter, we addressed the optimization of building a training set for any relevant CE-based model Hamiltonian while limiting the number of first-principles calculation. The benefits of such approach are multiple: not only the final model Hamiltonian is better defined (reduction of the covariance of the parameters), but the computational time required for its construction is reduced, and the building process itself is automated.

From the definition of an adequate training set, an active learning framework is designed, based on Reinforcement Learning and classical variance-based Design of Experiments, to tackle the challenge imposed by the limited number of first-principles calculation.

From preliminary results, it seems that the chemical relevance property is not as important as homogeneously sampling the pattern space in order to train a model Hamiltonian composed of highly dispersed contributions. Therefore, in such a case, a classical variance-based Design of Experiments approach on top of a random generation seems to provide the best results. However, if only the most significant contributions are of interest or if the relevant patterns have similar contributions (and no synergistic correlations with respect to the energy deviation) or if the total lateral interaction is an explicit target, then a pure Reinforcement Sampling approach may prove more suited.

For cases in between, additional efforts ought to be made to find a compromise between the pure Reinforcement Sampling approach and the pure pre-explored random approach, benefiting from both the preferential sampling of relevant patterns and the broadly homogeneous and efficient sampling within that subspace.

Bibliography

- [1] Ma, Z.; Zaera, F. *Encyclopedia of Inorganic Chemistry*; American Cancer Society, 2006.
- [2] Rothenberg, G. *Catalysis: concepts and green applications*; John Wiley & Sons, 2017.

- [3] Liu, J.-C.; Wang, Y.-G.; Li, J. Toward rational design of oxide-supported single-atom catalysts: atomic dispersion of gold on ceria. *Journal of the American Chemical Society* **2017**, *139*, 6190–6199.
- [4] Kumar, K. V.; Porkodi, K.; Rocha, F. Langmuir–Hinshelwood kinetics—a theoretical study. *Catalysis Communications* **2008**, *9*, 82–84.
- [5] Baxter, R.; Hu, P. Insight into why the Langmuir–Hinshelwood mechanism is generally preferred. *The Journal of chemical physics* **2002**, *116*, 4379–4381.
- [6] Farkas, A.; Hess, F.; Over, H. Experiment-based kinetic monte carlo simulations: CO oxidation over RuO₂ (110). *The Journal of Physical Chemistry C* **2012**, *116*, 581–591.
- [7] Sutton, J. E.; Lorenzi, J. M.; Krogel, J. T.; Xiong, Q.; Pannala, S.; Matera, S.; Savara, A. Electrons to reactors multiscale modeling: catalytic CO oxidation over RuO₂. *ACS Catalysis* **2018**, *8*, 5002–5016.
- [8] Hess, F.; Over, H. Rate-determining step or rate-determining configuration? the Deacon reaction over RuO₂ (110) studied by DFT-based KMC simulations. *ACS Catalysis* **2017**, *7*, 128–138.
- [9] Piccinin, S.; Stamatakis, M. CO Oxidation on Pd(111): A First-Principles-Based Kinetic Monte Carlo Study. *ACS Catal.* **2014**, *4*, 2143–2152.
- [10] Stamatakis, M. Kinetic modelling of heterogeneous catalytic systems. *Journal of Physics: Condensed Matter* **2014**, *27*, 013001.
- [11] Andersen, M.; Panosetti, C.; Reuter, K. A Practical Guide to Surface Kinetic Monte Carlo Simulations. *Frontiers in Chemistry* **2019**, *7*, 202.
- [12] Evans, M.; Polanyi, M. Inertia and driving force of chemical reactions. *Transactions of the Faraday Society* **1938**, *34*, 11–24.
- [13] Nørskov, J. K.; Bligaard, T.; Logadottir, A.; Bahn, S.; Hansen, L. B.; Bollinger, M.; Bengaard, H.; Hammer, B.; Sljivancanin, Z.; Mavrikakis, M., et al. Universality in heterogeneous catalysis. *Journal of catalysis* **2002**, *209*, 275–278.
- [14] Michaelides, A.; Liu, Z.-P.; Zhang, C.; Alavi, A.; King, D. A.; Hu, P. Identification of general linear relationships between activation energies and enthalpy changes for dissociation reactions at surfaces. *Journal of the American Chemical Society* **2003**, *125*, 3704–3705.
- [15] Chen, W.; Schmidt, D.; Schneider, W. F.; Wolverton, C. Ordering and Oxygen Adsorption in Au–Pt/Pt(111) Surface Alloys. *J. Phys. Chem. C* **2011**, *115*, 17915–17924.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

- [16] Zhu, B.; Creuze, J.; Mottet, C.; Legrand, B.; Guesmi, H. CO Adsorption-Induced Surface Segregation and Formation of Pd Chains on AuPd(100) Alloy: Density Functional Theory Based Ising Model and Monte Carlo Simulations. *J. Phys. Chem. C* **2016**, *120*, 350–359.
- [17] Vignola, E.; Steinmann, S. N.; Le Mapihan, K.; Vandegehuchte, B. D.; Curulla, D.; Sautet, P. Acetylene Adsorption on Pd–Ag Alloys: Evidence for Limited Island Formation and Strong Reverse Segregation from Monte Carlo Simulations. *J. Phys. Chem. C* **2018**, *122*, 15456–15463.
- [18] Cao, L.; Li, C.; Mueller, T. The use of cluster expansions to predict the structures and properties of surfaces and nanostructured materials. *Journal of chemical information and modeling* **2018**, *58*, 2401–2413.
- [19] Sanchez, J. M.; Ducastelle, F.; Gratias, D. Generalized cluster description of multicomponent systems. *Physica A: Statistical Mechanics and its Applications* **1984**, *128*, 334–350.
- [20] Lerch, D.; Wieckhorst, O.; Hammer, L.; Heinz, K.; Müller, S. Adsorbate cluster expansion for an arbitrary number of inequivalent sites. *Physical Review B* **2008**, *78*, 121405.
- [21] Stampfl, C. Predicting surface phase transitions from ab initio based statistical mechanics and thermodynamics. *Phase Transitions* **2007**, *80*, 311–332.
- [22] Schmidt, D. J.; Chen, W.; Wolverton, C.; Schneider, W. F. Performance of cluster expansions of coverage-dependent adsorption of atomic oxygen on Pt (111). *Journal of chemical theory and computation* **2012**, *8*, 264–273.
- [23] van de Walle, A.; Ceder, G. Automating first-principles phase diagram calculations. *Journal of Phase Equilibria* **2002**, *23*, 348.
- [24] Hess, F. Efficient Implementation of Cluster Expansion Models in Surface Kinetic Monte Carlo Simulations with Lateral Interactions: Subtraction Schemes, Supersites, and the Supercluster Contraction. *Journal of Computational Chemistry* **2019**, *40*, 2664–2676.
- [25] Vignola, E.; Steinmann, S. N.; Vandegehuchte, B. D.; Curulla, D.; Stamatakis, M.; Sautet, P. A machine learning approach to graph-theoretical cluster expansions of the energy of adsorbate layers. *The Journal of Chemical Physics* **2017**, *147*, 054106.
- [26] Blum, V.; Hart, G. L.; Walorski, M. J.; Zunger, A. Using genetic algorithms to map first-principles results to model Hamiltonians: Application to the generalized Ising model for alloys. *Physical Review B* **2005**, *72*, 165113.
- [27] Hart, G. L.; Blum, V.; Walorski, M. J.; Zunger, A. Evolutionary approach for determining first-principles hamiltonians. *Nature materials* **2005**, *4*, 391–394.
- [28] Bajpai, A.; Frey, K.; Schneider, W. F. Binary Approach to Ternary Cluster Expansions: NO–O–Vacancy System on Pt (111). *The Journal of Physical Chemistry C* **2017**, *121*, 7344–7354.

-
- [29] Seko, A.; Koyama, Y.; Tanaka, I. Cluster expansion method for multicomponent systems based on optimal selection of structures for density-functional theory calculations. *Physical Review B* **2009**, *80*, 165122.
- [30] Herder, L. M.; Bray, J. M.; Schneider, W. F. Comparison of cluster expansion fitting algorithms for interactions at surfaces. *Surface Science* **2015**, *640*, 104–111.
- [31] Settles, B. *Active learning literature survey*; 2009.
- [32] Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M., et al. Mastering the game of Go with deep neural networks and tree search. *nature* **2016**, *529*, 484–489.
- [33] Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A., et al. Mastering the game of go without human knowledge. *nature* **2017**, *550*, 354–359.
- [34] Greene, W. *Econometric Analysis*; Pearson, 2018.
- [35] Auer, P.; Cesa-Bianchi, N.; Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning* **2002**, *47*, 235–256.
- [36] Audibert, J.-Y.; Munos, R.; Szepesvári, C. Exploration-Exploitation Tradeoff Using Variance Estimates in Multi-Armed Bandits. *Theor. Comput. Sci.* **2009**, *410*, 1876–1902.
- [37] Audibert, J.-Y.; Bubeck, S. Regret Bounds and Minimax Policies under Partial Monitoring. *Journal of Machine Learning Research* **2010**, *11*, 2785–2836.
- [38] Honda, J.; Takemura, A. An asymptotically optimal bandit algorithm for bounded support models. In Proceedings of the Twenty-third Conference on Learning Theory (COLT 2010. 2010; pp 67–79.
- [39] Garivier, A.; Cappé, O. The KL-UCB algorithm for bounded stochastic bandits and beyond. Proceedings of the 24th annual conference on learning theory. 2011; pp 359–376.
- [40] Schadd, M. P.; Winands, M. H.; Van Den Herik, H. J.; Chaslot, G. M.-B.; Uiterwijk, J. W. Single-player monte-carlo tree search. International Conference on Computers and Games. 2008; pp 1–12.
- [41] Barth, D.; David, O.; Quessette, F.; Reinhard, V.; Strozecki, Y.; Vial, S. Efficient Generation of Stable Planar Cages for Chemistry. 2015.
- [42] Saito, J.; Winands, M.; Uiterwijk, J.; Van Den Herik, H. Grouping Nodes for Monte-Carlo Tree Search. 19th Belgian-Dutch Conference on Artificial Intelligence. 2007; pp 276–283.

Chapter 6. UCT-based active learning applied to cluster expansion Hamiltonians: Reinforcement Sampling

- [43] Childs, B. E.; Brodeur, J. H.; Kocsis, L. Transpositions and move groups in Monte Carlo tree search. 2008 IEEE Symposium On Computational Intelligence and Games. 2008; pp 389–395.
- [44] Bosc, G.; Boulicaut, J.-F.; Raïssi, C.; Kaytoue, M. Anytime Discovery of a Diverse Set of Patterns with Monte Carlo Tree Search. 2016.
- [45] Drake, P.; Uurtamo, S. Move ordering vs heavy playouts: Where should heuristics be applied in Monte Carlo Go,”. Proceedings of the 3rd North American Game-On Conference. 2007; pp 171–175.
- [46] James, S.; Rosman, B. S.; Konidaris, G. An investigation into the effectiveness of heavy rollouts in UCT. **2016**,
- [47] Toyoda, T.; Kotani, Y. Monte Carlo Go Using Previous Simulation Results. Proceedings of the 2010 International Conference on Technologies and Applications of Artificial Intelligence. USA, 2010; p 182–186.
- [48] Loth, M.; Sebag, M.; Hamadi, Y.; Schoenauer, M. Bandit-based search for constraint programming. International Conference on Principles and Practice of Constraint Programming. 2013; pp 464–480.

7 Conclusions and Outlook

As stated by its title, this thesis has been focused on the development of numerical tools for the understanding and quantification of chemical interactions of adsorbates on metallic surfaces. Let us review this statement:

Indeed, this thesis explored the chemical interaction in its many forms, depending on the level of abstraction considered:

- At a DFT level, the energy effects are addressed through an energy decomposition analysis, where the energy contribution of interactions is decomposed between charge transfer and polarization effects.
- At a classical level, the chemical bonds are addressed through geometry-based descriptors with the concept of nearest neighbours.
- Finally, their unification is addressed in the form of a model Hamiltonian, defining a linear relationship between the total energy of a system and the presence of geometry-based interactions.

Besides, methodological development and algorithms has been the focus of study, throughout this thesis:

- The ALMO formalism is the key tool at the core of any rigorous polarization vs. charge transfer EDA.
- The SANN algorithm is based on an elegant method for defining geometry-based nearest neighbors, in a fundamental and parameter-free fashion.

Chapter 7. Conclusions and Outlook

- RLS algorithms are essential and efficient engines for updating the minimum-norm least-squares solution of a linear system.
- Finally, Reinforcement Learning is a fundamental framework that is of relevance for powering active learning tools, which are themselves of critical importance for constructing an optimal training set.

All these developments are either already available publicly on GitHub or will be rendered public upon publication.

Finally, the work achieved during this thesis has been directed towards metallic surfaces applications:

- The ALMO formalism has been made compatible with electronic smearing, allowing for the description of metals, as exemplified by the application to water adlayers on noble metal surfaces.
- The SANN algorithm has been extended toward interfaces, by taking local anisotropy into account, ideally suited for nanoparticles and solvation shells of ions at interfaces.
- A reinforcement learning based active learning framework (including a custom RLS algorithm) has been designed to optimally train model Hamiltonians describing reactive surfaces for KMC applications.

Let us now review in more details each work separately:

- The ALMO formalism has been unified with mixed-state theory, using a mean-field approximation called S-ALMO. In this approximation, the interaction between orbitals are averaged (i.e. mean occupancy-rescaled) across all accessible quantum states at a given temperature (i.e. electronic smearing) and applied for all states independently. Therefore, S-ALMO allows for the quantification of charge transfer effects involving metals. In particular, S-ALMO has been applied to water adlayers adsorbed on surfaces of noble metals in order to examine the adsorption effects and their origins.
- The SANN algorithm has been adapted to tackle local anisotropy, through a simple yet efficient correction term, leading to the ASANN algorithm. Indeed, the barycenter of the SANN-based nearest neighbors allows to estimate the field of view that should be relevant for searching nearest neighbors. This effective field of view is then plugged back into the SANN algorithm through a correction term that, consequently, accounts for the local anisotropy. As a consequence, ASANN extends the SANN algorithm to interfaces

where the local anisotropy is strong (e.g. corrugated surfaces, edges, nanoparticles, ...) while remaining parameter-free.

- An RLS algorithm, called rank-Greville, has been developed by maintaining a general rank factorization. Indeed, a simple formula is derived for updating the pseudoinverse of a matrix $A \in \mathbb{R}^{n \times m}$ of rank r when a new row is added. This formula is similar to the Greville formula but relies on a rank factorization, which allows for solving two independent sub-systems separately and merging them with reduced overall computational cost. This formula has been implemented in the rank-Greville RLS solver to achieve an update of the linear least-squares solution in $O(mr)$ operations. Besides, its remarkable algebraic simplicity allows to use exact numerical representations of rationals, without the need for symbolic computing.
- Finally, an active learning framework, called Reinforcement Sampling, has been developed to automatically generate, in a reproducible way, an optimal training set for a CE-based model Hamiltonian (i.e. handling many-body effects). By reformulating the construction of a relevant input geometry as a strategy-based game, a reinforcement learning method is applied to the construction of an optimal training set. This UCT-based approach tackles the unusual properties of this strategy-based game by introducing various tweaks, including a carefully crafted pre-exploration step exploiting the very model Hamiltonian being trained. Consequently, the Reinforcement Sampling framework performs better than a random approach, effectively replacing the chemical intuition, typically used for building a curated training set, by reinforcement sampling.

Nonetheless, many challenges still remain in each topic addressed. In particular, many open questions/research topics can be formulated that are deeply linked to the work presented in this thesis.

- The unification between ALMO and mixed-states theory proposed in S-ALMO could be extended to the extensions of ALMO. In particular, (X)ALMO that have been recently developed¹ and allow the same basis function to be used in multiple fragment, but are currently incompatible with metals. These smearing-compatible extensions could then be applied to provide valuable insights on the origin of alloys effects on heterocatalytic processes. Finally, S-ALMO is not yet compatible with spin-polarized systems, which precludes the energy decomposition analysis for ferromagnetic surfaces.
- ASANN relies on a simple local anisotropy descriptor that is only sensitive to a single privileged direction. This is relevant for interfaces, but this description is no longer adapted for molecules, for example, where the bonding anisotropy is much more complex. Using a better description of local anisotropy that is able to handle non-trivial

anisotropy, by using a multipole expansion based on spherical harmonics for example, should make it possible to extend further the (A)SANN algorithm.

- Despite promising results, the Reinforcement Sampling framework could still be optimized in many ways. Indeed, this framework relies on numerous assumptions and approximations. In particular, tighter bounds for the payout distributions could accelerate the convergence of the UCT. Other metrics could also be considered for better performance. Finally, additional efforts could be made to better integrate the pre-exploration step within the UCT framework.
- The results presented in Chapter 6 are just a proof of concept. Additional tests with different Hamiltonians should be performed to confirm the results obtained. Finally, application of the algorithm with a true coupling to on-the-fly DFT geometry optimizations will make our tool truly valuable for the routine training of model Hamiltonians.
- Finally, the Reinforcement Sampling framework specifically addresses the optimal construction of an adequate training set, regardless of the final model Hamiltonian chosen (i.e. which energy contributions $\hat{\beta}_{rel}$ to fit), only assuming that such model will contain patterns associated with significant contributions. The optimisation of the relevant patterns to include in the final model Hamiltonian (i.e. dimensionality reduction) is a complementary yet crucial matter that has not been explored in this work²⁻⁶. In fact, the full model Hamiltonian (i.e. $\hat{\beta}_{full}$) is used internally to predict the deviation during the pre-exploration, but optimizing the form of the final model Hamiltonian during the training set building step (and using such a model internally) should help to design even better training sets, specifically designed for the final model Hamiltonian.

In fact, the Reinforcement Sampling framework has been mainly thought as a proof of concept for the broad application of reinforcement learning techniques to chemical challenges involving an exploitation/exploration dilemma. Indeed, reinforcement learning is a sub-field of machine learning that has yet to be exploited to solve chemical challenges (compared with the growing applications of supervised/unsupervised learning techniques in Chemistry), with a notable exception in retrosynthesis^{7,8}. Yet, I believe that reinforcement learning can prove to be an asset to plenty of other chemical challenges. In particular, the strategy-based game reformulation could prove beneficial to many areas of Chemistry involving optimization challenges, or at least, it would help suggest an unconventional yet appealing point of view.

Bibliography

- [1] Khaliullin, R. Z.; VandeVondele, J.; Hutter, J. Efficient Linear-Scaling Density Functional Theory for Molecular Systems. *J. Chem. Theory Comput.* **2013**, *9*, 4421–4427.

- [2] Hess, F. Efficient Implementation of Cluster Expansion Models in Surface Kinetic Monte Carlo Simulations with Lateral Interactions: Subtraction Schemes, Supersites, and the Supercluster Contraction. *Journal of Computational Chemistry* **2019**, *40*, 2664–2676.
- [3] Vignola, E.; Steinmann, S. N.; Vandegehuchte, B. D.; Curulla, D.; Stamatakis, M.; Sautet, P. A machine learning approach to graph-theoretical cluster expansions of the energy of adsorbate layers. *The Journal of Chemical Physics* **2017**, *147*, 054106.
- [4] Blum, V.; Hart, G. L.; Walorski, M. J.; Zunger, A. Using genetic algorithms to map first-principles results to model Hamiltonians: Application to the generalized Ising model for alloys. *Physical Review B* **2005**, *72*, 165113.
- [5] Hart, G. L.; Blum, V.; Walorski, M. J.; Zunger, A. Evolutionary approach for determining first-principles hamiltonians. *Nature materials* **2005**, *4*, 391–394.
- [6] Bajpai, A.; Frey, K.; Schneider, W. F. Binary Approach to Ternary Cluster Expansions: NO–O–Vacancy System on Pt (111). *The Journal of Physical Chemistry C* **2017**, *121*, 7344–7354.
- [7] Segler, M. H.; Preuss, M.; Waller, M. P. Planning chemical syntheses with deep neural networks and symbolic AI. *Nature* **2018**, *555*, 604–610.
- [8] Schreck, J. S.; Coley, C. W.; Bishop, K. J. Learning retrosynthetic planning through simulated experience. *ACS central science* **2019**, *5*, 970–981.

A Appendix to Chapter 2

A.1 Density recovery

Even if no proper references are available, a quality estimator can be conceived: the number of electrons recovered. Indeed, since we are constructing the density of a N-electrons system, then this density must account for exactly N electrons (i.e. the N electrons have to be recovered in the density).

The number of electrons recovered N_{calc} can be computed by integrating the density over the whole space. In matrix form, this calculation becomes:

$$N_{calc} = \sum_{\mu,\nu} R_{\mu\nu} S_{\mu\nu} = \sum_{\mu,\nu} R_{\mu\nu} S_{\nu\mu} = \sum_{\mu} (RS)_{\mu\mu} = \text{Tr}(RS) \quad (\text{A.1})$$

Since this estimator needs to be computed only once, at the end of a simulation, it can then be computed systematically with negligible additional cost.

Counterexample evaluation: As an example, the number of electron pairs recovered from a simple model system is evaluated. Let us investigate a case as far as possible from the domains of application where our construction is exact: infinite smearing with a 2-block system. Indeed, adding other blocks will not fundamentally change much the problem, since the real challenge here is in fact when we have blocks that interact strongly.

Let us therefore consider a 2-block system, composed of 1 doubly occupied orbital in the first block, and 2 orbitals of same energy in the second block (i.e. high degeneracy) with 1 pair of electron to share between the two orbitals of that block (i.e. infinite smearing in the second block). Let us consider an arbitrary basis set, and an acceptable orbital construction such that

Appendix A. Appendix to Chapter 2

there is a strong interaction between blocks. Such a system is studied here with the following description:

$$S = \left(\begin{array}{c|cc} 1 & 0.5 & 0.5 \\ \hline 0.5 & 1 & 0 \\ 0.5 & 0 & 1 \end{array} \right), T = \left(\begin{array}{c|cc} 1 & 0 & 0 \\ \hline 0 & 1 & 0 \\ 0 & 0 & 1 \end{array} \right), T' = \left(\begin{array}{c|cc} 1 & 0 & 0 \\ \hline 0 & \sqrt{\frac{1}{2}} & 0 \\ 0 & 0 & \sqrt{\frac{1}{2}} \end{array} \right) \quad (\text{A.2})$$

where S is the basis set function overlap matrix, and T' is the rescaled orbital coefficient matrix. Applying our mean-field approximation, we obtain a mean-field orbital overlap matrix ${}^\alpha\sigma$, and an associated density matrix ${}^\alpha R$ to be compared with the realⁱ density matrix R :

$${}^\alpha\sigma \approx \left(\begin{array}{c|cc} 1 & 0.35 & 0.35 \\ \hline 0.35 & 1 & 0 \\ 0.35 & 0 & 1 \end{array} \right), {}^\alpha R \approx \left(\begin{array}{c|cc} 1.33 & -0.33 & -0.33 \\ \hline -0.33 & 0.58 & 0.08 \\ -0.33 & 0.08 & 0.58 \end{array} \right) \neq R \approx \left(\begin{array}{c|cc} 1.33 & -0.33 & -0.33 \\ \hline -0.33 & 0.67 & 0 \\ -0.33 & 0 & 0.67 \end{array} \right) \quad (\text{A.3})$$

The density matrix ${}^\alpha R$ obtained is different from the real density matrix R . Indeed, this simplified model system exhibits a strong occupation-state dependency, coupled with strong inter-blocs interactions.

As expected from a quality estimator, the 2-electron pairs counterexample studied above is $N_{calc} \approx 1.83 \neq 2$. Therefore, 8.33% of the total density was lost.

However, this estimator provides a test that is only a necessary, not sufficient condition. Thus, a total recovery of the density would not necessarily mean that the S-ALMO result is exact.

As can be seen in the above example, the errors are largest when partially occupied orbitals overlap with orbitals from other blocks. This is typically the case when adsorbates with diffuse densities closely approach the metal surface or when the electronic smearing temperature is high. For example, the electron loss reaches 0.04 electrons for H_2S at 2000 K, while it is only 0.02 electrons at 300 K.

ⁱThe real density matrix R was obtained with the exact ensemble BLW formulation (i.e. in this case: $R = \frac{1}{2}R_{10} + \frac{1}{2}R_{01}$, where R_{10} (resp. R_{01}) is the system where the electron pair is on the first (resp. second) orbital of the second block).

A.2. Comparison of results with two different Pt pseudo potentials and results from VASP

A.2 Comparison of results with two different Pt pseudo potentials and results from VASP

A.3 Further analysis of the deformation energy

A.4 Influence of the electronic temperature

The difference between 300 K and 2000 K in Figure A.4 can be understood recalling that the higher electronic temperature increases the occupation of higher lying orbitals (or rather bands for periodic systems). Hence, the total interaction energy is less stabilizing as also slightly anti-bonding bands are occupied. On the other hand, since the electron density is more diffuse, the attractive interaction stretches out to longer ranges. Regarding the individual terms, $E_{BSS E}$ and E_{deform} do not change significantly. ΔE_{frozen} indicates that the metal surface becomes less repulsive at 2000 K due to the increased diffuseness of the electron density. ΔE_{pol} is roughly of the same magnitude but with opposite sign, suggesting that at 2000 K the polarization cannot achieve the same stabilization as at 300 K. This is due to the fact that at 2000 K the low lying bands (the ones that can easily be polarized) are already partially occupied, so that in the presence of the adsorbate only a smaller energy gain can be achieved. Last but not least, ΔE_{CT} is also negative, identifying this term as the responsible for the loss in overall stability seen in the graph on the left. Similarly to polarization, the charge-transfer stabilization is reduced due to the partial occupation of the orbitals that are most favorable for accepting electron density from the other fragment.

A.5 Expanded Figures 2, 3 and 4

A.6 Effect of K-point sampling

Appendix A. Appendix to Chapter 2

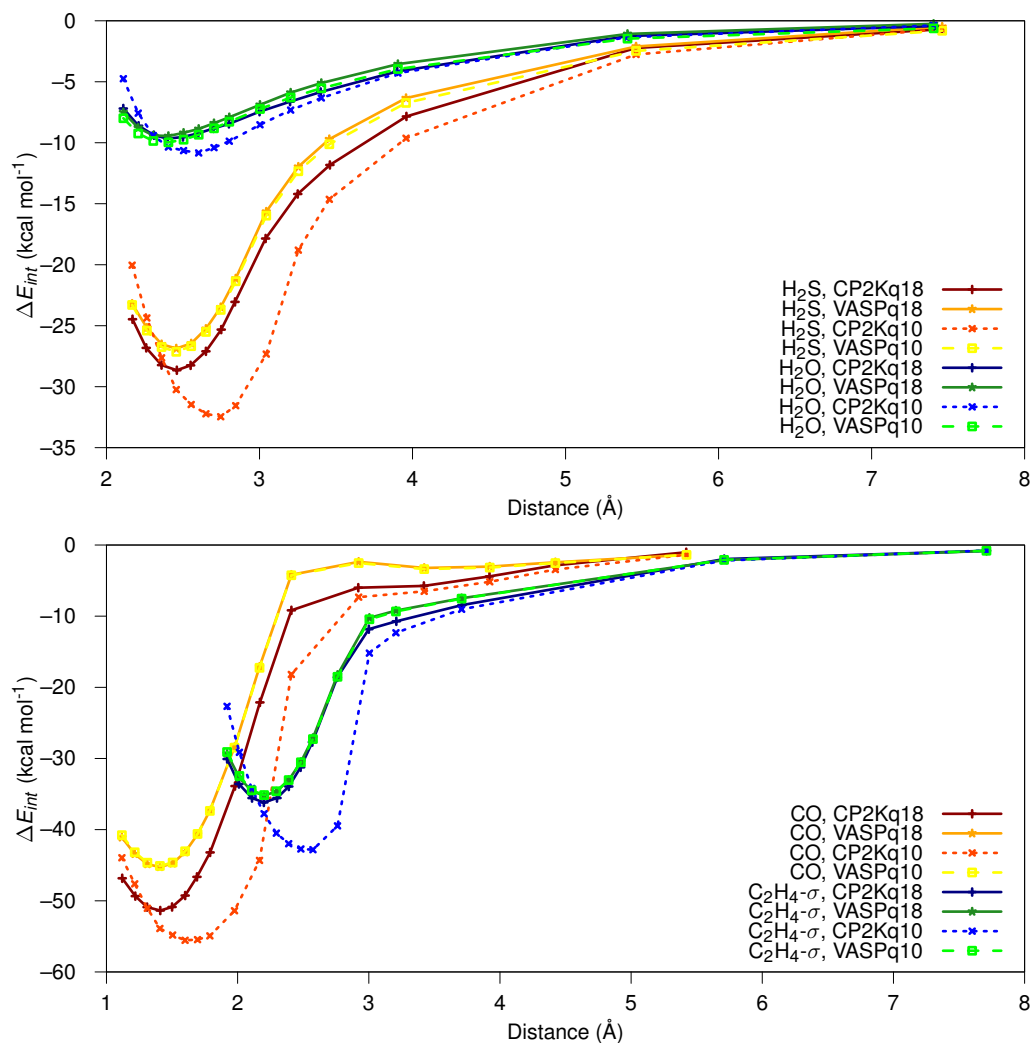


Figure A.1 – Top: H₂S and H₂O top adsorption on Pt(111). Bottom: Ethylene adsorption in a di- σ adsorption mode and CO on the fcc site. VASPq10 refers to standard VASP settings and pseudo-potentials, with a 400 eV energy cutoff plane-wave basis set. VASPq18 refers to a small-core pseudo-potential for Pt and the latest pseudo-potentials, recommended for GW computations, for all other elements. A refined FFT grid is used by using “accurate” precision settings and the plane-wave energy cutoff is set to 500 eV. CP2Kq10 refers to the GTH-PBE-q10 pseudo-potential, in conjunction with the DZVP-A5-Q10-323-MOL-T1-DERIVED_SET-0 basis set for Pt, while CP2Kq18 refers to the GTH-PBE-q18 small core pseudo-potential with DZVP-MOLOPT-SR-GTH as basis set.

A.6. Effect of K-point sampling

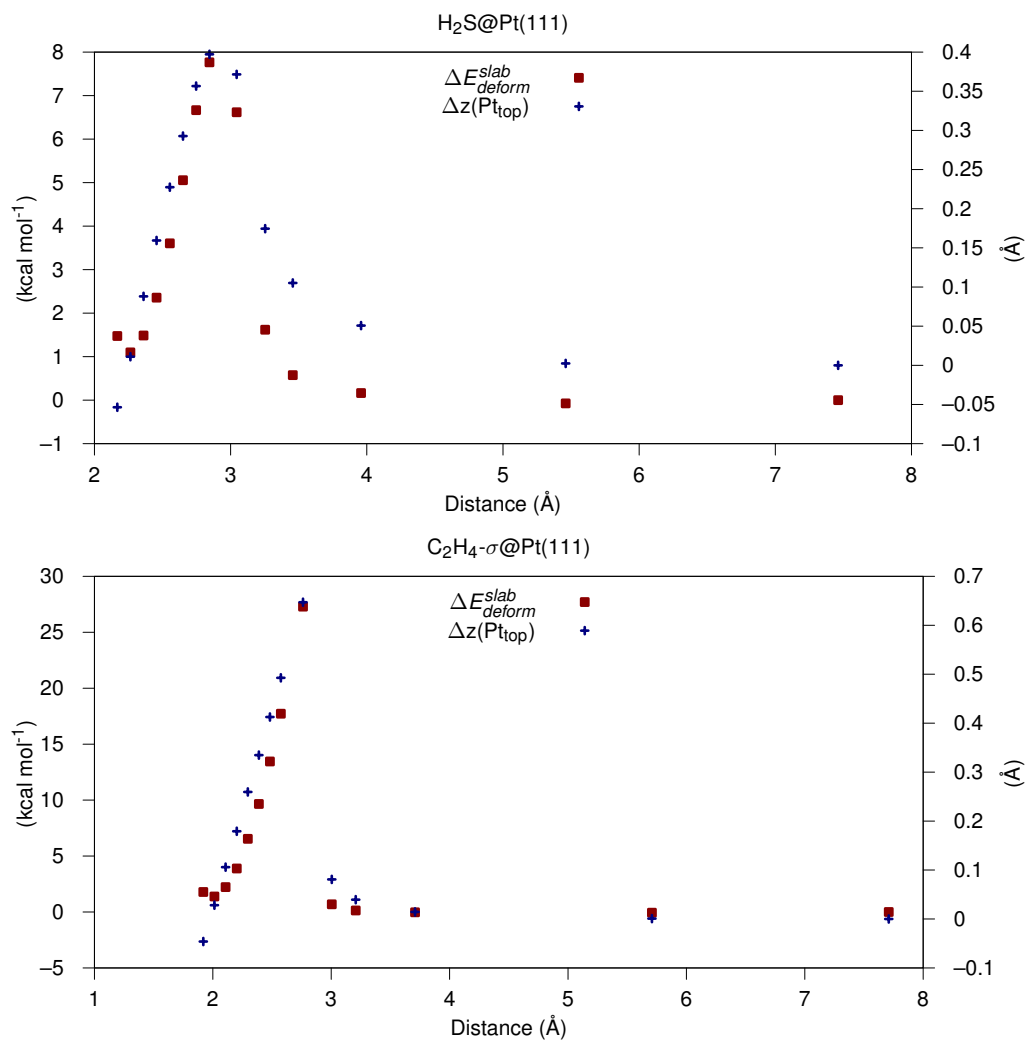


Figure A.2 – Top: H₂S top adsorption on Pt(111). Bottom: Ethylene adsorption in a di-σ adsorption mode. The left y-axis refers to the deformation energy in kcal/mol of the slab (red squares) and the right y-axis to the out-of-plane movement of the Pt atom with respect to the longest distance in Å (blue crosses).

Appendix A. Appendix to Chapter 2

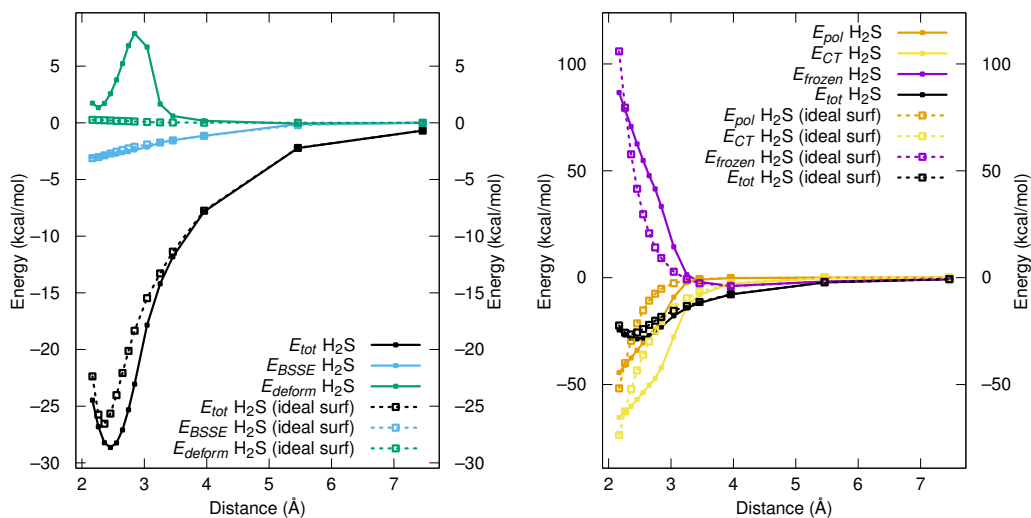


Figure A.3 – H₂S top adsorption on Pt(111). Full lines: Relaxed Pt(111) surface. Broken lines: ideal Pt(111) surface.

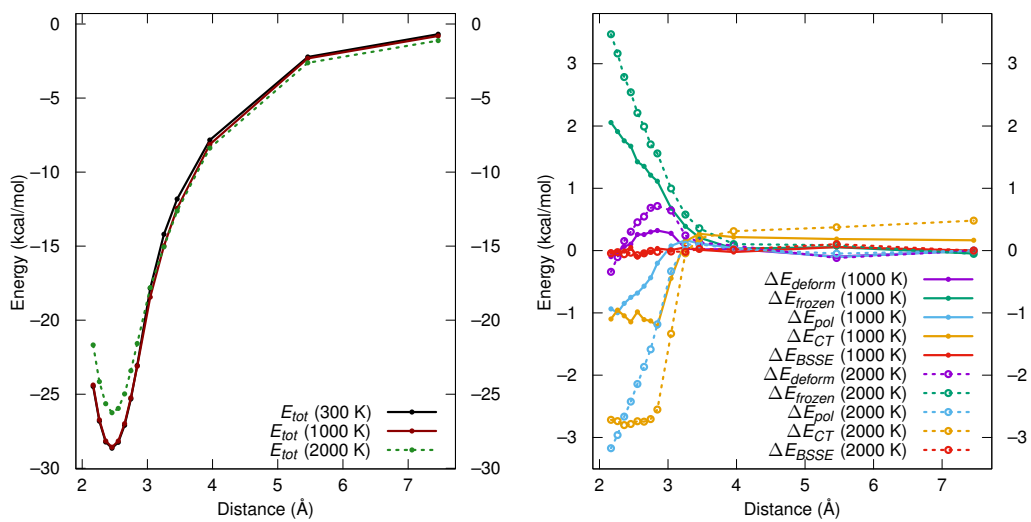


Figure A.4 – Left: Total adsorption energy of H₂S on Pt(111). Right: Difference in interaction energy components at 1000 K and 2000 K compared to 300 K, e.g., CT(300 K) - CT(2000 K)

A.6. Effect of K-point sampling

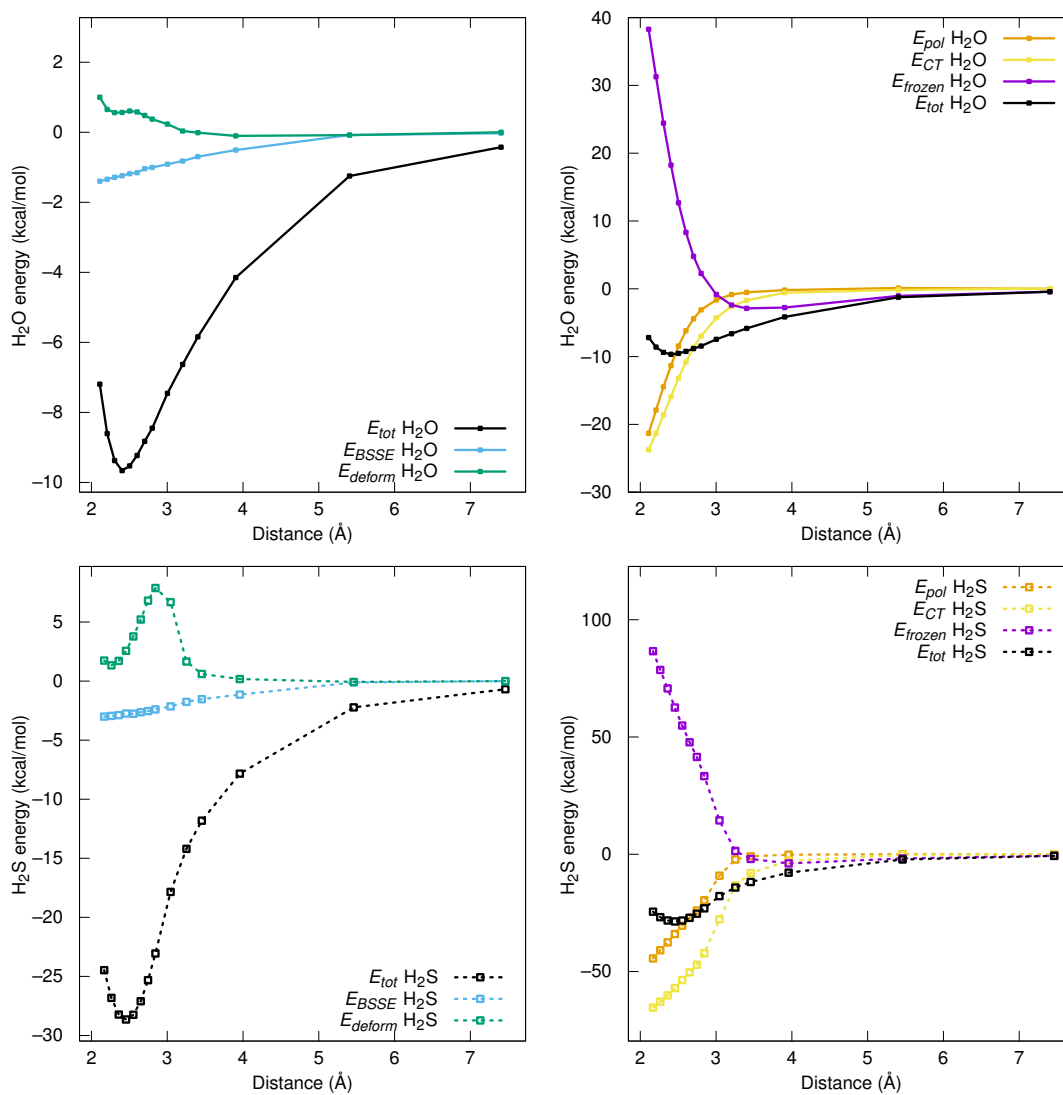


Figure A.5 – Energy decomposition analysis for H₂O (full lines, bottom) and H₂S (broken lines, top). The total interaction energy is compared to the BSSE and the deformation energy on the left, while the right quantifies the frozen monomers, polarization and charge transfer interaction energy. Note that the y-scales are aligned in such a way that the equilibrium interaction energy and the zero interaction energy are aligned for both systems.

Appendix A. Appendix to Chapter 2

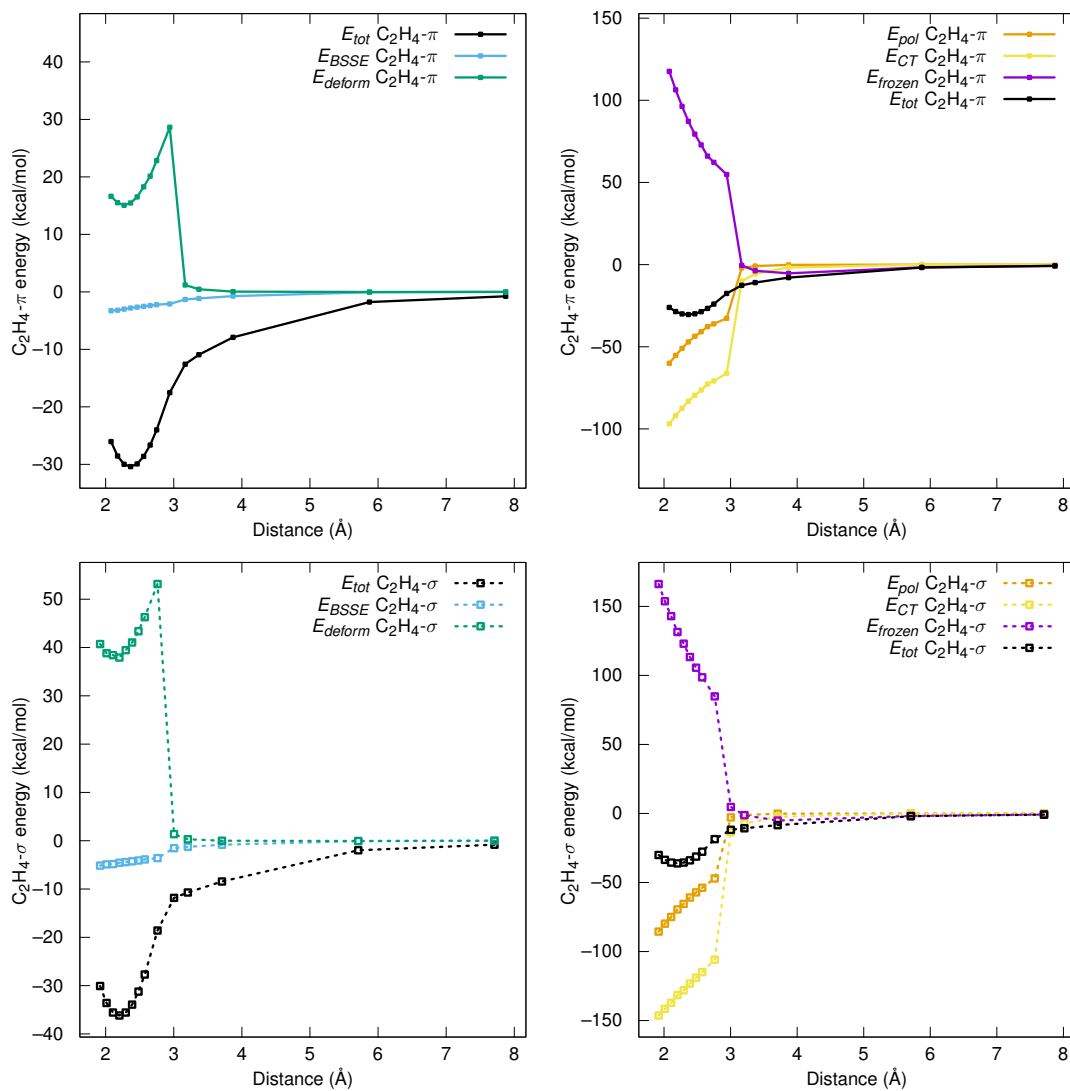


Figure A.6 – Energy decomposition analysis for di- σ (broken lines, bottom) and π (full lines, top) C_2H_4 . The total interaction energy is compared to the BSSE and the deformation energy on the left, while the right quantifies the frozen monomers, polarization and charge transfer interaction energy.

A.6. Effect of K-point sampling

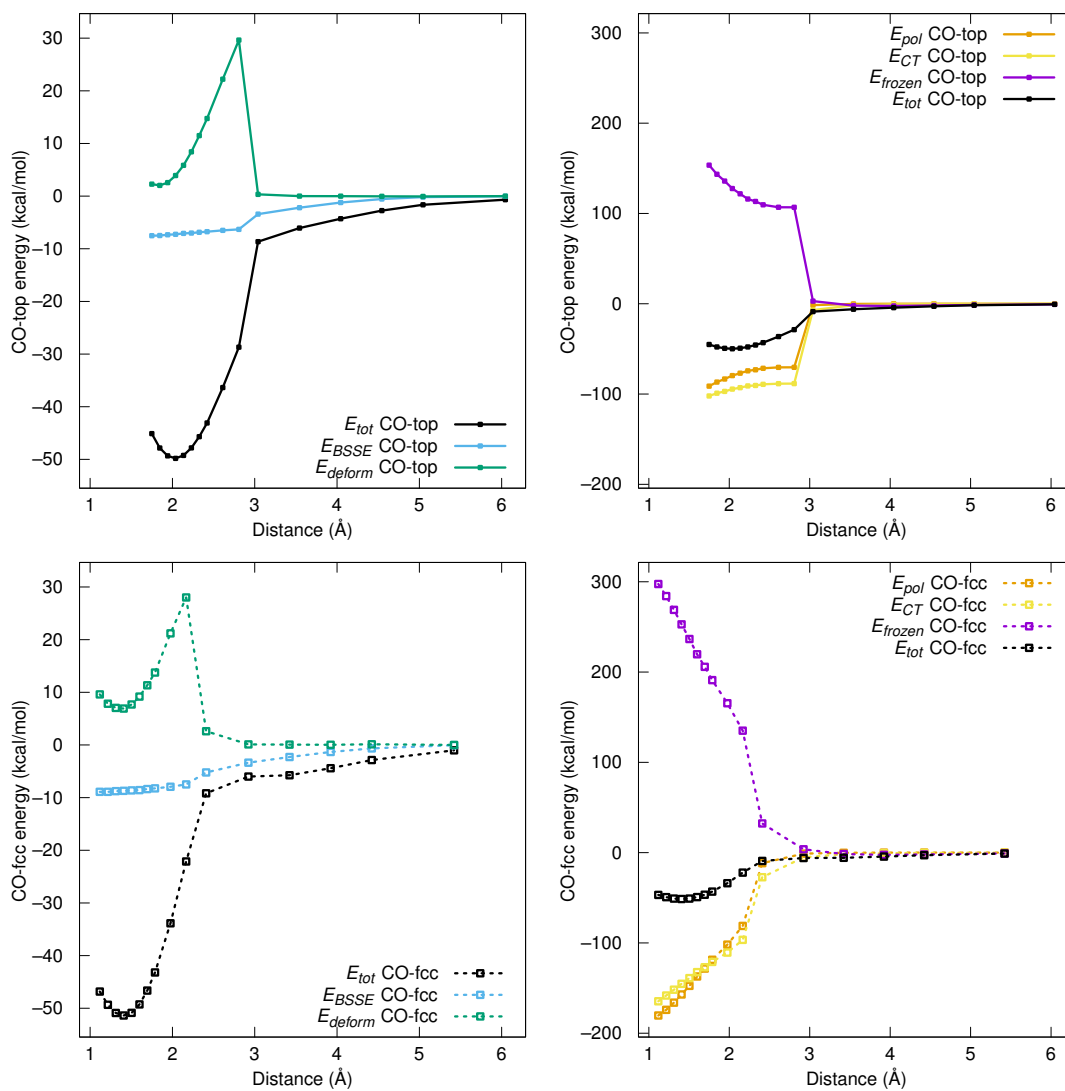


Figure A.7 – Energy decomposition analysis for top (full lines, bottom) and fcc (broken lines, top) CO on Pt(111). The total interaction energy is compared to the BSSE and the deformation energy on the left, while the right quantifies the frozen monomers, polarization and charge transfer interaction energy.

Appendix A. Appendix to Chapter 2

Table A.1 – Effect of K-point sampling on optimum geometry and adsorption energy (ΔE_{tot} in kcal mol⁻¹), as obtained with VASP. $\delta\Delta E_{tot}$ refers to the error compared to the optimized adsorption energy with a $3 \times 3 \times 1$ K-point Monkhorst-Pack K-point mesh and %Error is the corresponding error in percentage.

System	ΔE_{tot}^{Γ}	$\Delta E_{tot}^{3 \times 3 \times 1, SP}$	$\Delta E_{tot}^{3 \times 3 \times 1}$	$\delta\Delta E_{tot}^{\Gamma}$	$\delta\Delta E_{tot}^{3 \times 3 \times 1, SP}$	%Error ^{Γ}	%Error ^{$3 \times 3 \times 1, SP$}
H2O_top	-9.97	-9.64	-10.18	0.21	0.54	-2.1	-5.3
H2S_top	-27.14	-24.13	-24.35	-2.79	0.23	11.4	-0.9
C2H4_sigma	-35.14	-31.96	-31.98	-3.15	0.02	9.9	-0.1
C2H4_pi	-29.52	-26.02	-25.95	-3.57	-0.07	13.7	0.3
C6H6_hcp0	-43.56	-38.00	-37.66	-5.90	-0.35	15.7	0.9
C6H6_bridge	-47.78	-42.67	-42.41	-5.37	-0.26	12.6	0.6
CO_top	-42.40	-38.71	-38.65	-3.75	-0.05	9.7	0.1
CO_hcp	-43.94	-39.58	-39.17	-4.77	-0.40	12.2	1.0
CO_fcc	-45.11	-40.89	-40.59	-4.52	-0.30	11.1	0.7

B Appendix to Chapter 3

B.1 Additional Analysis

B.1.1 Water–Water interaction

The water–water interaction as assessed in the absence of the metal surfaces amounts to -4 (Pt) to -5 (Au) kcal/mol per water molecule, indicating that the chemisorption on Pt(111) imposes stronger constraints on the water geometries than the weaker interaction with Au(111).

In other words, the chemisorption on Pt leads to water layers that adopt less favorable geometries than on Au, with the deformation energy at the single molecule level being negligible. Interestingly, Cu is, in this respect, closer to the behavior of Pt than of Au. This is likely a coincidence, with the lattice mismatch on Cu(111) playing the role of the chemisorption on Pt(111).

The water–water interaction is, in magnitude, dominated by the ΔE_{CT} , which contributes -4.0 kcal/mol per water molecule over Pt(111), close to the -4.2 kcal/mol average total stabilization.

The water–water interaction contributes -10 kcal/mol per water molecule, more than twice as much as on average, to the stability of the H^{up} adlayer. In contrast, the $\sqrt{37}$ and $\sqrt{39}$ structures benefit “only” from a stabilization of about -6 kcal/mol per molecule, once again pointing to the fact that the H^{up} layer is “peculiar” and not a typical case for the interaction of water with metal surfaces.

B.2 Additional Tables and Figures

For chemisorbed acceptors (with a distance to the surface $z < 3.0 \text{ \AA}$), $E_{ads}^{acceptor} = -5.4 \text{ kcal}\cdot\text{mol}^{-1}$ if $\theta < 60^\circ$ and $-3.0 \text{ kcal}\cdot\text{mol}^{-1}$ else.




Figure B.1 – Interaction energies obtained with CP2K plotted against the ones from VASP. The black line corresponds to the bisector.

$\theta(^{\circ})/\phi(^{\circ})$	[0 – 100]	[100 – 120]	[120 – 180]
[80 – 90]	-6.2	-7.5	-7.4
[70 – 80]	-8.2	-5.6	-8.5
[0 – 70]	-9.3	-9.6	-10.4

Table B.1 – $E_{ads}^{acceptor}$ (kcal·mol⁻¹) for acceptor with a distance to the surface $z > 3.0 \text{ \AA}$




Figure B.2 – Surface energies in kcal/mol/Å² of the most stable ice-like water layer as a function of the experimental metal/liquid tension, retrieved from Heinz et al.¹ The black line corresponds to the bisector.




Figure B.3 – The correlation between the surface dipole moment and the change in workfunction $\Delta\Phi$ upon adsorption of ice adlayers.




Figure B.4 – Visualization of the different discussed water oligomers. The cell and slab used were similar to those used for H^{up} and H^{down} water layers on each metal ($3\sqrt{3}\times 3\sqrt{3}$ (111)). (1^{st} part)




Figure B.5 – Visualization of the different discussed water oligomers. The cell and slab used were similar to those used for H^{up} and H^{down} water layers on each metal ($3\sqrt{3}\times 3\sqrt{3}$ (111)). (2^{nd} part)




Figure B.6 – Correlation between DZVP basis set based computation and TZVP computations for ΔE_{pol} and ΔE_{CT} of water layers on Au(111), Cu(111), and Pt(111). The black continuous line corresponds to the fitted trend for polarisation energy and the dashed one to the fitted trend of charge transfer energy.




Figure B.7 – Correlation between ΔE_{int}^{NonAdd} for oligomers on Pt(111) vs Au(111). The black line corresponds to the bisector.




Figure B.8 – Parity plot for the linear regression model based on the H-bond acceptor type for the 27 investigated systems. Eleven variables are used (see Table B.1), leading to a root mean squared error of 1.37 kcal/mol. The black line represent the parity line.

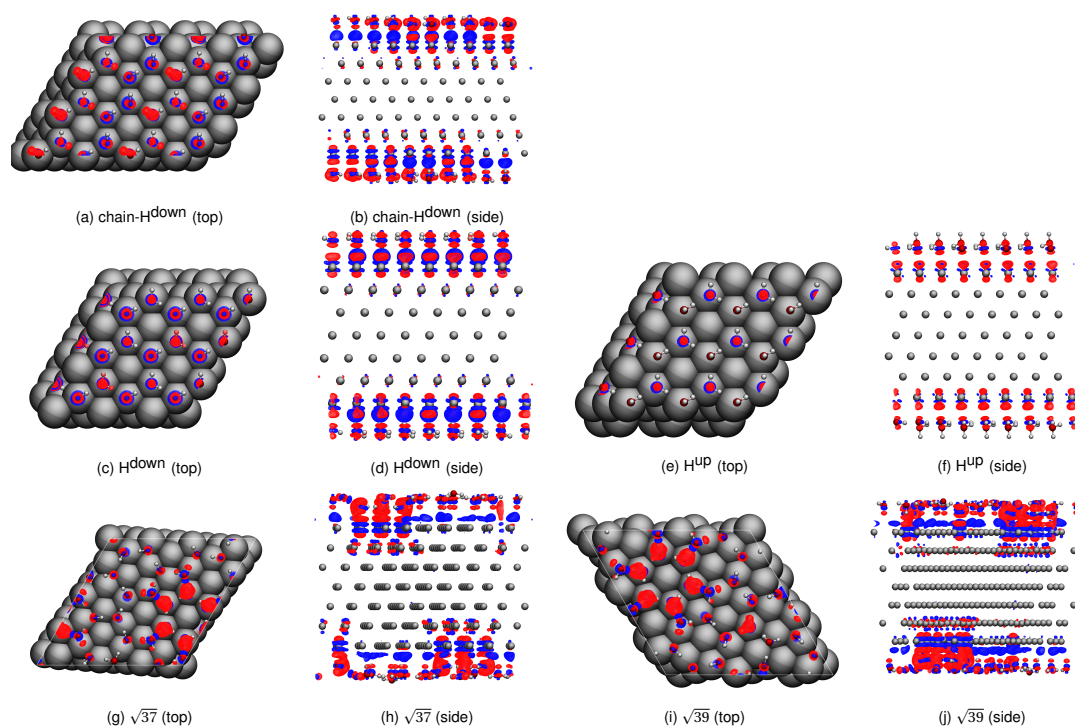


Figure B.9 – Visualization of the spatial distribution of the electronic density from water layers. Isosurfaces are plotted at -0.02 eV (?)

Bibliography

- [1] Heinz, H.; Vaia, R. A.; Farmer, B. L.; Naik, R. R. Accurate Simulation of Surfaces and Interfaces of Face-Centered Cubic Metals Using 12-6 and 9-6 Lennard-Jones Potentials. *J. Phys. Chem. C* **2008**, *112*, 17281–17290.




Figure B.10 – Averaged density difference ($\rho(\text{SCF}) - \rho(\cup\text{surface,adlayer})$) profiles along the out-of-plane unit-cell direction (1^{st} part). Density accumulations are positive, while density depletion with respect to the sum of the densities of the surface and the adlayer is negative. Grey vertical lines indicate the position of the Pt metal layers. Dark red and broken cyan lines indicate the top-most and lowest position of the oxygen and hydrogen atoms, respectively.

Appendix B. Appendix to Chapter 3




Figure B.11 – Averaged density difference ($\rho(\text{SCF}) - \rho(\cup\text{surface,adlayer})$) profiles along the out-of-plane unit-cell direction (2^{nd} part). Density accumulations are positive, while density depletion with respect to the sum of the densities of the surface and the adlayer is negative. Grey vertical lines indicate the position of the Pt metal layers. Dark red and broken cyan lines indicate the top-most and lowest position of the oxygen and hydrogen atoms, respectively.

C Appendix to Chapter 4

C.1 Angular correction term derivation

Here, we are interested in deriving a relation between the solid-angle $\Omega_{\mathcal{C}}$ of a spherical cap \mathcal{C} of radius $R_{\mathcal{C}}$, and its barycentre (more precisely, the distance $G_{\mathcal{C}}$ between its center and its barycentre).

We start from the relation between $G_{\mathcal{C}}$ and its height $h_{\mathcal{C}}$ ¹:

$$G_{\mathcal{C}} = \frac{3(2R_{\mathcal{C}} - h_{\mathcal{C}})^2}{4(3R_{\mathcal{C}} - h_{\mathcal{C}})} \quad (\text{C.1})$$

One can identify a second order polynomial in $h_{\mathcal{C}}$:

$$\begin{aligned} 12G_{\mathcal{C}}R_{\mathcal{C}} - 4G_{\mathcal{C}}h_{\mathcal{C}} &= 12R_{\mathcal{C}}^2 - 12R_{\mathcal{C}}h_{\mathcal{C}} + 3h_{\mathcal{C}}^2 \\ 3h_{\mathcal{C}}^2 + 4h_{\mathcal{C}}(G_{\mathcal{C}} - 3R_{\mathcal{C}}) + 12R_{\mathcal{C}}(R_{\mathcal{C}} - G_{\mathcal{C}}) &= 0 \end{aligned} \quad (\text{C.2})$$

Which solves for $h_{\mathcal{C}}$ as:

$$\begin{aligned} \Delta &= 16(G_{\mathcal{C}} - 3R_{\mathcal{C}})^2 - 12^2 R_{\mathcal{C}}(R_{\mathcal{C}} - G_{\mathcal{C}}) \\ &= 16(G_{\mathcal{C}}^2 - 6G_{\mathcal{C}}R_{\mathcal{C}} + 9R_{\mathcal{C}}^2 - 9R_{\mathcal{C}}^2 + 9R_{\mathcal{C}}G_{\mathcal{C}}) \\ &= 16G_{\mathcal{C}}(G_{\mathcal{C}} + 3R_{\mathcal{C}}) \end{aligned}$$

$$\begin{aligned} h_{\mathcal{C}} &= \frac{4(3R_{\mathcal{C}} - G_{\mathcal{C}}) - \sqrt{16G_{\mathcal{C}}(G_{\mathcal{C}} + 3R_{\mathcal{C}})}}{6} \\ &= \frac{2(3R_{\mathcal{C}} - G_{\mathcal{C}} - \sqrt{G_{\mathcal{C}}(G_{\mathcal{C}} + 3R_{\mathcal{C}})})}{3} \end{aligned} \quad (\text{C.3})$$

the other solution being unphysical (corresponding to $h_{\mathcal{C}} > 2R_{\mathcal{C}}$).

Introducing $\alpha = \frac{G_{\mathcal{C}}}{R_{\mathcal{C}}}$ and $\Omega_{\mathcal{C}} = 2\pi(1 - \frac{R_{\mathcal{C}} - h_{\mathcal{C}}}{R_{\mathcal{C}}}) = 2\pi \frac{h_{\mathcal{C}}}{R_{\mathcal{C}}}$ leads to a simplified final form:

$$\begin{aligned}\Omega_{\mathcal{C}} &= 2\pi \frac{h_{\mathcal{C}}}{R_{\mathcal{C}}} = 4\pi \frac{3R_{\mathcal{C}} - G_{\mathcal{C}} - \sqrt{G_{\mathcal{C}}(G_{\mathcal{C}} + 3R_{\mathcal{C}})}}{3R_{\mathcal{C}}} \\ &= 4\pi \left(1 - \frac{\alpha + \sqrt{\alpha(\alpha + 3)}}{3}\right) \\ &= 4\pi(1 - \gamma_{\mathcal{C}})\end{aligned}\tag{C.4}$$

where

$$\gamma_{\mathcal{C}} = \frac{\alpha + \sqrt{\alpha(\alpha + 3)}}{3}\tag{C.5}$$

is the angular correction term used in the fundamental ASANN inequality.

C.2 ASANN Theorems

The following theorems and properties make use of the previously introduced notations ($r_{i,m}$, $\Omega_{i,j}$, $R_i^{(m)}$, $R_i^{(m')}$, ...), and consider any particle distribution $\{r_{i,1}, \dots, r_{i,n}\}_{n \in \mathbb{N}^*}$ with n large enough.

Theorem 7. *For a given neighborhood distribution, the associated SANN-based coordination number m and the ASANN-based coordination number m' are related by the inequality:*

$$m' \leq m\tag{C.6}$$

Proof. Let us take an integer m° satisfying the definition domain of both the SANN and ASANN coordination radius m and m' (i.e. such that $R_i^{(m^\circ)}$ and $R_i^{(m^\circ)}$ are defined). By definition of m and m' , m° must respect the inequalities:

$$\begin{aligned}3 &\leq m^\circ \leq m \\ 1 &\leq \lfloor 2(1 - \gamma_{\mathcal{C}}) \rfloor + 1 \leq m^\circ \leq m'\end{aligned}\tag{C.7}$$

Such integer cannot be defined only if $m' < 3$, which would trivially leadⁱ to $m' \leq m$.

By definition of the angular correction term $\gamma_{\mathcal{C}}$:

$$0 \leq \gamma_{\mathcal{C}} < 1\tag{C.8}$$

ⁱIndeed, in such case: $m' < 3 \leq m$ by definition of m .

Therefore, using m° , one can write:

$$R_i^{(m^\circ)} = \frac{\sum_{j=1}^{m^\circ} r_{i,j}}{m^\circ - 2 \times (1 - \gamma_{\mathcal{E}})} \leq \frac{\sum_{j=1}^{m^\circ} r_{i,j}}{m^\circ - 2} = R_i^{(m^\circ)} \quad (\text{C.9})$$

By contradiction, let us suppose that $m < m'$. In that case, we can take $m^\circ = m$, and equation (4.3) of the main manuscript can be injected into equation (C.9) to give:

$$R_i^{(m)} \leq R_i^{(m)} < r_{i,m} \quad (\text{C.10})$$

This contradicts the definition of m' (i.e. the lowest acceptable integer such that the ASANN fundamental inequality is respected). Therefore, the SANN-based coordination number m cannot be strictly fewer than the associated ASANN-based coordination number m' , which proves the theorem. \square

This is in agreement with the intuitive idea that ASANN only focus on the relevant region of space with respect to coordination, and therefore less neighbors might cover enough of the ASANN coordination cap (even if they do not fill up the whole coordination sphere).

Theorem 8. *For any particle distribution, the ASANN algorithm is sure to converge to an acceptable solution. In other words, there exist a unique integer m' such that:*

$$r_{i,m'} \leq R_i^{(m')} = \frac{\sum_{j=1}^{m'} r_{i,j}}{m' - 2 \times (1 - \gamma_{\mathcal{E}})} < r_{i,m'+1} \quad (\text{C.11})$$

and this integer is the ASANN-based coordination number.

Proof. Theorem 7 ensures that the ASANN algorithm converge if the SANN algorithm converges, which can be proven². In other words, there exist an integer m' such that:

$$R_i^{(m')} = \frac{\sum_{j=1}^{m'} r_{i,j}}{m' - 2 \times (1 - \gamma_{\mathcal{E}})} < r_{i,m'+1} \quad (\text{C.12})$$

but

$$\forall m^\circ < m', \quad R_i^{(m^\circ)} \text{ is defined} \Leftrightarrow R_i^{(m^\circ)} \geq r_{i,m^\circ+1} \quad (\text{C.13})$$

Using this definition is becomes clear that such m' is unique, since it is the lowest acceptable integer respecting the fundamental ASANN inequality (4.10) of the main manuscript.

Appendix C. Appendix to Chapter 4

One must prove now that such integer m' yield an acceptable coordination number (i.e. the m' -th first neighbors distances are within the coordination radius). In other words, for such integer m' , the left-side of the stated inequality holds:

$$r_{i,m'} \leq R_i^{(m')} = \frac{\sum_{j=1}^{m'} r_{i,j}}{m' - 2 \times (1 - \gamma_{\mathcal{E}})} \quad (\text{C.14})$$

One can distinguish two cases whether $R_i^{(m'-1)}$ is defined or not:

- Let us rule out the case where m' is the minimal acceptable coordination number (i.e. $\lfloor 2(1 - \gamma_{\mathcal{E}}) \rfloor + 1 = m'$) which is handled by the definition of the lower bound $\lfloor 2(1 - \gamma_{\mathcal{E}}) \rfloor + 1$:

In such a case, if $R_i^{(m')} < r_{i,m'}$, then the m' -th nearest neighbor contributes to a negative value for the total solid angle, which means that the first $m' - 1$ neighbors contributions are more than enough to reach the total solid angle of $\Omega_{\mathcal{E}} = 4\pi(1 - \gamma_{\mathcal{E}})$. If $m' - 1 = 0$, this is absurd. Otherwise, this can be written as:

$$\sum_{j=1}^{m'-1} \Omega_{i,j} > \Omega_{\mathcal{E}} = 4\pi(1 - \gamma_{\mathcal{E}}) \quad (\text{C.15})$$

However, each single contribution cannot exceed 2π due to basic geometrical considerations. Therefore, one can write:

$$4\pi(1 - \gamma_{\mathcal{E}}) = \Omega_{\mathcal{E}} < \sum_{j=1}^{m'-1} \Omega_{i,j} = \sum_{j=1}^{\lfloor 2(1 - \gamma_{\mathcal{E}}) \rfloor} \Omega_{i,j} \leq 2\pi \times \lfloor 2(1 - \gamma_{\mathcal{E}}) \rfloor \quad (\text{C.16})$$

Which is absurd by definition of the floor operator $\lfloor \cdot \rfloor$, and therefore proves that $r_{i,m'} \leq R_i^{(m')}$ in such a case.

- If $R_i^{(m'-1)}$ is defined, by definition of m' , $R_i^{(m'-1)}$ do not respect the ASANN inequality:

$$R_i^{(m'-1)} = \frac{\sum_{j=1}^{m'-1} r_{i,j}}{m' - 1 - 2 \times (1 - \gamma_{\mathcal{E}})} \geq r_{i,m'} \quad (\text{C.17})$$

C.3. Discussion regarding a continuous version

Yet, by definition of $R_i^{(m')}$, and using the notation $\beta = 2 \times (1 - \gamma_{\mathcal{E}}) > 0$, one can write:

$$\begin{aligned}
 R_i^{(m')} &= \frac{\sum_{j=1}^{m'} r_{i,j}}{m' - \beta} = \frac{r_{i,m'} + \sum_{j=1}^{m'-1} r_{i,j}}{m' - \beta} \\
 &= \frac{r_{i,m'}}{m' - \beta} + R_i^{(m'-1)} \frac{m' - \beta - 1}{m' - \beta} \\
 &\geq \frac{r_{i,m'}}{m' - \beta} + r_{i,m'} \frac{m' - \beta - 1}{m' - \beta} \\
 &\geq r_{i,m'}
 \end{aligned} \tag{C.18}$$

In both cases, the inequality $r_{i,m'} \leq R_i^{(m')}$ holds, which finishes to prove the theorem. \square

C.3 Discussion regarding a continuous version

Coordination numbers act as topological descriptors (either local, or more global³), associated with the idea of bonding. As a consequence, they can be used as relevant variables (collective variables, reaction coordinates, ...).

A major application of these variables is in force fields⁴. However, force field parameters should be at least continuously differentiable with respect to the degrees of freedom involved (i.e. here, the atoms coordinates). Therefore, coordination numbers must be smoothed in order to be used in force fields.

Unfortunately, we could not find a satisfactory continuously differentiable version of the ASANN algorithmⁱⁱ.

In fact, one can ask about the relevance of such smoothed coordination numbers. To do so, we enumerate some desirable properties that a coordination number evaluator could verify, and check if these are compatible with the continuously differentiable condition.

Let us consider a local purely distance-based coordination number evaluator CN_i , providing a coordination number for an element V_i based on the distances d_{ij} between V_i and all the n others elements V_j . Mathematically, this reads as:

$$\begin{aligned}
 &CN_i: (\mathbb{R}^+)^n \rightarrow \mathbb{R}^+ \\
 &\mathbf{d} = d_{i1}, d_{i2}, \dots, d_{in} \mapsto \text{CN of } V_i
 \end{aligned}$$

ⁱⁱIn particular, the canonical rescaling of each neighbors contribution by $\frac{SA}{SA_{max}}$ (where SA is the solid angle of the corresponding neighbor, and SA_{max} is the solid angle of the nearest neighbor) leads to non continuously differentiable coordination numbers...

Appendix C. Appendix to Chapter 4

Supremum bound A fundamental property of CN is that if there $n + 1$ elements in total, Then the coordination number of V_i cannot be more than n :

$$\forall \mathbf{d} \in (\mathbb{R}^+)^n, \quad CN_i(\mathbf{d}) \leq n \quad (\text{C.19})$$

Ideal case There must exist an ideal configuration in which V_i is coordinated to all the n other elements (if this is only true for low enough n , then we consider such n for the rest of this proof). Such configuration (associated to $\mathbf{d} = \mathbf{d}_{\text{ideal}} = d_{i1}^\circ, d_{i2}^\circ, \dots, d_{in}^\circ$) should be associated with a coordination number of n for V_i :

$$\exists \mathbf{d}_{\text{ideal}} \in (\mathbb{R}^+)^n, \quad CN_i(\mathbf{d}_{\text{ideal}}) = n \quad (\text{C.20})$$

Sensitivity to small displacements The coordination number evaluator should not be less sensitive to small displacements near an ideal case, than far from it. In particular, a small displacement of an element V_j situated infinitely far from V_i should have less impact on CN_i than if V_j was closer to V_i . More generally, one can ask that the further V_j is from the ideal case, the fewer CN_i should be impacted by a small displacement of V_j .

This is coherent with the discrete version, where the most critical region for the coordination number (i.e. the frontier between coordinated/non-coordinated) is by definition located near an ideal case (i.e. a fully coordinated case). Not respecting this property means shifting the frontier to a partially coordinated region, and therefore, focusing on partial coordination instead of full coordination. This would require additional insights on the system studied, which is poorly compatible with a parameter-free method.

Let CN_i° be defined as:

$$\begin{aligned} CN_i^\circ: \mathbb{R}^+ &\rightarrow \mathbb{R}^+ \\ x &\mapsto CN_i(d_{i1}^\circ + x, d_{i2}^\circ, \dots, d_{in}^\circ) \end{aligned}$$

This property can then be written, in terms of CN_i° :

$$\forall \epsilon_1, \epsilon_2 \geq 0, \quad CN_i^\circ(0) - CN_i^\circ(\epsilon_2) \geq CN_i^\circ(\epsilon_1) - CN_i^\circ(\epsilon_1 + \epsilon_2) \quad (\text{C.21})$$

Definition: A coordination number evaluator satisfying the three properties above (i.e. (C.19), (C.20) and (C.21)) will be denoted in the following as *topologically coherent*.

Theorem 9. Any local purely distance-based coordination number evaluator is either non topologically coherent or non continuously differentiable. In other words, there cannot exist a

C.5. Model Hamiltonian parameters

Algorithm	Γ_{Cu}^{NP}	Γ_{Au}^{NP}	$\Gamma_{Cu,Cu}^{NP}$	$\Gamma_{Au,Cu}^{NP}$	$\Gamma_{Au,Au}^{NP}$
SANN	-1.97	-2.62	-0.24	-0.14	0.003
ASANN	-2.26	-2.62	-0.18	-0.12	0.003

Table C.1 – Model Hamiltonian parameters, using only nanoparticles in the training set (in eV).

local purely distance-based coordination number evaluator that is both topologically coherent and continuously differentiable.

Proof. Applying (C.20) to CN_i° , we have:

$$CN_i^\circ(0) = n \tag{C.22}$$

Therefore, according to (C.19), it comes that CN_i° admits a maximum at 0.

If CN_i is continuously differentiable, then CN_i° is also continuously differentiable:

$$CN_i \in C^1 \Rightarrow CN_i^\circ \in C^1 \tag{C.23}$$

Then, condition (C.21) can be rewritten:

$$\forall a, b \geq 0, \quad a \leq b \Rightarrow \frac{dCN_i^\circ}{dx}(a) \leq \frac{dCN_i^\circ}{dx}(b) \tag{C.24}$$

However, since CN_i° admits a maximum at 0, we have:

$$CN_i \in C^1 \Rightarrow \frac{dCN_i^\circ}{dx}(0) = 0 \tag{C.25}$$

Combining (C.25) and (C.24) gives that CN_i° is increasing on \mathbb{R}^+ . Applying (C.19) to this result, finally gives:

$$CN_i \in C^1 \Rightarrow \forall x \geq 0, CN_i^\circ(x) = n \tag{C.26}$$

It comes finally that such a coordination number evaluator would consider that an infinitely far neighbor is still fully bonded. Which is absurd, and therefore proves our theorem. \square

Appendix C. Appendix to Chapter 4




Figure C.1 – Parity plot for the total energy predicted by a 2-body model Hamiltonian fitted using the SANN vs. ASANN coordination algorithm. The set of 27 Au-Cu bimetallic nanoparticles was used for training and test set.

Algorithm	Γ_{Cu}^{all}	Γ_{Au}^{all}	$\Gamma_{Cu,Cu}^{all}$	$\Gamma_{Au,Cu}^{all}$	$\Gamma_{Au,Au}^{all}$
SANN	-1.38	-1.82	-0.39	-0.29	-0.23
ASANN	-1.68	-1.91	-0.34	-0.27	-0.22

Table C.2 – Model Hamiltonian parameters, using nanoparticles and bulks in the training set (in eV).

Algorithm	Co(100)	Co(110)	Co(111)	Co(211)
ASANN	6.67	5.83	7.5	8.75/7.33/5.5
SANN	6.67	5.83	7.5	8.75/7.33/5.5
Reference	$\frac{20}{3}$	$\frac{35}{6}$	$\frac{30}{4}$	$\frac{35}{4} / \frac{22}{3} / \frac{11}{2}$

Algorithm	Co(0001)	Co(10-10)	Co(10-11)	Co(10-12)	Co(11-20)
ASANN	7.5	8.33/6.67	7.5/6.58	9.17/6.75/5.67	9.17/5.83
SANN	7.5	8.33/6.67	7.5/6.58	9.17/6.75/5.67	9.83/8.17
Reference	$\frac{15}{2}$	$\frac{25}{3} / \frac{20}{3}$	$\frac{15}{2} / \frac{79}{12}$	$\frac{55}{6} / \frac{27}{4} / \frac{17}{3}$	$\frac{55}{6} / \frac{35}{6}$

Table C.3 – Generalized coordination numbers found by ASANN and SANN for various Cobalt surfaces.

C.4 Additional Figures

C.5 Model Hamiltonian parameters

The distance between the two set of parameters is:

$$\begin{aligned} \sqrt{\sum_{\chi} (\Gamma_{\chi}^{all} - \Gamma_{\chi}^{NP})^2} &= 1.04 \text{ eV}^2 \text{ for SANN} \\ &= 0.97 \text{ eV}^2 \text{ for ASANN} \end{aligned} \quad (C.27)$$

C.6 Generalized Coordination Numbers

In addition to the determination of conventional coordination numbers, we also applied our algorithm to the determination of generalized coordination numbers³, as this measure is believed to provide with a more realistic coordination picture, which is especially useful for adsorption properties prediction. Generalized coordination numbers are a first-neighbors extended generalisation of coordination defined by the formula:

$$GCN(i) = \overline{CN}(i) = \sum_{j=1}^{CN(i)} \frac{CN(j)}{CN_{max}} \quad (C.28)$$

where $CN(i)$ is the conventional coordination number of i , j represents a nearest neighbor of i , and CN_{max} is the bulk coordination number. Our Python 3 implementation of ASANN and SANN includes an optional generalized coordination number feature. The relevant generalized coordination numbers computed for the top sites of highlighted surfaces are presented in table C.3.

Appendix C. Appendix to Chapter 4

Atom type	ASANN (ref) CNs
Cs	8 (8)
Cl	8 (8)

Table C.4 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of CsCl (mp-22865).

Atom type	ASANN (ref) CNs
K	6 (6)
F	6 (6)

Table C.5 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of KF (mp-463)

C.7 Ionic crystals

Atom type	ASANN (ref) CNs
Na	6 (6)
Cl	6 (6)

Table C.6 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of NaCl (mp-22862)

C.7. Ionic crystals

Atom type	ASANN (ref) CNs
Cd	6 (6)
I	3 (3)

Table C.7 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of CdI₂ (mp-567259)

Atom type	ASANN (ref) CNs
Ca	7,9 (3)
Si(II)	4 (2)
Si(III)	4 (3)
O(I)	5-7 (1)
O(II)	7 (2)

Table C.8 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of CaSiO₃ (mp-4429)

Atom type	ASANN (ref) CNs
Zn	16 (4)
S	16 (4)

Table C.9 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of ZnS (mp-10695)

Atom type	ASANN (ref) CNs
Ti	6 (1)
Pb	8 (4)
O(I)	6 (1)
O(II)	6 (2)

Table C.10 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of TiPbO₃ (mp-20459)

Atom type	ASANN (ref) CNs
Ca	9 (3)
Ti	6 (6)
O	6 (2)

Table C.11 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of CaTiO₃ (mp-4019)

Atom type	ASANN (ref) CNs
Zn	5 (4)
S	5 (4)

Table C.12 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of ZnS (mp-560588)

Appendix C. Appendix to Chapter 4

Atom type	ASANN (ref) CNs
Ni	6 (3)
S	5 (3)

Table C.13 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of NiS (mp-1547)

Atom type	ASANN (ref) CNs
Ca	8 (8)
F	16 (4)

Table C.14 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of CaF₂ (mp-2741)

C.8 Molecules

Molecular structures are governed by atomic orbitals displaying a large higher-order anisotropy, compared to their first-order term (e.g. an atomic p-orbital has a strong anisotropy, yet its barycentre is on the nucleus). The ASANN description of such systems is therefore poorly improved over SANN, since ASANN only considers first-order anisotropy terms.

Atom type	ASANN (ref) CNs
Ti	6 (6)
O	7 (3)

Table C.15 – Coordination numbers found by ASANN (compared with reference) for crystallographic structure of TiO₂ (mp-2657).




Figure C.2 – Representation of a single EtOH molecule, with ASANN coordination numbers superimposed.




Figure C.3 – Representation of a single thiophene molecule, with ASANN coordination numbers superimposed.




Figure C.4 – Representation of a single benzene molecule, with ASANN coordination numbers superimposed.




Figure C.5 – Representation of a single graphene layer, with ASANN coordination numbers superimposed.

C.9 ROC data

Bibliography

- [1] Harris, J. W.; Stöcker, H. *Handbook of Mathematics and Computational Science*; Springer-Verlag: New York, 1998; p 107.
- [2] van Meel, J. A.; Filion, L.; Valeriani, C.; Frenkel, D. A parameter-free, solid-angle based, nearest-neighbor algorithm. *The Journal of Chemical Physics* **2012**, *136*, 234107.
- [3] Calle-Vallejo, E; Martínez, J. I.; García-Lastra, J. M.; Sautet, P; Loffreda, D. Fast Prediction of Adsorption Properties for Platinum Nanocatalysts with Generalized Coordination Numbers. *Angewandte Chemie International Edition* **2014**, *53*, 8316–8319.
- [4] Sprik, M. Coordination numbers as reaction coordinates in constrained molecular dynamics. *Faraday Discuss.* **1998**, *110*, 437–445.

Table C.16 – Raw data for classifying Cs ions as close enough to the surface (threshold 4.7 Å), based on their coordination numbers. Example of reading: 86% of Cs ions closer than 4.7 Å from the surface have ASANN coordination number below 9, while only 20% of the ions further from the surface have ASANN CNs below 9.

CN threshold	TPR	FPR	CN threshold	TPR	FPR	CN threshold	TPR	FPR
0	0%	0%	0	0%	0%	0	0%	0%
1	0%	0%	1	0%	0%	1	0%	0%
2	0%	0%	2	0%	0%	2	0%	0%
3	0%	0%	3	0%	0%	3	0%	0%
4	0%	0%	4	0%	0%	4	0%	0%
5	0%	0%	5	0%	0%	5	0%	0%
6	2%	0%	6	0%	0%	6	6%	0%
7	19%	0%	7	0%	0%	7	31%	2%
8	56%	3%	8	5%	0%	8	65%	11%
9	86%	20%	9	31%	3%	9	89%	35%
10	97%	53%	10	68%	19%	10	98%	66%
11	100%	82%	11	91%	51%	11	100%	89%
12	100%	96%	12	98%	80%	12	100%	98%
13	100%	99%	13	100%	95%	13	100%	100%
14	100%	100%	14	100%	99%	14	100%	100%
			15	100%	100%			

(a) ASANN algorithm.

(b) SANN algorithm.

(c) Fixed cut-off (4.0 Å).



D Appendix to Chapter 6

D.1 Linear regression models

The classical linear regression model $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$ relies on multiple hypothesis and assumptions. Yet, ordinary least squares estimators are commonly applied to the cluster expansion model. Here, we discuss whether the linear regression model hypotheses are indeed satisfied by the cluster expansion model.

D.1.1 Common assumptions

Linearity

Hypothesis: The model assumes a linear relationship between the regressand \mathbf{y} and the regressors \mathbf{X} .

Analysis: This is exactly the hypothesis behind the cluster expansion model. Therefore, this assumption is verified as long as the cluster expansion model is valid.

Full rank

Hypothesis: There should not be any collinearity within the regressors (i.e. there should be no exact linear relationship among any of the independent variables in the model).

Analysis: Using the rank-Greville recursive least squares scheme, we ensure that our sub-system is purely overdetermined (i.e. full column rank).

Exogeneity

Hypothesis: The residuals have zero conditional mean: $E[\epsilon|\mathbf{X}] = 0$. This assumption implies $E[\mathbf{y}|\mathbf{X}] = \mathbf{X}\beta$.

Analysis (non-truncated cluster expansion): In the full cluster expansion method, the residual is only due to the precision of the underlying DFT formalism. One can satisfy this assumption by considering a cluster expansion model for the average DFT energy, instead of the exact energy. The validity of the cluster expansion model for average DFT energy is a crucial hypothesis for the justification of using the DFT formalism. In our case, we must assume this hypothesis as higher quality computations are not reasonably available.

Analysis (truncated cluster expansion): However, in the truncated cluster expansion model, the residuals are coming from the DFT precision and the higher-order terms. Clearly, the higher-order terms are correlated with lower-order terms (imagine there is no 1-body water term, then higher-order terms involving water will be zero as well), meaning that the residuals are correlated with the regressors, which could indicate that a constant conditional mean is unlikely. However, this zero conditional mean assumption on the residuals is in practice not really constraining. Indeed, if the conditional average contribution of higher-order terms $E[\delta y_i | \mathbf{x}_i]$ associated with regressor \mathbf{x}_i is linear in \mathbf{x}_i , then those terms contribution is to be incorporated into the model's coefficients, providing with zero conditional mean residuals.

In the cluster expansion formalism, the (N+1)-order terms occurrences \mathbf{x}_i^{N+1} have convenient bounds. The lower bound is simply $\min_{\mathbf{x}_i^{n \leq N}}(\mathbf{x}_i^{N+1}) = \mathbf{0}$, since it is always possible to find a configuration such that the cluster expansion has negligible terms up to order N (e.g. considering a N-body configuration with low coverage). The higher bound is in practice linear with respect to the lower-order terms $\max_{\mathbf{x}_i^{n \leq N}}(\mathbf{x}_i^{N+1}) = \alpha \mathbf{x}_i^N$ where α contains geometry-constrained data.

Indeed, this bound is usually achieved at highest coverage, and the linear relationship arise from the maximum coordination number achievable. A justification of this can be found in the following: the occurrences in \mathbf{x}_i^{N+1} occurs from combinatorics, with maximum bound in $\binom{nb_e}{N+1} = O(nb_e^{N+1})$ for large enough number of entities nb_e in the system. However, due to geometric constrains and lattice-related considerations, each N-body pattern can only accept a fixed number of acceptable neighbouring entities, hence a linear maximal number of (N+1)-body terms.

As a consequence, the minimum and maximum contribution bounds are compatible with zero conditional mean of the residuals. Yet, we are interested in the average higher-order contributions, which can therefore be linear in the regressors. This is, at the end, dependant of the sampling method, and as such, hard to predict. We will assume that our sampling method

is compatible with the exogeneity assumption in the case of truncated cluster expansion, since it should focus on configurations where apparently strongly stabilizing or destabilizing terms are maximized, whereas apparently negligible terms are minimized.

Homoscedasticity and nonautocorrelation

Hypothesis Each residual ϵ_i has the same finite variance σ_i , and is uncorrelated with every other residual ϵ_j .

Analysis (non-truncated cluster expansion): In the full cluster expansion method, the residuals are only due to the precision of the underlying DFT formalism. This assumes that the DFT formalism does not describe some structures better than others, and does not consistently favour or disfavour some configurations (i.e. DFT-based errors are not correlated between each other). Therefore, this assumption should be satisfied as long as the DFT computations are performed at the same level of precision, and assuming DFT is similarly valid for all considered configurations.

Analysis (truncated cluster expansion): However, in the truncated cluster expansion model, the residuals are coming from the DFT precision and the higher-order terms. Clearly, the higher-order terms contribution range differs between lower-order grouped configurations. Therefore, hidden higher-order terms must provide with heteroscedastic residuals. Besides, those higher-order terms are likely to be shared among multiple configurations, leading to autocorrelation.

D.1.2 Ordinary least squares

If we consider a non-truncated cluster expansion model, the ordinary least square solution $\hat{\boldsymbol{\beta}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$ is unbiased, consistent and efficient (according to Gauss-Markov theorem).

Bias : The estimator $\hat{\boldsymbol{\beta}}$ is unbiasedⁱ, meaning that $E[\hat{\boldsymbol{\beta}}] = \boldsymbol{\beta}$.

Consistency : Assuming that $\frac{\mathbf{X}^T \mathbf{X}}{n}$ converges in probability to a positive definite matrix \mathbf{Q} (i.e. $\frac{\mathbf{X}^T \mathbf{X}}{n} \xrightarrow{p} \mathbf{Q}$), then the estimator $\hat{\boldsymbol{\beta}}$ is consistentⁱⁱ, meaning that it converges in probability to the true parameters of the model (i.e. $\hat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$).

ⁱSection 4.3.1 of Econometric Analysis (8th edition)

ⁱⁱSection 4.4.1 of Econometric Analysis (8th edition)

Efficiency According to the Gauss-Markov theorem, the estimator $\hat{\boldsymbol{\beta}}$ is efficientⁱⁱⁱ, meaning that its covariance matrix $\text{Var}(\hat{\boldsymbol{\beta}}) = \sigma^2(\mathbf{X}^T\mathbf{X})^{-1}$ (where σ^2 is the residual variance, and can be asymptotically estimated^{iv} by $s^2 = \frac{err^T \cdot err}{n-K}$) is the lowest among the unbiased estimators of $\boldsymbol{\beta}$.

D.1.3 Generalized least squares

However, the ordinary least squares estimator is no longer (the best estimator) adapted^v for describing the truncated cluster expansion model. Indeed, the ordinary least square is still unbiased^{vi} and consistent^{vii}, but no longer efficient^{viii}. In presence of heteroscedasticity and autocorrelation (i.e. if $E[\boldsymbol{\epsilon}^T\boldsymbol{\epsilon}] = \sigma^2\boldsymbol{\Omega} \neq \sigma^2\mathbf{I}$), the generalized least squares $\hat{\boldsymbol{\beta}} = (\mathbf{X}^T\boldsymbol{\Omega}^{-1}\mathbf{X})^{-1}\mathbf{X}^T\boldsymbol{\Omega}^{-1}\mathbf{y}$ is unbiased, consistent and efficient^{ix} with $\text{Var}(\hat{\boldsymbol{\beta}}|\mathbf{X}) = \sigma^2(\mathbf{X}^T\boldsymbol{\Omega}^{-1}\mathbf{X})^{-1}$.

D.1.4 Conclusion

An ordinary least squares is indeed well-adapted for the non-truncated expansion model, since we can assume that all the hypotheses of the linear regressions model are satisfied.

For the truncated version however, a generalized least squares would be required, since the homoscedasticity and nonautocorrelation hypothesis is violated by the presence of higher-order terms. Nevertheless, if we assume that the discarded higher-order terms are indeed negligible, then an ordinary least squares could be applied rigorously.

Consequently, the ordinary least squares estimator is adapted for describing the cluster expansion model, as long as the eventual truncation is performed properly, discarding only effectively negligible contributions.

D.2 Score range for UCB

For evaluating the score of an arm, most UCB algorithms only need the range of its associated rewards¹. Some UCB variants also use the variance (or its estimation via the empirical variance) of the rewards^{2;3}. It is expected that “variance-aware” UCB variants will perform better than UCB-like algorithms that are not aware of the reward variance³.

Unfortunately, the energy range per adsorbate in a subtree is not known beforehand. Therefore,

ⁱⁱⁱSection 4.3.5 of Econometric Analysis (8th edition)

^{iv}Section 4.3.4 and 4.4.2 of Econometric Analysis (8th edition)

^vSection 9.3 of Econometric Analysis (8th edition)

^{vi}Section 4.3.1 of Econometric Analysis (8th edition)

^{vii}Section 4.5.1 of Econometric Analysis (8th edition)

^{viii}Section 4.5.2 of Econometric Analysis (8th edition)

^{ix}Section 9.4.1 of Econometric Analysis (8th edition)

we consider in section D.2.1 an approximation using the empirical range.

In the case of the predicted energy, two major issues occur:

- Since the model parameters x will certainly change during the run, one should not store the predicted energies $E_{pred} = \Gamma^\top x_{model}$, but the observed configurations Γ . Storing all seen configurations for each arm (and sub-arm, i.e. node) would probably require too much space and recomputing the all the predicted energies seen would also be computationally expensive. An estimator of the model-dependent minimum and maximum predicted energy is introduced in section D.2.1. This estimator requires only the storage of two configurations per node.
- The linear model comes with uncertainties that should be taken into account. Indeed, the variance of the energy prediction associated with Γ is $\text{Var}(E_{pred}) = \text{Var}(\Gamma^\top x_{model} + err) = \Gamma^\top \text{Var}(x_{model})\Gamma + \text{Var}(err)$. The implication of such uncertainties is considered in sections D.2.2 and D.2.4.

D.2.1 Range without uncertainty

Sample minimum and maximum

In the case of real DFT energies per adsorbate, storing and maintaining the minimum and maximum energy observed for each arm is straightforward.

However, in the case of predicted energies $E_{pred} = \Gamma^\top x$. Maintaining a model-dependent minimum/maximum would require to store all observed configurations and compute on-the-fly:

$$\min_{exact}(E_{pred}) = \min_{\Gamma^{seen}} \Gamma^\top x \quad (\text{D.1})$$

This would quickly require too much storage and, in a lesser extent, would be slightly computationally expensive.

An approximation can be found in the natural bound:

$$\begin{aligned} \min_{approx}(E_{pred}) &= \sum_i x_i \times \begin{cases} \min_{\Gamma^{seen}} \Gamma_i, & \text{if } x_i \geq 0 \\ \max_{\Gamma^{seen}} \Gamma_i, & \text{else} \end{cases} \\ &= x \cdot \min_{\Gamma} / \max_{\Gamma}(\Gamma) \end{aligned} \quad (\text{D.2})$$

where we only need to store the minimum/maximum of each components of Γ seen (simply called \min_{Γ} and \max_{Γ}), which is straightforward to maintain and only represents the storage

of $2m$ integers^x.

Similar equations can be derived for the estimated sample maximum. These approximations are particularly good when the components of Γ are independent. Figure D.1 is an illustration of this approximation.




Figure D.1 – The min/max bound corresponds to the minimum/maximum over all configurations within the blue box (independently of those effectively seen in this box), whose frontiers are the empirical min/max for each component.

Estimated range

Let X_1, X_2, \dots, X_n be a sample from the uniform distribution on $[a, b]$. Estimating the range $b - a$ from the sample range $\max_{i=1, \dots, n} X_i - \min_{i=1, \dots, n} X_i$ is related to the German tank problem, and can be solved using the relation:

$$\mathbb{E}_{\Gamma \in \mathbb{N}^m} \left(\max_{i=1, \dots, n} X_i - \min_{i=1, \dots, n} X_i \right) = (b - a) \frac{n - 1}{n + 1} \tag{D.3}$$

Proof.

$$\begin{aligned}
 P\left(\max_{i=1,\dots,n} X_i < x\right) &= \left(\frac{x-a}{b-a}\right)^n \\
 \mathbb{E}\left(\max_{i=1,\dots,n} X_i\right) &= \int_a^b xn \frac{(x-a)^{n-1}}{(b-a)^n} dx \\
 &= \frac{a+bn}{n+1} \\
 P\left(\min_{i=1,\dots,n} X_i < x\right) &= 1 - \left(\frac{b-x}{b-a}\right)^n \\
 \mathbb{E}\left(\min_{i=1,\dots,n} X_i\right) &= \int_a^b xn \frac{(b-x)^{n-1}}{(b-a)^n} dx \\
 &= \frac{an+b}{n+1} \\
 \mathbb{E}\left(\max_{i=1,\dots,n} X_i - \min_{i=1,\dots,n} X_i\right) &= \frac{bn-b-(an-a)}{n+1} \\
 &= (b-a) \frac{n-1}{n+1}
 \end{aligned} \tag{D.4}$$

□

Therefore, assuming a uniform distribution for the rewards of a given arm, the reward range can be estimated from the sample range $range_{sample}$ using the (unbiased) estimator:

$$range_{approx} = range_{sample} \times \frac{n+1}{n-1} \tag{D.5}$$

D.2.2 Range with uncertainty

For real DFT energies per adsorbate, one can directly apply equation D.5 using the sample range:

$$range_{sample} = \max(E_{DFT}) - \min(E_{DFT}) \tag{D.6}$$

However, for predicted energies $E_{pred} = \Gamma^\top x$, one should take into account the uncertainty due to the model. Indeed, one can assume that the real energy associated with configuration Γ is within:

$$\begin{aligned}
 \Gamma^\top x - C\sqrt{\Gamma^\top \text{Var}(x)\Gamma + \text{Var}(err)} &\leq E_{real} \leq \Gamma^\top x + C\sqrt{\text{Var}(\Gamma^\top x) + \text{Var}(err)} \\
 &= \Gamma^\top x + C\sqrt{\Gamma^\top \text{Var}(x)\Gamma + \text{Var}(err)}
 \end{aligned} \tag{D.7}$$

where C is the confidence level desired.

Appendix D. Appendix to Chapter 6

Therefore, the uncertainty-corrected sample minimum for predicted energies should be:

$$\min_{exact}^{corrected}(E_{pred}) = \min_{\Gamma} \left(\Gamma^{\top} x - C \sqrt{\Gamma^{\top} \text{Var}(x) \Gamma + \text{Var}(err)} \right) \quad (\text{D.8})$$

This would require the storage of all observed configurations for each node, which will quickly become problematic. Furthermore, in a lesser extent, this computation would become slightly computationally expensive.

A bound approximation can be found in bounding the estimated sample minimum predicted energy with the largest confidence interval observed:

$$\begin{aligned} \min_{exact}^{corrected}(E_{pred}) &\geq \min_{approx}(E_{pred}) - C \sqrt{\max_{\Gamma} (\Gamma^{\top} \text{Var}(x) \Gamma) + \text{Var}(err)} \\ &= \min_{approx}(E_{pred}) - C \sqrt{\max_{\Gamma} (\sum (\Gamma \Gamma^{\top} \times \text{Var}(x))) + \text{Var}(err)} \\ &\geq \min_{bound}^{corrected}(E_{pred}) = \min_{approx}(E_{pred}) - C \sqrt{\sum \left(\max_{\Gamma} / \min_{\Gamma} (\Gamma \Gamma^{\top}) \times \text{Var}(x) \right) + \text{Var}(err)} \end{aligned} \quad (\text{D.9})$$

$$\text{where } \max_{\Gamma} / \min_{\Gamma} (\Gamma \Gamma^{\top})_{ij} = \begin{cases} \max_{\Gamma_{seen}} (\Gamma \Gamma^{\top})_{ij}, & \text{if } \text{Var}(x)_{ij} \geq 0 \\ \min_{\Gamma_{seen}} (\Gamma \Gamma^{\top})_{ij}, & \text{else} \end{cases}.$$

This bound requires to store $\max_{\Gamma_{seen}} (\Gamma \Gamma^{\top})$ and $\min_{\Gamma_{seen}} (\Gamma \Gamma^{\top})$ ($O(m^2)$ bits), which is likely to be problematic.

A more memory-efficient approximation (but even looser bound) can be found in:

$$\begin{aligned} \min_{approx}^{corrected}(E_{pred}) &= \min_{approx}(E_{pred}) - C \sqrt{\sum \left(\max_{\Gamma}^2 / \min_{\Gamma}^2 (\Gamma) \times \text{Var}(x) \right) + \text{Var}(err)} \\ &\leq \min_{bound}^{corrected}(E_{pred}) \leq \min_{exact}^{corrected}(E_{pred}) \end{aligned} \quad (\text{D.10})$$

$$\text{where } \max_{\Gamma}^2 / \min_{\Gamma}^2 (\Gamma)_{ij} = \begin{cases} \max_{\Gamma_{seen}} (\Gamma_i) \times \max_{\Gamma_{seen}} (\Gamma_j), & \text{if } \text{Var}(x)_{ij} \geq 0 \\ \min_{\Gamma_{seen}} (\Gamma_i) \times \min_{\Gamma_{seen}} (\Gamma_j), & \text{else} \end{cases}, \text{ because } \Gamma \text{ has positive components.}$$

Note that this bound only requires to store the minimum/maximum of each components of Γ seen, which is already stored from section D.2.1.

For these reasons, $\min_{approx}^{corrected}(E_{pred})$ was chosen in our implementation as the sample minimum bound.

A tighter bound can be found in using only the borders of the multidimensional blue box defined figure D.1:

$$\begin{aligned} \min_{borders}^{corrected}(E_{pred}) &= \min_{approx}(E_{pred}) - C \sqrt{\max_{\Gamma \in box} (\Gamma^\top \text{Var}(x)\Gamma) + \text{Var}(err)} \\ &= \min_{approx}(E_{pred}) - C \sqrt{\max_{\Gamma \in borders} (\Gamma^\top \text{Var}(x)\Gamma) + \text{Var}(err)} \end{aligned} \quad (D.11)$$

where

$$box = \left\{ \Gamma \mid \forall i, \min_{\Gamma_{seen}}(\Gamma_i) \leq \Gamma_i \leq \max_{\Gamma_{seen}}(\Gamma_i) \right\}$$

and

$$borders = \left\{ \Gamma \mid \forall i, \Gamma_i = \min_{\Gamma_{seen}}(\Gamma_i) \vee \Gamma_i = \max_{\Gamma_{seen}}(\Gamma_i) \right\}$$

But even with borders only, a naive implementation would need to test $O(2^m)$ configurations.

Finally a much simpler approximation comes from simply using the bounds for the approximate sample minimum/maximum:

$$\min_{basic}^{corrected}(E_{pred}) = \min_{approx}(E_{pred}) - C \sqrt{\min_{\Gamma} / \max_{\Gamma}(\Gamma)^\top \cdot \text{Var}(x) \cdot \min_{\Gamma} / \max_{\Gamma}(\Gamma) + \text{Var}(err)} \quad (D.12)$$

However, this approximation does not correspond systematically to a bound, and could therefore be difficult to interpret in the general case.

With trivial changes to the equations, one can obtain similar bounds for the maximum.

D.2.3 Variance without uncertainty

Here we are interested in evaluating the empiric variance for the reward of a given arm.

Model independent rewards

For the real DFT energies per adsorbate, these rewards are independent from the model. Therefore, estimating the reward variance is straightforward using the (unbiased) estimator:

$$\begin{aligned} \text{Var}(E_{DFT})_{empiric} &= \mathbb{E}((E_{DFT} - \mathbb{E}(E_{DFT})_{empiric})^2)_{empiric} \\ &= \mathbb{E}(E_{DFT}^2)_{empiric} - \mathbb{E}(E_{DFT})_{empiric}^2 \end{aligned} \quad (D.13)$$

Appendix D. Appendix to Chapter 6

which requires the storage of $\mathbb{E}(E_{DFT}^2)_{empiric}$ on top of $\mathbb{E}(E_{DFT})_{empiric}$.

Model dependent rewards

On the other hand, the predicted energies $E_{pred} = \Gamma^\top x$ are dependent on the model parameters x , which are likely to change during the run. Keeping a decomposed form for E_{pred} leads to the empirical variance formula:

$$\begin{aligned}
 \text{Var}(E_{pred})_{empiric} &= \mathbb{E}_\Gamma((\Gamma^\top x - \mathbb{E}_\Gamma(\Gamma^\top x)_{empiric})^2)_{empiric} \\
 &= \mathbb{E}_\Gamma((\Gamma^\top x)^2)_{empiric} - \mathbb{E}_\Gamma(\Gamma^\top x)_{empiric}^2 \\
 &= \mathbb{E}_\Gamma(x^\top \Gamma \Gamma^\top x)_{empiric} - (\mathbb{E}_\Gamma(\Gamma)_{empiric}^\top \cdot x)^2 \\
 &= x^\top \cdot \mathbb{E}_\Gamma(\Gamma \Gamma^\top)_{empiric} \cdot x - (\mathbb{E}_\Gamma(\Gamma)_{empiric}^\top \cdot x)^2
 \end{aligned} \tag{D.14}$$

This requires to store $\mathbb{E}_\Gamma(\Gamma \Gamma^\top)_{empiric}$ ($O(m^2)$) on top of $\mathbb{E}_\Gamma(\Gamma)_{empiric}$ ($O(m)$), which is likely to be problematic.

Instead, similarly to the range estimates, one can estimate a maximum bound for the variance with the formula:

$$\begin{aligned}
 \text{Var}(E_{pred})_{bound} &= (\max_\Gamma / \min_\Gamma(\Gamma)^\top \cdot x)^2 - (\mathbb{E}_\Gamma(\Gamma)_{empiric}^\top \cdot x)^2 \\
 &\geq \sum (xx^\top \times \max_\Gamma / \min_\Gamma(\Gamma \Gamma^\top)) - (\mathbb{E}_\Gamma(\Gamma)_{empiric}^\top \cdot x)^2 \\
 &= x^\top \cdot \max_\Gamma / \min_\Gamma(\Gamma \Gamma^\top) \cdot x - (\mathbb{E}_\Gamma(\Gamma)_{empiric}^\top \cdot x)^2 \\
 &\geq \text{Var}(E_{pred})_{empiric}
 \end{aligned} \tag{D.15}$$

$$\text{where } \max_\Gamma / \min_\Gamma(\Gamma \Gamma^\top)_{ij} = \begin{cases} \max_{\Gamma_{seen}}(\Gamma \Gamma^\top)_{ij}, & \text{if } (xx^\top)_{ij} \geq 0 \\ \min_{\Gamma_{seen}}(\Gamma \Gamma^\top)_{ij}, & \text{else} \end{cases}.$$

Note that this bound only requires to store the minimum/maximum of each components of Γ seen, which is already stored from section D.2.1.

This variance bound is expected to be very loose, and should therefore be capped with $\min(\frac{range^2}{4}, Var)$, to ensure that the variance estimate is not worse than the maximal variance acceptable for a bounded distribution. Note that this feature is explicitly used in UCB-Tuned².

D.2.4 Variance with uncertainty

Here, we are interested in evaluating the empirical variance of the predicted energy corrected with the uncertainties^{xi} from the model $\Gamma^\top x + \delta E_\Gamma$, where $\mathbb{E}(\delta E_\Gamma) = 0$ and $\text{Var}(\delta E_\Gamma) = \Gamma^\top \text{Var}(x)\Gamma + \text{Var}(err)$:

$$\begin{aligned}
 \text{Var}(E_{pred})_{empiric}^{corrected} &= \mathbb{E}_{\Gamma, \delta E_\Gamma} \left(\left(\Gamma^\top x + \delta E_\Gamma - \mathbb{E}_{\Gamma, \delta E_\Gamma} (\Gamma^\top x + \delta E_\Gamma)_{empiric} \right)^2 \right)_{empiric} \\
 &= \mathbb{E}_{\Gamma, \delta E_\Gamma} \left(\left(\Gamma^\top x - \mathbb{E}_\Gamma (\Gamma^\top x)_{empiric} + \delta E_\Gamma - \mathbb{E}_{\delta E_\Gamma} (\delta E_\Gamma) \right)^2 \right)_{empiric} \\
 &= \mathbb{E}_\Gamma \left(\mathbb{E}_{\delta E_\Gamma} \left(\left(\Gamma^\top x - \mathbb{E}_\Gamma (\Gamma^\top x)_{empiric} \right)^2 + (\delta E_\Gamma - \mathbb{E}_{\delta E_\Gamma} (\delta E_\Gamma))^2 \right. \right. \\
 &\quad \left. \left. + 2 \left(\Gamma^\top x - \mathbb{E}_\Gamma (\Gamma^\top x)_{empiric} \right) (\delta E_\Gamma - \mathbb{E}_{\delta E_\Gamma} (\delta E_\Gamma)) \right) \right)_{empiric} \quad (D.16) \\
 &= \text{Var}(E_{pred})_{empiric} + \mathbb{E}_\Gamma (\text{Var}(\delta E_\Gamma))_{empiric}
 \end{aligned}$$

Since $\text{Var}(E_{pred})_{empiric}^{corrected}$ contains the term $\text{Var}(E_{pred})_{empiric}$, then we can deduce that finding an approximation here is at least as problematic as in section D.2.3.

D.2.5 Conclusion

Simple yet reasonable approximations can be formulated for estimating the range at low memory cost, even when taking uncertainties into account.

However, the similar memory efficient approximations that were found for the empirical variance are likely to be unreasonable. Therefore, unless better approximations are found, variance-based UCB-like algorithms could only be applied with a large memory cost.

Despite the performance of some non variance-based UCB-like algorithms¹, one could expect a variance-based UCB variant to be particularly beneficial here, as the estimated range is very likely to be overestimated.

D.3 Heuristic score justification

D.3.1 A first guess: total adsorption energy

A natural and intuitive idea is to preferentially focus on chemically stable geometries for inclusion in a training set. “The more stable the better” is the typical way to go for creating

^{xi}This problem is similar to the more tangible one: estimating the variance of a die roll chosen at random from a finite set of different dice.

geometries that describe observed/realistic processes.

One could even devise a sort of rationale behind this choice: a model should better describe geometries that are more likely to occur (in other words, unrealistic geometries are almost never observed, so it is less useful to properly describe them). Such rationale would make sense for designing models aimed at handling only geometries that are typically observed. But KMC simulations do not handle only realistic geometries. Indeed, in a typical KMC simulation, geometries are generated at random. Even though the selection of a stable geometry is more likely than an unstable one, unrealistic geometries are constantly being generated and evaluated (and likely excluded). As a matter of fact, misjudging unrealistic geometries could lead to them being accepted in a disproportionate amount and ultimately result in unrealistic kinetic simulations.

For example, focusing on stable geometries would not provide much many-body insights if all lateral interactions were destabilizing. In such settings, a model Hamiltonian mainly trained on the most stable geometries would not be much better than the non-interacting (i.e. one-body) model.

D.3.2 Relevant heuristic: absolute deviation from non-interacting model

The one-body (i.e. non-interacting) model Hamiltonian is usually already known, as it is relatively easy to obtain (by computing the adsorption energy of each single molecular adsorbate at each possible active site) and essential for a preliminary study. So why not focusing on learning the lateral interactions only?

In this section we will justify that the quick learning of lateral interactions can benefit from focusing on geometries displaying the largest deviations from the non-interacting model.

Before diving into mathematical considerations, a simple intuition can be formulated in this very direction: including geometries with only negligible lateral interactions does not seem relevant to learn many-body effects.

Justification

We consider here a linear model (solved by ordinary least squares) where the regressors/features are the occurrences of each lateral interaction, the regressand/target is the total energy deviation from the non-interacting model (i.e. the total energy contribution due to lateral interactions) and the parameters to learn are the individual energy contributions of each lateral interactions.

Because of the sheer number of features and the relatively low number of observations (quick

learning), the corresponding set of linear equations is largely under-determined. Furthermore, the Reinforcement Sampling approach makes use of this linear model when it is being partially trained, which naturally amplifies its under-determination.

As a consequence, not only the interpolative power (i.e. accurate description within the subspace spanned by observed features), but also the extrapolative power (i.e. relatively accurate description outside, but still near, the subspace spanned by observed features) are critical for powering a heuristic that impacts the convergence rate of the whole Reinforcement Sampling procedure.

The extrapolative properties of the ordinary least squares come from the balanced distribution of the target among the observed features^{xii}, so that when learning with a novel geometry, associated many-body effects are typically attributed in a proportionate amount to the lateral interactions that are present in this new geometry. In other words, the ordinary least squares algorithm has a tendency to *a priori* provide uniform lateral interaction contributions (since it provides the set of best fitting parameters having minimum norm).

Therefore, learning from geometries displaying lateral interactions whose energy contributions are mostly similar contributes to provide a linear model with better extrapolative power i.e., having a relative accuracy even outside of the explored subspace despite being largely under-determined. On the contrary, learning from geometries displaying lateral interactions whose energy contributions are highly different would not benefit much the extrapolative power of the linear model. Indeed, a strongly positive contribution c , for example, is expected to “poison” the extrapolation of a negative contribution a (if seen together) by canceling it, and therefore, inducing the attribution of an overestimated positive value to this contribution a (and vice versa).

A trivial example can illustrate this phenomenon: assume 2 strongly stabilizing lateral interactions whose energy contributions are $a = b = -1$ eV and 2 strongly destabilizing lateral interactions with associated energy contributions $c = d = 1$ eV. Learning these 4 parameters with only 2 geometries is equivalent to solving an under-determined set of linear equations, and therefore, the resulting model will display a large extrapolated component. From this simple example, it is easy to illustrate how the extrapolative power is impacted by learning with similarly/dissimilarly combined parameters:

- Learning from observations where combined parameters are similar:

$$\begin{aligned} a + b &= -2 \text{ eV} \\ c + d &= 2 \text{ eV} \end{aligned} \tag{D.17}$$

^{xii}To be precise, the difference between the observed target and the linear prediction (the *a priori* error) is distributed proportionally to the previously unexplored component of the newly observed features, see theorem 4 for more details.

will lead to a set of fitted parameters $a' = a$, $b' = b$, $c' = c$ and $d' = d$ that perfectly matches the true contributions, and can therefore be applied accurately across the whole feature space.

- While learning from observations where combined parameters are dissimilar:

$$\begin{aligned} a + c &= 0 \text{ eV} \\ b + d &= 0 \text{ eV} \end{aligned} \tag{D.18}$$

will naturally lead to a set of fitted parameters $a' = b' = c' = d' = 0$ that cannot be much extrapolated outside of the feature subspace upon which the model was trained.

Even though learning with geometries where energetically similar interactions are combined is optimal, accurately finding such geometries (if they exist) is almost impossible *a priori*. Indeed, one would theoretically need to know beforehand the energy contributions of the lateral interactions involved, which are the very parameters being learned^{xiii}.

Nevertheless, focusing on geometries with large associated target in absolute value (i.e. absolute energy deviation from the non-interacting model) naturally introduces a bias toward geometries where lateral interactions with significant energy contributions (and same sign) are found simultaneously. At least, this bias is privileging geometries where large energy contributions are not canceled out too much by the other interactions being present.

As a consequence, using such heuristic (i.e. using the absolute energy deviation from the non-interacting model as a heuristic) for choosing geometries to learn from is expected to produce models with better extrapolative power on the subspace of significant non-canceled lateral interactions, on average. One should note that this gain in extrapolative power on one specific subspace is made at the cost of extrapolative power for describing the other lateral interactions present in the geometries (i.e. the more a set of contributions is dominating the whole many-body effects, the stronger is the “extrapolation poisoning” toward the remaining contributions).

In practice, however, the whole model Hamiltonian cannot reasonably be constructed. The raw feature space is combinatorially large, and would require far too much observations to be properly sampled, resulting in an unreasonable amount of computational resources that would likely cancel out the computational gain from such model Hamiltonian. Instead, only a truncated version of such many-body model Hamiltonian can reasonably be learned quickly.

Fast learning (i.e. with only a few number of observations/geometries) means that only a few number of lateral interaction contributions can be properly defined (i.e. without extrapo-

^{xiii}Furthermore, any prediction purely based on the few geometries already evaluated would be highly biased without external knowledge.

lation), due to rank considerations. Therefore, a truncated many-body model Hamiltonian should try to recover most of the many-body effects with the fewest number of lateral interactions. Fortunately, a cluster expansion based model is particularly prone to truncation as most many-body terms (especially long-range and/or high order correction terms involving more than 2-3 entities) are expected to be negligible. As a consequence, such truncated cluster expansion based model Hamiltonian typically focuses on the few strongest lateral interactions, and discards negligible ones.

Last but not least, let us consider the sets of strong interactions effectively canceling each others that are not found in geometries displaying a deviation from the non-interacting model. These interactions are therefore always found canceling each others, despite actively looking for geometries where they are not (except with geometries where they are canceled out by already known interactions). Consequently, such sets of interactions globally seem to be strongly correlated (in term of presence) in practice, and could therefore be relatively safely treated as negligible, especially for the truncation.

Note that if the heuristic-guided exploration is successful, the eventual completely unexplored features should correspond to lateral interactions with negligible energy contribution to the many-body effects. Therefore, this heuristic brings coherence to the default value of a least square solution (i.e. 0, due to the minimum norm constraint) for such contributions.

Bibliography

- [1] Garivier, A.; Cappé, O. The KL-UCB algorithm for bounded stochastic bandits and beyond. *Proceedings of the 24th annual conference on learning theory*. 2011; pp 359–376.
- [2] Auer, P.; Cesa-Bianchi, N.; Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning* **2002**, *47*, 235–256.
- [3] Audibert, J.-Y.; Munos, R.; Szepesvári, C. Exploration-Exploitation Tradeoff Using Variance Estimates in Multi-Armed Bandits. *Theor. Comput. Sci.* **2009**, *410*, 1876–1902.